On Construction of k-wise Independent Random Variables

Howard Karloff * Yishay Mansour †

Abstract

A *0-1 probability space* is a probability space $(\Omega, 2^\Omega, P)$, where the sample space $\Omega \subseteq \{0,1\}^n$ for some n. A probability space is *k-wise independent* if, when Y_i is defined to be the ith coordinate of the random n-vector, then any subset of k of the Y_i's is (mutually) independent, and it is said to be a probability space *for* $p_1, p_2, ..., p_n$ if $P[Y_i = 1] = p_i$.

We study constructions of k-wise independent 0-1 probability spaces in which the p_i's are arbitrary. It was known that for any $p_1, p_2, ..., p_n$, a k-wise independent probability space of size $m(n,k) = \binom{n}{k} + \binom{n}{k-1} + \binom{n}{k-2} + \cdots + \binom{n}{0}$ always exists. We prove that for some $p_1, p_2, ..., p_n \in [0,1]$, $m(n,k)$ is a lower bound on the size of any k-wise independent 0-1 probability space. For each fixed k, we prove that every k-wise independent 0-1 probability space when each $p_i = k/n$ has size $\Omega(n^k)$. For a very large degree of independence—$k = \lfloor \alpha n \rfloor$, for $\alpha > 1/2$—and all $p_i = 1/2$, we prove a lower bound on the size of $2^n(1 - \frac{1}{2\alpha})$. We also give explicit constructions of k-wise independent 0-1 probability spaces.

1 Introduction

A *0-1 probability space* $\mathcal{H}$ is a probability space $(\Omega, 2^\Omega, P)$ whose sample space, Ω, is a set of 0-1 strings of length n (for some n). The size of a probability space $\mathcal{H} = (\Omega, 2^\Omega, P)$ is $|\Omega|$. $\mathcal{H}$ is *k-wise independent* if, when Y_i is defined to be the ith coordinate of the random n-vector, then any subset of k of the Y_i's is (mutually) independent, and it is said to be a probability space *for* $p_1, p_2, ..., p_n$ if $P[Y_i = 1] = p_i$.

A great deal of research in theoretical computer science recently has dealt with the problem of constructing small probability spaces satisfying certain independence constraints, most commonly k-wise independence [Joffe74, Lub86, ABI86, Lub88]. Aside from its interest as a fundamental issue, such small probability spaces can often be used for derandomizing randomized algorithms [ABI86, Lub86, KM93, Sch92, BR89, MNN89]: if one has a randomized algorithm which needs to be run on a k-wise independent 0-1 probability space, then it can be derandomized by running it deterministically on all the points in such a space (in parallel, if possible). This deterministic algorithm will be efficient if the size of the sample space is small, so that not too much time (or processors, in the parallel case) will be needed, and if the algorithm runs quickly on each point of the sample space.

*College of Computing, Georgia Institute of Technology, Atlanta, GA 30332-0280. This author was supported in part by NSF grant CCR 9107349.

†Computer Science Dept., Tel-Aviv University. This research was supported in part by the Israel Science Foundation administered by the Israel Academy of Science and Humanities and by a grant of the Israeli Ministry of Science and Technology.

Almost all previous work has dealt with the construction of *uniform* k-wise independent 0-1 probability spaces. These are spaces Ω in which all points $\omega \in \Omega$ have the same probability. For derandomization, nonuniform spaces are as good as uniform ones. A sample space Ω whose $|\Omega|$ points have many different probabilities will serve just as well. Suppose, for example, that a probabilistic construction needs a pairwise independent 0-1 probability space for $p_1, p_2, ..., p_n$. The conclusion of the existence proof is that some structure exists with positive probability. Then if we have a "small" pairwise independent 0-1 probability space Ω for $p_1, p_2, ..., p_n$, then we can try each of the points of Ω, and we will be guaranteed to find one for which the structure exists.

We start with lower bounds. Define

$$m(n,k) = \binom{n}{k} + \binom{n}{k-1} + \binom{n}{k-2} + \cdots + \binom{n}{0}.$$

For constant k, $m(n,k)$ is $\Theta(n^k)$. The best general lower bound is from [ABI86, CGHFRS85]; cf. [AS92].

Theorem 1. [ABI86, CGHFRS85] If $\mathcal{H} = (\Omega, 2^\Omega, P)$ is a k-wise independent 0-1 probability space for $p_1, p_2, ..., p_n$ and all $p_i \in (0,1)$, then $|\Omega| \geq m(n, \lfloor k/2 \rfloor)$.

In fact, the theorem of [ABI86, CGHFRS85] is actually stronger, since it applies to random variables with arbitrary real range. We show that this lower bound is not always tight. We prove a lower bound of $m(n,k)$ for certain specified probabilities:

Theorem 2. Let n be a positive integer and let $1 \leq k \leq n$. Then there exist $p_1, p_2, ..., p_n \in (0,1)$ such that the size of every k-wise independent 0-1 probability space for $p_1, p_2, ..., p_n$ is at least $m(n,k)$.

Even when all the probabilities are identical, we can prove a lower bound on the size that is $\Omega(n^k)$, provided that k is fixed:

Theorem 3. Fix $k \geq 2$. The size of every k-wise independent 0-1 probability space when every $p_i = k/n$ is $\Omega(n^k)$.

In fact, we show that the result holds even if each p_i is sufficiently close to k/n.

Note that neither Theorem 2 nor Theorem 3 can be improved by allowing the probabilities to be arbitrary, for there *is* a k-wise independent 0-1 probability space where all $p_i = 1/2$ of size at most $2(2n)^{\lfloor k/2 \rfloor}$ [ABI86].

How close to optimal are these theorems? It is implicit in the work of D. Koller and N. Megiddo [KM93], who seem to have initiated the systematic study of nonuniform spaces in theoretical computer science, that there is always a space of size at most $m(n,k)$:

Theorem 4. [KM93] Let $n \geq 1$, let $1 \leq d \leq n$, and let $p_1, p_2, ..., p_n \in [0,1]$. Then there is a k-wise independent 0-1 probability space for $p_1, ..., p_n$ of size at most $m(n,k)$.

This implies that Theorem 2 is exactly tight and the bound of Theorem 3 is tight to within a constant factor (which depends on k but not n).

We also study k-wise independent 0-1 probability spaces when k is very large, say, $k = \lfloor \alpha n \rfloor$ for $\alpha \in (1/2, 1]$. In this case the best lower bound on size was $m(n, \lfloor k/2 \rfloor)$ (from Theorem 1). This quantity is bounded above by β^n, where $\beta < 2$ is a function only of α. Based on the techniques of [Fri92] we improve this bound as well.

