

Derandomizing Chernoff Bound with Union Bound with an Application to k -wise Independent Sets

Nader H. Bshouty
Technion, Haifa, Israel
bshouty@ca.technion.ac.il

May 23, 2016

Abstract

Derandomization of Chernoff bound with union bound is already proven in many papers. We here give another explicit version of it that obtains a construction of size that is arbitrary close to the probabilistic nonconstructive size.

We apply this to give a new simple polynomial time constructions of almost k -wise independent sets. We also give almost tight lower bounds for the size of k -wise independent sets.

1 Introduction

Derandomization of Chernoff bound with union bound is already proven in many papers. See for example [22, 6, 9]. We here give another explicit version of it that obtains a construction of a size that is arbitrary close to the size of the probabilistic nonconstructive size.

We then show that, for some construction problems, one can combine this method with the method of conditional probabilities to get a derandomization that runs in time polylogarithmic in the sample size.

In this paper we give the following application of this result:

For $p \in \mathbb{R}$, the *distance in the L_p -norm between two probability distribution D and Q over a sample space S* is

$$\|D - Q\|_p = \left(\sum_{s \in S} |D(s) - Q(s)|^p \right)^{\frac{1}{p}}$$

and for $p = \infty$ is $\|D - Q\|_p = \max_{s \in S} |D(s) - Q(s)|$.

Let $S = \{0, 1\}^n$. A *uniform distribution U over S* is a distribution where $U(s) = 1/2^n$ for all $s \in S$. For $I = (i_1, \dots, i_k)$ where $1 \leq i_1 < \dots < i_k \leq n$ the *distribution D_I restricted to I over $\{0, 1\}^k$* is $D_I(\sigma_1, \dots, \sigma_k) = \Pr_{s \sim D}[s_{i_1} = \sigma_1 \wedge \dots \wedge s_{i_k} = \sigma_k]$. A distribution Q over $\{0, 1\}^n$ is called *ϵ -almost k -wise independent in the L_p -norm* if for any $I = (i_1, \dots, i_k)$ we have $\|Q_I - U_I\|_p \leq \epsilon$.

The goal is to construct $S' \subset \{0, 1\}^n$ of small size such that the uniform distribution on S' is ϵ -almost k -wise independent in the L_p -norm. We will just say that S' is *ϵ -almost k -wise independent set* in the L_p -norm. The following table summarizes the results from the literature and our results. The table shows the sizes without the small terms $\log \log n$, $\log(1/\epsilon)$ and k . See the exact sizes in the table in Section 3.3 and the theorems in Subsection 3.4 and Section 4.

Construction Time	Reference	Size for L_∞	Size for L_1
Poly. time	[7]	$\frac{\log^2 n}{\epsilon^2}$	$\frac{2^k \log^2 n}{\epsilon^2}$
Poly. time	AGC+Hadamard	$\frac{\log n}{\epsilon^3}$	$\frac{2^{3k/2} \log n}{\epsilon^3}$
Poly. time	[13]	$\frac{\log^{5/4} n}{\epsilon^{2.5}}$	$\frac{2^{5k/4} \log^{5/4} n}{\epsilon^{2.5}}$
Poly. time	Ours	$\frac{\log n}{2^k \epsilon^3}$	$\frac{2^k + \log n}{\epsilon^3}$
$n^{O(k)}$ time	Ours	$\frac{\log n}{2^k \epsilon^2}$	$\frac{2^k + \log n}{\epsilon^2}$
Lower Bound	Ours	$\frac{\log n}{2^k \epsilon^2}$	$\frac{\log n}{\epsilon^2}$

With the techniques used in this paper, all the results in this table can be easily generalized to any product distribution and any alphabet.

Our construction is very simple. We first give a derandomization of Chernoff bound with union bound. For this we use the pessimistic method with a pessimistic estimator (potential function) that gives constructions of size that are arbitrary close to the size of the probabilistic nonconstructive size. Those constructions are polynomial in the space size that is exponential in the dimension of the problem. We then use the conditional probability method that reduces the complexity to polylogarithmic time in the sample size. Those are used to construct a dense perfect hash family and an ϵ -almost k -wise independent set of small dimension. We then combine both constructions to get the final construction. We also give lower bounds that are almost tight to the non-constructive constructions. Our constructions have sizes that are within a factor of $1/\epsilon$ from the lower bounds.

Our construction can be easily generalized to any product distribution and any alphabet (not necessarily alphabet of size power of prime) and can be used for other dense and balance constructions. See some other techniques for deterministic and randomized dense, balance and non-dense constructions in [5, 6, 10, 11, 12, 15, 16, 20, 21] and references within.

This paper is organized as follows. In Section 2 we give the main two theorems of the derandomization and show how to use the method of conditional probabilities to reduce its time complexity to polylogarithmic in the size of the sample space. In Section 3 we give an exponential time constructions of small size for a code that achieves the Gilbert-Varshamov bound, ϵ -balance error-correcting code and ϵ -bias sample space. In Section 3.3 we give all the constructions in the above table. Then in Section 5 we give the lower bounds.

2 The Derandomization

In this section we give the derandomization of Chernoff bound with union bound

The q -ary entropy function is

$$H_q(p) = p \log_q \frac{q-1}{p} + (1-p) \log_q \frac{1}{1-p}.$$

The Kullback-Leibler divergence between Bernoulli distributed random variables with parameter λ and η is

$$D(\lambda||\eta) = \lambda \ln \frac{\lambda}{\eta} + (1-\lambda) \ln \left(\frac{1-\lambda}{1-\eta} \right).$$

For two integers n and m we denote $[n] = \{1, 2, \dots, n\}$ and $(n, m] = \{n+1, n+2, \dots, m\}$. For a finite multiset of objects S we denote by U_S the uniform distribution over S .

We prove

Theorem 1. Let S be a finite sample space with a probability distribution D . Let X_1, \dots, X_N be random variables over S that take values from $\{0, 1\}$. Let $N' \leq N$ and $\{\lambda_i\}_{i \in [N]}$ be such that $0 < \lambda_i < p_i \leq \mathbf{E}_{s \sim D}[X_i]$ and $1 > \lambda_j > p_j \geq \mathbf{E}_{s \sim D}[X_j]$ for all $i \in [N']$ and $j \in (N', N]$. Let m be such that

$$P_0 := \sum_{i=1}^N e^{-D(\lambda_i \| p_i) \cdot m} \leq 1.$$

There is a multiset $S' = \{s_1, s_2, \dots, s_m\} \subseteq S$ such that for all $i \in [N']$ and $j \in (N', N]$

$$\mathbf{E}_{s \sim U_{S'}} [X_i(s)] \geq \lambda_i, \quad \mathbf{E}_{s \sim U_{S'}} [X_j(s)] \leq \lambda_j.$$

In particular, the result follows for

$$m \geq \max_{i \in [N]} \frac{\ln N}{D(\lambda_i \| p_i)} = \max_{i \in [N]} \frac{\log_{q_i} N}{1 - H_{q_i}(\lambda_i)}$$

where $q_i = 1/(1 - p_i)$.

