

Robust Characterizations of k -wise Independence over Product Spaces and Related Testing Results*

Ronitt Rubinfeld[†]
MIT and Tel Aviv University
ronitt@csail.mit.edu

Ning Xie[‡]
MIT
ningxie@csail.mit.edu

Abstract

A discrete distribution D over $\Sigma_1 \times \cdots \times \Sigma_n$ is called *(non-uniform) k -wise independent* if for any subset of k indices $\{i_1, \dots, i_k\}$ and for any $z_1 \in \Sigma_{i_1}, \dots, z_k \in \Sigma_{i_k}$, $\Pr_{\mathbf{X} \sim D}[X_{i_1} \cdots X_{i_k} = z_1 \cdots z_k] = \Pr_{\mathbf{X} \sim D}[X_{i_1} = z_1] \cdots \Pr_{\mathbf{X} \sim D}[X_{i_k} = z_k]$. We study the problem of testing (non-uniform) k -wise independent distributions over product spaces. For the uniform case we show an upper bound on the distance between a distribution D from k -wise independent distributions in terms of the sum of Fourier coefficients of D at vectors of weight at most k . Such a bound was previously known only when the underlying domain is $\{0, 1\}^n$. For the non-uniform case, we give a new characterization of distributions being k -wise independent and further show that such a characterization is robust based on our results for the uniform case. These results greatly generalize those of Alon et al. [STOC'07, pp. 496–505] on uniform k -wise independence over the Boolean cubes to non-uniform k -wise independence over product spaces. Our results yield natural testing algorithms for k -wise independence with time and sample complexity sublinear in terms of the support size of the distribution when k is a constant. The main technical tools employed include discrete Fourier transform and the theory of linear systems of congruences.

*A preliminary version of this work, titled “Testing Non-uniform k -wise Independent Distributions over Product Spaces”, appeared in the Proceedings of ICALP 2010.

[†]Research supported by NSF grants 0514771, 0728645, 0732334, Marie-Curie International Reintegration Grant PIRG03-GA-2008-231077 and Israel Science Foundation Grant Numbers 1147/09 and 1675/09.

[‡]Research supported in part by an Akamai Presidential Fellowship and NSF grants 0514771, 0728645 and 0732334.

1 Introduction

Nowadays we are both blessed and cursed by the colossal amount of data available for processing. In many situations, simply scanning the whole data set once can be a daunting task. It is then natural to ask what we can do in *sublinear time*. For many computational questions (for example, determining if a given graph is 3-colorable or bipartite, determining if an input Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ is a low-degree polynomial, determining if a given list of integers is in monotone order, etc), if instead of asking the decision version of the problems, one can relax the questions and consider the analogous property testing problems, then sublinear algorithms are often possible. See the survey articles [18, 35, 27, 13].

Property testing algorithms [36, 19] are usually based on *robust characterizations* of the objects being tested. For instance, the linearity test introduced in [11] is based on the characterization that a function is linear if and only if the linearity test (for uniformly and randomly chosen x and y , check if $f(x) + f(y) = f(x + y)$) has acceptance probability 1. Moreover, the characterization is *robust* in the sense that if the linearity test accepts a function with probability close to 1, then the function must be also close to some linear function. Property testing often leads to a new understanding of well-studied problems and sheds insight on related problems.

In this work, we show robust characterizations of k -wise independent distributions over discrete product spaces and give sublinear-time testing algorithms based on these robust characterizations. Note that distributions over product spaces are in general not *product distributions*, which by definition are n -wise independent distributions (see below for definition).

The k -wise Independent Distributions. For a finite set Σ , a discrete probability distribution D over Σ^n is (non-uniform) *k -wise independent* if for any set of k indices $\{i_1, \dots, i_k\}$ and for all $z_1, \dots, z_k \in \Sigma$, $\Pr_{\mathbf{X} \sim D}[X_{i_1} \cdots X_{i_k} = z_1 \cdots z_k] = \Pr_{\mathbf{X} \sim D}[X_{i_1} = z_1] \cdots \Pr_{\mathbf{X} \sim D}[X_{i_k} = z_k]$. That is, restricting D to any k coordinates gives rise to a fully independent distribution. For the special case when $\Pr_{\mathbf{X} \sim D}[X_i = z] = \frac{1}{|\Sigma|}$ for every index i and every letter z in the alphabet, we refer to the distribution as *uniform k -wise independent*¹. A distribution is *almost k -wise independent* if its restriction to any k coordinates is very close to some independent distribution. k -wise independent distributions look independent “locally” to any observer of only k coordinates, even though they may be far from the fully independent distributions “globally”. Furthermore, k -wise independent distributions can be constructed with exponentially smaller support sizes than fully independent distributions. Because of these useful properties, k -wise independent distributions have many applications in both probability theory and computational complexity theory [23, 25, 28, 31].

Given samples drawn from a distribution, it is natural to ask, how many samples are required to tell whether the distribution is k -wise independent or far from k -wise independent, where by “far from k -wise independent” we mean that the distribution has a large statistical distance from *any* k -wise independent distribution. Usually the time and query complexity of distribution testing algorithms are measured against the domain sizes of the distributions. For example, algorithms that test distributions over $\{0, 1\}^n$ with time complexity $o(2^n)$ are said to be sublinear-time testing algorithms.

Alon, Goldreich and Mansour [4] implicitly gave the first robust characterization of k -wise independence. Alon et al. [1] improved the bounds in [4] and also gave efficient testing algorithms. All of these results consider only uniform distributions over $\text{GF}(2)$. Our work generalizes previous results in two ways: to distributions over arbitrary finite product spaces and to non-uniform k -wise independent distributions.

¹In literature the term “ k -wise independence” usually refers to uniform k -wise independence.

Generic Algorithm for Testing Uniform k -wise Independence

1. Sample D independently M times
2. Use these samples to estimate all the Fourier coefficients of weight at most k
3. **Accept** if the magnitudes of *all* the estimated Fourier coefficients are at most δ

Figure 1: A *Generic Algorithm* for testing uniform k -wise independence.

1.1 Our Results

Let $\Sigma = \{0, 1, \dots, q - 1\}$ be the alphabet² and let $D : \Sigma^n \rightarrow [0, 1]$ be the distribution to be tested. For any vector $\mathbf{a} \in \Sigma^n$, the Fourier coefficient of distribution D at \mathbf{a} is $\hat{D}(\mathbf{a}) = \sum_{\mathbf{x} \in \Sigma^n} D(\mathbf{x}) e^{\frac{2\pi i}{q} \sum_{j=1}^n a_j x_j} = \mathbf{E}_{\mathbf{X} \sim D} \left[e^{\frac{2\pi i}{q} \sum_{j=1}^n a_j X_j} \right]$. The *weight* of \mathbf{a} is the number of non-zero entries in \mathbf{a} . It is a folklore fact that a distribution D is uniform k -wise independent if and only if for all non-zero vectors \mathbf{a} of weight at most k , $\hat{D}(\mathbf{a}) = 0$. A natural test for k -wise independence is thus the *Generic Algorithm* described in Fig. 1.

However, in order to prove that the *Generic Algorithm* works, one needs to show that the simple characterization of k -wise independence is *robust*. Here, *robustness* means that for any distribution D if all its Fourier coefficients at vectors of weight at most k are at most δ (in magnitude), then D is $\epsilon(\delta)$ -close to some uniform k -wise independent distribution, where the closeness parameter ϵ is in general a function of the error parameter δ , domain size and k . Consequently, the query and time complexity of the *Generic Algorithm* will depend on the underlying distance upper bound between D and k -wise independence.

Our first main result is the following robust characterization of uniform k -wise independence.

Theorem 1.1 (First Main Theorem). *Let $\Sigma = \{0, 1, \dots, q - 1\}$ and D be a distribution over Σ^n . Let $\Delta(D, \mathcal{D}_{\text{kwi}})$ denote the distance between D and the set of (uniform) k -wise independent distributions over Σ^n , then*

$$\Delta(D, \mathcal{D}_{\text{kwi}}) \leq \sum_{0 < \text{wt}(\mathbf{a}) \leq k} \left| \hat{D}(\mathbf{a}) \right|.$$

As it turns out, the sample complexity of our testing algorithm is $\tilde{O}\left(\frac{n^{2k}(q-1)^{2k}q^2}{\epsilon^2}\right)$ and the time complexity is $\tilde{O}\left(\frac{n^{3k}(q-1)^{3k}q^2}{\epsilon^2}\right)$, which are both sublinear when $k = O(1)$ and $q \leq \text{poly}(n)$. We further generalize this result to uniform k -wise independent distributions over product spaces, i.e., distributions over $\Sigma_1 \times \dots \times \Sigma_n$, where $\Sigma_1, \dots, \Sigma_n$ are (different) finite sets.

Our second main result is a robust characterization of *non-uniform* k -wise independent distributions over Σ^n .

Theorem 1.2 (Second Main Theorem). *Let $\Sigma = \{0, 1, \dots, q - 1\}$ and D be a distribution over Σ^n , then*

$$\Delta(D, \mathcal{D}_{\text{kwi}}) \leq \text{poly}(n, q) \max_{\mathbf{a}: 0 < \text{wt}(\mathbf{a}) \leq k} \left| \hat{D}^{\text{non}}(\mathbf{a}) \right|,$$

where the exponent in $\text{poly}(n, q)$ is a function of k only and $\{\hat{D}^{\text{non}}(\mathbf{a})\}_{\mathbf{a} \in \Sigma^n}$ are a set of non-uniform Fourier coefficients to be defined later (see Section 5.1 for details).

²This is without loss of generality, since we are not assuming any field or ring structure of the underlying alphabet of the distribution. All the properties about distributions considered in this paper are invariant under permutations of the symbols in the alphabet.

As we show in Sections 5.4 and 5.5, if all the marginal probabilities $\Pr_{\mathbf{X} \sim D}[X_i = z]$, $1 \leq i \leq n$ and $z \in \Sigma$, are bounded away from both zero and one, then Theorem 1.2 also implies a testing algorithm for non-uniform k -wise independence whose sample and time complexity are polynomial in n and q when k is a constant.

We remark that our result on non-uniform k -wise independent distributions also generalizes to distributions over product spaces.

To the best of our knowledge, there is no lower bound result for testing k -wise independence over general domains except [1] which is for the binary field case. It will be interesting to get good lower bounds for general domains as well.

Another related problem, namely testing *almost k -wise independence* over product spaces (see Section 6 for relevant definitions), admits a straightforward generalization of the testing algorithm given in [1], which was only proved there for the (uniform) binary case. We include these results in Section 6.

Our results add a new understanding of the structures underlying (non-uniform) k -wise independent distributions and it is hoped that one may find other applications of these robust characterizations.

As is often the case, commutative rings demonstrate different algebraic structures from those of prime fields. For example, the recent improved construction [16] of 3-query locally decodable codes of Yekhanin [42] relies crucially on the construction of set systems of superpolynomial size [21] such that the sizes of each set as well as all the pairwise intersections satisfy certain congruence relations modulo composite numbers (there is a polynomial upper bound when the moduli are primes). Generalizing results in the binary field (or prime fields) to commutative rings often poses new technical challenges and requires additional new ideas. We hope our results may find future applications in generalizing other results from the Boolean domains to general domains.

1.2 Techniques

Previous Techniques. Given a distribution D over the binary field which is not k -wise independent, a k -wise independent distribution was constructed in [4] by mixing³ D with a series of carefully chosen distributions in order to zero-out all the Fourier coefficients over subsets of size at most k . The total weight of the distributions used for mixing is an upper bound on the distance of D from k -wise independence. The distributions used for mixing are indexed by subsets $S \subset \{1, 2, \dots, n\}$ of size at most k . For a given such subset S , the added distribution U_S is picked such that, on the one hand it corrects the Fourier coefficient over S ; on the other hand, U_S 's Fourier coefficient over *any* other subset is zero. Using the orthogonality property of Hadamard matrices, one chooses U_S to be the uniform distribution over all strings whose parity over S is 1 (or -1 , depending on the sign of the distribution's bias over S). Although one can generalize it to work for prime fields, this construction breaks down when the alphabet size is a composite number.

For binary field a better bound is obtained in [1]. This is achieved by first working in the Fourier domain to remove all the first k -level Fourier coefficients of the input distribution. Such an operation ends up with a so-called “pseudo-distribution”, since at some points the resulting function may assume negative values. Then a series of carefully chosen k -wise independent distributions are added to the pseudo-distribution to fix the negative points. This approach does not admit a direct generalization to the non-Boolean cases because, for larger domains, the pseudo-distributions are in general complex-valued. Nevertheless⁴, one may use generalized Fourier expansion of real-valued functions to overcome this difficulty. We present this approach in Section 3. However, the bound obtained from this approach is weaker than our main results for the

³Here “mixing” means replacing the distribution D with a convex combination of D and some other distribution.

⁴We thank an anonymous referee for pointing this out.

uniform case which we discuss shortly. Moreover, the proof is “non-constructive” in the sense that we do not know exactly what distributions should we mix with the input distribution to make it k -wise independent. This drawback seems make it hard to generalize the approach to handle the non-uniform case. In contrast, our results on non-uniform k -wise independence relies crucially on the fact that the correction procedure for the uniform case is explicit and all the distributions used for mixing have some special structure (that is, they are uniform over all but at most k coordinates in the domain).

Uniform Distributions. Our results on uniform k -wise independent distributions extend the framework in [4]. As noted before, the key property used to mend a distribution into k -wise independent is the *orthogonality* relation between any pair of vectors. We first observe that all prime fields also enjoy this nice feature after some slight modifications. More specifically, for any two non-zero vectors \mathbf{a} and \mathbf{b} in \mathbb{Z}_p^n that are *linearly independent*, the set of strings with $\sum_{i=1}^n a_i x_i \equiv j \pmod{p}$ are *uniformly* distributed over the sets $S_{\mathbf{b},\ell} \stackrel{\text{def}}{=} \{\mathbf{x} : \sum_{i=1}^n b_i x_i \equiv \ell \pmod{p}\}$ for every $0 \leq \ell \leq p-1$. We call this the *strong orthogonality* between vectors \mathbf{a} and \mathbf{b} . The case when $q = |\Sigma|$ is not a prime is less straightforward. The main difficulty is that the strong orthogonality between pairs of vectors no longer holds, even when they are linearly independent⁵.

Suppose we wish to use some distribution $U_{\mathbf{a}}$ to correct the bias over \mathbf{a} . A simple but important observation is that we only need that $U_{\mathbf{a}}$'s Fourier coefficient at \mathbf{b} to be zero, which is a much weaker requirement than the property of being strongly orthogonal between \mathbf{a} and \mathbf{b} . Using a classical result in linear systems of congruences due to Smith [39], we are able to show that when \mathbf{a} satisfies $\gcd(a_1, \dots, a_n) = 1$ and \mathbf{b} is not a multiple of \mathbf{a} , the set of strings with $\sum_{i=1}^n a_i x_i \equiv j \pmod{q}$ are *uniformly* distributed over $S_{\mathbf{b},\ell}$ for ℓ 's that lie in a *subgroup* of \mathbb{Z}_q (compared with the uniform distribution over the whole group \mathbb{Z}_p for the prime field case). We refer to this as the *weak orthogonality* between vectors \mathbf{a} and \mathbf{b} . To zero-out the Fourier coefficient at \mathbf{a} , we instead bundle the Fourier coefficient at \mathbf{a} with the Fourier coefficients at $\ell\mathbf{a}$ for every $\ell = 2, \dots, q-1$, and think of them as the Fourier coefficients of some function over the one-dimensional space \mathbb{Z}_q . This allows us to upper bound the total weight required to simultaneously correct *all* the Fourier coefficients at \mathbf{a} and its multiples using only $U_{\mathbf{a}}$. We also generalize the result to product spaces $\Omega = \Sigma_1 \times \dots \times \Sigma_n$, which in general have different alphabets at different coordinates.

Non-uniform Distributions. One possible way of extending the upper bounds of the uniform case to the non-uniform case would be to map non-uniform probabilities to uniform probabilities over a larger domain. For example, consider when $q = 2$ a distribution D with $\Pr_D[x_i = 0] = 0.501$ and $\Pr_D[x_i = 1] = 0.499$. We could map $x_i = 0$ and $x_i = 1$ uniformly to $\{1, \dots, 501\}$ and $\{502, \dots, 1000\}$, respectively and test if the transformed distribution D' over $\{1, \dots, 1000\}$ is k -wise independent. Unfortunately, this approach produces a huge overhead on the distance upper bound, due to the fact that the alphabet size (and hence the distance bound) blowup depends on the closeness of marginal probabilities over different letters in the alphabet. However, in the previous example we should expect D behaves very much like the uniform case rather than with an additional factor of 1000 blowup in the alphabet size.

Instead we take the following approach. Consider a compressing/stretching factor for each marginal probability $\Pr_D[x_i = z]$, where $z \in \Sigma$ and $1 \leq i \leq n$. Specifically, let $\theta_i(z) \stackrel{\text{def}}{=} \frac{1}{q \Pr_D[x_i = z]}$ so that $\theta_i(z) \Pr_D[x_i = z] = \frac{1}{q}$, the probability that $x_i = z$ in the uniform distribution. If we multiply $D(\mathbf{x})$ for each \mathbf{x} in the domain by a product of n such factors, one for each coordinate, the non-uniform k -wise independent distribution will be transformed into a uniform one. The hope is that distributions *close to*

⁵We say two non-zero vectors \mathbf{a} and \mathbf{b} in \mathbb{Z}_q^n are linearly dependent if there exist two non-zero integers s and t in \mathbb{Z}_q such that $sa_i \equiv tbi \pmod{q}$ for every $1 \leq i \leq n$, and linearly independent if they are not linearly dependent

non-uniform k -wise independent will also be transformed into distributions that are *close to* uniform k -wise independent. However, this could give rise to exponentially large distribution weight at some points in the domain, making the task of estimating the relevant Fourier coefficients intractable. Intuitively, for testing k -wise independence purposes, all we need to know are the “local” weight distributions. To be more specific, for a vector $\mathbf{a} \in \Sigma^n$, the *support set* or simply *support* of \mathbf{a} is $\text{supp}(\mathbf{a}) \stackrel{\text{def}}{=} \{i \in [n] : a_i \neq 0\}$. For every non-zero vector \mathbf{a} of weight at most k , we define a new *non-uniform Fourier coefficient* at \mathbf{a} in the following steps:

1. Project D to $\text{supp}(\mathbf{a})$ to get $D_{\text{supp}(\mathbf{a})}$;
2. For every point in the support of $D_{\text{supp}(\mathbf{a})}$, multiply the marginal probability by the product of a sequence of compressing/stretching factors, one for each coordinate in $\text{supp}(\mathbf{a})$. Denote this transformed distribution by $D'_{\text{supp}(\mathbf{a})}$;
3. Define the non-uniform Fourier coefficient of D at \mathbf{a} to be the (uniform) Fourier coefficient of $D'_{\text{supp}(\mathbf{a})}$ at \mathbf{a} .

We then show a new characterization that D is non-uniform k -wise independent *if and only if* all the first k levels non-zero non-uniform Fourier coefficients of D are zero. This enables us to apply the Fourier coefficient correcting approach developed for the uniform case to the non-uniform case. Loosely speaking, for any vector \mathbf{a} , we can find a (small-weight) distribution $\mathcal{W}_{\mathbf{a}}$ such that mixing $D'_{\text{supp}(\mathbf{a})}$ with $\mathcal{W}_{\mathbf{a}}$ zeroes-out the (uniform) Fourier coefficient at \mathbf{a} , which is, by definition, the non-uniform Fourier coefficient of D at \mathbf{a} . But this $\mathcal{W}_{\mathbf{a}}$ is the distribution to mix with the “transformed” distribution, i.e., $D'_{\text{supp}(\mathbf{a})}$. To find out the distribution works for D , we apply an *inverse* compressing/stretching transformation to $\mathcal{W}_{\mathbf{a}}$ to get $\tilde{\mathcal{W}}_{\mathbf{a}}$. It turns out that mixing $\tilde{\mathcal{W}}_{\mathbf{a}}$ with the original distribution D not only corrects D ’s non-uniform Fourier coefficient at \mathbf{a} but also does not increase D ’s non-uniform Fourier coefficients at any other vectors except those vectors whose supports are strictly contained in $\text{supp}(\mathbf{a})$. Moreover, transforming from $\mathcal{W}_{\mathbf{a}}$ to $\tilde{\mathcal{W}}_{\mathbf{a}}$ incurs at most a constant (independent of n) blowup in the total weight. Therefore we can start from vectors of weight k and correct the non-uniform Fourier coefficients from level k to lower levels. This process terminates after we finish correcting all vectors of weight 1 and thus obtain a k -wise independent distribution. Bounding the total weight added during this process gives an upper bound on the distance between D and non-uniform k -wise independence. We hope that the notion of non-uniform Fourier coefficients may find other applications when non-uniform independence is involved.

1.3 Other Related Research

There are many works on k -wise independence, most focus on various *constructions* of k -wise independence or distributions that approximate k -wise independence. k -wise independent random variables were first studied in probability theory [23] and then in complexity theory [12, 2, 28, 29] mainly for derandomization purposes. Constructions of almost k -wise independent distributions were studied in [31, 3, 5, 17, 9]. Construction results of non-uniform k -wise independent distributions were given in [24, 26].

There has been much activity on property testing of distributions. Some examples include testing uniformity [20, 7], independence [6], monotonicity and being unimodal [8], estimating the support sizes [34] and testing a weaker notion than k -wise independence, namely, “almost k -wise independence” [1].

Many other techniques have also been developed to generalize results from Boolean domains to arbitrary domains [15, 30, 10].

1.4 Organization

We first give some necessary definitions and preliminary facts in Section 2. A proof of the first main theorem based on orthogonal polynomials is then given in Section 3. We present another proof of the robust characterization of the uniform k -wise independent distributions in Section 4. We then generalize the approach developed in Section 4 to the case of non-uniform k -wise independence and prove the second main theorem in Section 5. Finally in Section 6 we study the problem of testing almost k -wise independence over product spaces.

2 Preliminaries

Let n and m be two natural numbers with $m > n$. We write $[n]$ for the set $\{1, \dots, n\}$ and $[n, m]$ for the set $\{n, n+1, \dots, m\}$. For any integer $1 \leq k \leq n$, we write $\binom{[n]}{k}$ to denote the set of all k -subsets of $[n]$. Throughout this paper, Σ always stands for a finite set. Without loss of generality, we assume that $\Sigma = \{0, 1, \dots, q-1\}$, where $q = |\Sigma|$.

We use bold letters to denote vectors in Σ^n , for example, \mathbf{a} stands for the vector (a_1, \dots, a_n) with $a_i \in \Sigma$ being the i^{th} component of \mathbf{a} . For two vectors \mathbf{a} and \mathbf{b} in Σ^n , their inner product is $\mathbf{a} \cdot \mathbf{b} \stackrel{\text{def}}{=} \sum_{i=1}^n a_i b_i \pmod{q}$. The support of \mathbf{a} is the set of indices at which \mathbf{a} is non-zero. That is, $\text{supp}(\mathbf{a}) = \{i \in [n] : a_i \neq 0\}$. The weight of a vector \mathbf{a} is the cardinality of its support. Let $1 \leq k \leq n$ be an integer. We use $M(n, k, q) \stackrel{\text{def}}{=} \binom{n}{1}(q-1) + \dots + \binom{n}{k}(q-1)^k$ to denote the total number of non-zero vectors in Σ^n of weight at most k . Note that $M(n, k, q) = \Theta(n^k (q-1)^k)$ for $k = O(1)$.

We assume that there is an underlying probability distribution D from which we can receive independent, identically distributed (i.i.d) samples. The domain Ω of every distribution we consider in this paper will always be finite and in general is of the form $\Omega = \Sigma_1 \times \dots \times \Sigma_n$, where $\Sigma_1, \dots, \Sigma_n$ are finite sets. A point \mathbf{x} in Ω is said to be *in the support* of a distribution D if $D(\mathbf{x}) > 0$.