Theorem 5. Let $1/2 < \alpha \leq 1$. Then the size of any $\lfloor \alpha n \rfloor$-wise independent 0-1 probability space when all $p_i = 1/2$ is at least $2^n(1 - \frac{1}{2\alpha})$.

Since it is trivial to build an n-wise independent 0-1 probability space on $\Omega = \{0,1\}^n$ for any sequence $p_1, p_2, ..., p_n$, this theorem is also tight to within a constant factor. Furthermore, when $\alpha \in (0, 1/2)$, the upper bound of Theorem 4 implies that there is a k-wise independent 0-1 probability space of size at most γ^n for some $\gamma < 2$, so one cannot give an $\Omega(2^n)$ lower bound for any fixed $\alpha < 1/2$.

We now turn to construction of small k-wise independent 0-1 probability spaces. We first give a general construction for arbitrary p_i. Namely, given $p_1, p_2, ..., p_n \in [0,1]$, which are possibly irrational, we give an explicit construction of a k-wise independent 0-1 probability space for $p_1, ..., p_n$. The size of the sample space is at most $(8n)^{2k}$. Note that for a fixed k the size of the sample space is bounded by a polynomial in the lower bound.

(In fact, the Koller-Megiddo technique shows how to generate in polynomial time a k-wise independent 0-1 probability space of size at most $m(n,k)$, for any $p_1, ..., p_n$, but it is not a "construction." One has to execute n times an algorithm that takes a feasible solution to a linear program and produces a basic feasible solution. On the other hand, their method is more general in that it is not restricted to k-wise independence.)

An interesting and well-studied special case of k-wise independent 0-1 random variables is the case in which $k = 2$ and every $p_i = p$. For this case we give a construction with a sample space of size $O(n^2)$, which is the best one could hope for if p is to be arbitrary, by Theorem 3. For some specific ranges we show that linear-size constructions are possible.

The paper is organized as follows. In Section 2 we give the lower bound in the case that the probabilities are different. Section 3 derives a lower bound in the case that the probabilities are the same. The lower bound for a large degree of independence appears in Section 4. The general construction with arbitrary probabilities is in Section 5. The construction for pairwise independent random variables is in Section 6.

2 Lower Bound for Different Probabilities

Throughout this section, let $0 < \gamma < 1/m(n,k)$ and let $p_i = \gamma^{2^{i-1}}$. We derive a tight lower bound of $m(n,k)$ on the size of any k-wise independent 0-1 probability space $\mathcal{H}$ for $p_1, p_2, ..., p_n$.

First we define a lexicographic ordering between sets.

Definition. Let $S = \{s_1, \ldots, s_l\}$ and $T = \{t_1, \ldots, t_j\}$ be distinct subsets of $\{1, 2, ..., n\}$, where $s_1 > s_2 > \cdots > s_l$ and $t_1 > t_2 > \cdots > t_j$. Then $S \succ T$ if either (1) S is a proper superset of T, or (2) there exists an m such that $s_i = t_i$ for $i = 1, 2, ..., m$ and $s_{m+1} > t_{m+1}$. (Notice that either $S \succ T$ or $T \succ S$ if $S \neq T$.)

For a set $S \subseteq \{1, 2, ..., n\}$ and for any $y \in \{0,1\}^n$, with ith coordinate y_i, let $I_S(y) = \prod_{i \in S} y_i$. Since the distribution is k-wise independent, for any specific set S of size at most k the probability of the event $[I_S(y) = 1]$ is exactly $\prod_{i \in S} p_i$, so let us use q_S to denote $\prod_{i \in S} p_i$. The following lemma shows that the q's of different sets are well separated.

Lemma 6. Let $S = \{s_1, \ldots, s_l\}$ and $T = \{t_1, \ldots, t_j\}$ be sets of size at most k such that $S \succ T$. Then

$$q_T/q_S = (p_{t_1} \cdots p_{t_j})/(p_{s_1} \cdots p_{s_l}) \geq 1/\gamma > m(n,k).$$

Proof. If T is a proper subset of S, then the claim follows from the fact that for any i, $1/p_i \geq 1/\gamma > m(n,k)$, so assume T is not a proper subset of S. Without loss of generality $t_1 > t_2 > \cdots > t_j$, $s_1 > s_2 > \cdots > s_l$, $s_i = t_i$ for $i \leq m$, $s_{m+1} > t_{m+1}$. By definition, $q_T = p_{t_1} p_{t_2} \cdots p_{t_j} = (\prod_{i \leq m} p_{t_i})(\prod_{i>m} p_{t_i})$, and $q_S = p_{s_1} p_{s_2} \cdots p_{s_l} = (\prod_{i \leq m} p_{s_i})(\prod_{i>m} p_{s_i})$. Thus $q_T/q_S = \prod_{i>m} p_{t_i} / \prod_{i>m} p_{s_i}$. Let $r = t_{m+1}$. Since $s_{m+1} \geq t_{m+1} + 1 = r + 1$, the numerator of the quotient is at least $\prod_{i=1}^{r} p_i = \prod_{i=1}^{r} \gamma^{2^{i-1}} = \gamma^{2^r - 1}$. The denominator is at most $p_{s_{m+1}} = \gamma^{2^{s_{m+1}-1}} \leq \gamma^{2^r}$. Thus $q_T/q_S \geq \gamma^{(2^r-1)-2^r} = 1/\gamma > m(n,k)$. ■

Lemma 7. Let $\mathcal{H}$ be a k-wise independent 0-1 probability space for the $p_1, p_2, ..., p_n$ defined above. For any subset T of $\{1, 2, ..., n\}$ of size at most k, there is a vector x_T assigned positive probability by $\mathcal{H}$ such that $I_T(x_T) = 1$, and $I_S(x_T) = 0$ for all S such that $S \succ T$ and $|S| \leq k$.