Proof. We give an algorithm that constructs S' . Let

$$\alpha_i = \frac{\lambda_i}{(q_i - 1)(1 - \lambda_i)} = \frac{(1 - p_i)\lambda_i}{p_i(1 - \lambda_i)} \quad \text{and} \quad \gamma_i = \frac{1}{\alpha_i p_i + (1 - p_i)} = \frac{1 - \lambda_i}{1 - p_i}.$$

Suppose that the algorithm has already chosen s_1, \dots, s_ℓ . Consider the potential function

$$P_j := \sum_{i=1}^N e^{-D(\lambda_i \| p_i)m} \gamma_i^j \alpha_i^{Z_{j,i}}$$

where $Z_{j,i} = X_i(s_1) + \dots + X_i(s_j)$ for $j = 0, 1, \dots, \ell$. Here, $Z_{0,i} = 0$. We now show how the algorithm chooses $s_{\ell+1}$. Consider the random variable

$$P'_{\ell+1}(s) = \sum_{i=1}^N e^{-D(\lambda_i \| p_i)m} \gamma_i^{\ell+1} \alpha_i^{Z_{\ell,i} + X_i(s)}.$$

Since $\alpha_i < 1$ for $i \in [N']$ and $\alpha_j > 1$ for $j \in (N', N]$, we have

$$\begin{aligned} \mathbf{E}_{s \sim D} [P'_{\ell+1}(s)] &= \sum_{i=1}^N e^{-D(\lambda_i \| p_i)m} \gamma_i^{\ell+1} \alpha_i^{Z_{\ell,i}} \mathbf{E}_{s \sim D} [\alpha_i^{X_i(s)}] \\ &\leq \sum_{i=1}^N e^{-D(\lambda_i \| p_i)m} \gamma_i^{\ell} \alpha_i^{Z_{\ell,i}} \\ &= P_\ell. \end{aligned}$$

Now the algorithm chooses $s_{\ell+1} \in S$ that satisfies $P'_{\ell+1}(s_{\ell+1}) \leq P_\ell$. Let $P_{\ell+1} = P'_{\ell+1}(s_{\ell+1})$. Then $P_0 \leq 1$ and $P_{\ell+1} \leq P_\ell$. Therefore

$$P_m = \sum_{i=1}^N e^{-D(\lambda_i \| p_i)m} \gamma_i^m \alpha_i^{Z_{m,i}} \leq 1.$$

In particular, for all $i \in [N]$

$$e^{-D(\lambda_i|p_i)m} \gamma_i^m \alpha_i^{Z_{m,i}} \leq 1.$$

Therefore, for all $i \in [N]$,

$$\alpha_i^{Z_{m,i}} \leq (\gamma_i^{-1} \cdot e^{D(\lambda_i|p_i)m})^m = \alpha_i^{\lambda_i m}.$$

Thus, for all $i \in [N]$ we have

$$\alpha_i^{Z_{m,i}} \leq \alpha_i^{\lambda_i m}.$$

Since $\alpha_i < 1$ for all $i \in [N']$ and $\alpha_j > 1$ for all $j \in (N', N]$ we have $Z_{m,i} \geq \lambda_i m$ for all $i \in [N']$ and $Z_{m,j} \leq \lambda_j m$ for all $j \in (N', N]$. Since $\mathbf{E}_{s \sim U_{S'}}[X_k(s)] = Z_{m,k}/m$ for all $k \in [N]$, the result follows. \square

We now give the bit-time complexity of the algorithm described in the proof of Theorem 1

Theorem 2. *Let all notation and assumptions be as in Theorem 1. Suppose that for every $s \in S$ all the values $X_1(s), \dots, X_N(s)$ can be computed in bit-time¹ $\tilde{O}(N)$. Let $\mu_i = D(\lambda_i|p_i)$, $\mu = \min(1, \min_i \mu_i)$. Let $\tau = \max_i \max(\alpha_i, \gamma_i)$ where*

$$\alpha_i = \frac{(1-p_i)\lambda_i}{p_i(1-\lambda_i)} \quad \text{and} \quad \gamma_i = \frac{1-\lambda_i}{1-p_i}.$$

There is an algorithm that runs in bit-time

$$T = \tilde{O}(|S| \cdot Nm \cdot (m \log \tau + \log(1/\mu)))$$

and outputs a multiset $S' = \{s_1, \dots, s_m, s_{m+1}\} \subseteq S$ such that for all $i \in [N']$ and $j \in (N', N]$

$$\mathbf{E}_{s \sim U_{S'}}[X_i(s)] \geq \lambda_i, \quad \mathbf{E}_{s \sim U_{S'}}[X_j(s)] \leq \lambda_j$$

where $U_{S'}$ is the uniform distribution on S' .

Proof. Consider the algorithm in the proof of Theorem 1. Notice that here the size of S' is $m+1$ and not m as in Theorem 1. So the potential function used in the algorithm is

$$P_j := \sum_{i=1}^N e^{-\mu_i(m+1)} \gamma_i^j \alpha_i^{Z_{j,i}}$$

but with the same assumption

$$\sum_{i=1}^N e^{-\mu_i m} \leq 1$$

as in Theorem 1. Therefore

$$P_0 = \sum_{i=1}^N e^{-\mu_i(m+1)} \leq e^{-\mu} \sum_{i=1}^N e^{-\mu_i m} \leq 1 - \frac{\mu}{2}.$$

¹Here $\tilde{O}(N)$ is $O(N \cdot \text{poly}(\log T))$ where T is the time complexity of the construction.

First, notice that $\alpha_k \leq \tau$ and $\gamma_k \leq \tau$ for all $k \in [N]$. Let B be a positive integer that will be determined later. If we use $\Delta := B + 1 + \lceil \log(\tau + 1) \rceil$ bits for the representation of $e^{-\mu_i m}$, α_k and γ_k , i.e., the absolute error is less than 2^{-B} , then the absolute error in computing $r_k := e^{-\mu_i m} \gamma_k^\ell \alpha_k^{Z_{\ell,k}}$ is at most $O(\ell \tau^{2\ell} 2^{-B})$ and in computing $P_\ell = \sum_k r_k$ is at most $O(N \ell \tau^{2\ell} 2^{-B}) = O(N m \tau^{2m} 2^{-B})$. This error is at most $\mu/(4m)$ when $B \geq 2m \log \tau + \log N + 2 \log m + \log(1/\mu) + 2$. Notice that, since the absolute error is less than $\mu/(4m)$, we have $P_{\ell+1} \leq P_\ell + \mu/(4m)$. Since $P_0 \leq 1 - \mu/2$, we get

$$P_{m+1} \leq 1 - \frac{\mu}{2} + \sum_{i=1}^{m+1} \frac{\mu}{4m} < 1$$

which, as shown in the proof of Theorem 1, gives the required bound.