Let D_1 and D_2 be two distributions over the same domain Ω . The statistical distance between D_1 and D_2 is $\Delta(D_1, D_2) = \frac{1}{2} \sum_{\mathbf{x} \in \Omega} |D_1(\mathbf{x}) - D_2(\mathbf{x})|$. One can check that statistical distance is a metric and in particular satisfies the triangle inequality. We use statistical distance as the main metric to measure closeness between distributions in this paper. For any $0 \leq \epsilon \leq 1$, one may define a new distribution D' as the convex combination of D_1 and D_2 : $D' = \frac{1}{1+\epsilon} D_1 + \frac{\epsilon}{1+\epsilon} D_2$. It then follows that $\Delta(D', D_1) \leq \frac{\epsilon}{1+\epsilon} \leq \epsilon$. Sometimes we abuse notation and call the non-negative function ϵD_1 a *weighted* distribution (in particular, a *small-weight distribution* when ϵ is small).

Let $S = \{i_1, \dots, i_k\} \subseteq [n]$ be an index set. Let \mathbf{x} be an n -dimensional vector. We write \mathbf{x}_S to denote the k -dimensional vector obtained from projecting \mathbf{x} to the indices in S . Similarly, the *projection distribution* of a discrete distribution D over Σ^n with respect to S , denoted by D_S , is the distribution obtained by restricting to the coordinates in S . Namely, $D_S : \Sigma^k \rightarrow [0, 1]$ is a distribution such that $D_S(z_{i_1} \dots z_{i_k}) = \sum_{\mathbf{x}_S = (z_{i_1}, \dots, z_{i_k})} D(\mathbf{x})$. For brevity, we sometimes write $D_S(\mathbf{z}_S)$ for $D_S(z_{i_1} \dots z_{i_k})$.

2.1 The k -wise Independent Distributions

Let $D : \Sigma_1 \times \dots \times \Sigma_n \rightarrow [0, 1]$ be a distribution. We say D is the *uniform* distribution if for every $\mathbf{x} \in \Sigma_1 \times \dots \times \Sigma_n$, $\Pr_{\mathbf{X} \sim D}[\mathbf{X} = \mathbf{x}] = \frac{1}{q_1 \dots q_n}$, where $q_i = |\Sigma_i|$. D is *k -wise independent* if for any set of k indices $\{i_1, \dots, i_k\}$ and for any $z_1 \dots z_k \in \Sigma_{i_1} \times \dots \times \Sigma_{i_k}$, $\Pr_{\mathbf{X} \sim D}[X_{i_1} \dots X_{i_k} = z_1 \dots z_k] = \Pr_{\mathbf{X} \sim D}[X_{i_1} = z_1] \times \dots \times \Pr_{\mathbf{X} \sim D}[X_{i_k} = z_k]$. D is *uniform k -wise independent* if, on top of the previous condition, we have $\Pr_{\mathbf{X} \sim D}[X_i = z_j] = \frac{1}{|\Sigma_i|}$ for every $1 \leq i \leq n$ and every $z_j \in \Sigma_i$. Let \mathcal{D}_{kwi} denote

the set of all uniform k -wise independent distributions. The distance between D and \mathcal{D}_{kwi} , denoted by $\Delta(D, \mathcal{D}_{\text{kwi}})$, is the minimum statistical distance between D and any uniform k -wise independent distribution, i.e., $\Delta(D, \mathcal{D}_{\text{kwi}}) \stackrel{\text{def.}}{=} \inf_{D' \in \mathcal{D}_{\text{kwi}}} \Delta(D, D')$.

2.2 Discrete Fourier Transform.

For background on the discrete Fourier transform in computer science, the reader is referred to [40, 41, 14]. Let $f : \Sigma_1 \times \cdots \times \Sigma_n \rightarrow \mathbb{C}$ be any function defined over the discrete product space, we define the Fourier transform of D to be, for every $\mathbf{a} \in \Sigma_1 \times \cdots \times \Sigma_n$,

$$\hat{f}(\mathbf{a}) = \sum_{\mathbf{x} \in \Sigma_1 \times \cdots \times \Sigma_n} f(\mathbf{x}) e^{2\pi i \left(\frac{a_1 x_1}{q_1} + \cdots + \frac{a_n x_n}{q_n} \right)}. \quad (1)$$

$\hat{f}(\mathbf{a})$ is called f 's Fourier coefficient at \mathbf{a} . If the weight of \mathbf{a} is k , we then refer to $\hat{f}(\mathbf{a})$ as a *degree- k or level- k Fourier coefficient*.

One can easily verify that the inverse Fourier transform is

$$f(\mathbf{x}) = \frac{1}{q_1 \cdots q_n} \sum_{\mathbf{a} \in \Sigma_1 \times \cdots \times \Sigma_n} \hat{f}(\mathbf{a}) e^{-2\pi i \left(\frac{a_1 x_1}{q_1} + \cdots + \frac{a_n x_n}{q_n} \right)}. \quad (2)$$

Note that if $\Sigma_i = \Sigma$ for every $1 \leq i \leq n$ (which is the main focus of this paper), then $\hat{f}(\mathbf{a}) = \sum_{\mathbf{x} \in \Sigma^n} f(\mathbf{x}) e^{\frac{2\pi i}{q} \mathbf{a} \cdot \mathbf{x}}$ and $f(\mathbf{x}) = \frac{1}{|\Sigma|^n} \sum_{\mathbf{a} \in \Sigma^n} \hat{f}(\mathbf{a}) e^{-\frac{2\pi i}{q} \mathbf{a} \cdot \mathbf{x}}$.

We will use the following two simple facts about discrete Fourier transform which are straightforward to prove. Note that Fact 2.1 is a special case of Fact 2.2.

Fact 2.1. For any integer ℓ which is not congruent to 0 modulo q , $\sum_{j=0}^{q-1} e^{\frac{2\pi i}{q} \ell j} = 0$.

Fact 2.2. Let d, ℓ_0 be integers such that $d|q$ and $0 \leq \ell_0 \leq d-1$. Then $\sum_{j=0}^{\frac{q}{d}-1} e^{\frac{2\pi i}{q} (\ell_0 + dj)} = 0$.

Proposition 2.3. Let D be a distribution over $\Sigma_1 \times \cdots \times \Sigma_n$. Then D is the uniform distribution if and only if for any non-zero vector $\mathbf{a} \in \Sigma_1 \times \cdots \times \Sigma_n$, $\hat{D}(\mathbf{a}) = 0$.

Proof. First note that $\hat{D}(\mathbf{0}) = \sum_{\mathbf{x}} D(\mathbf{x}) = 1$. Therefore, if $\hat{D}(\mathbf{a}) = 0$ for all non-zero \mathbf{a} , then by the inverse Fourier transform (2),

$$D(\mathbf{x}) = \frac{1}{q_1 \cdots q_n} \hat{D}(\mathbf{0}) = \frac{1}{q_1 \cdots q_n}.$$

For the converse, let \mathbf{a} be any non-zero vector. Without loss of generality, suppose $a_1 \neq 0$. Since $D(\mathbf{x}) = \frac{1}{q_1 \cdots q_n}$ for all \mathbf{x} , we have

$$\begin{aligned} \hat{D}(\mathbf{a}) &= \frac{1}{q_1 \cdots q_n} \sum_{\mathbf{x}} e^{2\pi i \left(\frac{a_1 x_1}{q_1} + \cdots + \frac{a_n x_n}{q_n} \right)} \\ &= \frac{1}{q_1 \cdots q_n} \sum_{x_2, \dots, x_n} e^{2\pi i \left(\frac{a_2 x_2}{q_2} + \cdots + \frac{a_n x_n}{q_n} \right)} \sum_{x_1=0}^{q_1-1} e^{\frac{2\pi i}{q_1} a_1 x_1} \\ &= 0. \end{aligned} \quad (\text{by Fact 2.1}) \quad \square$$

By applying Proposition 2.3 to distributions obtained from restricting D to any k indices and observing the fact that, by the definition of Fourier transform, $\hat{D}(\mathbf{a}) = \hat{D}_S(\mathbf{a})$ when $\text{supp}(\mathbf{a}) \subseteq S$, we have the following characterization of k -wise independent distributions over product spaces, which is the basis of all the testing algorithms in this paper.

Corollary 2.4. *A distribution D over $\Sigma_1 \times \cdots \times \Sigma_n$ is k -wise independent if and only if for all non-zero vectors \mathbf{a} of weight at most k , $\hat{D}(\mathbf{a}) = 0$.*

2.3 Some other Definitions and Notation

We are going to use the following notation extensively in this paper.

Definition 2.5. Let D be a distribution over Σ^n . For every $\mathbf{a} \in \Sigma^n$ and every $0 \leq j \leq q - 1$, let $P_{\mathbf{a},j}^D \stackrel{\text{def}}{=} \Pr_{\mathbf{X} \sim D}[\mathbf{a} \cdot \mathbf{X} \equiv j \pmod{q}]$. When the distribution D is clear from the context, we often omit the superscript D and simply write $P_{\mathbf{a},j}$.

The Fourier transform (1) can be rewritten as

$$\hat{D}(\mathbf{a}) = \sum_{j=0}^{q-1} \Pr_{\mathbf{X} \sim D}[\mathbf{a} \cdot \mathbf{X} \equiv j \pmod{q}] e^{\frac{2\pi i}{q} j} = \sum_{j=0}^{q-1} P_{\mathbf{a},j} e^{\frac{2\pi i}{q} j}. \quad (3)$$

For any non-zero vector $\mathbf{a} \in \Sigma^n$ and any integer $0 \leq j \leq q - 1$, let $S_{\mathbf{a},j} \stackrel{\text{def}}{=} \{\mathbf{x} \in \Sigma^n : \sum_{i=1}^n a_i x_i \equiv j \pmod{q}\}$. Finally we write $U_{\mathbf{a},j}$ for the uniform distribution over $S_{\mathbf{a},j}$.

2.4 Query and Time Complexity Analysis of the Generic Testing Algorithm

We now provide a detailed analysis of the query and time complexity analysis of the generic testing algorithm as shown in Fig. 1. The main technical tool is the following standard Chernoff bound.

Theorem 2.6 (Chernoff Bound). *Let X_1, \dots, X_m be i.i.d. 0-1 random variables with $\mathbf{E}[X_i] = \mu$. Let $\bar{\mu} = \frac{1}{m} \sum_{i=1}^m X_i$. Then for all $\gamma, 0 < \gamma < 1$, we have $\Pr[|\bar{\mu} - \mu| \geq \gamma\mu] \leq 2 \cdot e^{-\frac{\gamma^2 \mu m}{3}}$.*

Theorem 2.7. *Let D be a distribution over Σ^n where $|\Sigma| = q$ and A be a subset of vectors in Σ^n . Suppose the distance between D and the set of k -wise independent distributions satisfies the following conditions:*

- (completeness) *For any $0 \leq \delta \leq 1$, if $\Delta(D, \mathcal{D}_{\text{kwi}}) \leq \delta$, then $|\hat{D}(\mathbf{a})| \leq \kappa\delta$ for every \mathbf{a} in A ;*
- (soundness) $\Delta(D, \mathcal{D}_{\text{kwi}}) \leq K \max_{\mathbf{a} \in A} |\hat{D}(\mathbf{a})|$, where K is a function of n, k, q and A .

*Then for any $0 < \epsilon \leq 1$, the generic testing algorithm draws⁶ $m = O(\frac{q^2 K^2}{\epsilon^2} \log(q|A|))$ independent samples from D and runs in time $O(\frac{q^2 K^2 |A|}{\epsilon^2} \log(q|A|))$ and satisfies the followings: If $\Delta(D, \mathcal{D}_{\text{kwi}}) \leq \frac{\epsilon}{3\kappa K}$, then with probability at least $2/3$, it outputs “**Accept**”; if $\Delta(D, \mathcal{D}_{\text{kwi}}) > \epsilon$, then with probability at least $2/3$, it outputs “**Reject**”.*

⁶For all the cases studied in this paper, the size of A is much larger than q , therefore we omit the factor q in the logarithm in all the subsequent formulas.

Proof. The algorithm is to sample D independently m times and use these samples to estimate, for each $\mathbf{a} \in A$, the Fourier coefficient of D at \mathbf{a} . Then if $\max_{\mathbf{a} \in A} \left| \hat{D}(\mathbf{a}) \right| \leq \frac{2\epsilon}{3K}$, the algorithm accepts D ; otherwise it rejects D . The running time bound follows from the fact that we need to estimate $|A|$ Fourier coefficients using m samples.

For every $\mathbf{a} \in A$ and $0 \leq j \leq q-1$, define a 0-1 indicator variable $I_{\mathbf{a},j}(\mathbf{x})$, where $\mathbf{x} \in \Sigma^n$, which is 1 if $\mathbf{a} \cdot \mathbf{x} \equiv j \pmod{q}$ and 0 otherwise. Clearly $\bar{I}_{\mathbf{a},j} \stackrel{\text{def}}{=} \mathbf{E}[I_{\mathbf{a},j}] = P_{\mathbf{a},j}$. Let $\bar{P}_{\mathbf{a},j} = \frac{1}{m} \sum_{\mathbf{x} \in Q} I_{\mathbf{a},j}(\mathbf{x})$; that is, $\bar{P}_{\mathbf{a},j}$ is the empirical estimate of $P_{\mathbf{a},j}$. Since $P_{\mathbf{a},j} \leq 1$, by Chernoff bound, $\Pr[|\bar{P}_{\mathbf{a},j} - P_{\mathbf{a},j}| > \frac{\epsilon}{3qK}] < \frac{2}{3q|A|}$. By union bound, with probability at least $2/3$, for every vector \mathbf{a} in A and every $0 \leq j < q$, $|\bar{P}_{\mathbf{a},j} - P_{\mathbf{a},j}| \leq \frac{\epsilon}{3qK}$.

The following fact provides an upper bound of the error in estimating the Fourier coefficient at \mathbf{a} in terms of the errors from estimating $P_{\mathbf{a},j}$.

Fact 2.8. *Let $f, g : \{0, \dots, q-1\} \rightarrow \mathbb{R}$ with $|f(j) - g(j)| \leq \epsilon$ for every $0 \leq j \leq q-1$. Then $|\hat{f}(\ell) - \hat{g}(\ell)| \leq q\epsilon$ for all $0 \leq \ell \leq q-1$.*

Proof. Let $h = f - g$, then $|h(j)| \leq \epsilon$ for every j . Therefore,

$$\begin{aligned} & |\hat{f}(\ell) - \hat{g}(\ell)| \\ &= |\hat{h}(\ell)| = \left| \sum_{j=0}^{q-1} h(j) e^{\frac{2\pi i}{q} \ell j} \right| \\ &\leq \sum_{j=0}^{q-1} |h(j) e^{\frac{2\pi i}{q} \ell j}| = \sum_{j=0}^{q-1} |h(j)| \\ &\leq \sum_{j=0}^{q-1} \epsilon = q\epsilon. \quad \square \end{aligned}$$

Let $\tilde{D}(\mathbf{a})$ be the estimated Fourier coefficient computed from $\bar{P}_{\mathbf{a},j}$. Fact 2.8 and (3) then imply that with probability at least $2/3$, $|\tilde{D}(\mathbf{a}) - \hat{D}(\mathbf{a})| \leq \frac{\epsilon}{3K}$ for every \mathbf{a} in A .

Now if $\Delta(D, \mathcal{D}_{\text{kwi}}) \leq \frac{\epsilon}{3\kappa K}$, then by our completeness assumption, we have $\max_{\mathbf{a} \in A} \left| \hat{D}(\mathbf{a}) \right| \leq \frac{\epsilon}{3K}$. Taking the error from estimation into account, $\max_{\mathbf{a} \in A} \left| \tilde{D}(\mathbf{a}) \right| \leq \frac{2\epsilon}{3K}$ holds with probability at least $2/3$. Therefore with probability at least $2/3$, the algorithm returns “**Accept**”.

If $\Delta(D, \mathcal{D}_{\text{kwi}}) > \epsilon$, then by our soundness assumption, $\max_{\mathbf{a} \in A} \left| \hat{D}(\mathbf{a}) \right| > \frac{\epsilon}{K}$. Again with probability at least $2/3$, $\max_{\mathbf{a} \in A} \left| \tilde{D}(\mathbf{a}) \right| > \frac{2\epsilon}{3K}$ for every \mathbf{a} in A , so the algorithm returns “**Reject**”. \square

3 A Proof of Theorem 1.1 Based on Orthogonal Polynomials

In this section we give our first and conceptually simple proof of Theorem 1.1. The bound we prove here is somewhat weaker than stated in Theorem 1.1. The basic idea is to apply the “cut in the Fourier space and then mend in the function space” approach in [1] to Fourier expansions with discrete orthogonal real polynomials as the basis functions.

3.1 Generalized Fourier series

The discrete Fourier transform reviewed in Section 2 can be generalized to decompositions over any orthonormal basis of an inner product space. In particular, for the discrete function space $\mathbb{R}^{\{0, \dots, q-1\}}$, any orthonormal basis of real functions $\{g_0(x), \dots, g_{q-1}(x)\}$ with $g_0(x) = 1$ for every x (the identity function)⁷ can be used in place of the standard Fourier basis $\{1, e^{\frac{2\pi ix}{q}}, \dots, e^{\frac{2\pi i(q-1)x}{q}}\}$. In general, such a basis of functions may be constructed by the Gram-Schmidt process. For concreteness, we present an explicit construction based on *discrete Legendre orthogonal polynomials* [32], a special case of Hahn polynomials. An extensive treatment of discrete orthogonal polynomials may be found in [33]. We remark that our proof works for any set of complete orthonormal basis of real functions as long as one of the basis functions is the identity function.

For $n \geq 0$, we write $(x)_n := x(x-1) \cdots (x-n+1)$ for the n^{th} falling factorial of x . For any integer $q \geq 2$, the discrete Legendre orthogonal polynomials, $\{P_a(x; q)\}_{a=0}^{q-1}$, are defined as

$$P_a(x; q) = \sum_{j=0}^a (-1)^j \binom{a}{j} \binom{a+j}{j} \frac{(x)_j}{(q-1)_j},$$

$$P_a(0; q) = 1, \text{ for all } a = 0, 1, \dots, q-1.$$

These polynomials satisfy the following orthogonal properties (see, e.g., [32]):

$$\sum_{x=0}^{q-1} P_a(x; q) P_b(x; q) = \begin{cases} 0, & \text{if } a \neq b, \\ \frac{1}{2a+1} \frac{(q+a)_{a+1}}{(q-1)_a}, & \text{if } a = b. \end{cases}$$

Now we define⁸ a complete set of orthonormal functions $\{\chi_a^{\text{OF}}(x)\}_{a=0}^{q-1}$ by

$$\chi_a^{\text{OF}}(x) = \sqrt{\frac{(2a+1)(q)_{a+1}}{(q+a)_{a+1}}} P_a(x; q),$$

then they form a complete basis for the real functions space over $\{0, 1, \dots, q-1\}$ and satisfy the orthogonality relation

$$\sum_{x=0}^{q-1} \chi_a^{\text{OF}}(x) \chi_b^{\text{OF}}(x) = \begin{cases} 0, & \text{if } a \neq b, \\ q, & \text{if } a = b. \end{cases}$$

Because of the orthogonality relation $\sum_{x=0}^{q-1} |\chi_a^{\text{OF}}(x)|^2 = q$ for every a , we immediately have

Fact 3.1. For every $0 \leq a \leq q-1$ and every $x \in \{0, 1, \dots, q-1\}$, $|\chi_a^{\text{OF}}(x)| \leq \sqrt{q}$.

Due to the orthogonality and the completeness of the basis functions, any real function $f : \{0, 1, \dots, q-1\} \rightarrow \mathbb{R}$ can be uniquely expanded in terms of $\{\chi_a^{\text{OF}}(x)\}_{a=0}^{q-1}$ as:

$$f(x) = \frac{1}{q} \sum_{a=0}^{q-1} \hat{f}^{\text{OF}}(a) \chi_a^{\text{OF}}(x),$$

⁷Therefore the uniform distribution is proportional to g_0 and then by the orthogonality relation, all the non-zero Fourier coefficients of the uniform distribution are zero.

⁸We add the superscript OF (denoting *orthogonal functions*) to distinguish them from the standard real Fourier basis functions over $\{0, 1\}^n$.

with the inversion formula

$$\hat{f}^{\text{OF}}(a) = \sum_{x=0}^{q-1} f(x) \chi_a^{\text{OF}}(x).$$

We call the expansion coefficients $\{\hat{f}^{\text{OF}}(a)\}$ the *generalized Fourier coefficients* of f .

Generalizing this expansion to real functions over higher dimensional spaces is straightforward. Let $n \geq 1$ be an integer and let $f : \{0, 1, \dots, q-1\}^n \rightarrow \mathbb{R}$. The generalized Fourier expansion of f is simply

$$f(\mathbf{x}) = \frac{1}{q^n} \sum_{\mathbf{a}} \hat{f}^{\text{OF}}(\mathbf{a}) \chi_{\mathbf{a}}^{\text{OF}}(\mathbf{x}),$$

with the inversion formula

$$\hat{f}^{\text{OF}}(\mathbf{a}) = \sum_{\mathbf{x}} f(\mathbf{x}) \chi_{\mathbf{a}}^{\text{OF}}(\mathbf{x}),$$

where $\chi_{\mathbf{a}}^{\text{OF}}(\mathbf{x}) \stackrel{\text{def}}{=} \prod_{i=1}^n \chi_{a_i}^{\text{OF}}(x_i)$ and satisfy the orthogonality relation $\sum_{\mathbf{x}} \chi_{\mathbf{a}}^{\text{OF}}(\mathbf{x}) \chi_{\mathbf{b}}^{\text{OF}}(\mathbf{x}) = \begin{cases} 0, & \text{if } \mathbf{a} \neq \mathbf{b}, \\ q^n, & \text{if } \mathbf{a} = \mathbf{b}. \end{cases}$

A direct consequence of the orthogonality of the basis functions $\{\chi_{\mathbf{a}}^{\text{OF}}(\mathbf{x})\}$ is the following Parseval's equality

$$\sum_{\mathbf{x}} f^2(\mathbf{x}) = \frac{1}{q^n} \sum_{\mathbf{a}} \hat{f}^{\text{OF}}(\mathbf{a})^2.$$

It is easy to check that the following characterizations of the uniform distribution and k -wise independent distributions over $\{0, 1, \dots, q-1\}^n$ in terms of the generalized Fourier coefficients. The proofs follow directly from the orthogonality of $\{\chi_{\mathbf{a}}^{\text{OF}}(\mathbf{x})\}$ and the definition of k -wise independence, therefore we omit them here.

Proposition 3.2. *Let D be a distribution over $\{0, 1, \dots, q-1\}^n$. Then D is the uniform distribution if and only if for all non-zero vector $\mathbf{a} \in \{0, 1, \dots, q-1\}^n$, $\hat{D}^{\text{OF}}(\mathbf{a}) = 0$.*

Corollary 3.3. *A distribution D over $\{0, 1, \dots, q-1\}^n$ is k -wise independent if and only if for all non-zero vectors \mathbf{a} of weight at most k , $\hat{D}^{\text{OF}}(\mathbf{a}) = 0$.*

3.2 Proof of Theorem 1.1

The basic idea of [1] is the following. Given a distribution D , we first operate in the Fourier space to construct a ‘‘pseudo-distribution’’ D_1 by setting all the first k -level generalized Fourier coefficients (except for the trivial Fourier coefficient) to zero. All other generalized Fourier coefficients of D_1 are the same as D . Generally speaking, D_1 is not going to be a distribution because it may assume negative values at some points. We then correct all these negative points by mixing D_1 with the uniform distribution with some appropriate weight. That is, we set $D' = \frac{1}{1+w} D_1 + \frac{w}{1+w} U$, where U is the uniform distribution and $w > 0$ is the weight of the uniform distribution. After such an operation, since the uniform distribution clearly has all its first k -level generalized Fourier coefficients equal to zero and due to linearity of the generalized Fourier transform, we maintain that all the first k -level generalized Fourier coefficients of D' are still zero; on the other hand, we increase the weights at negative points so that they now assume non-negative values in D' . Bounding the total statistical distance between D and D' then offers an upper bound on the distance between D and k -wise independence.