Proof. Consider the following event:

$$[I_T(x) = 1 \text{ and } I_S(x) = 0 \text{ for all } S \text{ such that } S \succ T,\ |S| \leq k].$$

The probability of this event is bounded from below by

$$P[I_T(x) = 1] - \sum_{\substack{S \succ T, \\ |S| \leq k}} P[I_S(x) = 1].$$

Since the size of T is at most k and the size of each S is at most k, the probability of the event can be bounded from below via k-wise independence. The result is $q_T - \sum_{\substack{S \succ T, \\ |S| \leq k}} q_S$. By Lemma 6, for any $S \succ T$, $q_T \geq q_S/\gamma$. Since $m(n,k) < 1/\gamma$, we have $q_T - \sum_{\substack{S \succ T, \\ |S| \leq k}} q_S > 0$. Thus the probability of the event

$$[I_T(x) = 1 \text{ and } I_S(x) = 0 \text{ for all } S \text{ such that } |S| \leq k \text{ and } S \succ T]$$

is positive. It follows that there is some vector x_T assigned positive probability by $\mathcal{H}$ that has the required properties. ■

Now Theorem 2 follows from Theorem 8.

Theorem 8. Let $\mathcal{H} = (\Omega, 2^\Omega, P)$ be any k-wise independent 0-1 probability space for the $p_1, ..., p_n$ defined above. Then $|\Omega| \geq m(n,k)$.

Proof. By Lemma 7, for any set T of size at most k, there is a vector $x_T \in \{0,1\}^n$ having positive probability in $\mathcal{H}$ such that $I_T(x_T) = 1$ and also $I_S(x_T) = 0$ for any S of size at most k such that $S \succ T$. We need to show for any S and T, both of size at most k and satisfying $S \neq T$, that $x_S \neq x_T$. Without loss of generality assume that $S \succ T$. Hence, $I_S(x_T) = 0$. But by the definition of x_S, $I_S(x_S) = 1$. Therefore $x_S \neq x_T$. ■

Remark: It is easy to see that it is not necessary that each probability be exactly p_i; there is some slack in the proof. This means that for each n and i there is a subinterval of $[0,1]$ of positive length so that if p_i' is in the given interval for all i, then the size of any k-wise independent 0-1 probability space for $p_1', p_2', ..., p_n'$ is at least $m(n,k)$.

3 Lower Bound for Identical Probabilities

Fix $k \geq 2$. Let $p_1, p_2, ..., p_n$ be a sequence of probabilities, $2 \leq k \leq n$. Choose reals $M, M' \geq 0$ such that M is a lower bound on the product of the smallest $k-1$ probabilities and M' is an upper bound on the product of the k (not $k-1$) largest probabilities. In this section we study k-wise independent 0-1 probability spaces for $p_1, p_2, ..., p_n$. We prove an $\Omega(n^k)$ lower bound on the size for $p_1, p_2, ..., p_n$, provided that

1. $M > 1.8/\binom{n}{k-1}$, and
2. $M > M' \cdot 0.9n/k$.

We will see later that these conditions are satisfied if $(M, M') = ((k/n)^{k-1}, (k/n)^k)$, i.e., if $p_i = k/n$ for all i. But notice that since both conditions are strict inequalities, if they are satisfied by (M, M') as above, then they are also satisfied by $(M - \epsilon, M' + \epsilon)$, where $\epsilon > 0$ depends on n and k. Thus these conditions hold if each p_i is in a small neighborhood around k/n.

First, we need the following technical lemma. It will help us prove that there are many strings in Ω with few ones apiece.

Lemma 9. Let $\alpha, \beta > 0$. Suppose $S_1, S_2, ..., S_m \subseteq \{1, 2, ..., N\}$ and $x_i > 0$ is associated with S_i such that

1. For $j = 1, 2, ..., N$, $\sum_{i:S_i \ni j} x_i \leq \alpha$, and
2. $\sum_{i=1}^{m} x_i \geq \beta$.

Then for any integer $l \geq 2$, there are at least $[l\beta - N\alpha]/[l(l-1)\alpha]$ pairwise disjoint S_i's, each of size less than l.

Proof. Let $T_j = \{i : S_i \ni j\}$. From the assumptions about the x_i's, for $j = 1, 2, ..., N$, $\sum_{i \in T_j} x_i \leq \alpha$, and therefore $\sum_{j=1}^{N} \sum_{i \in T_j} x_i \leq N\alpha$. Order the sets so that $|S_1| \leq |S_2| \leq \cdots \leq |S_m|$. Choose the minimal p so that $|S_p| \geq l$, choosing $p = m + 1$ if all $|S_i| < l$. Then $l \sum_{i=p}^{m} x_i \leq \sum_{i=p}^{m} |S_i| x_i \leq \sum_{i=1}^{m} |S_i| x_i \leq N\alpha$. Therefore $\sum_{i=p}^{m} x_i \leq N\alpha/l$. Since the sum of all x_i's is bounded below by β, we have $\beta \leq \sum_{i=1}^{p-1} x_i + \sum_{i=p}^{m} x_i \leq \sum_{i=1}^{p-1} x_i + N\alpha/l$. Therefore $\sum_{i=1}^{p-1} x_i \geq \beta - N\alpha/l$.

We may assume that $\beta - N\alpha/l > 0$, since otherwise the lemma is trivially true. Let $I = \emptyset$. Let $\mathcal{S} = \{1, 2, ..., p-1\}$. All sets S_i for $i \in \mathcal{S}$ have size less than l. Since $\sum_{i \in \mathcal{S}} x_i \geq \beta - N\alpha/l > 0$, there is a set S_t of size less than l with $t \in \mathcal{S}$.

Add t to I, and eliminate from $\mathcal{S}$ all i such that $S_i \cap S_t \neq \emptyset$ (including t itself). The sum of the x's of the sets remaining in $\mathcal{S}$ is at least $(\beta - N\alpha/l) - (l-1)\alpha$, since the sum of x_i over all S_i with S_i intersecting S_t is at most $\alpha|S_t|$.

If $(\beta - N\alpha/l) - (l-1)\alpha > 0$, we repeat the process. This process can be repeated at least $[\beta - N\alpha/l]/[(l-1)\alpha]$ times. At the end, we have a set of at least $[\beta l - N\alpha]/[l(l-1)\alpha]$ pairwise disjoint sets, each of size less than l. ■

Definition. Identify $B \subseteq \{1, 2, ..., n\}$ with its characteristic n-vector. We use $P[B]$ to denote the probability assigned by $\mathcal{H}$ to the characteristic vector of B.

Given $p_1, p_2, ..., p_n$ and M, M' as above, and a k-wise independent 0-1 probability space $\mathcal{H}$, say a $(k-1)$-subset S of $\{1, 2, ..., n\}$ is *valid* if according to $\mathcal{H} = (\Omega, 2^\Omega, P)$,

$$P[S] < 2M/3.$$

We first show that at least one sixth of the sets of size $k-1$ are valid.

Lemma 10. For every k-wise independent 0-1 probability space $\mathcal{H}$ as above, with the given restrictions on M and M', at least one sixth of the sets of size $k-1$ are valid.