Now arithmetic computations with $\Delta = B + 1 + \lceil \log(\tau + 1) \rceil$ bit numbers take bit-time $\tilde{O}(\Delta) = \tilde{O}(m \log \tau + \log(1/\mu))$. Since the number of arithmetic operations in the algorithm is $O(|S| \cdot Nm)$, the result follows. \square

In particular we have

Corollary 3. *Let all the notation be as in Theorem 1 and $0 < \epsilon_i < 1$ for all $i \in [N]$. Let $\alpha = \min_i \min(1/(1 - p_i), 1/(1 - \lambda_i))$. For*

$$m \geq \frac{3 \ln N}{\min_i p_i \epsilon_i^2}$$

the algorithm runs in bit-time $T = \tilde{O}(|S| \cdot Nm^2 \log \alpha)$ and outputs $S' = \{s_1, \dots, s_{m+1}\} \subseteq S$ such that for all $i \in [N']$ and $j \in (N', N]$

$$\mathbf{E}_{s \sim U_{S'}} [X_i(s)] \geq \lambda_i := (1 - \epsilon_i)p_i, \quad \mathbf{E}_{s \sim U_{S'}} [X_j(s)] \leq \lambda_j := (1 + \epsilon_j)p_j.$$

Proof. Follows from the fact that if $\lambda_i = (1 + \epsilon_i)p_i$ then $D(\lambda_i || p_i) \geq p_i \epsilon_i^2 / 2.5887 \geq p_i \epsilon_i^2 / 3$ and if $\lambda_i = (1 - \epsilon_i)p_i$ then $D(\lambda_i || p_i) \geq p_i \epsilon_i^2 / 2 \geq p_i \epsilon_i^2 / 3$.

Since $\alpha_i \leq 1/p_i(1 - \lambda_i)$ and $\gamma_i \leq 1/(1 - p_i)$ we have $m \log \tau = \tilde{O}(m \log \alpha)$. Since $1/\mu \leq \min 3/(p_i \epsilon_i^2) \leq m^2$ we get $T = \tilde{O}(|S| \cdot Nm \cdot (m \log \tau + \log(1/\mu))) = \tilde{O}(|S| \cdot Nm^2 \log \alpha)$. \square

2.1 Combining with the Method of Conditional Probabilities

In the above constructions, the time complexity is linear in $|S|$ which may be exponentially large. In the following we get around this problem when S is of the form $S_1 \times \dots \times S_n$ and the expectation of some ‘‘intermediate’’ random variables can be efficiently computed.

We prove

Theorem 4. *Let all notation and assumptions be as in Theorem 1 and Corollary 3. Suppose $S = S_1 \times S_2 \times \dots \times S_n$. If any expectation of the form*

$$\mathbf{E}[X_i(x_1, \dots, x_n) \mid x_1 = \xi_1, \dots, x_j = \xi_j]$$

can be computed in bit-time T then the constructions in Theorem 1 and 2 can be performed in bit-time

$$\tilde{O}(T(|S_1| + \dots + |S_n|) \cdot Nm(m \log \tau + \log(1/\mu)))$$

and in Corollary 3 in bit-time

$$\tilde{O}(T(|S_1| + \dots + |S_n|) \cdot Nm^2 \log \alpha)$$

Proof. The first result follows from the fact that since

$$\mathbf{E}[P'_{\ell+1}(s)] = \mathbf{E}_{y_1 \sim S_1}[\mathbf{E}[P'_{\ell+1}(s) \mid x_1 = y_1]] \leq 1,$$

there is $\xi_1 \in S_1$ such that $\mathbf{E}[P'_{\ell+1}(s) \mid x_1 = \xi_1] \leq 1$. So we find such ξ_1 . Then recursively find ξ_2, \dots, ξ_n .

For this case we need the absolute error to be less than $\mu/(4mn)$ (rather than $\mu/(4m)$ as in Theorem 2). This adds a factor of $\log n$ to the time complexity that is swallowed by the \tilde{O} . \square

3 Constructions of Almost k -Wise Independent Sets

3.1 ϵ -Balance Error-Correcting Code

A *linear code* over the field \mathbb{F}_q is a linear subspace $C \subset \mathbb{F}_q^m$. Elements in the code are called *codewords*. A linear code C is called a $[m, k, d]_q$ *linear code* if $C \subset \mathbb{F}_q^m$ is a linear code, $|C| = q^k$ and for every two distinct codewords v and u in the code we have $\text{dist}(v, u) := |\{i \mid v_i \neq u_i\}| \geq d$. The latter is equivalent to: For every nonzero codeword w in C , we have $\text{wt}(w) := |\{i \mid w_i \neq 0\}| \geq d$.

A linear code C is called a $[m, k, (1 - 1/q)m]_q$ *ϵ -balance error-correcting linear code* if $C \subset \mathbb{F}_q^m$ is a linear code, $|C| = q^k$ and for every nonzero codeword $w \in C$ and $\xi \in \mathbb{F}_q$ we have

$$(1 - \epsilon) \frac{m}{q} \leq |\{w_i \mid w_i = \xi\}| \leq (1 + \epsilon) \frac{m}{q}.$$

We show

Lemma 1. *Let q be a prime power, m and k positive integers and $0 \leq \epsilon \leq 1/2$. For*

$$m \geq O\left(\frac{kq \log q}{\epsilon^2}\right)$$

there is an $[m, k, (1 - 1/q)m]_q$ ϵ -balance error-correcting linear code that can be constructed in bit-time complexity $\tilde{O}(q^{k+3}/\epsilon^4)$.

Proof. We use Corollary 3. Consider $S = \mathbb{F}_q^k$ with the uniform distribution. Define for every $v \in \mathbb{F}_q^k$ of the form $v = (v_1, \dots, v_j, 1, 0, \dots, 0)$, $j = 0, 1, \dots, k-1$, every $\xi \in \mathbb{F}_q$ and every $t \in \{1, 2\}$ a random variable $X_{v, \xi, t} : \mathbb{F}_q^k \rightarrow \{0, 1\}$ where $X_{v, \xi, t}(w) = [v_1 w_1 + \dots + v_k w_k = \xi]$. That is, $X_{v, \xi, t}(w) = 1$ if $v_1 w_1 + \dots + v_k w_k = \xi$ and zero otherwise. The number of random variables is $N = 2q(q^k - 1)/(q - 1)$ and $\mathbf{E}[X_{v, \xi, t}] = 1/q$ for all v, ξ and t . The random variables satisfy the condition in Theorem 4 with $S_i = \mathbb{F}_q$ for $i = 1, \dots, k$. Therefore an $S' = \{s_1, \dots, s_m\} \subseteq S$ of size

$$m \geq \frac{3q \ln N}{\epsilon^2} = O\left(\frac{kq \log q}{\epsilon^2}\right)$$

that satisfies $\mathbf{E}_{s \sim U_{S'}}[X_{v, \xi, 1}(s)] \geq (1 - \epsilon)/q$ and $\mathbf{E}_{s \sim U_{S'}}[X_{v, \xi, 2}(s)] \leq (1 + \epsilon)/q$ for all v and ξ , can be constructed in bit-time complexity

$$\tilde{O}\left((kq) \left(2q \frac{q^k - 1}{q - 1}\right) \frac{k^2 q^2 \log^2 q}{\epsilon^4}\right) = \tilde{O}\left(\frac{q^{k+3}}{\epsilon^4}\right).$$