Let $D : \{0, 1, \dots, q-1\}^n \rightarrow \mathbb{R}^{\geq 0}$ be a distribution, that is, $D(\mathbf{x}) \geq 0$ for all \mathbf{x} and $\sum_{\mathbf{x}} D(\mathbf{x}) = 1$. First we define a real function $D_1 : \{0, 1, \dots, q-1\}^n \rightarrow \mathbb{R}$ by explicitly specifying all its generalized Fourier coefficients:

$$\hat{D}_1^{\text{OF}}(\mathbf{a}) = \begin{cases} 0, & \text{if } 0 < \text{wt}(\mathbf{a}) \leq k \\ \hat{D}^{\text{OF}}(\mathbf{a}), & \text{otherwise.} \end{cases}$$

We call D_1 a ‘‘pseudo-distribution’’ because D_1 may assume negative values at some points in the domain, which are called the *holes* in D_1 . Note that since $\hat{D}_1^{\text{OF}}(\mathbf{0}) = \hat{D}^{\text{OF}}(\mathbf{0}) = 1$, we have $\sum_{\mathbf{x}} D_1(\mathbf{x}) = 1$. So the only difference between D_1 and a distribution is these holes. The following lemma bounds the maximum depth of the holes in D_1 .

Lemma 3.4. *Let h be the maximum depth of the holes in D_1 , then*

$$h \leq \frac{q^{k/2}}{q^n} \sum_{0 < \text{wt}(\mathbf{a}) \leq k} |\hat{D}^{\text{OF}}(\mathbf{a})|.$$

Proof. From the upper bound in Fact 3.1, it follows that $|\chi_{\mathbf{a}}^{\text{OF}}(\mathbf{x})| \leq q^{k/2}$ if the weight of \mathbf{a} is at most k . Now since $D(\mathbf{x}) \geq 0$ for every \mathbf{x} in the domain and D_1 is obtained by cutting off all the first k level generalized Fourier coefficients of D , by linearity of the generalized Fourier expansion,

$$D_1(\mathbf{x}) = D(\mathbf{x}) - \frac{1}{q^n} \sum_{0 < \text{wt}(\mathbf{a}) \leq k} \hat{D}^{\text{OF}}(\mathbf{a}) \chi_{\mathbf{a}}^{\text{OF}}(\mathbf{x}).$$

Therefore, for all \mathbf{x} with $D_1(\mathbf{x}) < 0$, $\frac{1}{q^n} \sum_{0 < \text{wt}(\mathbf{a}) \leq k} \hat{D}^{\text{OF}}(\mathbf{a}) \chi_{\mathbf{a}}^{\text{OF}}(\mathbf{x}) > 0$, so we can upper bound the depth of every hole as

$$\begin{aligned} |D_1(\mathbf{x})| &= \left| \frac{1}{q^n} \sum_{0 < \text{wt}(\mathbf{a}) \leq k} \hat{D}^{\text{OF}}(\mathbf{a}) \chi_{\mathbf{a}}^{\text{OF}}(\mathbf{x}) - D(\mathbf{x}) \right| \\ &\leq \left| \frac{1}{q^n} \sum_{0 < \text{wt}(\mathbf{a}) \leq k} \hat{D}^{\text{OF}}(\mathbf{a}) \chi_{\mathbf{a}}^{\text{OF}}(\mathbf{x}) \right| \\ &\leq \frac{q^{k/2}}{q^n} \sum_{0 < \text{wt}(\mathbf{a}) \leq k} |\hat{D}^{\text{OF}}(\mathbf{a})|. \quad \square \end{aligned}$$

The following lemma bounds the ℓ_1 -distance between a function and its convex combination with other distributions.

Lemma 3.5 ([1]). *Let f be a real function defined over $\{0, 1, \dots, q-1\}^n$ such that $\sum_{\mathbf{x}} f(\mathbf{x}) = 1$. Let D_1, \dots, D_ℓ be distributions over the same domain and suppose there exist non-negative real numbers w_1, \dots, w_ℓ such that $D' \stackrel{\text{def}}{=} \frac{1}{1 + \sum_{i=1}^\ell w_i} (f + \sum_{i=1}^\ell w_i D_i)$ is non-negative for all $\mathbf{x} \in \{0, 1, \dots, q-1\}^n$. Then $\sum_{\mathbf{x}} |f(\mathbf{x}) - D'(\mathbf{x})| \leq 2 \sum_{i=1}^\ell w_i$.*

Now we can mix D_1 with the uniform distribution U over $\{0, 1, \dots, q-1\}^n$ of weight $q^n h$ (recall that $U(\mathbf{x}) = 1/q^n$ for every \mathbf{x} in $\{0, 1, \dots, q-1\}^n$) to obtain a distribution D' , that is,

$$D' \stackrel{\text{def}}{=} \frac{1}{1 + q^n h} D_1 + \frac{q^n h}{1 + q^n h} U.$$

Then D' is non-negative at every point in the domain and D' has all its first k -level generalized Fourier coefficients equal to zero. Thus D' is a k -wise independent distribution by Corollary 3.3. Furthermore, by Lemma 3.5,

$$\sum_{\mathbf{x}} |D_1(\mathbf{x}) - D'(\mathbf{x})| \leq 2q^n h \leq 2q^{k/2} \sum_{0 < \text{wt}(\mathbf{a}) \leq k} |\hat{D}^{\text{OF}}(\mathbf{a})|.$$

By Parseval's equality, $\sum_{\mathbf{x}} |D(\mathbf{x}) - D_1(\mathbf{x})|^2 = \frac{1}{q^n} \sum_{0 < \text{wt}(\mathbf{a}) \leq k} |\hat{D}^{\text{OF}}(\mathbf{a})|^2$. Combining this with Cauchy-Schwarz inequality yields

$$\sum_{\mathbf{x}} |D(\mathbf{x}) - D_1(\mathbf{x})| \leq \sqrt{\sum_{0 < \text{wt}(\mathbf{a}) \leq k} |\hat{D}^{\text{OF}}(\mathbf{a})|^2}.$$

Now the distance between D and k -wise independence can be upper bounded as

$$\begin{aligned} \Delta(D, \mathcal{D}_{\text{kwi}}) &\leq \Delta(D, D') \\ &= \frac{1}{2} \sum_{\mathbf{x}} |D(\mathbf{x}) - D'(\mathbf{x})| \\ &\leq \frac{1}{2} \sum_{\mathbf{x}} |D(\mathbf{x}) - D_1(\mathbf{x})| + \frac{1}{2} \sum_{\mathbf{x}} |D_1(\mathbf{x}) - D'(\mathbf{x})| \quad (\text{by the triangle inequality}) \\ &\leq \frac{1}{2} \sqrt{\sum_{0 < \text{wt}(\mathbf{a}) \leq k} |\hat{D}^{\text{OF}}(\mathbf{a})|^2} + q^{k/2} \sum_{0 < \text{wt}(\mathbf{a}) \leq k} |\hat{D}^{\text{OF}}(\mathbf{a})| \\ &= O(q^{k/2}) \sum_{0 < \text{wt}(\mathbf{a}) \leq k} |\hat{D}^{\text{OF}}(\mathbf{a})|. \end{aligned}$$

We thus prove the following theorem

Theorem 3.6. *Let D be a distribution over $\{0, 1, \dots, q-1\}^n$, then*

$$\Delta(D, \mathcal{D}_{\text{kwi}}) \leq O(q^{k/2}) \sum_{0 < \text{wt}(\mathbf{a}) \leq k} |\hat{D}^{\text{OF}}(\mathbf{a})|. \quad (4)$$

In particular,

$$\Delta(D, \mathcal{D}_{\text{kwi}}) \leq O(q^{k/2}) M(n, k, q) \max_{0 < \text{wt}(\mathbf{a}) \leq k} |\hat{D}^{\text{OF}}(\mathbf{a})|.$$

Remark 3.7. One may try to generalize the approach of discrete orthogonal polynomials to the non-uniform k -wise independence as well. However, this seems to require some additional new ideas and we leave it as an interesting open problem. To see the obstacle, consider the simplest one-dimensional case and let $p(x)$, for every $x \in \{0, 1, \dots, q-1\}$, be the non-uniform marginal probabilities. We need to find a complete set of orthonormal functions $\{\chi_a^{\text{OF}}(x)\}_{a=0}^{q-1}$. On the one hand, the constraint $\hat{D}^{\text{OF}}(0) = 1$ for every distribution D (so that the ‘‘cut and paste’’ method may apply) requires that $\chi_0^{\text{OF}}(x) = 1$ for every $x \in \{0, 1, \dots, q-1\}$; on the other hand, if we stick to the characterization that $D = p$ if and only if all the non-zero Fourier coefficients of D vanish, then combining this with the orthonormality of $\{\chi_a^{\text{OF}}(x)\}_{a=0}^{q-1}$ yields that $\chi_0^{\text{OF}}(x) = qp(x)$ for every x . Clearly only the uniform distribution $p(x) = 1/q$ can satisfy both conditions.

3.3 Testing algorithm analysis

Since the bound in Theorem 3.6 is slightly weaker than the bound in Theorem 1.1, we will not give a detailed analysis of the testing algorithm based on orthogonal polynomials. In fact, by combining Fact 3.1 with the proof of Fact 2.8, it is easy to see that for any $0 \leq \delta \leq 1$ and any non-zero vector \mathbf{a} of weight at most k , if $\Delta(D, \mathcal{D}_{\text{kwi}}) \leq \delta$, then $\left| \hat{D}^{\text{OF}}(\mathbf{a}) \right| \leq q^{3/2} \delta$. We thus have the following theorem

Theorem 3.8. *There is an algorithm that tests the k -wise independence over $\{0, 1, \dots, q-1\}^n$ with query complexity $O\left(\frac{q^{k+2}M(n,k,q)^2}{\epsilon^2} \log(M(n,k,q))\right)$ and time complexity $O\left(\frac{q^{k+2}M(n,k,q)^3}{\epsilon^2} \log(M(n,k,q))\right)$ and satisfies the following: for any distribution D over Σ^n , if $\Delta(D, \mathcal{D}_{\text{kwi}}) \leq \frac{\epsilon}{3q^{(k+3)/2}M(n,k,q)}$, then with probability at least $2/3$, the algorithm accepts; if $\Delta(D, \mathcal{D}_{\text{kwi}}) > \epsilon$, then with probability at least $2/3$, the algorithm rejects.*

4 Uniform k -wise Independence

We now give another proof of Theorem 1.1 based on the standard Fourier transform. The advantage of this approach is twofold: first it gives slightly better bound; second and more importantly, the construction of a k -wise independent distribution from an input distribution is explicit and this enables us to generalize it the non-uniform case. For ease of exposition, we start from the simplest case: when the domain is a prime field.

4.1 Warm-up: Distributions over \mathbb{Z}_p^n

We begin our study with testing k -wise independent distributions when the alphabet size is a prime. Our main result is that in this case the distance between a distribution and k -wise independence can be upper bounded by the sum of the biases (to be defined later) of the distribution, slightly generalizing an idea of Alon, Goldreich and Mansour [4] that they applied to the binary field case.

Let D be a discrete distribution over \mathbb{Z}_p^n , where p is a prime number.

Definition 4.1. Let $\mathbf{a} \in \mathbb{Z}_p^n$ be a non-zero vector. We say D is *unbiased* over \mathbf{a} if $P_{\mathbf{a},\ell}^D = 1/p$ for every $0 \leq \ell \leq p-1$. The $\text{MaxBias}_D(\mathbf{a})$ of a distribution D is defined to be $\text{MaxBias}_D(\mathbf{a}) \stackrel{\text{def}}{=} \max_{0 \leq j < p} P_{\mathbf{a},j}^D - \frac{1}{p}$.

Note that the MaxBias is non-negative for any distribution. It is well-known that when p is prime, the Fourier coefficient $\hat{D}(\mathbf{a})$ of a distribution D over \mathbb{Z}_p^n as defined by (3) is zero if and only if $P_{\mathbf{a},j} = 1/p$ for every $0 \leq j \leq p-1$. Combining this with the fact that D is unbiased over \mathbf{a} if and only if $\text{MaxBias}_D(\mathbf{a})$ is zero, we thus have the following simple characterization of k -wise independence in terms of MaxBias .

Proposition 4.2. *D is k -wise independent if and only if for all non-zero $\mathbf{a} \in \mathbb{Z}_p^n$ with $\text{wt}(\mathbf{a}) \leq k$, $\text{MaxBias}_D(\mathbf{a}) = 0$.*

We say two non-zero vectors \mathbf{a} and \mathbf{b} are *linearly dependent* if there exists some $c \in \mathbb{Z}_p^*$ such that $\mathbf{b} = c\mathbf{a}$ and *linearly independent* if they are not linearly dependent.

Claim 4.3. *If \mathbf{a} and \mathbf{b} are linearly dependent, then $\text{MaxBias}_D(\mathbf{a}) = \text{MaxBias}_D(\mathbf{b})$.*

Proof. Suppose $\text{MaxBias}_D(\mathbf{a})$ is attained at j , i.e., $\text{MaxBias}_D(\mathbf{a}) = P_{\mathbf{a},j} - \frac{1}{p}$. Then $\text{MaxBias}_D(\mathbf{b}) \geq P_{\mathbf{b},cj \pmod p} - \frac{1}{p} = P_{\mathbf{a},j} - \frac{1}{p} = \text{MaxBias}_D(\mathbf{a})$. Similarly, since c^{-1} exists, we also have $\text{MaxBias}_D(\mathbf{a}) \geq \text{MaxBias}_D(\mathbf{b})$. It follows that $\text{MaxBias}_D(\mathbf{a}) = \text{MaxBias}_D(\mathbf{b})$. \square

For each $\mathbf{a} \in \mathbb{Z}_p^n$ there are $p - 2$ other vectors (namely, by taking $c = 2, \dots, p - 1$) that are linearly dependent with \mathbf{a} .

Lemma 4.4. *Let $\mathbf{a}, \mathbf{b} \in \mathbb{Z}_p^n$ be two non-zero, linearly independent vectors, then for any $0 \leq r_a, r_b \leq p - 1$,*

$$\Pr_{\mathbf{x} \in \mathbb{Z}_p^n} \left[\sum_{i=1}^n a_i x_i \equiv r_a \pmod{p} \wedge \sum_{i=1}^n b_i x_i \equiv r_b \pmod{p} \right] = \frac{1}{p^2}$$

Proof. This follows from the well-known fact that the number of solutions to a system of 2 linearly independent linear equations over \mathbb{Z}_p in n variables is p^{n-2} , independent of the vectors of free coefficients. \square

Definition 4.5 (Strong Orthogonality). Let \mathbf{a} and \mathbf{b} be two non-zero vectors in \mathbb{Z}_p^n . We say \mathbf{a} is *strongly orthogonal* to \mathbf{b} if $U_{\mathbf{a},j}$ is *unbiased* over \mathbf{b} for every $0 \leq j \leq p - 1$. That is, $\Pr_{\mathbf{X} \sim U_{\mathbf{a},j}}[\mathbf{b} \cdot \mathbf{X} \equiv \ell \pmod{p}] = 1/p$, for all $0 \leq j, \ell \leq p - 1$.

Corollary 4.6. *Let \mathbf{a} be a non-zero vector in \mathbb{Z}_p^n and \mathbf{b} be another non-zero vector that is linearly independent of \mathbf{a} . Then \mathbf{a} is strongly orthogonal to \mathbf{b} .*

Proof. Clearly we have $|S_{\mathbf{a},j}| = p^{n-1}$ for all non-zero \mathbf{a} and all j . Then by Lemma 4.4, the p^{n-1} points in $S_{\mathbf{a},j}$ are uniformly distributed over each of the p sets $S_{\mathbf{b},\ell}$, $0 \leq \ell \leq p - 1$. \square

Now we are ready to prove the following main result of this section.

Theorem 4.7. *Let D be a distribution over \mathbb{Z}_p^n . Then $\Delta(D, \mathcal{D}_{kwi}) \leq \frac{p}{p-1} \sum_{0 < \text{wt}(\mathbf{a}) \leq k} \text{MaxBias}_D(\mathbf{a})$.*

Note that this generalizes the result of [4] for $\text{GF}(2)$ to $\text{GF}(p)$ for any prime p . When $p = 2$, we recover the same (implicit) bound there (our MaxBias is exactly half of their “Bias”).

We first give a brief overview of the proof. We are going to prove Theorem 4.7 by constructing a k -wise independent distribution that is close to D . Generalizing the approach in [4], we start from D , step by step, zeroing-out $\text{MaxBias}_D(\mathbf{a})$ for every non-zero vector \mathbf{a} of weight at most k . By Proposition 4.2, the resulting distribution will be a k -wise independent one. At each step, we pick any \mathbf{a} with $\text{MaxBias}_D(\mathbf{a}) > 0$. To zero-out $\text{MaxBias}_D(\mathbf{a})$, we apply a convex combination between the old distribution and some carefully chosen distribution to get a new distribution. By the strong orthogonality between linearly independent vectors (c.f. Corollary 4.6), if for every $0 \leq j \leq q - 1$, we mix with D the uniform distribution over all strings in $S_{\mathbf{a},j}$ with some appropriate weight (this weight can be zero), we will not only zero-out the MaxBias at \mathbf{a} but also guarantee that for any \mathbf{b} that is linearly independent from \mathbf{a} , $\text{MaxBias}_D(\mathbf{b})$ is not going to increase (therefore the MaxBias of all zeroed-out vectors will remain zero throughout the correcting steps). This enables us to repeat the zeroing-out process for all other vectors of weight at most k and finally obtain a k -wise independent distribution.

Proof of Theorem 4.7. First we partition all the non-zero vectors of weight at most k into families of linearly dependent vectors, say F_1, F_2, \dots , etc. Pick any vector \mathbf{a} from F_1 . If $\text{MaxBias}_D(\mathbf{a}) = 0$, we move on to the next family of vectors. Now suppose $\text{MaxBias}_D(\mathbf{a}) > 0$, and without loss of generality, assume that $P_{\mathbf{a},0} \leq P_{\mathbf{a},1} \leq \dots \leq P_{\mathbf{a},p-1}$. Let $\epsilon_j = P_{\mathbf{a},j} - \frac{1}{p}$. Since $\sum_{j=0}^{p-1} P_{\mathbf{a},j} = 1$, we have $\epsilon_0 + \dots + \epsilon_{p-1} = 0$. Also note that $\text{MaxBias}_D(\mathbf{a}) = \epsilon_{p-1}$.

Now we define a new distribution D' as

$$D' = \frac{1}{1 + \epsilon} D + \frac{\epsilon_{p-1} - \epsilon_0}{1 + \epsilon} U_{\mathbf{a},0} + \dots + \frac{\epsilon_{p-1} - \epsilon_{p-2}}{1 + \epsilon} U_{\mathbf{a},p-2},$$

where $\epsilon = (\epsilon_{p-1} - \epsilon_0) + \dots + (\epsilon_{p-1} - \epsilon_{p-2})$. Now by the triangle inequality,

$$\begin{aligned}\Delta(D, D') &\leq \epsilon = (\epsilon_{p-1} - \epsilon_0) + \dots + (\epsilon_{p-1} - \epsilon_{p-2}) \\ &= p\epsilon_{p-1} = p\text{MaxBias}_D(\mathbf{a}).\end{aligned}$$

It is easy to check that $\text{MaxBias}_{D'}(\mathbf{a}) = 0$, since for every $0 \leq j \leq p-1$,

$$\begin{aligned}P_{\mathbf{a},j}^{D'} &= \frac{1}{1+\epsilon} P_{\mathbf{a},j}^D + \frac{\epsilon_{p-1} - \epsilon_j}{1+\epsilon} \\ &= \frac{1}{1+\epsilon} (P_{\mathbf{a},j}^D + \epsilon_{p-1} - \epsilon_j) \\ &= \frac{1}{1+\epsilon} (\epsilon_{p-1} + \frac{1}{p}) \\ &= \frac{1}{p} \quad (\text{because } \epsilon = p\epsilon_{p-1}).\end{aligned}$$

Moreover, due to Corollary 4.6 and the fact that $U_{\mathbf{a},j}$ is unbiased over \mathbf{b} for every $0 \leq j < p$, we have for any vector \mathbf{b} that is not in the same family with \mathbf{a} (i.e., in F_2, \dots , etc.),

$$\text{MaxBias}_{D'}(\mathbf{b}) = \frac{1}{1+\epsilon} \text{MaxBias}_D(\mathbf{b}) \leq \text{MaxBias}_D(\mathbf{b}).$$

In particular, if $\text{MaxBias}_D(\mathbf{b})$ is zero, then after zeroing-out the bias at \mathbf{a} , $\text{MaxBias}_{D'}(\mathbf{b})$ remains zero.

Note that once we zero-out the MaxBias over \mathbf{a} , then by Claim 4.3, the biases over all other $p-2$ vectors in F_1 vanish as well (that is, we only need to perform one zeroing-out for the $p-1$ vectors in the same family). Repeating this process for all other families of vectors, we reach a distribution D_f that is unbiased over all vectors of weight at most k . By Proposition 4.2 D_f is k -wise independent and the distance between D_f and D is at most as claimed in the theorem. \square

4.2 Distributions over \mathbb{Z}_q^n

We now address the main problem of this section, that is, robust characterization of k -wise independent distributions over domains of the form \mathbb{Z}_q^n when q is composite. A straightforward application of the method for the prime fields case breaks down for general commutative rings because the strongly orthogonal condition in Corollary 4.6 does not hold, even if the two vectors are linearly independent. Recall that a distribution D over \mathbb{Z}_q^n is k -wise independent if and only if for all non-zero vectors \mathbf{a} of weight at most k , $\hat{D}(\mathbf{a}) = 0$. Our main technical result in this section is to show, analogous to the prime field case, for a distribution D over the general domain \mathbb{Z}_q^n , the following holds: for every non-zero vector \mathbf{a} of weight at most k , there exists a (small-weight) distribution such that mixing it with D zeroes-out the Fourier coefficient at \mathbf{a} and does not increase the Fourier coefficient at any other vector.

Unless stated otherwise, all arithmetic operations in this section are performed modulo q ; for instance, we write $\mathbf{a} = \mathbf{b}$ to mean that $a_i \equiv b_i \pmod{q}$ for each $1 \leq i \leq n$.

Definition 4.8 (Prime Vectors). Let $\mathbf{a} = (a_1, \dots, a_n)$ be a non-zero vector in \mathbb{Z}_q^n . \mathbf{a} is called a *prime vector* if $\gcd(a_1, \dots, a_n) = 1$. If \mathbf{a} is a prime vector, then we refer to the set of vectors $\{2\mathbf{a}, \dots, (q-1)\mathbf{a}\}$ (note that all these vectors are distinct) as the *multiples* of \mathbf{a} . A prime vector and its *multiples* are collectively referred to as a *family of vectors*.

Note that families of vectors do *not* form a partition of the set of all the vectors. For example when $n = 2$ and $q = 6$, vector $(4, 0)$ is a *multiple* of both $(1, 0)$ and $(2, 3)$, but the latter two are not *multiples* of each other. Furthermore, there can be more than one prime vector in a family of vectors, e.g., for $q = 6$ again, $(2, 3)$ and $(4, 3)$ are *multiples* while they are both prime vectors.

Recall that we use $S_{\mathbf{a},j}$ to denote the set $\{\mathbf{x} \in \mathbb{Z}_q^n : \sum_{i=1}^n a_i x_i \equiv j \pmod{q}\}$.

Proposition 4.9. *If \mathbf{a} is a prime vector, then $|S_{\mathbf{a},j}| = q^{n-1}$ for any $0 \leq j \leq q - 1$.*

Proof. Since $\gcd(a_1, \dots, a_n) = 1$, there exist integers z_1, \dots, z_n such that $a_1 z_1 + \dots + a_n z_n = 1$. Note that for any $\mathbf{z} \in \mathbb{Z}_q^n$ the map $h_{\mathbf{z}}(\mathbf{x}) = \mathbf{x} + \mathbf{z}$ is injective. Now if $\mathbf{x} \in S_{\mathbf{a},0}$, then $h_{\mathbf{z}}(\mathbf{x}) = (x_1 + z_1, \dots, x_n + z_n) \in S_{\mathbf{a},1}$. Therefore $|S_{\mathbf{a},0}| \leq |S_{\mathbf{a},1}|$. Similarly we have $|S_{\mathbf{a},1}| \leq |S_{\mathbf{a},2}| \leq \dots \leq |S_{\mathbf{a},q-1}| \leq |S_{\mathbf{a},0}|$. Since the sets $S_{\mathbf{a},0}, \dots, S_{\mathbf{a},q-1}$ form a partition of \mathbb{Z}_q^n , it follows that $|S_{\mathbf{a},0}| = |S_{\mathbf{a},1}| = \dots = |S_{\mathbf{a},q-1}| = q^{n-1}$. \square

4.2.1 Linear Systems of Congruences

A *linear system of congruences* is a set of linear modular arithmetic equations in some variables. We will be particularly interested in the case when all modular arithmetic equations are modulo q . If the number of variables is k , then a solution to the system of congruences is a vector in \mathbb{Z}_q^k . Two solutions \mathbf{x}, \mathbf{x}' in \mathbb{Z}_q^k are *congruent* to each other if $\mathbf{x} = \mathbf{x}'$ (i.e. $x_i \equiv x'_i \pmod{q}$ for every $1 \leq i \leq k$) and *incongruent* otherwise.