Proof. If S is not valid, then the characteristic vector of S must be assigned probability at least $2M/3$. Clearly, the sum of the probabilities assigned by $\mathcal{H}$ to characteristic vectors of $(k-1)$-sets is at most 1. Therefore the number of invalid $(k-1)$-sets is at most $3/(2M)$. Because $M > 1.8/\binom{n}{k-1}$, we have $3/(2M) < (5/6)\binom{n}{k-1}$. ■

Theorem 11. Every k-wise independent 0-1 probability space $\mathcal{H}$ for $p_1, p_2, ..., p_n$ as above has size at least $\frac{1}{6}\binom{n}{k-1}\frac{n}{45k^2}/\binom{7k}{k-1}$. If $n \geq 2k$, then this is at least $c \cdot m(n,k)/(k^3 2^{7k})$ for a universal constant $c > 0$.

Proof. Suppose $\{1, 2, ..., k-1\}$ is valid. Let $U_1, U_2, ..., U_m$ be the proper supersets of $\{1, 2, ..., k-1\}$ such that $P[U_i] > 0$. Let $x_i = P[U_i]$ and $S_i = U_i - \{1, 2, ..., k-1\}$. Because $\mathcal{H}$ is a k-wise independent 0-1 probability space for $p_1, p_2, ..., p_n$, the chance of having ones simultaneously in positions $1, 2, ..., k-1$ and j (for $j \geq k$) is at most M'. But this probability is exactly $\sum_{i:S_i \ni j} x_i$, so $\sum_{i:S_i \ni j} x_i \leq M'$. We have

1. $S_i \subseteq \{k, k+1, ..., n\}$.
2. For $j = k, k+1, ..., n$, $\sum_{i:S_i \ni j} x_i \leq M'$, and
3. $\sum_{i=1}^m x_i > M/3$ (because $\{1, 2, ..., k-1\}$ is valid).

Apply Lemma 9 with $N = n-(k-1)$, $\alpha = M'$, $\beta = M/3$. By the lemma, for all $l \geq 2$, there are at least

$$\frac{\frac{1}{3}lM - (n-k+1)M'}{l(l-1)M'} = \frac{\frac{1}{3}M/M' - (n-k+1)/l}{l-1}$$

pairwise disjoint S_i's, each of size at most $l-1$. Now set $l = 6k+1$. Because $M/M' \geq 0.9n/k$, we have at least

$$\frac{1}{6k}\left[\frac{1}{3}\frac{M}{M'} - \frac{n}{6k}\right] \geq \frac{n}{45k^2}$$

pairwise disjoint S_i's, each of size at most $6k-1$. Thus there are at least $\frac{n}{45k^2}$ strings having positive probability in $\mathcal{H}$ with ones in positions $1, 2, ..., k-1$ and at most $(k-1)+(l-1) = 7k-1$ ones in total.

Now this argument obviously goes through for each of the $\frac{1}{6}\binom{n}{k-1}$ or more valid sets. This means that, *counting multiplicities*, there are at least $\frac{1}{6}\binom{n}{k-1} \cdot \frac{n}{45k^2}$ strings with at most $7k-1$ ones each. Since each such string contains at most $\binom{7k-1}{k-1} < \binom{7k}{k-1}$ subsets of size $k-1$, each such string is generated fewer than $\binom{7k}{k-1}$ times by different valid sets. Thus the number of *different* strings of positive probability exceeds $\frac{1}{6}\binom{n}{k-1}\frac{n}{45k^2}/\binom{7k}{k-1}$.

Proving that this last quantity is at least $c \cdot \frac{m(n,k)}{k^3 2^{7k}}$ for a universal constant $c > 0$ if $k \leq n/2$ is not difficult, if one uses facts such as $(k+1)\binom{n}{k} \geq m(n,k)$ if $k \leq n/2$, $\binom{n}{k-1}n \geq \binom{n}{k}$, and $\binom{7k}{k-1} \leq 2^{7k}$. ■

Now the reader can verify that the two conditions above on M and M' are satisfied if $M' = (k/n)^k$ and $M = (k/n)^{k-1}$, using inequalities such as $\binom{a}{b} \geq (a/b)^b$ if $a \geq b \geq 1$ and $[k/(k-1)]^{k-1} \geq 2$ for all $k \geq 2$. Theorem 3 follows.

4 Lower Bounds for Large Independence

In this section we show how to derive a lower bound for the case in which the degree of independence is very large. The main tool in proving our result is the following theorem, in which $\mathcal{C}_d$ denotes the set of all d-dimensional subcubes of the n-dimensional cube $\mathcal{C} = \{1, -1\}^n$.

Theorem 12. Let $n \geq 1$. Let g be a 2^n-dimensional probability vector viewed as a function on the vertices of the n-dimensional cube $\mathcal{C}$. Let $1 \leq d \leq n$. Let $N = |\{v \in \mathcal{C} : g(v) > 0\}|$. Suppose that there is an α such that for all $C \in \mathcal{C}_d$, $\alpha = \sum_{v \in C} g(v)$. Then $N \geq 2^{n-1}(n+2-2d)/(n+1-d)$.

The proof of the above theorem uses the Fourier analysis techniques of [Fri92]. Before giving the proof, we derive an interesting consequence.

Corollary 13. Let $\mathcal{H} = (\Omega, 2^\Omega, P)$ be a k-wise independent 0-1 probability space where all $p_i = 1/2$. Then $|\Omega| \geq 2^n(\frac{2k+2-n}{2k+2}) = 2^n(1 - \frac{n}{2k+2})$.

Proof. For all $(n-k)$-dimensional subcubes C, $\sum_{v \in H} g(v) = 2^{-k}$. So let $d = n-k$, $\alpha = 2^{-k}$. Then $N \geq 2^{n-1}(n+2-2(n-k))/(n+1-(n-k)) = 2^n(2k-n+2)/(2+2k)$. ■

Corollary 14. Let $0 < \alpha \leq 1$ and $\lfloor \alpha n \rfloor \leq k \leq n$, and suppose that $\mathcal{H} = (\Omega, 2^\Omega, P)$ is a k-wise independent 0-1 probability space where all $p_i = 1/2$. Then $|\Omega| \geq (1 - 1/(2\alpha))\, 2^n$.

Proof. Because $k \geq \lfloor \alpha n \rfloor$, $1 - \frac{n}{2(k+1)} \geq 1 - 1/(2\alpha)$. By Corollary 13, we have $|\Omega| \geq 2^n(1 - n/(2k+2)) \geq 2^n(1 - 1/(2\alpha))$. ■

Corollary 14 implies Theorem 5.