Now, $C = \{(us_1, \dots, us_m) \mid u \in \mathbb{F}_q^k\}$ is the code. \square

3.2 ϵ -Bias Sample Space

Let D be a probability distribution over \mathbb{F}_2^n . The *bias* of D with respect to a set of indices $I \subseteq [n]$ is defined as

$$\text{bias}_I(D) = \left| \Pr_{x \sim D} \left(\sum_{i \in I} x_i = 0 \right) - \Pr_{x \sim D} \left(\sum_{i \in I} x_i = 1 \right) \right|.$$

We say that D is ϵ -*bias sample space* if $\text{bias}_I(D) \leq \epsilon$ for all non-empty subset $I \subseteq [n]$. If D is the uniform distribution over a multiset $S \subseteq \mathbb{F}_2^n$ then we call S an ϵ -*bias set*. The goal is to construct a small ϵ -bias set in polynomial time in n/ϵ . The following constructions are known from the literature

Reference	Size $ S = O(\cdot)$
Alon et. al. [7]	$\frac{n^2}{\epsilon^2 \log^2(n/\epsilon)}$
AGC+Hadamard code ²	$\frac{n}{\epsilon^3 \log(1/\epsilon)}$
Ben-Aroya and Ta-Shma [13]	$\frac{n^{5/4}}{\epsilon^{2.5} \log^{5/4}(1/\epsilon)}$

The best lower bound for the size of ϵ -bias set is [7, 1]

$$\Omega \left(\frac{n}{\epsilon^2 \log(1/\epsilon)} \right).$$

Let C be an ϵ -balance error-correcting linear code $[m, n, m/2]_2$ over \mathbb{F}_2 with a $m \times n$ generator matrix A . It is easy to see that the set of rows of A is ϵ -bias set of size m . Therefore, by Lemma 1, for $q = 2$, we have

Lemma 2. *An ϵ -bias set $S \subseteq \mathbb{F}_2^n$ of size*

$$O \left(\frac{n}{\epsilon^2} \right)$$

can be constructed in time $O(2^n/\epsilon^4)$.

Remark: Using the powering construction in [7] with $b_{ij} = (\text{bin}(v_j x^i), \text{bin}(y))$ where $\{y\}$ is an ϵ -bias set $S' \subseteq \mathbb{F}_2^m$ (rather than all the elements of \mathbb{F}_2^m) gives a polynomial time construction of an ϵ -bias set $S \subseteq \mathbb{F}_2^n$ of size $O(n/\epsilon^3)$.

3.3 k -wise Approximating Distributions in Time $O(n^k)$

The *distance in the L_p -norm between two probability distribution D and Q over the sample space S is*

$$\|D - Q\|_p = \left(\sum_{s \in S} |D(s) - Q(s)|^p \right)^{\frac{1}{p}}$$

for $p \in \mathbb{R}$ and $\|D - Q\|_p = \max_{s \in S} |D(s) - Q(s)|$ for $p = \infty$.

Let $S = \Sigma^n$. A *uniform distribution U over S* is a distribution where $U(s) = 1/|\Sigma|^n$ for all $s \in S$. A *product distribution D over S* is a distribution where $D(s_1, \dots, s_n) = p_{1,s_1} \cdots p_{n,s_n}$ where $0 \leq p_{i,s_i} \leq 1$ for all $i \in [n]$ and $s_i \in \Sigma$.

²See the construction in [13]

For $I = (i_1, \dots, i_k)$ where $1 \leq i_1 < \dots < i_k \leq n$ the distribution D_I restricted to I over Σ^k is $D_I(\sigma_1, \dots, \sigma_k) = \Pr_{s \sim D}[s_{i_1} = \sigma_1 \wedge \dots \wedge s_{i_k} = \sigma_k]$. Two distributions D and Q over Σ^n are called k -wise ϵ -close in the L_p -norm, if for any $I = (i_1, \dots, i_k)$, $\|D_I - Q_I\|_p \leq \epsilon$. If D is the uniform distribution then Q is called ϵ -almost k -wise independent in the L_p -norm.

When $\epsilon = 0$ then S is called k -wise independent set. It is known that the size of any k -wise independent set is $n^{\Theta(k)}$, [3, 14]. See also [8].

The goal is: given a distribution D . Construct $S' \subset S$ of small size such that the uniform distribution on S' is k -wise ϵ -close to D in the L_p -norm. We will just say that S' is k -wise ϵ -close to D in the L_p -norm and if D is the uniform distribution we say that S is ϵ -almost k -wise independent in the L_p -norm.

For $\Sigma = \{0, 1\}$, Naor and Naor proved

Lemma 3. [18]. *Let $k < n$ be an odd integer, t is a power of 2 and*

$$n \leq 2^{\lfloor 2^{(m-1)/(k-1)} \rfloor} - 1.$$

Given an ϵ -bias set $S \subseteq \{0, 1\}^m$ of size t , one can, in polynomial time, construct a set $R \subseteq \{0, 1\}^n$ of size t that is ϵ -almost k -wise independent in the L_∞ -norm and $2^{k/2}\epsilon$ -almost k -wise independent in the L_1 -norm.

For $\Sigma = \{0, 1\}$, the following are the best known polynomial time constructions of sets that are ϵ -almost k -wise independent in the L_∞ -norm and L_1 -norm. The constructions use the ϵ -bias sets in Section 3.2 with Lemma 3. See also the sizes without the small terms k , $\log(1/\epsilon)$ and $\log \log n$ in the first three rows of the table in the introduction.

Reference	Size for L_∞	Size for L_1
[7]	$\frac{k^2 \log^2 n}{\epsilon^2 (\log^2 k + (\log \log n)^2 + \log^2(1/\epsilon))}$	$\frac{k^2 2^k \log^2 n}{\epsilon^2 (k^2 + (\log \log n)^2 + \log^2(1/\epsilon))}$
AGC+Hadamard	$\frac{k \log n}{\epsilon^3 \log(1/\epsilon)}$	$\frac{k 2^{3k/2} \log n}{\epsilon^3 (k + \log(1/\epsilon))}$
[13]	$\frac{k^{5/4} \log^{5/4} n}{\epsilon^{2.5} \log^{5/4}(1/\epsilon)}$	$\frac{k^{5/4} 2^{5k/4} \log^{5/4} n}{\epsilon^{2.5} (k^{5/4} + \log^{5/4}(1/\epsilon))}$

By Lemma 2 and Lemma 3 we have

Lemma 4. *An ϵ -almost k -wise independent set in the L_∞ -norm of size*

$$O\left(\frac{k \log n}{\epsilon^2}\right)$$

can be constructed in time $O(n^{(k-1)/2}/\epsilon^4)$.