We record some useful results on linear systems of congruences in this section. For more on this, the interested reader is referred to [22] and [39]. These results will be used in the next section to show some important orthogonality properties of vectors in \mathbb{Z}_q^n . In this section, all matrices are integer-valued. Let M be a $k \times n$ matrix with $k \leq n$. The *greatest divisor* of M is the greatest common divisor (gcd) of the determinants of all $k \times k$ sub-matrices of M . M is a *prime matrix* if the greatest divisor of M is 1.

Lemma 4.10 ([39]). *Let M be a $(k + 1) \times n$ matrix. If the sub-matrix consisting of the first k rows of M is a prime matrix and M has greatest divisor d , then there exist integers u_1, \dots, u_k such that for every $1 \leq j \leq n$,*

$$u_1 M_{1,j} + u_2 M_{2,j} + \dots + u_k M_{k,j} \equiv M_{k+1,j} \pmod{d}.$$

Consider the following system of linear congruent equations:

$$\begin{cases} M_{1,1}x_1 + M_{1,2}x_2 + \dots + M_{1,n}x_n \equiv M_{1,n+1} \pmod{q} \\ \vdots \\ M_{k,1}x_1 + M_{k,2}x_2 + \dots + M_{k,n}x_n \equiv M_{k,n+1} \pmod{q}. \end{cases} \quad (5)$$

Let M denote the $k \times n$ matrix consisting of the coefficients of the linear system of equations and let \tilde{M} denote the corresponding augmented matrix of M , that is, the $k \times (n + 1)$ matrix with one extra column consisting of the free coefficients.

Definition 4.11. Let M be the coefficient matrix of (5) and \tilde{M} be the augmented matrix of M . Suppose $k < n$ so that system (5) is a defective system of equations. Define Y_k, Y_{k-1}, \dots, Y_1 , respectively, to be the greatest common divisors of the determinants of all the $k \times k, (k - 1) \times (k - 1), \dots, 1 \times 1$, respectively, sub-matrices of M . Analogously define Z_k, Z_{k-1}, \dots, Z_1 for the augmented matrix \tilde{M} . Also we set $Y_0 = 1$ and $Z_0 = 1$. Finally let $s = \prod_{j=1}^k \gcd(q, \frac{Y_j}{Y_{j-1}})$ and $t = \prod_{j=1}^k \gcd(q, \frac{Z_j}{Z_{j-1}})$.

The following Theorem of Smith gives the necessary and sufficient conditions for a system of congruent equations to have solutions.

Theorem 4.12 ([39]). *If $k < n$, then the (defective) linear system of congruences (5) has solutions if and only if $s = t$. Moreover, if this condition is met, the number of incongruent solutions is sq^{n-k} .*

4.2.2 Weak Orthogonality between Families of Vectors

To generalize the proof idea of the GF(2) case (and also the prime field case studied in Section 4.1) to commutative rings \mathbb{Z}_q for arbitrary q , it seems crucial to relax the requirement that linearly independent vectors are strongly orthogonal. Rather, we introduce the notion of weak orthogonality between a pair of vectors.

Definition 4.13 (Weak Orthogonality). Let \mathbf{a} and \mathbf{b} be two non-zero vectors in \mathbb{Z}_q^n . We say \mathbf{a} is *weakly orthogonal* to \mathbf{b} if for all $0 \leq j \leq q-1$, $\hat{U}_{\mathbf{a},j}(\mathbf{b}) = 0$.

Remark 4.14. Note that in general weak orthogonality is not a symmetric relation, that is, \mathbf{a} is weakly orthogonal to \mathbf{b} does not necessarily imply that \mathbf{b} is weakly orthogonal to \mathbf{a} . Also note that strong orthogonality implies weak orthogonality while the converse is not necessarily true. In particular, strong orthogonality does not hold in general for linearly independent vectors in \mathbb{Z}_q^n . However, for our purpose of constructing k -wise independent distributions, weak orthogonality between pairs of vectors suffices.

The following lemma is the basis of our upper bound on the distance between a distribution and k -wise independence. This lemma enables us to construct a small-weight distribution using an appropriate convex combination of $\{U_{\mathbf{a},j}\}_{j=0}^{q-1}$, which on the one hand zeroes-out all the Fourier coefficients at \mathbf{a} and its *multiple* vectors, on the other hand has zero Fourier coefficient at all other vectors. The proof of the Lemma 4.15 relies crucially on the results in Section 4.2.1 about linear system of congruences.

Lemma 4.15. *Let \mathbf{a} be a non-zero prime vector and \mathbf{b} any non-zero vector that is not a multiple of \mathbf{a} . Then \mathbf{a} is weakly orthogonal to \mathbf{b} .*

Proof. Consider the following system of linear congruences:

$$\begin{cases} a_1x_1 + a_2x_2 + \cdots + a_nx_n \equiv a_0 \pmod{q} \\ b_1x_1 + b_2x_2 + \cdots + b_nx_n \equiv b_0 \pmod{q}. \end{cases} \quad (6)$$

Following our previous notation, let $M = \begin{bmatrix} a_1 & a_2 & \cdots & a_n \\ b_1 & b_2 & \cdots & b_n \end{bmatrix}$ and $\tilde{M} = \begin{bmatrix} a_1 & a_2 & \cdots & a_n & a_0 \\ b_1 & b_2 & \cdots & b_n & b_0 \end{bmatrix}$. Since \mathbf{a} is a prime vector, $Y_1 = Z_1 = 1$. We next show that if \mathbf{b} is not a *multiple* of \mathbf{a} , then Y_2 can not be a multiple of q .

Claim 4.16. *Let \mathbf{a} be a prime vector and let $M = \begin{bmatrix} a_1 & a_2 & \cdots & a_n \\ b_1 & b_2 & \cdots & b_n \end{bmatrix}$. The determinants of all 2×2 sub-matrices of M are congruent to 0 modulo q if and only if \mathbf{a} and \mathbf{b} are multiple vectors.*

Proof. If \mathbf{a} and \mathbf{b} are *multiple* vectors, then it is clear that the determinants of all the sub-matrices are congruent to 0 modulo q . For the *only if* direction, all we need to prove is that $\mathbf{b} = c\mathbf{a}$ for some integer c . First suppose that the determinants of all 2×2 sub-matrices of M are 0. Then it follows that $\frac{b_1}{a_1} = \cdots = \frac{b_n}{a_n} = c$. If c is an integer then we are done. If c is not an integer, then $c = \frac{u}{v}$, where u, v are integers and $\gcd(u, v) = 1$. But this implies $v|a_i$ for every $1 \leq i \leq n$, contradicting our assumption that \mathbf{a} is a prime vector. Now if not all of the determinants are 0, it must be the case that the greatest common divisor of the determinants of all 2×2 sub-matrices, say d' , is a multiple of q . By Lemma 4.10, there is an integer c such that $ca_i \equiv b_i \pmod{d'}$ for every $1 \leq i \leq n$. Consequently, $b_i \equiv ca_i \pmod{q}$ for every i and hence \mathbf{b} is a *multiple* of \mathbf{a} . \square

Let $d = \gcd(q, Y_2)$. Clearly $1 \leq d \leq q$ and according to Claim 4.16, $d \neq q$ so $d|q$. Applying Theorem 4.12 with $k = 2$ to (6), the two linear congruences are solvable if and only if $d = \gcd(q, Y_2) = \gcd(q, Z_2)$. If this is the case, the total number of incongruent solutions is dq^{n-2} . Furthermore, if we let h denote the greatest common divisor of the determinants of all 2×2 sub-matrices of \tilde{M} , then $d|h$. By Lemma 4.10, there is an integer u such that $b_0 \equiv ua_0 \pmod{h}$. It follows that $d|(b_0 - ua_0)$. Let us consider a fixed a_0 and write $\ell_0 = ua_0 \pmod{d}$. Since \mathbf{a} is a prime vector, by Proposition 4.9, there are in total q^{n-1} solutions to (6). But for any fixed b_0 that has solutions to (6), there must be dq^{n-2} solutions to (6) and in addition $d|q$. Since there are exactly q/d b_0 's in $\{0, \dots, q-1\}$, we conclude that (6) has solutions for b_0 if and only if $b_0 = \ell_0 + d\ell$, where ℓ_0 is some constant and $\ell = 0, \dots, \frac{q}{d} - 1$. Finally we have

$$\begin{aligned} \hat{U}_{\mathbf{a},j}(\mathbf{b}) &= \sum_{\mathbf{x} \in \mathbb{Z}_q^n} U_{\mathbf{a},j}(\mathbf{x}) e^{\frac{2\pi i}{q} \mathbf{b} \cdot \mathbf{x}} = \frac{1}{q^{n-1}} \sum_{\mathbf{a} \cdot \mathbf{x} \equiv j \pmod{q}} e^{\frac{2\pi i}{q} \mathbf{b} \cdot \mathbf{x}} \\ &= \frac{d}{q} \sum_{b_0: b_0 = \ell_0 + d\ell} e^{\frac{2\pi i}{q} b_0} = 0. \end{aligned} \quad (\text{by Fact 2.2})$$

This finishes the proof of Lemma 4.15. \square

4.2.3 Correcting the Fourier Coefficients of Multiple Vectors

Now we show how to zero-out a distribution's Fourier coefficient at every vector in a family. Let D be a distribution over \mathbb{Z}_q^n . By (3), for every $1 \leq \ell \leq q-1$, the Fourier coefficient of a vector $\ell \mathbf{a}$ can be rewritten as $\hat{D}(\ell \mathbf{a}) = \sum_{j=0}^{q-1} P_{\mathbf{a},j} e^{\frac{2\pi i}{q} \ell j}$. Recall that $\text{MaxBias}(\mathbf{a}) = \max_{0 \leq j \leq q-1} P_{\mathbf{a},j} - \frac{1}{q}$.

Claim 4.17. *We have that $\text{MaxBias}(\mathbf{a}) \leq \frac{1}{q} \sum_{\ell=1}^{q-1} |\hat{D}(\ell \mathbf{a})|$.*

Proof. Since $\hat{D}(\ell \mathbf{a}) = \sum_{j=0}^{q-1} P_{\mathbf{a},j} e^{\frac{2\pi i}{q} \ell j}$, by the inverse Fourier transform (2), for every $0 \leq j \leq q-1$,

$$P_{\mathbf{a},j} = \frac{1}{q} \sum_{\ell=0}^{q-1} \hat{D}(\ell \mathbf{a}) e^{-\frac{2\pi i}{q} \ell j}.$$

Since $\hat{D}(0) = 1$, we have for every $0 \leq j \leq q-1$,

$$\begin{aligned} \left| P_{\mathbf{a},j} - \frac{1}{q} \right| &= \frac{1}{q} \left| \sum_{\ell=1}^{q-1} \hat{D}(\ell \mathbf{a}) e^{-\frac{2\pi i}{q} \ell j} \right| \\ &\leq \frac{1}{q} \sum_{\ell=1}^{q-1} |\hat{D}(\ell \mathbf{a}) e^{-\frac{2\pi i}{q} \ell j}| \leq \frac{1}{q} \sum_{\ell=1}^{q-1} |\hat{D}(\ell \mathbf{a})|. \end{aligned} \quad \square$$

Now we are ready to prove the main theorem of this section.

Theorem 1.1. *Let D be a distribution over \mathbb{Z}_q^n , then ⁹*

$$\Delta(D, \mathcal{D}_{\text{kwi}}) \leq \sum_{0 < \text{wt}(\mathbf{a}) \leq k} |\hat{D}(\mathbf{a})|.$$

⁹ It is easy to verify that the same bound holds for prime field case if we transform the bound in MaxBias there into a bound in terms of Fourier coefficients. Conversely we can equivalently write the bound of the distance from k -wise independence in terms of MaxBias at *prime vectors*. However, we believe that stating the bound in terms of Fourier coefficients is more natural and generalizes more easily.

Test-Uniform-KWI(D, k, q, ϵ)

1. Sample D independently $O\left(\frac{q^2 M(n, k, q)^2}{\epsilon^2} \log(M(n, k, q))\right)$ times to obtain a set Q
2. For every non-zero vector \mathbf{a} of weight at most k , use Q to estimate $\hat{D}(\mathbf{a})$
3. If $\max_{\mathbf{a}} \left| \hat{D}(\mathbf{a}) \right| \leq \frac{2\epsilon}{3M(n, k, q)}$, return **“Accept”**; else return **“Reject”**

Figure 2: Algorithm for testing if a distribution D over Σ^n is uniform k -wise independent.

In particular, $\Delta(D, \mathcal{D}_{\text{kwi}}) \leq M(n, k, q) \max_{0 < \text{wt}(\mathbf{a}) \leq k} \left| \hat{D}(\mathbf{a}) \right|$.

Proof. Let \mathbf{a} be a prime vector and $\hat{D}(\mathbf{a}), \hat{D}(2\mathbf{a}), \dots, \hat{D}((q-1)\mathbf{a})$ be the Fourier coefficients of \mathbf{a} and all the multiples of \mathbf{a} . Now construct a new distribution D' over \mathbb{Z}_q^n as

$$D' = \frac{1}{1+\epsilon} D + \frac{1}{1+\epsilon} \sum_{j=0}^{q-1} v(j) U_{\mathbf{a}, j},$$

where $\epsilon = \sum_{j=0}^{q-1} v(j)$ and $\{v(j)\}_{j=0}^{q-1}$ are a set of non-negative real numbers that will be specified later. It is easy to check that D' is indeed a distribution. Moreover, by Lemma 4.15 and linearity of the Fourier transform, for every \mathbf{b} that is not a multiple of \mathbf{a} ,

$$\left| \hat{D}'(\mathbf{b}) \right| = \frac{1}{1+\epsilon} \left| \hat{D}(\mathbf{b}) \right| \leq \left| \hat{D}(\mathbf{b}) \right|.$$

Without loss of generality, assume that $P_{\mathbf{a}, 0} \leq \dots \leq P_{\mathbf{a}, q-1}$. That is, $\text{MaxBias}(\mathbf{a}) = P_{\mathbf{a}, q-1} - \frac{1}{q}$. If we choose $v(j) = P_{\mathbf{a}, q-1} - P_{\mathbf{a}, j}$, then clearly $v(j)$ is non-negative for every $0 \leq j \leq q-1$. Furthermore, by our construction $P_{\mathbf{a}, j}^{D'} = \frac{1}{q}$ for every j . Therefore by Fact 2.1, $\hat{D}'(\ell\mathbf{a}) = 0$ for every $1 \leq \ell \leq q-1$. Since $\sum_{j=0}^{q-1} P_{\mathbf{a}, j} = 1$, it follows that $\sum_{j=0}^{q-1} v(j) = q \text{MaxBias}(\mathbf{a})$. By Claim 4.17,

$$\Delta(D, D') \leq \epsilon = \sum_{j=0}^{q-1} v(j) \leq \sum_{\ell=1}^{q-1} \left| \hat{D}(\ell\mathbf{a}) \right|. \quad (7)$$

Finally observe that although some vectors are multiples of more than one prime vector (thus they belong to more than one family and appear more than once in (7)), because the distance bound in (7) is the sum of magnitudes of all the Fourier coefficients in the family, once a vector's Fourier coefficient is zeroed-out, it will not contribute to the distance bound at any later stage. This completes the proof of the theorem. \square

4.2.4 Testing Algorithm and its Analysis

We are now ready to prove the following result on testing k -wise independence over \mathbb{Z}_q^n .

Theorem 4.18. *There is an algorithm that tests the k -wise independence over Σ^n with query complexity $\tilde{O}\left(\frac{n^{2k}(q-1)^{2k}q^2}{\epsilon^2}\right)$ and time complexity $\tilde{O}\left(\frac{n^{3k}(q-1)^{3k}q^2}{\epsilon^2}\right)$ and satisfies the following: for any distribution D over Σ^n , if $\Delta(D, \mathcal{D}_{\text{kwi}}) \leq \frac{\epsilon}{3qM(n, k, q)}$, then with probability at least $2/3$, the algorithm accepts; if $\Delta(D, \mathcal{D}_{\text{kwi}}) > \epsilon$, then with probability at least $2/3$, the algorithm rejects.*

Proof. Our testing algorithm simply plugs the upper bound on distance to k -wise independence in Theorem 1.1 into the *Generic Algorithm* as shown in Fig. 1. The algorithm is described in Figure 2. For the analysis of **Test-Uniform-KWI**(D, k, q, ϵ), we simply apply Theorem 2.7 with $K = M(n, k, q)$, $A = \{\mathbf{a} \in \Sigma^n : 0 < \text{wt}(\mathbf{a}) \leq k\}$ and $\kappa = q$. To see $\kappa = q$, note that $P_{\mathbf{a},j} = 1/q$ holds for every \mathbf{a} in A and $0 \leq j \leq q - 1$ for any k -wise independent distribution. Since no (randomized) algorithm can increase the statistical difference between two distributions [37], by Fact 2.8 (more precisely, the proof of Fact 2.8), if $\Delta(D, \mathcal{D}_{\text{kwi}}) \leq \delta$, then we have $|\hat{D}(\mathbf{a})| \leq q\delta$ for every $\mathbf{a} \in A$. \square

4.3 Distributions over Product Spaces

Now we generalize the underlying domains from \mathbb{Z}_q^n to product spaces. Let $\Sigma_1, \dots, \Sigma_n$ be n finite sets. Without loss of generality, let $\Sigma_i = \{0, 1, \dots, q_i - 1\}$. In this section, we consider distributions over the product space $\Omega = \Sigma_1 \times \dots \times \Sigma_n$. For a set of integers $\{q_1, \dots, q_n\}$, denote their *least common multiple* (lcm) by $\text{lcm}(q_1, \dots, q_n)$. Let $Q \stackrel{\text{def}}{=} \text{lcm}(q_1, \dots, q_n)$ and in addition, for every $1 \leq i \leq n$, set $M_i = \frac{Q}{q_i}$. Then we can rewrite the Fourier coefficient defined in (1) as

$$\begin{aligned} \hat{D}(\mathbf{a}) &= \sum_{\mathbf{x} \in \Sigma_1 \times \dots \times \Sigma_n} D(\mathbf{x}) e^{\frac{2\pi i}{Q}(M_1 a_1 x_1 + \dots + M_n a_n x_n)} \\ &= \sum_{\mathbf{x} \in \Sigma_1 \times \dots \times \Sigma_n} D(\mathbf{x}) e^{\frac{2\pi i}{Q}(a'_1 x_1 + \dots + a'_n x_n)}, \end{aligned}$$

where $a'_i \equiv M_i a_i \pmod{Q}$ for every $1 \leq i \leq n$. This suggests that we may view D as a distribution over Σ^n with *effective alphabet size* $|\Sigma| = Q = \text{lcm}(q_1, \dots, q_n)$ and consider the following map from vectors in $\Sigma_1 \times \dots \times \Sigma_n$ to vectors in \mathbb{Z}_Q^n :

$$\mathcal{H} : (a_1, \dots, a_n) \mapsto (M_1 a_1 \pmod{Q}, \dots, M_n a_n \pmod{Q}). \quad (8)$$

Then we only need to consider the Fourier coefficients at vectors $\mathbf{a}' \stackrel{\text{def}}{=} \mathcal{H}(\mathbf{a}) = (a'_1, \dots, a'_n) \in \mathbb{Z}_Q^n$ (that is, vectors in \mathbb{Z}_Q^n whose i^{th} component is a multiple of M_i for every i). Note that in general $M = \text{lcm}(q_1, \dots, q_n)$ could be an exponentially large number and is therefore not easy to handle in practice¹⁰. However, this difficulty can be overcome by observing the following simple fact. Since we are only concerned with vectors of weight at most k , we may take different effective alphabet sizes for different index subsets of size k . For example, consider a k -subset $S = \{i_1, \dots, i_k\}$. Then the effective alphabet size of S is $|\Sigma_S| = \text{lcm}(q_{i_1}, \dots, q_{i_k})$, which is at most $\text{poly}(n)$ if we assume k is a constant and each q_i is polynomially bounded.

Our main result for distributions over product spaces is the following theorem.

Theorem 4.19. *Let D be a distribution over $\Sigma_1 \times \dots \times \Sigma_n$. Then $\Delta(D, \mathcal{D}_{\text{kwi}}) \leq \sum_{0 < \text{wt}(\mathbf{a}) \leq k} |\hat{D}(\mathbf{a})|$.*

We now sketch the proof of Theorem 4.19.

A vector $\mathbf{a} \in \Sigma_1 \times \dots \times \Sigma_n$ is a *prime vector* if $\text{gcd}(a_1, \dots, a_n) = 1$. For any integer $\ell > 0$, the ℓ -multiple of \mathbf{a} is $\ell \mathbf{a} \stackrel{\text{def}}{=} (\ell a_1 \pmod{q_1}, \dots, \ell a_n \pmod{q_n})$. Let \mathbf{a} be a prime vector. Then vectors in the set $\{2\mathbf{a}, \dots, (Q - 1)\mathbf{a}\}$ are called the *multiple vectors* of \mathbf{a} . Note that these $Q - 1$ vectors may not be all distinct.

¹⁰Recall that the testing algorithm requires estimating all the low-degree Fourier coefficients, where each Fourier coefficient is an exponential sum with M as the denominator.

The main difficulty in applying our result for distributions over \mathbb{Z}_q^n to distributions over product spaces is that the mapping in (8) is not surjective. In particular, after the mapping some families of vectors may have no prime vector in it. To handle this problem, we slightly generalize the result of weak orthogonality in Lemma 4.15 to non-prime vectors. Specifically, we say a non-zero vector \mathbf{a} (not necessarily prime) is *weakly orthogonal* to vector \mathbf{b} if $\hat{U}_{\mathbf{a},\ell}(\mathbf{b}) = 0$ for all ℓ such that $S_{\mathbf{a},\ell}$ is non-empty.

Lemma 4.20. *Let \mathbf{a} and \mathbf{b} be two vectors in \mathbb{Z}_q^n . If \mathbf{b} is not a multiple of \mathbf{a} , then vector \mathbf{a} is weakly orthogonal to \mathbf{b} .*

Proof. Clearly we only need to prove the case when \mathbf{a} is not a prime vector. Let $\tilde{\mathbf{a}}$ be any prime vector that is a multiple of \mathbf{a} and suppose $\mathbf{a} = d\tilde{\mathbf{a}}$. Now $S_{\mathbf{a},\ell}$ is non-empty only if $\ell \equiv \ell'd \pmod{q}$ for some integer ℓ' . Note that $S_{\mathbf{a},\ell'd} = \cup_{j:jd \equiv \ell'd \pmod{q}} S_{\tilde{\mathbf{a}},j}$. Since the sets $\{S_{\tilde{\mathbf{a}},j}\}_{j=0}^{q-1}$ are pairwise disjoint, it follows that $U_{\mathbf{a},\ell'd} = \frac{1}{\gcd(d,q)} \sum_{j:jd \equiv \ell'd \pmod{q}} U_{\tilde{\mathbf{a}},j}$, where $\gcd(d,q)$ is the number of incongruent j 's satisfying $jd \equiv \ell'd \pmod{q}$. Now by Lemma 4.15, if \mathbf{b} is not a multiple of $\tilde{\mathbf{a}}$, then $\hat{U}_{\tilde{\mathbf{a}},j}(\mathbf{b}) = 0$ for every j . It follows that $\hat{U}_{\mathbf{a},\ell d}(\mathbf{b}) = 0$. \square

Note that for any integer $\ell > 0$ and every $1 \leq i \leq n$, $\ell a_i \equiv b_i \pmod{q_i}$ if and only if $\ell a_i m_i \equiv b_i m_i \pmod{Q}$, hence the map \mathcal{H} preserves the *multiple* relationship between vectors. Now Lemma 4.20 implies that if we map the vectors in $\Sigma_1 \times \cdots \times \Sigma_n$ to vectors in \mathbb{Z}_Q^n as defined in (8), then we can perform the same zeroing-out process as before: for each family of vectors, zero-out all the Fourier coefficients at the vectors in this family using a mixture of uniform distributions without increasing the magnitudes of the Fourier coefficients everywhere else. This will end up with a k -wise independent distribution over the product space $\Sigma_1 \times \cdots \times \Sigma_n$.