Now we do the proof of Theorem 12. We start with the following simple technical lemma, which follows from the convexity of the function $f(x) = x^2$.

Lemma 15. Let g be a probability vector. Suppose that $N = |\{i : g(i) > 0\}|$. Then $|g|^2 = \sum_v g^2(v) \geq 1/N$.

We use A to denote the $2^n \times 2^n$ adjacency matrix of $\mathcal{C}$, i.e., the (u, v) entry $A(u, v)$ of A is 1 if u and v differ in exactly one position, and 0 otherwise. For each $T \subseteq \{1, 2, ..., n\}$, let χ_T be a vector, viewed as a function on $\mathcal{C}$, whose vth coordinate $\chi_T(v) = \prod_{i \in T} v_i$, where $v \in \{+1, -1\}^n$. If $f, h : \mathcal{C} \to \mathbb{R}$, let $< f, h >= 2^{-n} \sum_{v \in \mathcal{C}} f(v)h(v)$.

Lemma 16. Let $0 \leq r \leq n$. Then for each $T \subseteq \{1, 2, ..., n\}$ of size r, χ_T is an eigenvector of A, corresponding to eigenvalue $n - 2r$.

Proof. Let $T \subseteq \{1, 2, ..., n\}$ with $|T| = r$. Choose any $u \in \mathcal{C}$. The uth entry of $A\chi_T$ is

$$\sum_{v \in \mathcal{C}} A(u, v)\chi_T(v) = \sum_{v \in N(u)} \chi_T(v),$$

where $N(u)$ denotes those $v \in \mathcal{C}$ differing from u in exactly one position. Of the n v's in $N(u)$, $n-r$ v's differ from u in a bit position not in T. These v's satisfy $\chi_T(v) = \chi_T(u)$. The remaining r v's differ from u in one position in T and satisfy $\chi_T(v) = -\chi_T(u)$. It follows that

$$\begin{aligned} \textstyle\sum_{v \in N(u)} \chi_T(v) &= (n-r)\chi_T(u) - r\chi_T(u) \\ &= (n-2r)\chi_T(u). \end{aligned}$$

Since u was arbitrary, we have $A\chi_T = (n-2r)\chi_T$. ■

It is easy to see that if $T \neq U$, then $< \chi_T, \chi_U >= 0$, and that $< \chi_T, \chi_T >= 1$ for all T. This means that the χ_T's are orthogonal and hence $\{\chi_T : T \subseteq \{1, 2, ..., n\}\}$ is linearly independent. Thus any 2^n-dimensional vector g can be written as a linear combination $g = \sum_T \beta_T \chi_T$ for some reals β_T.

Proof of Theorem 12. As the theorem is vacuous if $d = n$, we assume without loss of generality that $d < n$. Fixing g, let us set $\hat{g}(U) =< g, \chi_U >$ for any U. It is easy to see, from the orthogonality of the χ_T's, that $g(v) = \sum_T \hat{g}(T)\chi_T(v)$.

We are given that $\sum_{v \in C} g(v) = \alpha$ for all $C \in \mathcal{C}_d$. Choose any $U \subseteq \{1, 2, ..., n\}$ such that $1 \leq |U| \leq n - d$.

$$\begin{aligned} \hat{g}(U) &= < g, \chi_U > \\ &= 2^{-n} \textstyle\sum_v g(v)\chi_U(v). \end{aligned}$$

Let us use "$v|_U = s$" to mean that the restriction of v to the bit positions in U is s.

$$\begin{aligned} 2^n \hat{g}(U) &= \textstyle\sum_v g(v)\chi_U(v) \\ &= \textstyle\sum_{s \in \{1,-1\}^{|U|}} \sum_{v:v|_U = s} g(v)\chi_U(v) \\ &= \textstyle\sum_{s \in \{1,-1\}^{|U|}} (\prod_{i=1}^{|U|} s_i)(\sum_{v:v|_U = s} g(v)). \end{aligned}$$

Now $n - |U| \geq d$ and $\{v : v|_U = s\}$ is a cube of dimension $n - |U|$. This $(n - |U|)$-dimensional cube can be partitioned into $2^{n-|U|}/2^d$ d-dimensional subcubes. As the sum of $g(v)$ over any d-dimensional subcube is α,

$$\sum_{v:v|_U = s} g(v) = \frac{2^{n-|U|}}{2^d}\alpha.$$

Thus

$$2^n \hat{g}(U) = \sum_{s \in \{1,-1\}^{|U|}} \left(\prod_{i=1}^{|U|} s_i\right)\left(\sum_{v:v|_U = s} g(v)\right) = \left(\frac{2^{n-|U|}}{2^d}\alpha\right) \sum_{s \in \{1,-1\}^{|U|}} \prod_{i=1}^{|U|} s_i.$$

Now

$$\sum_{s \in \{1,-1\}^{|U|}} \prod_{i=1}^{|U|} s_i = [1 + (-1)]^{|U|} = 0,$$

since $|U| \geq 1$. Therefore $\hat{g}(U) = 0$ for any U such that $1 \leq |U| \leq n - d$.

Let us define $h(v) = (Ag)_v = \sum_{u \in N(v)} g(u)$. Now

$$h(v) = (Ag)_v = \left(A\left[\sum_T \hat{g}(T)\chi_T\right]\right)_v = \sum_T \hat{g}(T)(A\chi_T)_v = \sum_T (n - 2|T|)\hat{g}(T)\chi_T(v),$$

where the third equality follows from the linearity of A and the last equality follows from Lemma 16. Above we showed that for any T with $1 \leq |T| \leq n - d$, $\hat{g}(T) = 0$. However, $\hat{g}(\emptyset) = 1/2^n$, since g is a probability vector.