An ϵ -almost k -wise independent in the L_1 -norm of size

$$O\left(\frac{k 2^k \log n}{\epsilon^2}\right)$$

can be constructed in time $O(2^{2k} n^{(k-1)/2}/\epsilon^4)$.

We now prove

Theorem 5. An ϵ -almost k -wise independent set in the L_∞ -norm of size

$$O\left(\frac{k \log n}{2^k \epsilon^2}\right)$$

can be constructed in time $\tilde{O}(n^k/\epsilon^4)$.

Proof. Consider $S = \{0, 1\}^n$ with the uniform distribution. Define for every $1 \leq i_1 < i_2 < \dots < i_k \leq n$ and $(\xi_1, \dots, \xi_k) \in \{0, 1\}^k$ the random variable $X_{i_1, \dots, i_k, \xi_1, \dots, \xi_k}(s)$ that is equal to 1 if and only if $s_{i_j} = \xi_j$ for all $j = 1, \dots, k$. Now the result follows from Corollary 3 and Theorem 4. \square

This gives an ϵ -almost k -wise independent set in the L_1 -norm of size $O(k2^k \log n/\epsilon^2)$. We now give a better bound. We first prove

Lemma 5. Let $0 \leq r < k$. An ϵ -almost k -wise independent set in the L_1 -norm of size

$$m = O\left(\frac{2^k + k2^r \log n}{\epsilon^2}\right)$$

can be constructed in time $\tilde{O}(2^{2^{k-r}} n^k/\epsilon^4)$.

Proof. Consider $S = \{0, 1\}^n$ with the uniform distribution. For every $a \in \{0, 1\}^r$ and $B \subseteq \{0, 1\}^{k-r}$ we define the random variable $Z_{i_1, \dots, i_k, a, B}(s) = 1$ if and only if $(s_{i_1}, \dots, s_{i_k}) \in \{a\} \times B$. Let $S' \subset S$ and suppose for every $I = (i_1, \dots, i_k)$, $a \in \{0, 1\}^r$ and $B \subseteq \{0, 1\}^{k-r}$ we have

$$|\mathbf{E}_{s \sim U_{S'}}[Z_{I, a, B}] - \mathbf{E}_{s \sim U_S}[Z_{I, a, B}]| = \left| \mathbf{E}_{s \sim U_{S'}}[Z_{I, a, B}] - \frac{|B|}{2^k} \right| \leq \frac{\epsilon}{2^{r+1}}.$$

Then (here \mathbf{E} is $\mathbf{E}_{s \sim U_{S'}}$)

$$\begin{aligned} \sum_{a \in \{0, 1\}^r, b \in \{0, 1\}^{k-r}} \left| \mathbf{E}[Z_{I, a, \{b\}}] - \frac{1}{2^k} \right| &= \sum_{a \in \{0, 1\}^r} \sum_{b \in \{0, 1\}^{k-r}} \left| \mathbf{E}[Z_{I, a, \{b\}}] - \frac{1}{2^k} \right| \\ &= \sum_{a \in \{0, 1\}^r} \max_{B \subseteq \{0, 1\}^{k-r}} \left(\mathbf{E}[Z_{I, a, B}] - \frac{|B|}{2^k} \right) + \left(\frac{|\bar{B}|}{2^k} - \mathbf{E}[Z_{I, a, \bar{B}}] \right) \\ &\leq \sum_{a \in \{0, 1\}^r} \frac{\epsilon}{2^r} = \epsilon, \end{aligned}$$

and therefore, S' is an ϵ -almost k -wise independent in the L_1 norm.

Now to construct such a set S' we use Corollary 3 and Theorem 4. We have $N = 2^{2^{k-r}+1} 2^r \binom{n}{k}$ and define for each $Z_{I, a, B}$, $\epsilon_{I, a, B} = \epsilon 2^{k-r-1}/|B|$ and $p_{I, a, B} = |B|/2^k$. By Theorem 1,

$$m \geq \frac{\ln 2^{2^{k-r}+1} 2^r \binom{n}{k}}{\min p_{I, a, B} \epsilon_{I, a, B}^2} = O\left(\frac{2^k + k2^r \log n}{\epsilon^2}\right).$$

\square

We now prove

Theorem 6. Let $d > 0$ be any real number. An ϵ -almost k -wise independent set in the L_1 -norm of size

$$m = O\left(\frac{2^k + (k/d)2^k + k \log n}{\epsilon^2}\right)$$

can be constructed in time $\tilde{O}(n^{k+d}/\epsilon^4)$.

In particular, an ϵ -almost k -wise independent set in the L_1 -norm of size

$$m = O\left(\frac{2^k + k \log n}{\epsilon^2}\right)$$

can be constructed in time $\tilde{O}(n^{2k}/\epsilon^4)$.

Proof. Follows from Lemma 6 with $r = \max(\lceil k - \log \log n - \log d \rceil, 0)$. □

The above results can be extended to any product distribution over any alphabet.

3.4 Efficient Construction for Any k

In this subsection, we give a construction that is efficient for any k . We will give the results for the uniform distribution. Similar results can be obtained for the product distribution.

We first define the dense perfect hash family. We say that $H \subseteq [q]^n$ is a $(1 - \epsilon)$ -dense (n, q, k) -perfect hash family if for every $1 \leq i_1 < i_2 < \dots < i_k \leq n$ there are at least $(1 - \epsilon)|H|$ elements $h \in H$ such that h_{i_1}, \dots, h_{i_k} are distinct.

We will use the following lemma. The proof is in [11] for a power of prime q . See also [17]. Here we give the proof for any q .

Lemma 6. If $n > q > 4k^2/\epsilon$ then there is a $(1 - \epsilon)$ -dense (n, q, k) -PHF of size

$$O\left(\frac{k^2 \log n}{\epsilon \log(\epsilon q/k^2)}\right)$$

that can be constructed in polynomial time.

Proof. We use Theorem 1. The sample space is $S = [q]^n$ with the uniform distribution. The random variables are $X_{i,j}(s) = I[s_i = s_j]$ for all $1 \leq i < j \leq n$. That is, $X_{i,j}(s) = 1$ if $s_i = s_j$ and $X_{i,j}(s) = 0$ otherwise. We have $p = \mathbf{E}[X_{i,j}] = 1/q$. The number of such random variables is $N = \binom{n}{2}$. Let $h = \lceil k^2/\epsilon \rceil$ and $\lambda = 1/h > p$. Then

$$\begin{aligned} D(\lambda||p) &= \frac{1}{h} \ln\left(\frac{1/h}{1/q}\right) + \left(1 - \frac{1}{h}\right) \ln \frac{1 - 1/h}{1 - 1/q} \\ &\geq \frac{1}{h} \ln q - \frac{1}{h} \ln h + \left(1 - \frac{1}{h}\right) \ln\left(1 - \frac{1}{h}\right) \\ &\geq \frac{\ln q - \ln h - 1}{h}. \end{aligned}$$

By Theorem 1, in polynomial time, we can find a multiset $S' = \{s_1, \dots, s_m\} \subset S$ where

$$m = \frac{\ln N}{D(\lambda||p)} = O\left(\frac{k^2 \log n}{\epsilon \log(\epsilon q/k^2)}\right)$$

such that for all $1 \leq i \leq j \leq n$ we have $\Pr_{s \in U_{S'}}[s_i = s_j] \leq 1/h$. Then for any $1 \leq i_1 < i_2 < \dots < i_k \leq n$ we have

$$\Pr_{s \in U_{S'}}[(\exists 1 \leq j_1 < j_2 \leq k) s_{i_{j_1}} = s_{i_{j_2}}] \leq \frac{\binom{k}{2}}{h} \leq \epsilon.$$

□

We are now ready to give our main three results.