Next we bound the total weight required to zero-out a family of vectors. Let S be any k -subset of $[n]$. Without loss of generality, we may take $S = [k]$. Let $q_S = \text{lcm}(q_1, \dots, q_k)$ and let $m_i = \frac{q_S}{q_i}$ for each $1 \leq i \leq k$. Let $\mathbf{a} \in \Sigma_1 \times \cdots \times \Sigma_n$ be a prime vector whose support is contained in $[k]$. Then

$$\begin{aligned} \hat{D}(\mathbf{a}) &= \sum_{\mathbf{x} \in \Sigma_1 \times \cdots \times \Sigma_k} D_S(\mathbf{x}) e^{2\pi i \left(\frac{a_1 x_1}{q_1} + \cdots + \frac{a_k x_k}{q_k} \right)} \\ &= \sum_{\mathbf{x} \in \Sigma_1 \times \cdots \times \Sigma_k} D_S(\mathbf{x}) e^{\frac{2\pi i}{q_S} (m_1 a_1 x_1 + \cdots + m_k a_k x_k)} \\ &= \sum_{\mathbf{x} \in \Sigma_1 \times \cdots \times \Sigma_k} D_S(\mathbf{x}) e^{\frac{2\pi i}{q_S} (a'_1 x_1 + \cdots + a'_k x_k)}, \end{aligned}$$

where, as before, we define $\mathbf{a}' = (a'_1, \dots, a'_k)$ with $a'_i = m_i a_i \pmod{q_S}$ for $1 \leq i \leq k$.

Let $d = \gcd(m_1 a_1 \pmod{q_S}, \dots, m_k a_k \pmod{q_S}) = \gcd(a'_1, \dots, a'_k)$ and set $S_{\mathbf{a}',j} = \{\mathbf{x} \in \Sigma_1 \times \cdots \times \Sigma_k : a'_1 x_1 + \cdots + a'_k x_k \equiv j \pmod{q_S}\}$. Clearly $S_{\mathbf{a}',j}$ is non-empty only if $d|j$.

Claim 4.21. *Let \mathbf{a} be a vector in $\Sigma_1 \times \cdots \times \Sigma_k$ with $d = \gcd(a'_1, \dots, a'_k)$. Then $|S_{\mathbf{a}',\ell d}| = \frac{dq_1 \cdots q_k}{q_S}$ for every $0 \leq \ell \leq \frac{q_S}{d} - 1$.*

Proof. Since $d = \gcd(a'_1, \dots, a'_k)$, if we let $b_i = \frac{a'_i}{d}$ for each $1 \leq i \leq k$, then $\gcd(b_1, \dots, b_k) = 1$. Now applying the same argument as in the proof of Proposition 4.9 gives the desired result. \square

Now for every $1 \leq \ell \leq \frac{q_S}{d} - 1$ and put $q^* \stackrel{\text{def}}{=} \frac{q_S}{d}$, we have

$$\begin{aligned}
\hat{D}(\ell \mathbf{a}) &= \sum_{\mathbf{x} \in \Sigma_1 \times \dots \times \Sigma_k} D_S(\mathbf{x}) e^{2\pi i \left(\frac{\ell a_1 x_1}{q_1} + \dots + \frac{\ell a_k x_k}{q_k} \right)} \\
&= \sum_{\mathbf{x} \in \Sigma_1 \times \dots \times \Sigma_k} D_S(\mathbf{x}) e^{\frac{2\pi i}{q_S} \ell \mathbf{a}' \cdot \mathbf{x}} = \sum_{j=0}^{\frac{q_S}{d} - 1} \Pr_{\mathbf{X} \sim D} [\mathbf{a}' \cdot \mathbf{X} \equiv jd \pmod{q_S}] e^{\frac{2\pi i}{q_S} \ell jd} \\
&= \sum_{j=0}^{\frac{q_S}{d} - 1} w(j) e^{\frac{2\pi i}{q_S} \ell jd} = \sum_{j=0}^{q^* - 1} w(j) e^{\frac{2\pi i}{q^*} \ell j},
\end{aligned}$$

where $w(j) \stackrel{\text{def}}{=} P_{\mathbf{a}', jd}$. That is, each of the Fourier coefficients $\hat{D}(\mathbf{a}), \hat{D}(2\mathbf{a}), \dots, \hat{D}((q^* - 1)\mathbf{a})$ can be written as a one-dimensional Fourier transform of a function (namely, $w(j)$) over \mathbb{Z}_{q^*} . Then following the same proofs as those in Sec. 4.2.3, we have that the total weight to zero-out the Fourier coefficients at \mathbf{a} and its multiples is at most $\sum_{\ell=1}^{\frac{q_S}{d} - 1} |\hat{D}(\ell \mathbf{a})|$. This in turn gives the upper bound stated in Theorem 4.19 on the distance between D and k -wise independence over product spaces.

4.3.1 Testing Algorithm and its Analysis

We study the problem of testing k -wise independence over the product space $\Sigma_1 \times \dots \times \Sigma_n$ in this section.

To simplify notation, in the following we write

$$M^{\text{prod}} = \sum_{\ell=1}^k \sum_{I \in \binom{[n]}{\ell}} \prod_{i \in I} (q_i - 1)$$

for the total number of non-zero Fourier coefficients of weight at most k , and

$$q_{\max} = \max_{S \in \binom{[n]}{k}} \text{lcm}(q_i : i \in S)$$

for the maximum effective alphabet size of any index subset of size k .

Note that a simple corollary of Theorem 4.19 is

$$\Delta(D, \mathcal{D}_{\text{kwi}}) \leq M^{\text{prod}} \max_{0 < \text{wt}(\mathbf{a}) \leq k} |\hat{D}(\mathbf{a})|,$$

which gives the soundness condition for the distance bound. For the completeness condition, it is easy to see that for any $0 \leq \delta \leq 1$ and any non-zero vector \mathbf{a} of weight at most k , if $\Delta(D, \mathcal{D}_{\text{kwi}}) \leq \delta$, then $|\hat{D}(\mathbf{a})| \leq q_{\max} \delta$. The following theorem can now be proved easily by plugging these two conditions into Theorem 2.7. We omit the proof.

Theorem 4.22. *There is an algorithm that tests the k -wise independence over the product space $\Sigma_1 \times \dots \times \Sigma_n$ (as shown in Fig 3) with query complexity $O\left(\frac{q_{\max}^2 M^{\text{prod}}(n, k, q)^2}{\epsilon^2} \log(M^{\text{prod}}(n, k, q))\right)$ and time complexity $O\left(\frac{q_{\max}^2 M^{\text{prod}}(n, k, q)^3}{\epsilon^2} \log(M^{\text{prod}}(n, k, q))\right)$ and satisfies the following: for any distribution D over Σ^n , if $\Delta(D, \mathcal{D}_{\text{kwi}}) \leq \frac{\epsilon}{3q_{\max} M^{\text{prod}}(n, k, q)}$, then with probability at least $2/3$, the algorithm accepts; if $\Delta(D, \mathcal{D}_{\text{kwi}}) > \epsilon$, then with probability at least $2/3$, the algorithm rejects.*

Test-Product-KWI(D, k, q, ϵ)

1. Sample D independently $O\left(\frac{q_{\max}^2 M^{\text{prod}}(n, k, q)^2}{\epsilon^2} \log(M^{\text{prod}}(n, k, q))\right)$ times to obtain a set Q
2. For every non-zero vector \mathbf{a} of weight at most k , use Q to estimate $\hat{D}(\mathbf{a})$
3. If $\max_{\mathbf{a}} \left| \hat{D}(\mathbf{a}) \right| \leq \frac{2\epsilon}{3M^{\text{prod}}(n, k, q)}$, return “**Accept**”; else return “**Reject**”

Figure 3: Algorithm for testing uniform k -wise independence over product spaces.

5 Non-uniform k -wise Independence

In this section we seek a robust characterization of non-uniform k -wise independent distributions. For ease of exposition, we present our results only for the case when the underlying domain is $\{0, 1, \dots, q-1\}^n$. Our approach can be generalized easily to handle distributions over product spaces.

Recall that a distribution $D : \Sigma^n \rightarrow [0, 1]$ is k -wise independent if for any index subset $S \subset [n]$ of size k , $S = \{i_1, \dots, i_k\}$, and for any $z_1 \dots z_k \in \Sigma^k$, $D_S(z_1 \dots z_k) = \Pr_D[X_{i_1} = z_1] \dots \Pr_D[X_{i_k} = z_k]$. Our strategy of showing an upper bound on the distance between D and non-uniform k -wise independence is to reduce the non-uniform problem to the uniform case and then apply Theorem 1.1.

5.1 Non-uniform Fourier Coefficients

In the following we define a set of factors which are used to transform non-uniform k -wise independent distributions into uniform ones. Let $p_i(z) \stackrel{\text{def}}{=} \Pr_D[X_i = z]$. We assume that $0 < p_i(z) < 1$ for every $i \in [n]$ and every $z \in \Sigma$ (this is without loss of generality since if some $p_i(z)$'s are zero, then it reduces to the case of distributions over product spaces). Let $\theta_i(z) \stackrel{\text{def}}{=} \frac{1}{qp_i(z)}$. Intuitively, one may think of the $\theta_i(z)$'s as a set of compressing/stretching factors which transform a non-uniform k -wise distribution into a uniform one. For convenience of notation, if $S = \{i_1, \dots, i_\ell\}$ and $\mathbf{z} = z_{i_1} \dots z_{i_\ell}$, we write $\theta_S(\mathbf{z})$ for the product $\theta_{i_1}(z_{i_1}) \dots \theta_{i_\ell}(z_{i_\ell})$.

Definition 5.1 (Non-uniform Fourier Coefficients). Let D be a distribution over Σ^n . Let \mathbf{a} be a non-zero vector in Σ^n and $\text{supp}(\mathbf{a})$ to be its support. Let $D_{\text{supp}(\mathbf{a})}$ be the projection of D to coordinates in $\text{supp}(\mathbf{a})$. For every \mathbf{z} in the support of $D_{\text{supp}(\mathbf{a})}$, define $D'_{\text{supp}(\mathbf{a})}(\mathbf{z}) = \theta_{\text{supp}(\mathbf{a})}(\mathbf{z}) D_{\text{supp}(\mathbf{a})}(\mathbf{z})$, which is the transformed distribution¹¹ of the projected distribution $D_{\text{supp}(\mathbf{a})}$. The *non-uniform Fourier coefficient* of D at \mathbf{a} , denoted $\hat{D}^{\text{non}}(\mathbf{a})$, is defined by

$$\hat{D}^{\text{non}}(\mathbf{a}) \stackrel{\text{def}}{=} \hat{D}'_{\text{supp}(\mathbf{a})}(\mathbf{a}) = \sum_{\mathbf{z} \in \Sigma^{|\text{supp}(\mathbf{a})|}} D'_{\text{supp}(\mathbf{a})}(\mathbf{z}) e^{\frac{2\pi i}{q} \mathbf{a} \cdot \mathbf{z}}. \quad (9)$$

Remark 5.2. In the following we always refer to \hat{D}^{non} collectively as a set of (complex) numbers that will be used to indicate the distance between distribution D and the non-uniform k -wise independence. Strictly speaking, \hat{D}^{non} are not Fourier coefficients since in general there is no distribution whose (low degree) Fourier coefficients are exactly \hat{D}^{non} .

¹¹Note that in general $D'_{\text{supp}(\mathbf{a})}$ is not a distribution: it is non-negative everywhere but $\sum_{\mathbf{x}} D'_{\text{supp}(\mathbf{a})}(\mathbf{x}) = 1$ may not hold.

To summarize, let us define a function

$$\mathcal{F} : (\mathbb{R}^{\geq 0})^{\Sigma^n} \times \left(\binom{[n]}{k} \times \Sigma^k \right) \rightarrow (\mathbb{R}^{\geq 0})^{\Sigma^k}$$

which maps a distribution D over Σ^n and a vector $\mathbf{a} \in \Sigma^n$ of weight k to a non-negative function over $\Sigma^{|\text{supp}(\mathbf{a})|}$. That is, for every $\mathbf{z} \in \Sigma^k$,

$$\mathcal{F}(D, \mathbf{a})(\mathbf{z}) = D_{\text{supp}(\mathbf{a})}(\mathbf{z})\theta_{\text{supp}(\mathbf{a})}(\mathbf{z}). \quad (10)$$

Then the non-uniform Fourier coefficient of D at \mathbf{a} is simply the ordinary uniform Fourier coefficient of $\hat{\mathcal{F}}(D, \mathbf{a})$ at \mathbf{a} :

$$\hat{D}^{\text{non}}(\mathbf{a}) = \hat{\mathcal{F}}(D, \mathbf{a})(\mathbf{a}).$$

The idea of defining $D'_{\text{supp}(\mathbf{a})}$ is that if D is non-uniform k -wise independent, then $D'_{\text{supp}(\mathbf{a})}$ will be a uniform distribution over the coordinates in $\text{supp}(\mathbf{a})$. Indeed, our main result in this section is to show a connection between the non-uniform Fourier coefficients of D and the property that distribution D is non-uniform k -wise independent. In particular we have the following simple characterization of the non-uniform k -wise independence.

Theorem 5.3. *A distribution D over Σ^n is non-uniform k -wise independent if and only if for every non-zero vector $\mathbf{a} \in \Sigma^n$ of weight at most k , $\hat{D}^{\text{non}}(\mathbf{a}) = 0$.*

5.2 New Characterization of Non-uniform k -wise Independence

We prove Theorem 5.3 in this section. It is straightforward to show that if D is a non-uniform k -wise independent distribution, then all the non-zero non-uniform Fourier coefficients of degree at most k are zero. However, the proof of the converse is more involved. The key observation is that if we write the non-uniform Fourier transform as a linear transformation, the non-uniform Fourier transform matrix, like the uniform Fourier transform matrix, can be expressed as a tensor product of a set of heterogeneous DFT (discrete Fourier transform) matrices (as opposed to homogeneous DFT matrices in the uniform case). This enables us to show that the non-uniform Fourier transform is invertible. Combined with the condition that all the non-trivial non-uniform Fourier coefficients are zero, this invertibility property implies that D must be a non-uniform k -wise independent distribution.

Recall that our new characterization of non-uniform k -wise independent distributions is:

Theorem 5.3. *A distribution D over Σ^n is k -wise independent if and only if for every non-zero vector $\mathbf{a} \in \Sigma^k$ with $\text{wt}(\mathbf{a}) \leq k$, $\hat{D}^{\text{non}}(\mathbf{a}) = 0$.*

Proof. Suppose D is a non-uniform k -wise independent distribution. Then it is easy to see that for any non-empty $T \subset [n]$ of size at most k (not just for subsets whose sizes are exactly k),

$$D_T(\mathbf{z}_T) = \prod_{i \in T} p_i(z_i).$$

Indeed, if $|T| = k$ then this follows directly from the definition of non-uniform k -wise independent distributions. If $|T| < k$, let $S \supset T$ be any index set of size k , then

$$D_T(\mathbf{z}_T) = \sum_{\mathbf{z}_S: \mathbf{z}_T = \mathbf{z}_S|_T} D_S(\mathbf{z}_S)$$

$$\begin{aligned}
&= \sum_{z_j: j \in S \setminus T} \prod_{\ell \in S} p_\ell(z_\ell) \\
&= \prod_{i \in T} p_i(z_i) \sum_{z_j: j \in S \setminus T} \prod_{j \in S \setminus T} p_j(z_j) \\
&= \prod_{i \in T} p_i(z_i) \prod_{j \in S \setminus T} \left(\sum_{z_j \in \Sigma} p_j(z_j) \right) \\
&= \prod_{i \in T} p_i(z_i),
\end{aligned}$$

as $\sum_{z_j \in \Sigma} p_j(z_j) = 1$ for every $1 \leq j \leq n$.

Let \mathbf{a} be any non-zero vector of weight $\ell \leq k$ whose support set is $\text{supp}(\mathbf{a})$. Now we show that $D'_{\text{supp}(\mathbf{a})}$ is a uniform distribution and consequently all the non-uniform Fourier coefficients whose support sets are $\text{supp}(\mathbf{a})$ must be zero. Indeed, by the definition of D' ,

$$\begin{aligned}
D'_{\text{supp}(\mathbf{a})}(\mathbf{z}_{\text{supp}(\mathbf{a})}) &= D_{\text{supp}(\mathbf{a})}(\mathbf{z}_{\text{supp}(\mathbf{a})}) \prod_{i \in \text{supp}(\mathbf{a})} \theta_i(z_i) \\
&= \prod_{i \in \text{supp}(\mathbf{a})} p_i(z_i) \prod_{i \in \text{supp}(\mathbf{a})} \frac{1}{qp_i(z_i)} \\
&= \frac{1}{q^\ell}
\end{aligned}$$

for every $\mathbf{z}_{\text{supp}(\mathbf{a})} \in \{0, 1, \dots, q-1\}^\ell$. Hence $\hat{D}^{\text{non}}(\mathbf{a}) = \hat{D}'_{\text{supp}(\mathbf{a})}(\mathbf{a}) = 0$ by Corollary 2.4.

The converse direction will follow directly from Lemma 5.4 below by setting $E = D_S$ in the statement. \square

Lemma 5.4. *Let $E : \Sigma^k \rightarrow \mathbb{R}^{\geq 0}$ be a distribution. For any index set $T \subseteq [k]$, let $E_T(\mathbf{z})$, $E'_T(\mathbf{z})$ and $\hat{E}^{\text{non}}(\mathbf{a})$ be defined analogously to those of $D_T(\mathbf{z})$, $D'_T(\mathbf{z})$ and $\hat{D}^{\text{non}}(\mathbf{a})$, respectively. If $\hat{E}^{\text{non}}(\mathbf{a}) = 0$ for every non-zero vector \mathbf{a} , then E is a non-uniform independent distribution, i.e. $E'_{[k]}$ is the uniform distribution and consequently E is a product distribution.*

One may think of Lemma 5.4 as the non-uniform version of Proposition 2.3.

Proof. For notational simplicity we write $S = [k]$. Let T be a subset of S of size $k-1$, and without loss of generality, we assume that $T = \{1, \dots, k-1\}$. We first observe the following relation between $E'_S(\mathbf{z})$ and $E'_T(\mathbf{z}_T)$.

$$\begin{aligned}
E'_T(z_1, \dots, z_{k-1}) &= E_T(z_1, \dots, z_{k-1}) \theta_1(z_1) \cdots \theta_{k-1}(z_{k-1}) \\
&= \sum_{z_k} E_S(z_1, \dots, z_{k-1}, z_k) \theta_1(z_1) \cdots \theta_{k-1}(z_{k-1}) \\
&= \sum_{z_k} \frac{1}{\theta_k(z_k)} E'_S(z_1, \dots, z_k) \\
&= \sum_{z_k} qp_k(z_k) E'_S(z_1, \dots, z_k).
\end{aligned}$$

By induction, we have in general, for any $T \subset S$,

$$E'_T(\mathbf{z}_T) = \sum_{z_j: j \in S \setminus T} E'_S(z_1, \dots, z_k) \prod_{j \in S \setminus T} (qp_j(z_j)). \quad (11)$$

Next we use (11) to eliminate the intermediate projection distributions E'_T , and write the non-uniform Fourier transform of E as a linear transform of $\{E'_S(\mathbf{z})\}_{\mathbf{z} \in \Sigma^k}$. Let \mathbf{a} be a vector whose support set is T , then

$$\begin{aligned} \hat{E}^{\text{non}}(\mathbf{a}) &= \hat{E}'_T(\mathbf{a}) \\ &= \sum_{z_i: i \in T} E'_T(\mathbf{z}_T) e^{\frac{2\pi i}{q} \sum_{i \in T} a_i z_i} \\ &= \sum_{z_i: i \in T} \sum_{z_j: j \in S \setminus T} E'_S(\mathbf{z}) e^{\frac{2\pi i}{q} \sum_{i \in T} a_i z_i} \prod_{j \in S \setminus T} (qp_j(z_j)) \\ &= \sum_{\mathbf{z} \in \Sigma^k} E'_S(\mathbf{z}) \prod_{i \in T} e^{\frac{2\pi i}{q} a_i z_i} \prod_{j \in S \setminus T} (qp_j(z_j)) \\ &= \sum_{\mathbf{z} \in \Sigma^k} E'_S(\mathbf{z}) \prod_{i \in \text{supp}(\mathbf{a})} e^{\frac{2\pi i}{q} a_i z_i} \prod_{j \in S \setminus \text{supp}(\mathbf{a})} (qp_j(z_j)). \end{aligned} \quad (12)$$

Define a q^k -dimensional column vector \mathbf{E}' with entries $E'_S(\mathbf{z})$ (we will specify the order of the entries later). Similarly define another q^k -dimensional column vector whose entries are the non-uniform Fourier coefficients $\hat{\mathbf{E}}^{\text{non}}$. Then we may write (12) more compactly as

$$\hat{\mathbf{E}}^{\text{non}} = \tilde{\mathbf{F}} \mathbf{E}'. \quad (13)$$

In what follows, we will show that $\tilde{\mathbf{F}}$ can be written nicely as a tensor product of k matrices. This in turn enables us to show that $\tilde{\mathbf{F}}$ is non-singular.

Let $\omega = e^{\frac{2\pi i}{q}}$ be a primitive q^{th} root of unity. The q -point discrete Fourier transform (DFT) matrix is given by

$$\mathbf{F} = \begin{bmatrix} 1 & 1 & 1 & 1 & \cdots & 1 \\ 1 & \omega & \omega^2 & \omega^3 & \cdots & \omega^{q-1} \\ 1 & \omega^2 & \omega^4 & \omega^6 & \cdots & \omega^{2(q-1)} \\ 1 & \omega^3 & \omega^6 & \omega^9 & \cdots & \omega^{3(q-1)} \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & \omega^{q-1} & \omega^{2(q-1)} & \omega^{3(q-1)} & \cdots & \omega^{(q-1)(q-1)} \end{bmatrix}$$

Note that a DFT matrix is also a Vandermonde matrix and therefore $\det(\mathbf{F}) \neq 0$.