We would like to derive an upper bound on $|g|^2$. To do this we consider $< h, g >$. Clearly, $< h, g > \geq 0$. On the other hand,

$$< h, g >= 2^{-n} \sum_v g(v)h(v) = 2^{-n} \sum_v \sum_U \sum_T [(n - 2|T|)\hat{g}(T)\chi_T(v)] [\hat{g}(U)\chi_U(v)]$$

$$= 2^{-n} \sum_T \sum_U \hat{g}(U)(n - 2|T|)\hat{g}(T) \sum_v \chi_T(v)\chi_U(v).$$

The last summation equals $2^n < \chi_T, \chi_U >$, which is 0 if $U \neq T$ and 2^n if $U = T$. Therefore

$$< h, g >= \sum_T (n - 2|T|)\hat{g}^2(T) = \frac{n}{4^n} + \sum_{T:|T| \geq n-d+1} (n - 2|T|)\hat{g}^2(T).$$

We can bound $< h, g >$ as follows.

$$0 \leq \ < h, g > \ \leq \frac{n}{4^n} + (2d - 2 - n) \sum_{T:|T| \geq n-d+1} \hat{g}^2(T).$$

By Parseval's identity, $||g||^2 = \sum_T \hat{g}^2(T)$, where $||g||^2 =< g, g >$. Hence

$$< g, g >= 2^{-n}|g|^2 = \frac{1}{4^n} + \sum_{T:|T| \geq n-d+1} \hat{g}^2(T) \ .$$

Using this identity we have

$$0 \leq \frac{n}{4^n} + (2d - 2 - n) \left(\frac{|g|^2}{2^n} - \frac{1}{4^n} \right) = \frac{2n + 2 - 2d}{4^n} + \frac{2d - 2 - n}{2^n}|g|^2.$$

Therefore

$$\frac{1}{|g|^2} \geq \frac{2^n}{2} \left(\frac{n + 2 - 2d}{n + 1 - d} \right).$$

By Lemma 15, $|g|^2 \geq 1/N$. It follows that

$$N \geq \frac{1}{|g|^2} \geq \frac{2^n}{2} \left(\frac{n + 2 - 2d}{n + 1 - d} \right). \blacksquare$$

5 General Construction

In this section we give a general construction for k-wise independent 0-1 random variables. The construction receives as input the desired $p_1, p_2, ..., p_n$.

We first need a simple lemma, which is an easy consequence of the Moebius inversion formula.

Lemma 17. Let $\mathcal{H} = (\Omega, 2^\Omega, P)$ be a 0-1 probability space. Define Y_i to be the ith bit of an n-vector chosen according to $\mathcal{H}$. Let $1 \leq k \leq n$. Suppose that there are reals $p_1, p_2, ..., p_n$ such that for all subsets S of $\{1, 2, ..., n\}$ of size at most k, $P[\cap_{i \in S}[Y_i = 1]] = \prod_{i \in S} p_i$. Then $\mathcal{H}$ is a k-wise independent 0-1 probability space for $p_1, p_2, ..., p_n$.

We omit the simple proof of the following corollary.

Corollary 18. Suppose $\mathcal{H}$ is a k-wise independent 0-1 probability space for $p_1, p_2, ..., p_n$. Suppose that $p'_i = 1 - p_i$ if $i \in S$, $p'_i = p_i$ otherwise. A k-wise independent 0-1 probability space for $p'_1, p'_2, ..., p'_n$ can be obtained from $\mathcal{H}$ by interchanging zeroes and ones in every position $i \in S$ in each point of the probability space.

Our general construction for $p_1, p_2, ..., p_n$ consists of three parts: (1) a construction for probabilities t_i very close to 1, (2) a construction for $s_1, s_2, ..., s_n$ such that each s_i either equals p_i or is slightly larger, and (3) a method of combining the two probability spaces so as to get $p_1, ..., p_n$ exactly.

We must give a construction for very large probabilities. Since it is easier to understand what's going on when the probabilities are very small, we will assume $t_i \leq 1/(2n)$ for all i, and then invoke Corollary 18 to get the space for $t_i \geq 1 - 1/(2n)$.

Lemma 19. Let $\epsilon = 1/(2n)$. Suppose that $0 \leq t_1, t_2, ..., t_n \leq \epsilon$. There is a k-wise independent 0-1 probability space $\mathcal{H}$ for $t_1, t_2, ..., t_n$ whose points are the elements of $\{0,1\}^n$ with at most k ones. Its size is at most $m(n,k)$.

Proof. We assign a probability to each 0-1 vector as follows. Identify $B \subseteq \{1, 2, ..., n\}$ with its characteristic n-vector. We use $P[B]$ to denote the probability that $\mathcal{H}$ assigns to the characteristic vector of B. Set $P[B] = 0$ if $|B| > k$. We will create probabilities $P[B]$ for each $B \subseteq \{1, 2, ..., n\}$, $|B| \leq k$, such that for all $A \subseteq \{1, 2, ..., n\}$ such that $|A| \leq k$, $\sum_{B:B \supseteq A} P[B] = q_A = \prod_{i \in A} t_i$. In $\mathcal{H}$, where Y_i denotes the random variable which is the ith bit of the random n-vector, we have $P[\cap_{i \in A}[Y_i = 1]] = \sum_{B:B \supseteq A} P[B] = q_A$. By Lemma 17, this is enough.

We assign probabilities $P[B]$ by downward induction on $|B|$. As each $P[B]$ is defined, we prove inductively that $0 \leq P[B] \leq q_B$. If $|B| = k$, set $P[B] = q_B$. (Notice that $0 \leq P[B] \leq q_B$.)

Now suppose that $|B| = l$, $l < k$, and $P[C]$ has been defined for all C of size exceeding l. Set

$$P[B] = q_B - \sum_{\substack{C:C \supseteq B \\ \neq}} P[C] = q_B - \sum_{\substack{C:C \supseteq B, |C| \leq k \\ \neq}} P[C].$$

It is obvious with this definition that $q_B = \sum_{C:C \supseteq B} P[C]$. Now if $0 \leq P[C] \leq q_C$ for all C of size exceeding l, then clearly $P[B] \leq q_B$. Also,

$$\begin{aligned} P[B] &\geq q_B - \textstyle\sum_{C:C \supseteq B, |C| \leq k} q_C \\ &\geq q_B - q_B \textstyle\sum_{i=1}^{k-l} \binom{n-l}{i} \epsilon^i \\ &\geq q_B [1 - \textstyle\sum_{i=1}^{\infty} (n\epsilon)^i] \\ &= q_B [1 - \textstyle\sum_{i=1}^{\infty} (\frac{1}{2})^i] \\ &= 0. \end{aligned}$$

That $m(n,k)$ is the number of 0-1 vectors with at most k ones is obvious. ■

Now we give a construction and lemma due to Joffe [Joffe74]. Let $r \geq n$ be a prime. Suppose that $s_1, s_2, ..., s_n \in [0,1]$ satisfy $s_i = j_i/r$ for all i, for some integers j_i. Define a probability space, with (at most) n^k points, as follows. Choose $a_0, a_1, ..., a_{k-1}$ uniformly and independently at random from $\mathbb{Z}_r$, and let $X_i = a_0 + a_1 i + a_2 i^2 + \cdots + a_{k-1} i^{k-1} \bmod r$, for $1 \leq i \leq n$. For each point $< a_0, a_1, ..., a_{k-1} >$ in the sample space, let $Y_i = 1$ if $X_i \in \{0, 1, ..., j_i - 1\}$ and let $Y_i = 0$ otherwise, $1 \leq i \leq n$. Associate with $< a_0, a_1, ..., a_{k-1} >$ the vector $< Y_1, Y_2, ..., Y_n >$ and give each of the r^k r-vectors probability $1/r^k$ (and then combine identical r-vectors, if necessary).