We first show

Theorem 7. *An ϵ -almost k -wise independent set in the L_∞ -norm of size*

$$O\left(\frac{k^3 \log n}{2^k \epsilon^3}\right)$$

can be constructed in polynomial time.

Proof. In this proof, c_i , $i = 1, 2, 3, \dots$, are some constants. Obviously, $\epsilon < 1/2^k$. For $n \leq (k/\epsilon)^{3k}$, the size of the ϵ -almost k -wise independent set in [7] (see the table) is at most

$$\frac{c_1 k^2 \log^2 n}{\epsilon^2 \log^2(1/\epsilon)} \leq \frac{c_2 k^3 \log n}{\epsilon^2 \log(1/\epsilon)} \leq \frac{c_2 k^3 \log n}{2^k \epsilon^3}.$$

Now let $n \geq (k/\epsilon)^{3k}$. Let $q = \lceil (k/\epsilon)^3 \rceil$ and $H \subseteq [q]^n$ be a $(1 - \epsilon/4)$ -dense (n, q, k) -PHF of size

$$|H| = \frac{c_3 k^2 \log n}{\epsilon \log(\epsilon q/k^2)} \leq \frac{c_4 k^2 \log n}{\epsilon \log(1/\epsilon)}.$$

Such an H exists by Lemma 6 and can be constructed in polynomial time. By Theorem 5, an $(\epsilon/4)$ -almost k -wise independent set $R \subseteq \{0, 1\}^q$ in the L_∞ -norm of size

$$|R| = \frac{c_5 k \log q}{2^k \epsilon^2} \leq \frac{c_6 k \log(1/\epsilon)}{2^k \epsilon^2}$$

can be constructed in time $\tilde{O}(q^k/\epsilon^4) = \text{poly}(n, 1/\epsilon)$. Define

$$S = \{(v_{u_1}, \dots, v_{u_n}) \mid u \in H, v \in R\}.$$

The size of S is

$$|R| \cdot |H| = \frac{c_7 k^3 \log n}{2^k \epsilon^3}.$$

Now for a random uniform $s \in S$, any $1 \leq i_1 < \dots < i_k \leq n$ and $\xi \in \{0, 1\}^k$,

$$\begin{aligned} \Pr_{s \in S}[(s_{i_1}, \dots, s_{i_k}) = \xi] &= \Pr_{v \in R, u \in H}[(v_{u_{i_1}}, \dots, v_{u_{i_k}}) = \xi] \\ &\leq \Pr_{v \in R, u \in H}[(v_{u_{i_1}}, \dots, v_{u_{i_k}}) = \xi \mid u_{i_1}, \dots, u_{i_k} \text{ are distinct}] + \\ &\quad \Pr_{v \in R, u \in H}[u_{i_1}, \dots, u_{i_k} \text{ are not distinct}] \\ &\leq \left(\frac{1}{2^k} + \frac{\epsilon}{4}\right) + \frac{\epsilon}{4} \leq \frac{1}{2^k} + \epsilon \end{aligned}$$

and

$$\begin{aligned}
\Pr_{s \in S}[(s_{i_1}, \dots, s_{i_k}) = \xi] &= \Pr_{v \in R, u \in H}[(v_{u_{i_1}}, \dots, v_{u_{i_k}}) = \xi] \\
&\geq \Pr_{v \in R, u \in H}[(v_{u_{i_1}}, \dots, v_{u_{i_k}}) = \xi \mid u_{i_1}, \dots, u_{i_k} \text{ are distinct}] \cdot \\
&\quad \Pr_{v \in R, u \in H}[u_{i_1}, \dots, u_{i_k} \text{ are distinct}] \\
&\geq \left(\frac{1}{2^k} - \frac{\epsilon}{4}\right) \left(1 - \frac{\epsilon}{4}\right) \geq \frac{1}{2^k} - \epsilon.
\end{aligned}$$

Therefore, S is ϵ -almost k -wise independent set in L_∞ -norm. \square

For the L_1 -norm we prove

Theorem 8. *An ϵ -almost k -wise independent set in the L_1 -norm of size*

$$\tilde{O}\left(\frac{2^k + \log n}{\epsilon^3}\right)$$

can be constructed in polynomial time.

This theorem follows from the following two results

Lemma 7. *For $n > 2^{2^k}$, an ϵ -almost k -wise independent set in the L_1 -norm of size*

$$O\left(\frac{k^3 \log^{1/2}(1/\epsilon) \log n}{\epsilon^3}\right) = \tilde{O}\left(\frac{\log n}{\epsilon^3}\right)$$

can be constructed in polynomial time.

Proof. Let $n \geq 2^{2^k}$. If $\epsilon < 1/2^{2^k/k}$ then the AGC+Hadamard construction is of size (see the table)

$$\frac{c_1 k 2^{3k/2} \log n}{\epsilon^3 (k + \log(1/\epsilon))} \leq \frac{c_2 k^{2.5} \log^{1/2}(1/\epsilon) \log n}{\epsilon^3}.$$

Now let $\epsilon \geq 1/2^{2^k/k}$. Let $q = \lceil 2^{2^k/k} k^2 / \epsilon \rceil$ and $H \subseteq [q]^n$ be a $(1 - \epsilon/4)$ -dense (n, q, k) -PHF of size

$$\frac{c_3 k^2 \log n}{\epsilon \log(\epsilon q / k^2)} \leq \frac{c_4 k^3 \log n}{\epsilon 2^k}.$$