Definition 5.5 (Tensor Product of Vectors and Matrices). Let \mathbf{A} be an $m \times n$ matrix and \mathbf{B} be a $p \times q$ matrix. Then the tensor product (a.k.a. Kronecker product) $\mathbf{A} \otimes \mathbf{B}$ is an $mp \times nq$ block matrix given by

$$\mathbf{A} \otimes \mathbf{B} = \begin{bmatrix} a_{00} \mathbf{B} & \cdots & a_{0,n-1} \mathbf{B} \\ \vdots & \ddots & \vdots \\ a_{m-1,0} \mathbf{B} & \cdots & a_{m-1,n-1} \mathbf{B} \end{bmatrix}$$

$$= \begin{bmatrix} a_{00}b_{00} & \cdots & a_{00}b_{0,q-1} & \cdots & \cdots & a_{0,n-1}b_{00} & \cdots & a_{0,n-1}b_{0,q-1} \\ \vdots & \ddots & \vdots & & & \vdots & \ddots & \vdots \\ a_{00}b_{p-1,0} & \cdots & a_{00}b_{p-1,q-1} & \cdots & \cdots & a_{0,n-1}b_{p-1,0} & \cdots & a_{0,n-1}b_{p-1,q-1} \\ \vdots & & \vdots & \ddots & & \vdots & & \vdots \\ \vdots & & \vdots & & \ddots & \vdots & & \vdots \\ a_{m-1,0}b_{00} & \cdots & a_{m-1,0}b_{0,q-1} & \cdots & \cdots & a_{m-1,n-1}b_{00} & \cdots & a_{m-1,n-1}b_{0,q-1} \\ \vdots & \ddots & \vdots & & & \vdots & \ddots & \vdots \\ a_{m-1,0}b_{p-1,0} & \cdots & a_{m-1,0}b_{p-1,q-1} & \cdots & \cdots & a_{m-1,n-1}b_{p-1,0} & \cdots & a_{m-1,n-1}b_{p-1,q-1} \end{bmatrix}.$$

Let \mathbf{a} be an m -dimensional column vector in \mathbb{R}^m and \mathbf{b} be a p -dimensional column vector in \mathbb{R}^p . Then the tensor product $\mathbf{a} \otimes \mathbf{b}$ is an mp -dimensional column vector in \mathbb{R}^{mp} and its entries are given by

$$\mathbf{a} \otimes \mathbf{b} = \begin{bmatrix} a_0 \\ \vdots \\ a_{m-1} \end{bmatrix} \otimes \begin{bmatrix} b_0 \\ \vdots \\ b_{p-1} \end{bmatrix} = \begin{bmatrix} a_0 b_0 \\ \vdots \\ a_0 b_{p-1} \\ \vdots \\ \vdots \\ a_{m-1} b_0 \\ \vdots \\ a_{m-1} b_{p-1} \end{bmatrix}.$$

Let $q \geq 2$ be an integer. The q -ary representation of a natural number r is an ordered tuple (b_k, \dots, b_1, b_0) such that $0 \leq b_i \leq q-1$ for every $0 \leq i \leq k$ and $r = b_0 + b_1 \cdot q + \cdots + b_k \cdot q^k$. The following simple while useful fact about the tensor product of matrices can be proved easily by induction on the number of matrices in the product.

Fact 5.6. Let $F^{(1)}, \dots, F^{(k)}$ be a set of $q \times q$ matrices where the (i, j) th entry of $F^{(\ell)}$ is denoted by $F_{i,j}^{(\ell)}$, $0 \leq i, j \leq q-1$. Let $G = F^{(1)} \otimes \cdots \otimes F^{(k)}$. For $0 \leq I, J \leq q^k - 1$, let the q -ary representations of I and J be $I = (i_1, \dots, i_k)$ and $J = (j_1, \dots, j_k)$, respectively. Then

$$G_{I,J} = F_{i_1, j_1}^{(1)} \cdots F_{i_k, j_k}^{(k)}.$$

Let's first consider the simple case when E is a one-dimensional distribution. Let \mathbf{E} be the column vector whose entries are values of E at $\{0, 1, \dots, q-1\}$. Similarly let $\hat{\mathbf{E}}$ be the column vector of E 's Fourier transform. If we arrange the entries of \mathbf{E} and $\hat{\mathbf{E}}$ in increasing order, then the one-dimensional (uniform) Fourier transform can be written in the matrix multiplication form as

$$\hat{\mathbf{E}} = \begin{bmatrix} \hat{E}(0) \\ \vdots \\ \hat{E}(q-1) \end{bmatrix} = \mathbf{F} \begin{bmatrix} E(0) \\ \vdots \\ E(q-1) \end{bmatrix} = \mathbf{F}\mathbf{E}. \quad (14)$$

For the general case in which E is a distribution over $\{0, 1, \dots, q-1\}^k$, we may view every k -dimensional point (x_1, \dots, x_k) in $E(x_1, \dots, x_k)$ as the representation of a natural number X in the q -ary representation: $X = x_1 \cdot q^{k-1} + \cdots + x_{k-1} \cdot q + x_k$. Then this provides a *natural order* of the entries in any column

vector defined over $\{0, 1, \dots, q-1\}^k$: view each vector (x_1, \dots, x_k) as a natural number X in the q -ary representation and arrange them in the increasing order. By tensor product and arranging the entries in \mathbf{E} and $\hat{\mathbf{E}}$ in the natural order, the k -dimensional Fourier transform can be written as

$$\hat{\mathbf{E}} = \begin{bmatrix} \hat{E}(0, 0, \dots, 0) \\ \vdots \\ \hat{E}(q-1, q-1, \dots, q-1) \end{bmatrix} = \underbrace{\mathbf{F} \otimes \dots \otimes \mathbf{F}}_{k \text{ times}} \begin{bmatrix} E(0, 0, \dots, 0) \\ \vdots \\ E(q-1, q-1, \dots, q-1) \end{bmatrix} = \left(\underbrace{\mathbf{F} \otimes \dots \otimes \mathbf{F}}_{k \text{ times}} \right) \mathbf{E}. \quad (15)$$

Definition 5.7 (Non-uniform DFT Matrices). For every $1 \leq i \leq k$, define (recall that $p_i(z)$'s are the marginal probabilities of E at coordinate i) the *non-uniform DFT matrix* at coordinate i to be

$$\tilde{\mathbf{F}}_i = \begin{bmatrix} qp_i(0) & qp_i(1) & qp_i(2) & qp_i(3) & \dots & qp_i(q-1) \\ 1 & \omega & \omega^2 & \omega^3 & \dots & \omega^{q-1} \\ 1 & \omega^2 & \omega^4 & \omega^6 & \dots & \omega^{2(q-1)} \\ 1 & \omega^3 & \omega^6 & \omega^9 & \dots & \omega^{3(q-1)} \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & \omega^{q-1} & \omega^{2(q-1)} & \omega^{3(q-1)} & \dots & \omega^{(q-1)(q-1)} \end{bmatrix}$$

The following lemma follows directly from Fact 5.6 and (12).

Lemma 5.8. *If we arrange the entries in \mathbf{E}' and $\hat{\mathbf{E}}^{\text{non}}$ in the natural order, then the $q^k \times q^k$ matrix $\tilde{\mathbf{F}}$ in (13) is the tensor product of k non-uniform DFT matrices, i.e.,*

$$\tilde{\mathbf{F}} = \tilde{\mathbf{F}}_1 \otimes \dots \otimes \tilde{\mathbf{F}}_k,$$

and consequently

$$\hat{\mathbf{E}}^{\text{non}} = (\tilde{\mathbf{F}}_1 \otimes \dots \otimes \tilde{\mathbf{F}}_k) \mathbf{E}'.$$

The following is a well-known fact on the determinants of tensor product matrices, see e.g. [38] for an elementary proof.

Fact 5.9. *If \mathbf{A} is an $m \times m$ square matrix and \mathbf{B} is an $n \times n$ square matrix, then*

$$\det(\mathbf{A} \otimes \mathbf{B}) = (\det(\mathbf{A}))^n (\det(\mathbf{B}))^m.$$

Proposition 5.10. *The non-uniform DFT matrix is non-singular for every $1 \leq i \leq k$. In particular,*

$$\det(\tilde{\mathbf{F}}_i) = q(p_i(0) + \dots + p_i(q-1)) (-1)^{q-1} \prod_{1 \leq \ell < m \leq q-1} (\omega^m - \omega^\ell) = (-1)^{q-1} q \prod_{1 \leq \ell < m \leq q-1} (\omega^m - \omega^\ell) \neq 0.$$

Proof. By Laplace expansion along the first row, we have

$$\det(\tilde{\mathbf{F}}_i) = \sum_{j=0}^{q-1} (-1)^j qp_i(j) \det(\mathbf{M}_{1j}). \quad (16)$$

The determinant of the minor \mathbf{M}_{1j} is

$$\det(\mathbf{M}_{1j}) = \begin{vmatrix} 1 & \omega & \dots & \omega^{j-1} & \omega^{j+1} & \dots & \omega^{q-1} \\ 1 & \omega^2 & \dots & \omega^{2(j-1)} & \omega^{2(j+1)} & \dots & \omega^{2(q-1)} \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 1 & \omega^{q-1} & \dots & \omega^{(j-1)(q-1)} & \omega^{(j+1)(q-1)} & \dots & \omega^{(q-1)(q-1)} \end{vmatrix}$$

$$\begin{aligned}
&= \left(\prod_{\ell=0, \ell \neq j}^{q-1} \omega^\ell \right) \begin{vmatrix} 1 & 1 & \cdots & 1 & 1 & \cdots & 1 \\ 1 & \omega & \cdots & \omega^{j-1} & \omega^{j+1} & \cdots & \omega^{q-1} \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 1 & \omega^{q-2} & \cdots & \omega^{(j-1)(q-2)} & \omega^{(j+1)(q-2)} & \cdots & \omega^{(q-1)(q-2)} \end{vmatrix} \\
&= \prod_{\ell=0, \ell \neq j}^{q-1} \omega^\ell \prod_{\substack{0 \leq \ell < m \leq q-1 \\ \ell, m \neq j}} (\omega^m - \omega^\ell) \\
&= \frac{\prod_{\ell=0, \ell \neq j}^{q-1} \omega^\ell \prod_{0 \leq \ell < m \leq q-1} (\omega^m - \omega^\ell)}{\prod_{\ell=0}^{j-1} (\omega^j - \omega^\ell) \prod_{\ell=j+1}^{q-1} (\omega^\ell - \omega^j)},
\end{aligned}$$

since the matrix in the second step is a Vandermonde matrix.

Using the fact that $\omega^q = 1$, the denominator may be simplified as

$$\begin{aligned}
&\prod_{\ell=0}^{j-1} (\omega^j - \omega^\ell) \prod_{\ell=j+1}^{q-1} (\omega^\ell - \omega^j) \\
&= (-1)^j \prod_{\ell=0}^{j-1} \omega^\ell \prod_{\ell=1}^j (1 - \omega^\ell) \prod_{\ell=j+1}^{q-1} (\omega^\ell - \omega^j) \\
&= (-1)^j \prod_{\ell=0}^{j-1} \omega^\ell \prod_{\ell=1}^j (1 - \omega^\ell) \prod_{\ell=j+1}^{q-1} \omega^{\ell-q} (\omega^q - \omega^{q+j-\ell}) \\
&= (-1)^j \prod_{\ell=0}^{j-1} \omega^\ell \prod_{\ell=1}^j (1 - \omega^\ell) \prod_{\ell=j+1}^{q-1} \omega^\ell \prod_{\ell=j+1}^{q-1} (1 - \omega^{q+j-\ell}) \\
&= (-1)^j \prod_{\ell=0, \ell \neq j}^{q-1} \omega^\ell \prod_{\ell=1}^{q-1} (1 - \omega^\ell).
\end{aligned}$$

Therefore we have

$$\det(\mathbf{M}_{1j}) = (-1)^j (-1)^{q-1} \prod_{1 \leq \ell < m \leq q-1} (\omega^m - \omega^\ell).$$

Plugging $\det(\mathbf{M}_{1j})$ into (16) completes the proof. \square

Combining Fact 5.9 and Proposition 5.10 gives

Lemma 5.11. *We have that*

$$\det(\tilde{\mathbf{F}}) = \det(\tilde{\mathbf{F}}_1 \otimes \cdots \otimes \tilde{\mathbf{F}}_k) \neq 0.$$

Recall that we assume that all the non-zero Fourier coefficients $\hat{E}^{\text{non}}(\mathbf{a})$ are zero. Now to make the linear system of equations in (13) complete, we add another constraint that $\hat{E}^{\text{non}}(\mathbf{0}) = \sum_z E'(z) = cq^k$, where c is a constant which will be determined later. Since $\tilde{\mathbf{F}}$ is non-singular, there is a unique solution to this system of q^k linear equations. But we know the uniform distribution $E'(z) = c$ for every $z \in \Sigma^k$ is a solution (by the proof of the *only if* direction of Theorem 5.3), therefore this is the unique solution.

Now we have, for every $\mathbf{z} \in \Sigma^k$, $E(\mathbf{z})\theta_S(\mathbf{z}) = c$. Observe that $1/\theta_S(\mathbf{z}) = q^k p_1(z_1) \cdots p_k(z_k)$, and since $p_i(z)$'s are marginal probabilities, $\sum_{z \in \Sigma} p_i(z) = 1$ for every i , it follows that

$$\sum_{\mathbf{z} \in \Sigma^k} \frac{1}{\theta_S(\mathbf{z})} = q^k \sum_{\mathbf{z} \in \Sigma^k} p_1(z_1) \cdots p_k(z_k) = q^k.$$

Using the fact that $\sum_{\mathbf{z} \in \Sigma^k} E(\mathbf{z}) = 1$, we arrive at

$$1 = \sum_{\mathbf{z} \in \Sigma^k} E(\mathbf{z}) = c \sum_{\mathbf{z} \in \Sigma^k} \frac{1}{\theta_S(\mathbf{z})} = q^k c,$$

and therefore $c = \frac{1}{q^k}$ and $E(\mathbf{z}) = \frac{1}{q^k \theta_S(\mathbf{z})} = p_1(z_1) \cdots p_k(z_k)$ as desired. This completes the proof of Lemma 5.4. \square

5.3 Zeroing-out Non-uniform Fourier Coefficients

Given a distribution D which is not k -wise independent, what is the distance between D and the non-uniform k -wise independence? In the following we will, based on the approach that has been applied to the uniform case, try to find a set of small-weight distributions to mix with D in order to zero-out all the non-uniform Fourier coefficients of weight at most k . Moreover, we can bound the total weight added to the original distribution in this zeroing-out process in terms of the non-uniform Fourier coefficients of D . This will show that the characterization of the non-uniform k -wise independence given in Theorem 5.3 is robust.

A careful inspection of Theorem 1.1 and its proof shows that if we focus on the weights added to correct any fixed prime vector and its *multiples*, we actually prove the following.

Theorem 5.12. *Let E' be a non-negative function¹² defined over Σ^n , \mathbf{a} be a prime vector of weight at most k and $\hat{E}'(\mathbf{a}), \hat{E}'(2\mathbf{a}), \dots, \hat{E}'((q-1)\mathbf{a})$ be the Fourier coefficients at \mathbf{a} and its multiple vectors. Then there exist a set of non-negative real numbers $w_j, j = 0, 1, \dots, q-1$, such that the (small-weight) distribution¹³ $\mathcal{W}_{E', \mathbf{a}} \stackrel{\text{def}}{=} \sum_{j=0}^{q-1} w_j U_{\mathbf{a}, j}$ satisfies the following properties. The Fourier coefficients of $E' + \mathcal{W}_{E', \mathbf{a}}$ at $\mathbf{a}, 2\mathbf{a}, \dots, (q-1)\mathbf{a}$ all equal zero and $\hat{\mathcal{W}}_{E', \mathbf{a}}(\mathbf{b}) = 0$ for all non-zero vectors that are not multiples of \mathbf{a} . Moreover, the total weight of $\mathcal{W}_{E', \mathbf{a}}$ is at most $\sum_{j=0}^{q-1} w_j \leq \sum_{\ell=1}^{q-1} |\hat{E}'(\ell\mathbf{a})|$.*

Applying Theorem 5.12 with E' equal to $D'_{\text{supp}(\mathbf{a})}$ gives rise to a small-weight distribution $\mathcal{W}_{D'_{\text{supp}(\mathbf{a})}, \mathbf{a}}$ which, by abuse of notation, we denote by $\mathcal{W}_{\mathbf{a}}$. When we add $\mathcal{W}_{\mathbf{a}}$ to $D'_{\text{supp}(\mathbf{a})}$, the resulting non-negative function has zero Fourier coefficients at \mathbf{a} and all its *multiple* vectors. That is,

$$\hat{\mathcal{W}}_{\mathbf{a}}(\ell\mathbf{a}) = -\hat{D}'_{\text{supp}(\mathbf{a})}(\ell\mathbf{a}), \quad \text{for every } 1 \leq \ell \leq q-1. \quad (17)$$

$$= -\hat{D}^{\text{non}}(\ell'\mathbf{a}), \quad \text{for every } \ell' \text{ such that } \text{supp}(\ell'\mathbf{a}) = \text{supp}(\mathbf{a}). \quad (17')$$

and for any \mathbf{b} which is not a *multiple* of \mathbf{a} ,

$$\hat{\mathcal{W}}_{\mathbf{a}}(\mathbf{b}) = 0. \quad (18)$$

¹²In Theorem 1.1 we only prove this for the case when E' is a distribution. However it is easy to see that the result applies to non-negative functions as well.

¹³Recall that $U_{\mathbf{a}, j}$ is the uniform distribution over all strings $x \in \mathbb{Z}_q^n$ such that $\mathbf{a} \cdot \mathbf{x} \equiv j \pmod{q}$.

However, this small-weight distribution only works for the auxiliary function $D'_{\text{supp}(\mathbf{a})}$ but what we are looking for is a small-weight distribution that corrects the non-uniform Fourier coefficients of D at \mathbf{a} . To this end, we apply the reversed compressing/stretching factor to $\mathcal{W}_{\mathbf{a}}$ to get $\tilde{\mathcal{W}}_{\mathbf{a}}$,

$$\tilde{\mathcal{W}}_{\mathbf{a}}(\mathbf{x}) \stackrel{\text{def}}{=} \frac{\mathcal{W}_{\mathbf{a}}(\mathbf{x})}{\theta_{[n]}(\mathbf{x})}. \quad (19)$$

The following lemma shows that mixing D with $\tilde{\mathcal{W}}_{\mathbf{a}}$ results in a distribution whose non-uniform Fourier coefficients at \mathbf{a} as well as its *multiple* vectors are zero¹⁴. In addition, the mixing only adds a relatively small weight and may increase the magnitudes of the non-uniform Fourier coefficients only at vectors whose supports are completely contained in the support of \mathbf{a} .

Lemma 5.13. *Let D be a distribution over Σ^n and \mathbf{a} be a prime vector of weight at most k . Let $\text{supp}(\mathbf{a})$ be the support set of \mathbf{a} and $\tilde{\mathcal{W}}_{\mathbf{a}}$ be as defined in (19). Let the maximum factor over all possible compressing/stretching factors be denoted as $\gamma_k \stackrel{\text{def}}{=} \max_{S,z} \frac{1}{\theta_S(z)}$, where S ranges over all subsets of $[n]$ of size at most k and $z \in \Sigma^{|S|}$. Then $\tilde{\mathcal{W}}_{\mathbf{a}}$ satisfies the following properties:*

1. *The non-uniform Fourier coefficients of $D + \tilde{\mathcal{W}}_{\mathbf{a}}$ at \mathbf{a} as well as at the multiple vectors of \mathbf{a} whose support sets are also $\text{supp}(\mathbf{a})$ are all zero.¹⁵ Moreover, $\hat{\mathcal{W}}_{\mathbf{a}}^{\text{non}}(\mathbf{a}') = 0$ for every vector \mathbf{a}' whose support set is $\text{supp}(\mathbf{a})$ but is not a multiple vector of \mathbf{a} .*
2. *For any vector \mathbf{b} with $\text{supp}(\mathbf{b}) \not\subseteq \text{supp}(\mathbf{a})$, $\hat{\mathcal{W}}_{\mathbf{a}}^{\text{non}}(\mathbf{b}) = 0$.*
3. *The total weight of $\tilde{\mathcal{W}}_{\mathbf{a}}$ is at most $\gamma_k \sum_{\mathbf{x} \in \Sigma^n} \mathcal{W}_{\mathbf{a}}(\mathbf{x}) \leq \gamma_k \sum_{j=1}^{q-1} |\hat{D}^{\text{non}}(j\mathbf{a})|$.*
4. *For any non-zero vector \mathbf{c} with $\text{supp}(\mathbf{c}) \subset \text{supp}(\mathbf{a})$, $\hat{\mathcal{W}}_{\mathbf{a}}^{\text{non}}(\mathbf{c}) \leq \gamma_k \sum_{j=1}^{q-1} |\hat{D}^{\text{non}}(j\mathbf{a})|$.*

Proof. For simplicity, we assume that $\text{supp}(\mathbf{a}) = [k]$. Recall that $\mathcal{W}_{\mathbf{a}} = \sum_{j=0}^{q-1} w_j U_{\mathbf{a},j}$ and $U_{\mathbf{a},j}$ is the uniform distribution over the strings $\mathbf{x} \in \mathbb{Z}_q^n$ such that $\sum_{i=1}^n a_i x_i \equiv j \pmod{q}$. A simple while important observation is the following: since the support of \mathbf{a} is $[k]$, if $x_1 \cdots x_k$ satisfies the constraint $\sum_{i=1}^k a_i x_i \equiv j \pmod{q}$, then for any $y_{k+1} \cdots y_n \in \Sigma^{n-k}$, $x_1 \cdots x_k y_{k+1} \cdots y_n$ will satisfy the constraint and thus is in the support of the distribution.

Remark on notation. In the rest of this section, we always write \mathbf{x} for an n -bit vector in Σ^n and write \mathbf{z} for a k -bit vector in Σ^k .

Note that we may decompose $\mathcal{W}_{\mathbf{a}}$ (or any non-negative function) into a sum of q^k weighted distributions as $\mathcal{W}_{\mathbf{a}} = \sum_{\mathbf{z} \in \Sigma^k} w_{\mathbf{z}} \mathcal{U}_{\mathbf{z}}$, such that each of the distribution $\mathcal{U}_{\mathbf{z}}$ is supported on the $|\Sigma|^{n-k}$ strings whose k -bit prefixes are \mathbf{z} . That is,

$$w_{\mathbf{z}} \mathcal{U}_{\mathbf{z}}(\mathbf{x}) = \begin{cases} \mathcal{W}_{\mathbf{a}}(\mathbf{x}), & \text{if } \mathbf{x}_{[k]} = \mathbf{z}, \\ 0, & \text{otherwise.} \end{cases}$$

¹⁴In fact, the lemma only guarantees to zero-out the Fourier coefficients at \mathbf{a} and its *multiples* whose support sets are the same as that of \mathbf{a} . But that will not be a problem since we will perform the correction process in stages and will come to vectors with smaller support sets at some later stages.

¹⁵Note that if \mathbf{a} is a prime vector and \mathbf{a}' is a *multiple* vector of \mathbf{a} , then $\text{supp}(\mathbf{a}') \subseteq \text{supp}(\mathbf{a})$.

To make \mathcal{U}_z indeed a distribution, i.e., $\sum_{\mathbf{x}} \mathcal{U}_z(\mathbf{x}) = 1$, we simply set

$$w_z \stackrel{\text{def}}{=} (\mathcal{W}_{\mathbf{a}})_{[k]}(z). \quad (20)$$

That is, w_z equals the mass of the projected distribution $\mathcal{W}_{\mathbf{a}}$ at z . By Theorem 5.12 clearly we have

$$\sum_{z \in \Sigma^k} w_z \leq \sum_{j=1}^{q-1} \left| \hat{D}^{\text{non}}(j\mathbf{a}) \right|. \quad (21)$$

The aforementioned observation then implies that for every $z \in \Sigma^k$, \mathcal{U}_z is the uniform distribution over all $|\Sigma|^{n-k}$ strings whose k -bit prefixes are z . In other words, \mathcal{U}_z is uniform over the strings in its support. We will refer to these distributions as *atomic uniform distributions*. More explicitly,

$$\mathcal{U}_z(\mathbf{x}) = \begin{cases} \frac{1}{q^{n-k}}, & \text{if } \mathbf{x}_{[k]} = z, \\ 0, & \text{otherwise.} \end{cases} \quad (22)$$

After applying the compressing/stretching factor, \mathcal{U}_z is transformed into $\tilde{\mathcal{U}}_z$:

$$\tilde{\mathcal{U}}_z(\mathbf{x}) = \begin{cases} \frac{1}{q^{n-k}\theta_{[n]}(\mathbf{x})}, & \text{if } \mathbf{x}_{[k]} = z, \\ 0, & \text{otherwise.} \end{cases} \quad (23)$$

We call $\tilde{\mathcal{U}}_z$ a *transformed atomic uniform distribution*. Clearly we have

$$\tilde{\mathcal{W}}_{\mathbf{a}} = \sum_{z \in \Sigma^k} w_z \tilde{\mathcal{U}}_z.$$

We remark that both atomic uniform distributions and transformed atomic uniform distributions are introduced only for the sake of analysis; they play no role in the testing algorithm.