Lemma 20. The random variables $X_1, X_2, ..., X_n$ are k-wise independent and each X_i is uniform on $\mathbb{Z}_r$. It follows that the 0-1 probability space defined by the vectors $< Y_1, Y_2, ..., Y_n >$ is a k-wise independent 0-1 probability space for $s_1, s_2, ..., s_n$.

Lemma 20 follows easily from the fact that Van der Monde matrices are invertible over $\mathbb{Z}_r$.

We now present an almost trivial technique that takes a k-wise independent 0-1 probability space A for $s_1, s_2, ..., s_n$, and a k-wise independent 0-1 probability space A' for $t_1, t_2, ..., t_n$, and produces a k-wise independent 0-1 probability space $\mathcal{H}$ for $s_1t_1, s_2t_2, ..., s_nt_n$. To generate an element of $\mathcal{H}$, simply independently pick x from A and x' from A', and output the bit-wise AND $x \wedge x' \in \{0,1\}^n$. Clearly the size of the implied space $\mathcal{H}$ is at most $|A||A'|$.

Lemma 21. The above construction of $\mathcal{H}$ from A and A' guarantees that $\mathcal{H}$ is a k-wise independent 0-1 probability space for $s_1t_1, s_2t_2, ..., s_nt_n$.

Proof. From Lemma 17, it is sufficient to show that for any subset $S \subseteq \{1, 2, ..., n\}$ of size at most k, $P[\cap_{i \in S}[Y_i = 1]] = \prod_{i \in S}(s_it_i)$, where Y_i is the ith bit of the random n-vector $x \wedge x'$. But the ith bit of $x \wedge x'$ is 1 if and only if the ith bit of x and the ith bit of x' are both one. By independence, this occurs with probability $(\prod_{i \in S} s_i)(\prod_{i \in S} t_i)$. ■

The probability space for the given probabilities $p_1, p_2, ..., p_n$ is obtained routinely as follows.

1. Assume without loss of generality that every $p_i \geq 1/2$.
2. Build a probability space in which the expectation s_i of the ith variable satisfies $s_i \in [p_i, p_i + 1/(4n)]$. This is done via Lemma 20 with r being a prime in $[4n, 8n]$.
3. Build a space for $t_1, t_2, ..., t_n$, where $t_i = p_i/s_i$ and hence $t_i \in [1 - 1/(2n), 1]$. We do this via Lemma 19.
4. Use Lemma 21 to form a space for $p_1, ..., p_n$ where $p_i = t_is_i$.

We conclude with

Theorem 22. The new space is a k-wise independent 0-1 probability space for the given probabilities $p_1, p_2, ..., p_n$. Its size is at most $(8n)^k m(n,k) \leq (8n)^{2k}$.

6 Construction for Pairwise Independence

In this section we exhibit a space of size at most $4n^2$ for pairwise independent random variables in the case that all the probabilities are equal. From the lower bound of Theorem 11 it follows that this construction is optimal in size, to within a constant factor.

Definition. Let $\mathcal{H}$ be a 0-1 probability space. $\mathcal{H}$ is an $(a, b; n)$-*space* if the probability of a 1 in position i is a for all i, and the probability of simultaneous ones in positions i and j is b, for all $i < j$.

For a pairwise independent 0-1 probability space, we would want $b = a^2$. However, here we will need $b < a^2$.

The following lemma shows that given such an $(a, b; n)$-space, we can build a *pairwise independent* 0-1 probability space if all $p_i = p$, and if p is in a certain interval:

Lemma 23. Let $\mathcal{H} = (\Omega, 2^\Omega, P)$ be an $(a, b; n)$-space with $b < a^2$. Then there is a construction of a pairwise independent 0-1 probability space $\mathcal{H}$ if all $p_i = p$, on a sample space of size at most $|\Omega| + 2$, provided that $p \in [b/a, (a-b)/(1-a)]$.

Proof. The new probability space will be defined in the following way. Let $x, y \geq 0$ such that $x + y \leq 1$. With probability x, choose a point in Ω (according to the probabilities in P). With probability y, choose the all-ones row. With probability $1 - x - y$, choose the all-zeroes row. It is clear that the implied probability space has size at most $|\Omega| + 2$.

In order that the new space be pairwise independent for all $p_i = p$, it suffices to ensure that the probability of a 1 in position i is p and the probability of simultaneous ones in positions i and j is p^2 (and that the sum of the probabilities of the points is 1). In other words, it is sufficient to satisfy the following system:

$$\begin{aligned} x \cdot a + y \cdot 1 &= p \\ x \cdot b + y \cdot 1 &= p^2 \end{aligned}$$

and satisfy $x + y \leq 1$, $x, y \geq 0$. The unique solution to the system is

$$\begin{aligned} x &= (p - p^2)/(a - b), \\ y &= p(ap - b)/(a - b). \end{aligned}$$

We have $x + y \leq 1$, $x, y \geq 0$ if and only if $p \in [b/a, (a-b)/(1-a)]$. ■

Note that since $b < a^2$, $b/a < (a-b)/(1-a)$.

Based on the above lemma we can exhibit a quadratic size space for an arbitrary probability p.

Theorem 24. For any $n > 1$, for any $0 \leq p \leq 1$, there is a construction of a pairwise independent 0-1 probability space of size at most $4n^2$ if every $p_i = p$.

Proof. First we claim that if $q \geq n$ is a prime, then for any l, $0 \leq l \leq q$, there is an $(l/q, (l^2 - l)/(q^2 - q); n)$-space. The points of the sample space are the $q^2 - q$ pairs (c, d), $c, d \in \mathbb{Z}_q$, $d \neq 0$. Let $X_i(c, d) = c + di \bmod q$ and $Y_i(c, d) = 1$ if $X_i(c, d) \in \{0, 1, 2, ..., l - 1\}$ and $Y_i(c, d) = 0$ otherwise, $1 \leq i \leq n$. Build a space with at most $q^2 - q$ vectors by associating with point (c, d) the n-vector $(Y_1, Y_2, ..., Y_n)$, and then combining identical n-vectors.