Such an H exists by Lemma 6 and can be constructed in polynomial time. By Theorem 6, an $(\epsilon/4)$ -almost k -wise independent set $R \subseteq \{0, 1\}^q$ in the L_1 -norm of size

$$\frac{c_5 (2^k + k \log q)}{\epsilon^2} \leq \frac{c_6 2^k}{\epsilon^2}$$

can be constructed in time $\tilde{O}(q^{2^k}/\epsilon^4) = \text{poly}(n, 1/\epsilon)$. Define

$$S = \{(v_{u_1}, \dots, v_{u_n}) \mid u \in H, v \in R\}.$$

The size of S is

$$|R| \cdot |H| = \frac{c_7 k^3 \log n}{\epsilon^3}.$$

Now for a random uniform $s \in S$, any $1 \leq i_1 < \dots < i_k \leq n$ and any boolean function $f : \{0, 1\}^k \rightarrow \{0, 1\}$,

$$\begin{aligned} \Pr_{s \in S}[f(s_{i_1}, \dots, s_{i_k}) = 1] &= \Pr_{v \in R, u \in H}[f(v_{u_{i_1}}, \dots, v_{u_{i_k}}) = 1] \\ &\leq \Pr_{v \in R, u \in H}[f(v_{u_{i_1}}, \dots, v_{u_{i_k}}) = 1 \mid u_{i_1}, \dots, u_{i_k} \text{ are distinct}] + \\ &\quad \Pr_{v \in R, u \in H}[u_{i_1}, \dots, u_{i_k} \text{ are not distinct}] \\ &\leq \left(\Pr[f = 1] + \frac{\epsilon}{4} \right) + \frac{\epsilon}{4} \leq \Pr[f = 1] + \epsilon. \end{aligned}$$

Therefore S is ϵ -almost k -wise independent set in L_∞ -norm. \square

Lemma 8. For $n \leq 2^{2^k}$, an ϵ -almost k -wise independent set in the L_1 -norm of size

$$\tilde{O}\left(\frac{2^k}{\epsilon^3}\right)$$

can be constructed in polynomial time.

Proof. If $n < (4k^2/\epsilon)^{2k}$ then we use the bound in [7] (see the table) and get an ϵ -almost k -wise independent set in the L_1 -norm of size

$$O\left(\frac{k^2 \min(k^2, \log^2(1/\epsilon)) \cdot 2^k}{\epsilon^2}\right) = \tilde{O}\left(\frac{2^k}{\epsilon^2}\right).$$

Otherwise, we take $q = \lceil n^{1/k} \rceil \geq (4k^2/\epsilon)^2$ and use Theorem 6 to construct an $(\epsilon/4)$ -almost k -wise set in the L_1 -norm, $R \subseteq \{0, 1\}^q$, of size (notice that $\log n \leq 2^k$)

$$O\left(\frac{2^k}{\epsilon^2}\right)$$

in time $\tilde{O}(q^{2k}/\epsilon^4) = \text{poly}(n, 1/\epsilon)$. Composing this with the $(\epsilon/4)$ -dense (n, q, k) -PFH in Lemma 6 of size

$$O\left(\frac{k^2 \log n}{\epsilon \log(\epsilon q/k^2)}\right) = O\left(\frac{k^3}{\epsilon}\right)$$

gives the result. \square

4 Lower Bounds

In this section we give the following two lower bounds

Theorem 9. Let $1/\text{poly}(n) \leq \epsilon < 1/2^{k+1}$ and $k < n/2$. Any ϵ -almost k -wise independent set in the L_∞ -norm is of size

$$\Omega\left(\frac{\log n}{2^k \epsilon^2 \log(1/2^{k+1}\epsilon)}\right) = \tilde{\Omega}\left(\frac{\log n}{2^k \epsilon^2}\right).$$

Theorem 10. Let $\epsilon > 1/n^{k/5}$. Any ϵ -almost k -wise independent set in the L_1 -norm is of size

$$\Omega\left(\frac{k \log n}{\epsilon^2 \log(1/\epsilon)}\right) = \tilde{\Omega}\left(\frac{\log n}{\epsilon^2}\right).$$

The following is proved in [2, 1]. We give the proof for completeness.

Lemma 9. *Let $S \subset \{0, 1\}^n$ and $r < n$ be an even number. If for every distinct $i_1, \dots, i_r \in [n]$, $\alpha \in \{0, 1\}^r \setminus \{0^r\}$ and a random uniform $s \in S$ we have $1/2 + \epsilon \geq \Pr[\alpha_1 s_{i_1} + \dots + \alpha_r s_{i_r}] \geq 1/2 - \epsilon$ then*

$$|S| \geq \Omega \left(\min \left(\frac{r \log(n/r)}{\epsilon^2 \log(1/\epsilon)}, 2^{r/2} \binom{n}{r/2} \right) \right).$$

In particular, for $\epsilon > 1 / \left(2^{r/4} \binom{n}{r/2}^{1/2} \right)$

$$|S| \geq \Omega \left(\frac{r \log(n/r)}{\epsilon^2 \log(1/\epsilon)} \right).$$

Proof. Let $S = \{s^{(1)}, \dots, s^{(m)}\}$ and $I = \{(i_1, \dots, i_{r/2}) \mid 1 \leq i_1 < i_2 < \dots < i_{r/2} \leq 1\}$. Consider

$$C = \{(\beta_1 s_{i_1}^{(1)} + \dots + \beta_{r/2} s_{i_{r/2}}^{(1)}, \dots, \beta_1 s_{i_1}^{(m)} + \dots + \beta_{r/2} s_{i_{r/2}}^{(m)}) \mid i \in I, \beta \in \{0, 1\}^{r/2} \setminus \{0^{r/2}\}\}.$$

Then for two distinct $u, v \in C$ we have (for some i_1, \dots, i_r and $(\alpha_1, \dots, \alpha_r) \in \{0, 1\}^r \setminus \{0^r\}$)

$$\text{dist}(u, v) = wt(u + v) = m \cdot \Pr_{s \in U_S} [\alpha_1 s_{i_1} + \dots + \alpha_r s_{i_r}] \in \left[\left(\frac{1}{2} - \epsilon \right) m, \left(\frac{1}{2} + \epsilon \right) m \right].$$

Therefore C is an ϵ -balance error-correcting code size $(2^{r/2} - 1) \binom{n}{r/2}$. By MRRW bound,[19], for binary code with the results in Section 7 and (3) in [7] and the bound in [1], the result follows. \square

We now prove Theorem 9

Proof. Let S be a ϵ -almost k -wise set in the L_∞ norm. Let $r \geq 2$ be an even constant such that $\epsilon > 1 / \left(2^{r/4} \binom{n-k+r}{r/2}^{1/2} \right)$. For $\xi \in \{0, 1\}^{k-r}$ consider the sets $S_\xi = \{s \in S \mid (s_1, \dots, s_{k-r}) = \xi\}$. Obviously, $\{S_\xi\}_{\xi \in \{0, 1\}^{k-r}}$ is a partition of S .