Our plan is to show the following: on the one hand, $\{w_z \tilde{\mathcal{U}}_z\}_z$, the weighted transformed atomic uniform distributions, *collectively* zero-out the non-uniform Fourier coefficients of D at \mathbf{a} and all the *multiple* vectors of \mathbf{a} whose supports are the same as \mathbf{a} . On the other hand, *individually*, each transformed atomic uniform distribution $\tilde{\mathcal{U}}_z$ has zero non-uniform Fourier coefficient at any vector whose support is not a subset of $\text{supp}(\mathbf{a})$. Then by linearity of the Fourier transform, $\tilde{\mathcal{W}}_{\mathbf{a}}$ also has zero Fourier coefficients at these vectors.

We first show that if we project $\tilde{\mathcal{U}}_z$ to index set $[k]$ to obtain the distribution $(\tilde{\mathcal{U}}_z)_{[k]}$, then $(\tilde{\mathcal{U}}_z)_{[k]}$ is supported only on a single string (namely z) and has total weight $\frac{1}{\theta_{[k]}(z)}$, which is independent of the compressing/stretching factors applied to the last $n - k$ coordinates.

Remark on notation. To simplify notation, we will use Kronecker's delta function, $\delta(\mathbf{u}, \mathbf{v})$, in the following. By definition, $\delta(\mathbf{u}, \mathbf{v})$ equals 1 if $\mathbf{u} = \mathbf{v}$ and 0 otherwise. An important property of δ -function is $\sum_{\mathbf{u}'} f(\mathbf{u}')\delta(\mathbf{u}, \mathbf{u}') = f(\mathbf{u})$, where f is an arbitrary function.

Claim 5.14. *We have*

$$\left(\tilde{\mathcal{U}}_z \right)_{[k]}(z') = \frac{\delta(z', z)}{\theta_{[k]}(z)}, \quad (24)$$

and consequently

$$\sum_{\mathbf{x} \in \Sigma^n} \tilde{\mathcal{U}}_z(\mathbf{x}) = \frac{1}{\theta_{[k]}(z)}. \quad (24')$$

Proof. Note that $\tilde{\mathcal{U}}_{\mathbf{z}}(\mathbf{x})$ can be written as

$$\tilde{\mathcal{U}}_{\mathbf{z}}(\mathbf{x}) = \frac{\delta(\mathbf{x}_{[k]}, \mathbf{z})}{\theta_{[k]}(\mathbf{z})} \frac{1}{q^{n-k} \theta_{[k+1, n]}(\mathbf{x}_{[k+1, n]})} = \frac{\delta(\mathbf{x}_{[k]}, \mathbf{z})}{\theta_{[k]}(\mathbf{z})} \frac{1}{q^{n-k} \theta_{k+1}(x_{k+1}) \cdots \theta_n(x_n)}.$$

Then by simple calculation,

$$\begin{aligned} \left(\tilde{\mathcal{U}}_{\mathbf{z}}\right)_{[k]}(\mathbf{x}_{[k]}) &= \sum_{x_{k+1}, \dots, x_n} \tilde{\mathcal{U}}_{\mathbf{z}}(\mathbf{x}) = \frac{\delta(\mathbf{x}_{[k]}, \mathbf{z})}{\theta_{[k]}(\mathbf{z})} \sum_{x_{k+1}, \dots, x_n} \frac{1}{q^{n-k} \theta_{k+1}(x_{k+1}) \cdots \theta_n(x_n)} \\ &= \frac{\delta(\mathbf{x}_{[k]}, \mathbf{z})}{\theta_{[k]}(\mathbf{z})} \frac{1}{q^{n-k}} \sum_{x_{k+1}, \dots, x_n} q^{n-k} p_{k+1}(x_{k+1}) \cdots p_n(x_n) \\ &= \frac{\delta(\mathbf{x}_{[k]}, \mathbf{z})}{\theta_{[k]}(\mathbf{z})} \left(\sum_{x_{k+1}} p_{k+1}(x_{k+1}) \right) \cdots \left(\sum_{x_n} p_n(x_n) \right) \\ &= \frac{\delta(\mathbf{x}_{[k]}, \mathbf{z})}{\theta_{[k]}(\mathbf{z})}. \quad \square \end{aligned}$$

Note that (24) is exactly what we want, since to compute the non-uniform Fourier coefficient of $w_{\mathbf{z}} \tilde{\mathcal{U}}_{\mathbf{z}}(\mathbf{z}')$ at \mathbf{a} , we need to multiply the projected distribution by $\theta_{[k]}(\mathbf{z}')$. Specifically, denote the transformed function $\mathcal{F}(\tilde{\mathcal{W}}_{\mathbf{a}}, \mathbf{a})$ (as defined in (10)) by \mathcal{W}' and use (24), then for every $\mathbf{z}' \in \Sigma^k$,

$$\begin{aligned} \mathcal{W}'(\mathbf{z}') &= \left(\tilde{\mathcal{W}}_{\mathbf{a}}\right)_{[k]}(\mathbf{z}') \theta_{[k]}(\mathbf{z}') \\ &= \sum_{\mathbf{z}} w_{\mathbf{z}} \left(\tilde{\mathcal{U}}_{\mathbf{z}}\right)_{[k]}(\mathbf{z}') \theta_{[k]}(\mathbf{z}') \\ &= \sum_{\mathbf{z}} w_{\mathbf{z}} \frac{\delta(\mathbf{z}', \mathbf{z})}{\theta_{[k]}(\mathbf{z})} \theta_{[k]}(\mathbf{z}') \\ &= w_{\mathbf{z}'}. \end{aligned}$$

It follows that $\mathcal{W}' = \mathcal{W}_{\mathbf{a}}$ by (20). Therefore for any vector \mathbf{b} whose support set is $[k]$, we have $\hat{\mathcal{W}}_{\mathbf{a}}^{\text{non}}(\mathbf{b}) = \hat{\mathcal{W}}_{\mathbf{a}}(\mathbf{b})$. In particular, by (17') and (18), $\hat{\mathcal{W}}_{\mathbf{a}}^{\text{non}}(\ell' \mathbf{a}) = -\hat{D}^{\text{non}}(\ell' \mathbf{a})$ for every vector $\ell' \mathbf{a}$ such that $\text{supp}(\ell' \mathbf{a}) = \text{supp}(\mathbf{a})$ and $\hat{\mathcal{W}}_{\mathbf{a}}^{\text{non}}(\mathbf{b}) = 0$ for every vector \mathbf{b} which is not a *multiple* of \mathbf{a} and satisfies $\text{supp}(\mathbf{b}) = \text{supp}(\mathbf{a})$. This proves the first part of the Lemma 5.13.

Next we consider the non-uniform Fourier coefficient of $\tilde{\mathcal{U}}_{\mathbf{a}}$ at \mathbf{b} , where $\text{supp}(\mathbf{b}) \not\subseteq [k]$. Without loss of generality, assume that $\text{supp}(\mathbf{b}) = \{\ell + 1, \dots, k, k + 1, \dots, k + m\}$, where $\ell \leq k - 1$ and $m \geq 1$. Consider the non-uniform Fourier coefficient of any atomic uniform distribution $\tilde{\mathcal{U}}_{\mathbf{z}}$ at \mathbf{b} . By the form of $\tilde{\mathcal{U}}_{\mathbf{z}}(\mathbf{x})$ in (23),

$$\begin{aligned} \left(\tilde{\mathcal{U}}_{\mathbf{z}}\right)_{\text{supp}(\mathbf{b})}(x_{\ell+1}, \dots, x_{k+m}) &= \left(\tilde{\mathcal{U}}_{\mathbf{z}}\right)_{[\ell+1, k+m]}(x_{\ell+1}, \dots, x_{k+m}) \\ &= \sum_{x_1, \dots, x_{\ell}} \sum_{x_{k+m+1}, \dots, x_n} \tilde{\mathcal{U}}_{\mathbf{z}}(\mathbf{x}) \\ &= \frac{1}{q^{n-k}} \sum_{x_1, \dots, x_{\ell}} \frac{\delta(\mathbf{x}_{[k]}, \mathbf{z})}{\theta_{[k]}(\mathbf{z})} \sum_{x_{k+m+1}, \dots, x_n} \frac{1}{\theta_{k+1}(x_{k+1}) \cdots \theta_{k+m}(x_{k+m}) \theta_{k+m+1}(x_{k+m+1}) \cdots \theta_n(x_n)} \end{aligned}$$

$$\begin{aligned}
&= \frac{\delta(\mathbf{x}_{[\ell+1,k]}, \mathbf{z}_{[\ell+1,k]})}{q^{n-k}\theta_{[k]}(\mathbf{z})} \frac{q^{n-k-m}}{\theta_{k+1}(x_{k+1}) \cdots \theta_{k+m}(x_{k+m})} \left(\sum_{x_{k+m+1}} p_{k+m+1}(x_{k+m+1}) \right) \cdots \left(\sum_{x_n} p_n(x_n) \right) \\
&= \frac{1}{q^m \theta_{[k]}(\mathbf{z}) \theta_{k+1}(x_{k+1}) \cdots \theta_{k+m}(x_{k+m})} \delta(\mathbf{x}_{[\ell+1,k]}, \mathbf{z}_{[\ell+1,k]}).
\end{aligned}$$

Therefore, after applying the compressing/stretching transformation, $\tilde{\mathcal{U}}_{\mathbf{z}}$ is uniform over $[k+1, k+m]$. Consequently, its non-uniform Fourier coefficient at \mathbf{b} is

$$\begin{aligned}
\hat{\mathcal{U}}_{\mathbf{z}}^{\text{non}}(\mathbf{b}) &= \sum_{x_{\ell+1}, \dots, x_{k+m}} \frac{\delta(\mathbf{x}_{[\ell+1,k]}, \mathbf{z}_{[\ell+1,k]}) \theta_{\ell+1}(x_{\ell+1}) \cdots \theta_{k+m}(x_{k+m})}{q^m \theta_{[k]}(\mathbf{z}) \theta_{k+1}(x_{k+1}) \cdots \theta_{k+m}(x_{k+m})} e^{\frac{2\pi i}{q}(b_{\ell+1}x_{\ell+1} + \cdots + b_{k+m}x_{k+m})} \\
&= \frac{e^{\frac{2\pi i}{q}(b_{\ell+1}z_{\ell+1} + \cdots + b_k z_k)}}{q^m \theta_1(z_1) \cdots \theta_{\ell}(z_{\ell})} \sum_{x_{k+1}, \dots, x_{k+m}} e^{\frac{2\pi i}{q}(b_{k+1}x_{k+1} + \cdots + b_{k+m}x_{k+m})} \\
&= \frac{e^{\frac{2\pi i}{q}(b_{\ell+1}z_{\ell+1} + \cdots + b_k z_k)}}{q^m \theta_1(z_1) \cdots \theta_{\ell}(z_{\ell})} \sum_{x_{k+2}, \dots, x_{k+m}} e^{\frac{2\pi i}{q}(b_{k+2}x_{k+2} + \cdots + b_{k+m}x_{k+m})} \sum_{x_{k+1}} e^{\frac{2\pi i}{q}(b_{k+1}x_{k+1})} \\
&= 0,
\end{aligned}$$

where the last step follows from Fact 2.1 as b_{k+1} is non-zero. This proves the second part of the Lemma 5.13.

By (24') in Claim 5.14 the total weight added by a transformed atomic uniform distribution is $\frac{w_{\mathbf{z}}}{\theta_{[k]}(\mathbf{z})} \leq \gamma_k w_{\mathbf{z}}$. Adding the weights of all the atomic uniform distributions together and using the upper bound on total weights in (21) proves the third part of Lemma 5.13.

For the last part, assume $\text{supp}(\mathbf{c}) = T \subset [k]$. Now consider the contribution of a transformed atomic uniform distribution $w_{\mathbf{z}} \tilde{\mathcal{U}}_{\mathbf{z}}$ to the non-uniform Fourier coefficient at \mathbf{c} . The probability mass at \mathbf{z}'_T of the transformed atomic distribution is

$$\begin{aligned}
\mathcal{F}(w_{\mathbf{z}} \tilde{\mathcal{U}}_{\mathbf{z}}, \mathbf{c})(\mathbf{z}'_T) &= w_{\mathbf{z}} \left(\frac{\delta(\mathbf{z}'_T, \mathbf{z})}{\theta_{[k]}(\mathbf{z})} \right)_T \theta_T(\mathbf{z}'_T) \\
&= w_{\mathbf{z}} \frac{\theta_T(\mathbf{z}'_T)}{\theta_{[k]}(\mathbf{z})} \delta(\mathbf{z}'_T, \mathbf{z}_T).
\end{aligned}$$

Therefore we can upper bound its non-uniform Fourier coefficient at \mathbf{c} by

$$\begin{aligned}
\left| \hat{\mathcal{F}}(w_{\mathbf{z}} \tilde{\mathcal{U}}_{\mathbf{z}}, \mathbf{c})(\mathbf{c}) \right| &\leq \left| \sum_{\mathbf{z}'_T} \mathcal{F}(w_{\mathbf{z}} \tilde{\mathcal{U}}_{\mathbf{z}}, \mathbf{c})(\mathbf{z}'_T) \right| \\
&= w_{\mathbf{z}} \frac{\theta_T(\mathbf{z}'_T)}{\theta_{[k]}(\mathbf{z})} && \text{(since } \mathcal{F}(w_{\mathbf{z}} \tilde{\mathcal{U}}_{\mathbf{z}}, \mathbf{c}) \text{ is non-negative)} \\
&\leq w_{\mathbf{z}} \frac{1}{\theta_{[k]}(\mathbf{z})} && \text{(since } \theta_T(\mathbf{z}'_T) \leq 1) \\
&\leq \gamma_k w_{\mathbf{z}}.
\end{aligned}$$

Finally we add up the weights of all transformed atomic uniform distributions in $\tilde{\mathcal{W}}$ and apply (21) to prove the last part of Lemma 5.13. \square

Now for any prime vector \mathbf{a} of weight k , we can mix D with $\tilde{\mathcal{U}}_{\mathbf{a}}$ to zero-out the non-uniform Fourier coefficient at \mathbf{a} and all its *multiples* whose supports sets are $\text{supp}(\mathbf{a})$. By Lemma 5.13, the added small-weight distribution will only increase the magnitudes of the non-uniform Fourier coefficients at vectors whose supports are strict subsets of $\text{supp}(\mathbf{a})$. After doing this for all the prime vectors at level k , we obtain a distribution whose non-uniform Fourier coefficients at level k are all zero. We then recompute the non-uniform Fourier coefficients of the new distribution and repeat this process for prime vectors whose weights are $k - 1$. By iterating this process k times, we finally zero out all the non-uniform Fourier coefficients on the first level and obtain a non-uniform k -wise independent distribution.

Theorem 1.2. *Let D be a distribution over Σ^n , then*

$$\Delta(D, \mathcal{D}_{\text{kwi}}) \leq O\left(n^{\frac{k^2-k+2}{2}} q^{k(k+1)}\right) \max_{\mathbf{a}: 0 < \text{wt}(\mathbf{a}) \leq k} \left| \hat{D}^{\text{non}}(\mathbf{a}) \right|.$$

Proof. First observe that for every $1 \leq i \leq n$ and every $z \in \Sigma$, $\frac{1}{\theta_i(z)} = qp_i(z) < q$, so $\gamma_j < q^j$, for every $1 \leq j \leq k$.

We consider the zeroing-out processes in $k + 1$ stages. At stage 0 we have the initial distribution. Finally at stage k , we zero-out all the level-1 non-uniform Fourier coefficients and obtain a non-uniform k -wise independent distribution.

Let $f_{\max} = \max_{0 < \text{wt}(\mathbf{a}) \leq k} \left| \hat{D}^{\text{non}}(\mathbf{a}) \right|$. To simplify notation, we shall normalize by f_{\max} every bound on the magnitudes of the non-uniform Fourier coefficients as well as every bound on the total weight added in each stage. That is, we divide all the quantities by f_{\max} and work with the ratios.

Let $f^{(j)}$ denote the maximum magnitude, divided by f_{\max} , of all the non-uniform Fourier coefficients that have not been zeroed-out at stage j ; that is, the non-uniform Fourier coefficients at level i for $1 \leq i \leq k - j$. Clearly $f^{(0)} = 1$.

Now we consider the zeroing-out process at stage 1. There are $\binom{n}{k}(q-1)^k$ vectors at level k , and by part(3) of Lemma 5.13, correcting the non-uniform Fourier coefficient at each vector adds a weight at most $\gamma_k(q-1)f^{(0)}$. Therefore, the total weight added at stage 1 is at most $\binom{n}{k}(q-1)^k \gamma_k(q-1)f^{(0)} = O(n^k q^{2k+1})$. Next we calculate $f^{(1)}$, the maximal magnitude of the remaining non-uniform Fourier coefficients. For any vector \mathbf{c} at level i , $1 \leq i \leq k - 1$, there are $\binom{n-i}{k-i}(q-1)^{k-i}$ vectors at level k whose support sets are supersets of $\text{supp}(\mathbf{c})$. By part(4) of Lemma 5.13, zeroing-out the non-uniform Fourier coefficient at each such vector may increase $\left| \hat{D}^{\text{non}}(\mathbf{c}) \right|$ by $\gamma_k(q-1)f^{(0)}$. Therefore the magnitude of the non-uniform Fourier coefficient at \mathbf{c} is at most

$$f^{(0)} + \binom{n-i}{k-i} (q-1)^{k-i} \gamma_k(q-1) f^{(0)} = O\left(n^{k-i} q^{2k-i+1}\right).$$

Clearly the worst case happens when $i = 1$ and we thus have $f^{(1)} \leq O\left(n^{k-1} q^{2k}\right)$.

In general it is easy to see that at every stage, the maximum magnitude increases of the non-uniform Fourier coefficients always occur at level 1. At stage j , we need to zero-out the non-uniform Fourier coefficients at level $k - j + 1$. For a vector \mathbf{a} at level 1, there are $\binom{n-1}{k-j}(q-1)^{k-j}$ vectors at level $k - j + 1$ whose support sets are supersets of $\text{supp}(\mathbf{a})$, and the increase in magnitude of $\hat{D}^{\text{non}}(\mathbf{a})$ caused by each such level- $(k - j + 1)$ vector is at most $\gamma_{k-j+1}(q-1)f^{(j-1)}$. We thus have

$$f^{(j)} \leq \binom{n-1}{k-j} (q-1)^{k-j} \gamma_{k-j+1}(q-1) f^{(j-1)} \leq O\left(n^{k-j} q^{2(k-j+1)}\right) f^{(j-1)}, \quad \text{for } 1 \leq j \leq k - 1.$$

This in turn gives

$$f^{(j)} \leq O\left(n^{\frac{j(2k-j-1)}{2}} q^{j(2k-j+1)}\right), \quad \text{for } 1 \leq j \leq k-1.$$

It is easy to check that the weights added at stage k dominates the weights added at all previous stages, therefore the total weight added during all $k+1$ stages is at most

$$O\left(\binom{n}{1}(q-1)\gamma_1\right) f^{(k-1)} \leq O\left(n^{\frac{k^2-k+2}{2}} q^{k(k+1)}\right). \quad \square$$

5.4 Testing Algorithm and its Analysis

We now study the problem of testing non-uniform k -wise independence over \mathbb{Z}_q^n . Define

$$\theta_{\max} \stackrel{\text{def}}{=} \max_{S \subset [n], 0 < |S| \leq k, \mathbf{z} \in \Sigma^{|S|}} \theta_S(\mathbf{z})$$

to be the maximum compressing/stretching factor we ever apply when compute the non-uniform Fourier coefficients.

Claim 5.15. *For any $0 \leq \delta \leq 1$, if $\Delta(D, \mathcal{D}_{\text{kwi}}) \leq \delta$, then for any non-zero vector \mathbf{a} of weight at most k , $|\hat{D}^{\text{non}}(\mathbf{a})| \leq q\theta_{\max}\delta$.*

Proof. Recall that we compute the non-uniform Fourier coefficient of D at \mathbf{a} by first projecting D to $\text{supp}(\mathbf{a})$ and then apply a compressing/stretching factor to each marginal probability in $D_{\text{supp}(\mathbf{a})}$. Let D' be any k -wise independent distribution with $\Delta(D, D') \leq \delta$. For every $0 \leq j \leq q-1$, let $P_{\mathbf{a},j}^{\text{non}}$ and $P'_{\mathbf{a},j}^{\text{non}}$ be the total probability mass of points in D and D' that satisfy $\mathbf{a} \cdot \mathbf{z} \equiv j \pmod{q}$ after applying the compressing/stretching factors. By the definitions of statistical distance and θ_{\max} , we have

$$\begin{aligned} |P_{\mathbf{a},j}^D - 1/q| &= |P_{\mathbf{a},j}^{\text{non}} - P'_{\mathbf{a},j}^{\text{non}}| \\ &= \left| \sum_{\mathbf{a} \cdot \mathbf{z} \equiv j \pmod{q}} (D_{\text{supp}(\mathbf{a})}(\mathbf{z}) - D'_{\text{supp}(\mathbf{a})}(\mathbf{z})) \theta_{\text{supp}(\mathbf{a})}(\mathbf{z}) \right| \\ &\leq \sum_{\mathbf{a} \cdot \mathbf{z} \equiv j \pmod{q}} |(D_{\text{supp}(\mathbf{a})}(\mathbf{z}) - D'_{\text{supp}(\mathbf{a})}(\mathbf{z})) \theta_{\text{supp}(\mathbf{a})}(\mathbf{z})| \\ &\leq \theta_{\max} \sum_{\mathbf{a} \cdot \mathbf{z} \equiv j \pmod{q}} |D_{\text{supp}(\mathbf{a})}(\mathbf{z}) - D'_{\text{supp}(\mathbf{a})}(\mathbf{z})| \\ &\leq \theta_{\max} \delta. \end{aligned}$$

Now applying Fact 2.8 gives the claimed bound. □

For simplicity, in the following we use $M^{\text{non}}(n, k, q) \stackrel{\text{def}}{=} O\left(n^{\frac{k^2-k+2}{2}} q^{k(k+1)}\right)$ to denote the bound in Theorem 1.2.

Theorem 5.16. *There is an algorithm that tests the non-uniform k -wise independence over Σ^n with query complexity $\tilde{O}\left(\frac{\theta_{\max}^2 n^{(k^2-k+2)} q^{2(k^2+2k+1)}}{\epsilon^2}\right)$ and time complexity $\tilde{O}\left(\frac{\theta_{\max}^2 n^{(k^2+2)} q^{(2k^2+5k+2)}}{\epsilon^2}\right)$ and satisfies the following: for any distribution D over Σ^n , if $\Delta(D, \mathcal{D}_{\text{kwi}}) \leq \frac{\epsilon}{3q\theta_{\max}M^{\text{non}}(n,k,q)}$, then with probability at least $2/3$, the algorithm accepts; if $\Delta(D, \mathcal{D}_{\text{kwi}}) > \epsilon$, then with probability at least $2/3$, the algorithm rejects.*

Test-Non-Uniform-KWI(D, k, q, ϵ)

1. Sample D independently $m = O\left(\frac{q^{2(k+1)}\theta_{\max}^2 M^{\text{non}}(n, k, q)^2}{\epsilon^2} \log(M(n, k, q))\right)$ times
2. Use the samples to estimate, for each non-zero vector \mathbf{a} of weight at most k and each $\mathbf{z} \in \Sigma^{|\text{supp}(\mathbf{a})|}$, $D_{\text{supp}(\mathbf{a})}(\mathbf{z})$
 - Compute $D'_{\text{supp}(\mathbf{a})}(\mathbf{z}) = \theta_S(\mathbf{z})D_{\text{supp}(\mathbf{a})}(\mathbf{z})$
 - Compute $\hat{D}^{\text{non}}(\mathbf{a}) \stackrel{\text{def}}{=} \hat{D}'_{\text{supp}(\mathbf{a})}(\mathbf{a}) = \sum_{\mathbf{z}} D'_{\text{supp}(\mathbf{a})}(\mathbf{z})e^{\frac{2\pi i}{q}\mathbf{a}\cdot\mathbf{z}}$
3. If $\max_{\mathbf{a}} \left| \hat{D}^{\text{non}}(\mathbf{a}) \right| \leq \frac{2\epsilon}{3M^{\text{non}}(n, k, q)}$ return **“Accept”**; else return **“Reject”**

Figure 4: Algorithm for testing non-uniform k -wise independence.