That this is an $(l/q, (l^2 - l)/(q^2 - q); n)$-space can easily be verified by noting that the q omitted pairs $(c, 0)$ define l rows that are all 1 and $q - l$ rows that are all 0, and if the q omitted rows were added, we'd have a pairwise independent space in which each random variable is 1 with probability l/q.

Now for the construction of the pairwise independent probability space. Choose a prime q such that $n \leq q \leq 2n$. Choose an l, $1 \leq l \leq q - 1$, such that $p \in [(l-1)/(q-1), l/(q-1)]$. The construction of this proof gives us an $(l/q, (l^2 - l)/(q^2 - q); n)$-space of size at most $q^2 - q$. Lemma 23 now gives us a pairwise independent sample space, of size at most $(q^2 - q) + 2 \leq q^2 \leq (2n)^2 = 4n^2$, with $p_i = p$ for all i, provided that $p \in [b/a, (a-b)/(1-a)]$. A simple calculation with $a = l/q$, $b = (l^2 - l)/(q^2 - q)$ shows that $[b/a, (a-b)/(1-a)] = [(l-1)/(q-1), l/(q-1)]$, which contains p. ■

For certain ranges of p, better constructions are possible. Theorem 25 uses Lemma 23 to get a linear construction for $p \leq 1/(n-1)$, while Corollary 28 exhibits a linear construction for values in the neighborhood of 1/2.

Theorem 25. For $p \leq 1/(n-1)$ there is a pairwise independent 0-1 probability space of size at most $n + 2$ with $p_i = p$ for all i.

Proof. We exhibit a $(1/n, 0; n)$-space; this is enough for all $p \in [b/a, (a-b)/(1-a)] = [0, (1/n)/(1-1/n)] = [0, 1/(n-1)]$. This is simply the identity matrix in which each row gets probability $1/n$. By Lemma 23 there is a construction of size at most $n+2$. ■

We can also get linear size constructions around $1/2$.

Lemma 26. Let A be a normalized $m \times m$ Hadamard matrix (i.e., an $m \times m$ $(+1, -1)$ matrix satisfying $A^T A = mI$, whose first row and first column are all $+1$'s). Let A' be obtained from A by omitting the first row and first column and then replacing all -1's by 0's. The probability space obtained by giving each of the $m-1$ rows of A' probability $1/(m-1)$ is a $((1/2)(m-2)/(m-1), (1/4)(m-4)/(m-1); m-1)$-space.

The proof is simple and is omitted.

Theorem 27. Let $n > 1$ and let A be a normalized $m \times m$ Hadamard matrix with $m \geq n+1$. There is a pairwise independent 0-1 probability space (whose points are elements of $\{0,1\}^n$) of size at most $m+1$ with $p_i = p$ for all i if $p \in [1/2 - 1/(m-2), 1/2 + 1/(m-2)]$.

Proof. With $a = (1/2)(m-2)/(m-1)$, $b = (1/4)(m-4)/(m-1)$ (from Lemma 26), we have $b/a = 1/2 - 1/(m-2)$ and $(a-b)/(1-a) = 1/2$. This gives a construction for $p \in [1/2 - 1/(m-2), 1/2]$. Its size is at most $(m-1)+2 = m+1$. For $p \in [1/2, 1/2 + 1/(m-2)]$, we first build a space for $1-p$ and then exchange zeroes and ones. ■

Corollary 28. For any $n > 1$, for any $p \in [1/2 - 1/(2n-2), 1/2 + 1/(2n-2)]$, there is a pairwise independent 0-1 probability space of size at most $2n+1$ with all $p_i = p$.

Proof. For any t, there is a $2^t \times 2^t$ Hadamard matrix. Given n, take m to be the least power of two exceeding n; $m \leq 2n$. By Theorem 27, there is a construction of size at most $m+1 \leq 2n+1$ for any $p \in [1/2 - 1/(m-2), 1/2 + 1/(m-2)] \supseteq [1/2 - 1/(2n-2), 1/2 + 1/(2n-2)]$. ■

7 Acknowledgments

For their help in the early stages of our work, we are very grateful to D. Koller and N. Megiddo, whose paper [KM93], in opening up the area of nonuniform spaces in theoretical computer science, inspired this one. We thank Shi-Chun Tsai for pointing out a typo.

References

[ABI86] 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), pages 567–583.

[AS92] N. Alon and J. Spencer, *The Probabilistic Method,* Wiley, 1992.

[BR89] B. Berger and J. Rompel, "Simulating $(log^c n)$-wise Independence in NC," *Proc.* 30^{th} *IEEE Symposium on Foundations of Computer Science*, 1989, pages 2–7.

[CGHFRS85] B. Chor, O. Goldreich, J. Håstad, J. Friedman, S. Rudich, and R. Smolensky, "The Bit Extraction Problem or t-Resilient Functions," *Proc.* 26^{th} *IEEE Symposium on Foundations of Computer Science*, 1985, pages 396–407.

[Fri92] J. Friedman, "On the Bit Extraction Problem," *Proc. 33rd IEEE Symposium on Foundations of Computer Science*, 1992, pages 314–319.

[Joffe74] A. Joffe, "On a Set of Almost Deterministic k-Independent Random Variables," *Annals of Probability* 2 (1974), pages 161–162.

[KM93] D. Koller and N. Megiddo, Constructing Small Sample Spaces Satisfying Given Constraints," *Proc.* 25^{th} *ACM Symposium on Theory of Computing*, 1993, pages 268–277.

[Lub86] M. Luby, "A Simple Parallel Algorithm for the Maximal Independent Set Problem," *SIAM Journal on Computing* 15 (1986), pages 1036–1053.

[Lub88] M. Luby, "Removing Randomness in Parallel Computation Without a Processor Penalty," *Proc.* 29^{th} *IEEE Symposium on Foundations of Computer Science*, 1988, pages 162-173.

[MNN89] R. Motwani, J. Naor, and J. Naor, "The Probabilistic Method Yields Deterministic Parallel Algorithms," *Proc.* 30^{th} *IEEE Symposium on Foundations of Computer Science*, 1989, pages 8–13.

[Sch92] L. Schulman, "Sample Spaces Uniform on Neighborhoods," *Proc.* 24^{th} *ACM Symposium on Theory of Computing*, 1992, pages 17–25.