Let $I = \{k - r + 1, k - r + 2, \dots, n\}$. For distinct $i_1, \dots, i_r \in I$, $\alpha \in \{0, 1\}^r \setminus \{0^r\}$, $\xi = \xi_1, \dots, \xi_{k-r} \in \{0, 1\}^{k-r}$, and a random uniform $x \in S$ we have

$$\begin{aligned} \Pr[\alpha_1 x_{i_1} + \dots + \alpha_r x_{i_r} = 1 \mid x \in S_\xi] &= \frac{\Pr[\alpha_1 x_{i_1} + \dots + \alpha_r x_{i_r} = 1, (x_1, \dots, x_{k-r}) = \xi]}{\Pr[(x_1, \dots, x_{k-r}) = \xi]} \\ &= \frac{\Pr[\alpha_1 x_{i_1} + \dots + \alpha_r x_{i_r} = 1, (x_1, \dots, x_{k-r}) = \xi]}{\sum_{u \in \{0, 1\}^r} \Pr[(x_{i_1}, \dots, x_{i_r}) = u, (x_1, \dots, x_{k-r}) = \xi]} \\ &\geq \frac{2^{r-1} \left(\frac{1}{2^k} - \epsilon \right)}{2^r \left(\frac{1}{2^k} + \epsilon \right)} \geq \frac{1}{2} - 2^{k+1} \epsilon. \end{aligned}$$

In the same way

$$\Pr[\alpha_1 x_{i_1} + \dots + \alpha_r x_{i_r} = 1 \mid x \in S_\xi] \leq \frac{1}{2} + 2^{k+1} \epsilon.$$

Therefore, by Lemma 9, for $\epsilon > 1 / \left(2^{r/4} \binom{n-k+r}{r/2}^{1/2} \right)$,

$$|S| = \sum_{\xi \in \{0, 1\}^{k-r}} |S_\xi| = 2^{k-r} \cdot \Omega \left(\frac{r \log(|I|/r)}{(2^{k+1} \epsilon)^2 \log(1/(2^{k+1} \epsilon))} \right) = \Omega \left(\frac{\log n}{2^k \epsilon^2 \log(1/2^{k+1} \epsilon)} \right).$$

\square

We now prove Theorem 10

Proof. For every distinct $i_1, \dots, i_k \in [n]$ and $\alpha \in \{0, 1\}^k \setminus \{0^k\}$ and any random uniform $x \in S$

$$\begin{aligned} \left| \Pr[\alpha_1 x_{i_1} + \dots + \alpha_k x_{i_k} = 1] - \frac{1}{2} \right| &\leq \sum_{\alpha_1 \xi_1 + \dots + \alpha_k \xi_k = 1} \left| \Pr[x_{i_1} = \xi_1, \dots, x_{i_k} = \xi_k] - \frac{1}{2^k} \right| \\ &\leq \sum_{\xi \in \{0,1\}^k} \left| \Pr[x_{i_1} = \xi_1, \dots, x_{i_k} = \xi_k] - \frac{1}{2^k} \right| \leq \epsilon. \end{aligned}$$

Therefore, by Lemma 9, for $\epsilon > 1/n^{k/5}$, we have

$$|S| = \Omega\left(\frac{k \log n}{\epsilon^2 \log(1/\epsilon)}\right).$$

□

Acknowledgement. We would like to thank Noga Alon for pointing out the paper [1].

References

- [1] N. Alon. Perturbed Identity Matrices Have High Rank: Proof and Applications. *Combinatorics, Probability & Computing*. 18(1-2). pp. 3-15. (2009).
- [2] N. Alon, A. Andoni, T. Kaufman, K. Matulef, R. Rubinfeld, N. Xie. Testing k -wise and almost k -wise independence. STOC 2007. pp. 496–505. (2007).
- [3] N. Alon, L. Babai, A. Itai. A Fast and Simple Randomized Parallel Algorithm for the Maximal Independent Set Problem. *J. Algorithms* 7(4). pp. 567–583. (1986).
- [4] N. Alon, J. Bruck, J. Naor, M. Naor, R. M. Roth. Construction of asymptotically good low-rate error-correcting codes through pseudo-random graphs. *IEEE Transactions on Information Theory* 38(2). pp. 509-516. (1992)
- [5] N. Alon, S. Gutner. Balanced hashing, color coding and approximate counting. *Electronic Colloquium on Computational Complexity (ECCC)* 16: 12 (2009).
- [6] N. Alon, S. Gutner. Balanced families of perfect hash functions and their applications. *ACM Transactions on Algorithms*. 6(3). (2010).
- [7] N. Alon, O. Goldreich, J. Hastad, R. Peralta. Simple constructions of almost k -wise independent random variables. FOCS 1990. pp. 544–553. (1990).
- [8] N. Alon, O. Goldreich, Y. Mansour. Almost k -wise independence versus k -wise independence. *Inf. Process. Lett.* 88(3): 107-110 (2003)
- [9] S. Arora, E. Hazan, S. Kale. The multiplicative weights update method: A meta-algorithm and applications. *Theory of computing*, 8, pp. 121–164. (2012).

- [10] N. Alon, D. Moshkovitz, S. Safra. Algorithmic construction of sets for k -restrictions. *ACM Transactions on Algorithms* 2(2). pp. 153–177. (2006).
- [11] N. H. Bshouty. Linear time Constructions of some d -Restriction Problems. CoRR abs/1406.2108 (2014). CIAC 2015. pp. 74–88. (2015).
- [12] N. H. Bshouty, A. Gabizon. Almost Optimal Cover-Free Families. CoRR abs1507.07368. (2015).
- [13] A. Ben-Aroya, A. Ta-Shma. Constructing Small-Bias Sets from Algebraic-Geometric Codes. *Theory of Computing*. 9. pp. 253–272. (2013).
- [14] B. Chor, O. Goldreich, J. Håstad, J. Friedman, S. Rudich, R. Smolensky. The Bit Extraction Problem of t -Resilient Functions. *FOCS 1985*: pp. 396–407. (1985).
- [15] R. Impagliazzo, V. Kabanets. Constructive Proofs of Concentration Bounds *ECCC TR10-072*. (2010).
- [16] D. Koller, N. Megiddo. Constructing Small Sample Spaces Satisfying Given Constraints. *SIAM J. Discrete Math.* 7(2). pp. 260–274. (1994).
- [17] E. Porat, A. Rothschild. Explicit Nonadaptive Combinatorial Group Testing Schemes. *IEEE Transactions on Information Theory* 57(12), pp. 7982–7989 (2011).
- [18] J. Naor, M. Naor. Small-Bias Probability Spaces: Efficient Constructions and Applications. *SIAM J. Comput.* 22(4). pp. 838–856 (1993)
- [19] R. J. McEliece, E. R. Rodemich, H. Rumsey Jr., L. R. Welch. New upper bounds on the rate of a code via the Delsarte-MacWilliams inequalities. *IEEE Transactions on Information Theory* 23(2): 157–166 (1977)
- [20] E. Porat, A. Rothschild. Explicit Nonadaptive Combinatorial Group Testing Schemes. *IEEE Transactions on Information Theory*. 57(12), pp. 7982–7989. (2011).
- [21] M. Naor, L. J. Schulman, A. Srinivasan. Splitters and Near-Optimal Derandomization. *FOCS 1995*: 182–191. (1995).
- [22] A. Wigderson, D. Xiao. Derandomizing the Ahlswede-Winter matrix-valued Chernoff bound using pessimistic estimators, and applications. *Theory of Computing* 4(1): pp. 53–76. (2008).