We now briefly sketch the proof of Theorem 5.16. Instead of estimating $P_{\mathbf{a},j}$ as in the proof Theorem 2.7, we estimate $D_{\text{supp}(\mathbf{a})}(\mathbf{z})$ for every \mathbf{z} such that $\mathbf{a} \cdot \mathbf{z} \equiv j \pmod{q}$. Since each $P_{\mathbf{a},j}^{\text{non}}$ is the sum of at most q^k terms, where each term is some $D_{\text{supp}(\mathbf{a})}(\mathbf{z})$ multiplied by a factor at most θ_{\max} , it suffices to estimate each $D_{\text{supp}(\mathbf{a})}(\mathbf{z})$ within additive error $\epsilon/3qM^{\text{non}}(n, k, q)q^k\theta_{\max}$. The soundness part follows directly from Claim 5.15.

5.5 Testing Algorithm when the Marginal Probabilities are Unknown

If the one-dimensional marginal probabilities $p_i(z)$ are not known, we can first estimate these probabilities by sampling the distribution D and then plug these empirical estimates into the testing algorithm shown in Fig 4. The only difference between this case and the known probabilities case is that we need to deal with errors from two sources: apart from those in estimating $D_{\text{supp}(\mathbf{a})}(\mathbf{z})$ there are additional errors from estimating the compressing/stretching factors. It turns out that the query and time complexity are essentially the same when all the one-dimensional marginal probabilities are bounded away from zero.

In the following we write $p_{\min} = \min_{i,z} p_i(z)$ for the minimum one-dimensional marginal probability. Note that $\theta_{\max} \leq (qp_{\min})^{-k}$.

Theorem 5.17. *There is an algorithm that tests the non-uniform k -wise independence over Σ^n where the one-dimensional marginal probabilities of the given distribution are unknown. The algorithm has query complexity $\tilde{O}\left(\frac{\theta_{\max}^2 n^{(k^2-k+2)} q^{2(k^2+2k+1)}}{\epsilon^2 p_{\min}}\right)$ and time complexity $\tilde{O}\left(\frac{\theta_{\max}^2 n^{(k^2+2)} q^{(2k^2+5k+2)}}{\epsilon^2 p_{\min}}\right)$ and satisfies the following: for any distribution D over Σ^n , if $\Delta(D, \mathcal{D}_{\text{kwi}}) \leq \frac{\epsilon}{3q\theta_{\max}M^{\text{non}}(n, k, q)}$, then with probability at least $2/3$, the algorithm accepts; if $\Delta(D, \mathcal{D}_{\text{kwi}}) > \epsilon$, then with probability at least $2/3$, the algorithm rejects.*¹⁶

Proof. The algorithm is essentially the same as that of Theorem 5.16 shown in Fig 4. The only difference is that this new algorithm first uses $m = \tilde{O}\left(\frac{\theta_{\max}^2 n^{(k^2-k+2)} q^{2(k^2+2k+1)}}{\epsilon^2 p_{\min}}\right)$ samples to estimate, for each $1 \leq i \leq n$ and each $\mathbf{z} \in \Sigma$, the one-dimensional marginal probability $p_i(z)$. Denote the estimated marginal probabilities by $p'_i(z)$ and similarly the estimated compressing/stretching factors by $\theta'_S(z)$. After that, the algorithm

¹⁶Note that if p_{\min} is extremely small, the query and time complexity of the testing algorithm can be superpolynomial. One possible fix for this is to perform a “cutoff” on the marginal probabilities. That is, if any of the estimated marginal probabilities is too small, we simply treat it as zero. Then we test the input distribution against some k -wise independent distribution over a product space. We leave this as an open question for future investigation.

uses the same samples to estimate, for every non-zero \mathbf{a} of weight at most k and every \mathbf{z} , the projected probability $D_{\text{supp}(\mathbf{a})}(\mathbf{z})$. Then it uses these probabilities together with the estimated one-dimensional marginal probabilities to calculate $\hat{D}^{\text{non}}(\mathbf{a})$.

By Chernoff bound, for every $p'_i(\mathbf{z})$, with probability at least $1 - 1/6q^n$, we have $1 - \epsilon' \leq \frac{p'_i(\mathbf{z})}{p_i(\mathbf{z})} \leq 1 + \epsilon'$, where $\epsilon' = \epsilon/(12kq\theta_{\max}M^{\text{non}}(n, k, q))$. Therefore by union bound, with probability at least $5/6$, all the estimated one-dimensional marginal probabilities have at most $(1 \pm \epsilon')$ multiplicative errors.

It is easy to verify by Taylor's expansion that for any fixed integer $k > 1$, $(1 + y)^k \leq 1 + 2ky$ for all $0 \leq y \leq 1/(k - 1)$. Also by Bernoulli's inequality, $(1 - y)^k \geq 1 - ky$ for all $0 \leq y \leq 1$. Combining these two facts with the multiplicative error bound for $p'_i(\mathbf{z})$, we get that with probability at least $5/6$ all the estimated compressing/stretching factors have at most $(1 \pm 2k\epsilon')$ multiplicative errors, as every such factor is a product of at most k factors of the form $1/qp_i(\mathbf{z})$.

Also by Chernoff bound, we have with probability at least $5/6$,

$$|\bar{D}_{\text{supp}(\mathbf{a})}(\mathbf{z}) - D_{\text{supp}(\mathbf{a})}(\mathbf{z})| \leq \frac{\epsilon}{12qM^{\text{non}}(n, k, q)q^k\theta_{\max}}$$

for every \mathbf{a} and \mathbf{z} .

Define

$$P_{\mathbf{a},j}^{\text{non def}} = \sum_{\mathbf{a} \cdot \mathbf{z} \equiv j \pmod{q}} D_{\text{supp}(\mathbf{a})}(\mathbf{z})\theta_{\text{supp}(\mathbf{a})}(\mathbf{z})$$

as the ‘‘non-uniform’’ $P_{\mathbf{a},j}$.

Our estimated value of $P_{\mathbf{a},j}^{\text{non def}}$, denoted by $\bar{P}_{\mathbf{a},j}^{\text{non}}$, is in fact

$$\bar{P}_{\mathbf{a},j}^{\text{non}} = \sum_{\mathbf{a} \cdot \mathbf{z} \equiv j \pmod{q}} \bar{D}_{\text{supp}(\mathbf{a})}(\mathbf{z})\theta'_{\text{supp}(\mathbf{a})}(\mathbf{z}),$$

where $\bar{D}_{\text{supp}(\mathbf{a})}(\mathbf{z})$ denotes the empirical estimate of $D_{\text{supp}(\mathbf{a})}(\mathbf{z})$. To simplify notation, in the following we write $P(\mathbf{z}) = D_{\text{supp}(\mathbf{a})}(\mathbf{z})$, $\bar{P}(\mathbf{z}) = \bar{D}_{\text{supp}(\mathbf{a})}(\mathbf{z})$, $\theta(\mathbf{z}) = \theta_{\text{supp}(\mathbf{a})}(\mathbf{z})$ and $\theta'(\mathbf{z}) = \theta'_{\text{supp}(\mathbf{a})}(\mathbf{z})$.

Putting the two error estimates together, we have with probability at least $2/3$, for every \mathbf{a} and j

$$\begin{aligned} |\bar{P}_{\mathbf{a},j}^{\text{non}} - P_{\mathbf{a},j}^{\text{non def}}| &= \left| \sum_{\mathbf{a} \cdot \mathbf{z} \equiv j \pmod{q}} \bar{P}(\mathbf{z})\theta'(\mathbf{z}) - P(\mathbf{z})\theta(\mathbf{z}) \right| \\ &= \left| \sum_{\mathbf{a} \cdot \mathbf{z} \equiv j \pmod{q}} \bar{P}(\mathbf{z})\theta'(\mathbf{z}) - P(\mathbf{z})\theta'(\mathbf{z}) + P(\mathbf{z})\theta'(\mathbf{z}) - P(\mathbf{z})\theta(\mathbf{z}) \right| \\ &\leq \left| \sum_{\mathbf{a} \cdot \mathbf{z} \equiv j \pmod{q}} \bar{P}(\mathbf{z})\theta'(\mathbf{z}) - P(\mathbf{z})\theta'(\mathbf{z}) \right| + \left| \sum_{\mathbf{a} \cdot \mathbf{z} \equiv j \pmod{q}} P(\mathbf{z})\theta'(\mathbf{z}) - P(\mathbf{z})\theta(\mathbf{z}) \right| \\ &\leq \sum_{\mathbf{a} \cdot \mathbf{z} \equiv j \pmod{q}} \theta'(\mathbf{z})|\bar{P}(\mathbf{z}) - P(\mathbf{z})| + \sum_{\mathbf{a} \cdot \mathbf{z} \equiv j \pmod{q}} P(\mathbf{z})|\theta'(\mathbf{z}) - \theta(\mathbf{z})| \\ &\leq 2\theta_{\max} \sum_{\mathbf{a} \cdot \mathbf{z} \equiv j \pmod{q}} |\bar{P}(\mathbf{z}) - P(\mathbf{z})| + (2k\epsilon')\theta_{\max} \sum_{\mathbf{a} \cdot \mathbf{z} \equiv j \pmod{q}} P(\mathbf{z}) \end{aligned}$$

$$\begin{aligned} &\leq 2\theta_{\max} q^k \frac{\epsilon}{12qM^{\text{non}}(n, k, q)q^k\theta_{\max}} + 2k\theta_{\max} \frac{\epsilon}{12kq\theta_{\max}M^{\text{non}}(n, k, q)} \\ &= \frac{\epsilon}{3qM^{\text{non}}(n, k, q)}. \end{aligned}$$

The rest of the proof is similar to that of Theorem 2.7 so we omit the details. \square

6 Almost k -wise Independence over Product Spaces

We study the problem of testing the almost k -wise independent distributions over product spaces in this section. In the following we will follow [1] and define the almost k -wise independence in terms of max-norm.

Definition 6.1 (Uniform Almost k -wise Independence). Let Σ be a finite set with $|\Sigma| = q$. A discrete probability distribution D over Σ^n is (uniform) (ϵ, k) -wise independent if for any set of k indices $\{i_1, \dots, i_k\}$ and for all $z_1, \dots, z_k \in \Sigma$,

$$\left| \Pr_{\mathbf{X} \sim D} [X_{i_1} \cdots X_{i_k} = z_1 \cdots z_k] - 1/q^k \right| \leq \epsilon.$$

Generalizing this definition to non-uniform almost k -wise independence over product spaces is straightforward.

Definition 6.2 (Non-uniform Almost k -wise Independence over Product Spaces). Let $\Sigma_1, \dots, \Sigma_n$ be finite sets. A discrete probability distribution D over $\Sigma_1 \times \cdots \times \Sigma_n$ is (non-uniform) (ϵ, k) -wise independent if for any set of k indices $\{i_1, \dots, i_k\}$ and for all $z_{i_1} \in \Sigma_{i_1}, \dots, z_{i_k} \in \Sigma_{i_k}$,

$$\left| \Pr_{\mathbf{X} \sim D} [X_{i_1} \cdots X_{i_k} = z_{i_1} \cdots z_{i_k}] - \Pr_{\mathbf{X} \sim D} [X_{i_1} = z_{i_1}] \times \cdots \times \Pr_{\mathbf{X} \sim D} [X_{i_k} = z_{i_k}] \right| \leq \epsilon.$$

From now on we will work with the most general notion of the almost k -wise independence, that is the non-uniform almost k -wise independent distributions over product spaces. Let $\mathcal{D}_{(\epsilon, k)}$ denote the set of all (ϵ, k) -wise independent distributions. The distance between a distribution D and the set of (ϵ, k) -wise independent distributions is the minimum statistical distance between D and any distribution in $\mathcal{D}_{(\epsilon, k)}$, i.e., $\Delta(D, \mathcal{D}_{(\epsilon, k)}) = \inf_{D' \in \mathcal{D}_{(\epsilon, k)}} \Delta(D, D')$. D is said to be δ -far from (ϵ, k) -wise independence if $\Delta(D, \mathcal{D}_{(\epsilon, k)}) > \delta$. We write q_m for $\max_{1 \leq i \leq n} |\Sigma_i|$. To simplify notation, we use vectors $\mathbf{p}_1, \dots, \mathbf{p}_n$ of dimensions $|\Sigma_1|, \dots, |\Sigma_n|$, respectively to denote the marginal probabilities at each coordinates. That is, for every $z_j \in \Sigma_i$, the j^{th} component of \mathbf{p}_i is $\mathbf{p}_i(z_j) = \Pr_{\mathbf{X} \sim D} [X_i = z_j]$. Clearly we have $\sum_{z_j \in \Sigma_i} \mathbf{p}_i(z_j) = 1$ for every $1 \leq i \leq n$.

In the property testing setting, for a given distribution D , we would like to distinguish between the case that D is in $\mathcal{D}_{(\epsilon, k)}$ from the case that D is δ -far from $\mathcal{D}_{(\epsilon, k)}$.

Theorem 6.3. *Given a discrete distribution D over $\Sigma_1 \times \cdots \times \Sigma_n$, there is a testing algorithm with query complexity $O(\frac{k \log(nq_m)}{\epsilon^2 \delta^2})$ and time complexity $\tilde{O}(\frac{(nq_m)^k}{\epsilon^2 \delta^2})$ such that the following holds. If $D \in \mathcal{D}_{(\epsilon, k)}$, then the algorithm accepts with probability at least $2/3$; if D is δ -far from $\mathcal{D}_{(\epsilon, k)}$, then the algorithm rejects with probability at least $2/3$.*

To analyze the testing algorithm we will need the following lemma which, roughly speaking, states that the distance parameter δ can be translated into the error parameter ϵ (up to a factor of ϵ) in the definition of the almost k -wise independence.

Test-AKWI($D, k, \Sigma, \epsilon, \delta$)

1. Sample D independently $Q = O(\frac{k \log(nq_m)}{\epsilon^2 \delta^2})$ times
2. Use the samples to estimate, for every k -subset $I = \{i_1, \dots, i_k\}$ of $[n]$ and every $z_{i_1} \dots z_{i_k}$,
 $\bar{p}_I(z_{i_1} \dots z_{i_k}) \stackrel{\text{def}}{=} \Pr_{\mathbf{X} \sim D}[X_{i_1} \dots X_{i_k} = z_{i_1} \dots z_{i_k}]$
3. Let $p_I(z_{i_1} \dots z_{i_k}) \stackrel{\text{def}}{=} \Pr_{\mathbf{X} \sim D}[X_{i_1} = z_{i_1}] \times \dots \times \Pr_{\mathbf{X} \sim D}[X_{i_k} = z_{i_k}]$
4. If $\max_{I, \mathbf{z}} |\bar{p}_I(z_{i_1} \dots z_{i_k}) - p_I(z_{i_1} \dots z_{i_k})| > \epsilon + \epsilon\delta/2$, return **“Reject”**; else return **“Accept”**

Figure 5: Algorithm for testing almost k -wise independence over product spaces.

Lemma 6.4 ([1]). *Let D be a distribution over $\Sigma_1 \times \dots \times \Sigma_n$. If $\Delta(D, \mathcal{D}_{(\epsilon, k)}) > \delta$, then $D \notin \mathcal{D}_{(\epsilon + \epsilon\delta, k)}$. If $\Delta(D, \mathcal{D}_{(\epsilon, k)}) \leq \delta$, then $D \in \mathcal{D}_{(\epsilon + \epsilon\delta, k)}$.*

Proof. For the first part, suppose $D \in \mathcal{D}_{(\epsilon + \epsilon\delta, k)}$. Let $U_{\mathbf{p}_1, \dots, \mathbf{p}_n}$ denote the distribution that for every $z_1 \dots z_n \in \Sigma_1 \times \dots \times \Sigma_n$, $U_{\mathbf{p}_1, \dots, \mathbf{p}_n}(z_1 \dots z_n) = \mathbf{p}_1(z_1) \dots \mathbf{p}_n(z_n)$. It is easy to check that since $\sum_{z_i} \mathbf{p}_i = 1$, $U_{\mathbf{p}_1, \dots, \mathbf{p}_n}$ is indeed a distribution. Now define a new distribution D' as $D' = (1 - \delta)D + \delta U_{\mathbf{p}_1, \dots, \mathbf{p}_n}$, then one can easily verify that $D' \in \mathcal{D}_{(\epsilon, k)}$, therefore $\Delta(D, \mathcal{D}_{(\epsilon, k)}) \leq \delta$.

For the second part, recall that no randomized procedure can increase the statistical difference between two distributions [37], therefore to project distributions to any set of k coordinates and then look at the probability of finding any specific string of length k can not increase the statistical distance between D and any distribution in $\mathcal{D}_{(\epsilon, k)}$. It follows that when restricted to any k coordinates, the max-norm of D is at most $\epsilon + \delta$. \square

Proof of Theorem 6.3. The testing algorithm is illustrated in Fig. 5. The query complexity and time complexity of the testing algorithm are straightforward to check. Now we prove the correctness of the algorithm. As shown in Fig. 5, we write $\bar{p}_I(z_{i_1} \dots z_{i_k})$ for the estimated probability from the samples, $p_I^D(z_{i_1} \dots z_{i_k})$ for $\Pr_{\mathbf{X} \sim D}[X_{i_1} \dots X_{i_k} = z_{i_1} \dots z_{i_k}]$ and $p_I(z_{i_1} \dots z_{i_k})$ for $\Pr_{\mathbf{X} \sim D}[X_{i_1} = z_{i_1}] \times \dots \times \Pr_{\mathbf{X} \sim D}[X_{i_k} = z_{i_k}]$. Observe that $\mathbf{E}[\bar{p}_I(z_{i_1} \dots z_{i_k})] = p_I^D(z_{i_1} \dots z_{i_k})$. Since $\bar{p}_I(z_{i_1} \dots z_{i_k})$ is the average of Q independent 0/1 random variables, Chernoff bound gives

$$\Pr[|\bar{p}_I(z_{i_1} \dots z_{i_k}) - p_I^D(z_{i_1} \dots z_{i_k})| \geq \epsilon\delta/2] \leq \exp(-\Omega(\epsilon^2 \delta^2 Q)).$$

By setting $Q = C \frac{k \log(nq_m)}{\epsilon^2 \delta^2}$ for large enough constant C and applying a union bound argument to all k -subsets and all possible strings of length k , we get that with probability at least $2/3$, for every I and every z_{i_1}, \dots, z_{i_k} , $|\bar{p}_I(z_{i_1} \dots z_{i_k}) - p_I^D(z_{i_1} \dots z_{i_k})| < \epsilon\delta/2$.

Now if $D \in \mathcal{D}_{(\epsilon, k)}$, then with probability at least $2/3$, for all I and all z_{i_1}, \dots, z_{i_k} , $|p_I^D(z_{i_1} \dots z_{i_k}) - p_I(z_{i_1} \dots z_{i_k})| \leq \epsilon$, so by the triangle inequality $|\bar{p}_I(z_{i_1} \dots z_{i_k}) - p_I(z_{i_1} \dots z_{i_k})| \leq \epsilon + \epsilon\delta/2$. Therefore the algorithm accepts.

If D is δ -far from (ϵ, k) -wise independence, then by Lemma 6.4, $D \notin \mathcal{D}_{(\epsilon + \epsilon\delta, k)}$. That is, there are some I and z_{i_1}, \dots, z_{i_k} such that $|p_I^D(z_{i_1} \dots z_{i_k}) - p_I(z_{i_1} \dots z_{i_k})| > \epsilon + \epsilon\delta$. Then with probability at least $2/3$, $|\bar{p}_I(z_{i_1} \dots z_{i_k}) - p_I(z_{i_1} \dots z_{i_k})| > \epsilon + \epsilon\delta/2$. Therefore the algorithm rejects. \square

Acknowledgments

We would like to thank Tali Kaufman for useful discussions and suggestions, Elchanan Mossel for his enthusiasm in this problem and a helpful conversation and Per Austrin for correspondence on constructing orthogonal real functions. We are most grateful to several anonymous referees who pointed out a critical error in an earlier version of this paper, suggested the approach presented in Section 3, caught errors and omissions and offered numerous comments that greatly improve the presentation of the paper. Finally we thank an anonymous referee for suggesting the current title.

References

- [1] N. Alon, A. Andoni, T. Kaufman, K. Matulef, R. Rubinfeld, and N. Xie. Testing k -wise and almost k -wise independence. In *Proc. 39th Annual ACM Symposium on the Theory of Computing*, pages 496–505, 2007.
- [2] N. Alon, L. Babai, and A. Itai. A fast and simple randomized algorithm for the maximal independent set problem. *Journal of Algorithms*, 7:567–583, 1986.
- [3] N. Alon, O. Goldreich, J. Håstad, and R. Peralta. Simple constructions of almost k -wise independent random variables. *Random Structures and Algorithms*, 3(3):289–304, 1992. Earlier version in FOCS’90.
- [4] N. Alon, O. Goldreich, and Y. Mansour. Almost k -wise independence versus k -wise independence. *Information Processing Letters*, 88:107–110, 2003.
- [5] Y. Azar, J. Naor, and R. Motwani. Approximating probability distributions using small sample spaces. *Combinatorica*, 18(2):151–171, 1998.
- [6] T. Batu, E. Fischer, L. Fortnow, R. Kumar, R. Rubinfeld, and P. White. Testing random variables for independence and identity. In *Proc. 42nd Annual IEEE Symposium on Foundations of Computer Science*, pages 442–451, 2001.
- [7] T. Batu, L. Fortnow, R. Rubinfeld, W. D. Smith, and P. White. Testing that distributions are close. In *Proc. 41st Annual IEEE Symposium on Foundations of Computer Science*, pages 189–197, 2000.
- [8] T. Batu, R. Kumar, and R. Rubinfeld. Sublinear algorithms for testing monotone and unimodal distributions. In *Proc. 36th Annual ACM Symposium on the Theory of Computing*, pages 381–390, New York, NY, USA, 2004. ACM Press.
- [9] C. Bertram-Kretzberg and H. Lefmann. MOD_p -tests, almost independence and small probability spaces. *Random Structures and Algorithms*, 16(4):293–313, 2000.
- [10] E. Blais. Testing juntas nearly optimally. In *Proc. 41st Annual ACM Symposium on the Theory of Computing*, pages 151–158, 2009.
- [11] M. Blum, M. Luby, and R. Rubinfeld. Self-testing/correcting with applications to numerical problems. *Journal of Computer and System Sciences*, 47:549–595, 1993. Earlier version in STOC’90.
- [12] B. Chor and O. Goldreich. On the power of two-point based sampling. *Journal of Complexity*, 5(1):96–106, 1989.

- [13] A. Czumaj and C. Sohler. Sublinear-time algorithms. *Bulletin of the European Association for Theoretical Computer Science*, 89:23–47, 2006.
- [14] R. de Wolf. A brief introduction to Fourier analysis on the boolean cube. *Theory of Computing, Graduate Surveys*, TCGS 1:1–20, 2008.
- [15] I. Diakonikolas, H. Lee, K. Matulef, K. Onak, R. Rubinfeld, R. Servedio, and A. Wan. Testing for concise representations. In *Proc. 48th Annual IEEE Symposium on Foundations of Computer Science*, pages 549–558, 2007.
- [16] K. Efremenko. 3-query locally decodable codes of subexponential length. In *Proc. 41st Annual ACM Symposium on the Theory of Computing*, pages 39–44, 2009.
- [17] G. Even, O. Goldreich, M. Luby, N. Nisan, and B. Velickovic. Efficient approximation of product distributions. *Random Structures and Algorithms*, 13(1):1–16, 1998. Earlier version in STOC’92.
- [18] E. Fischer. The art of uninformed decisions: A primer to property testing. *Bulletin of the European Association for Theoretical Computer Science*, 75, 2001.
- [19] O. Goldreich, S. Goldwasser, and D. Ron. Property testing and its connection to learning and approximation. *Journal of the ACM*, 45:653–750, 1998.
- [20] O. Goldreich and D. Ron. On testing expansion in bounded-degree graphs. Technical Report TR00-020, Electronic Colloquium on Computational Complexity, 2000.
- [21] V. Grolmusz. Superpolynomial size set-systems with restricted intersections mod 6 and explicit ramsey graphs. *Combinatorica*, 20(1):71–86, 2000.
- [22] G. H. Hardy and E. M. Wright. *An Introduction to the Theory of Numbers*. Oxford University Press, 5th edition, 1980.
- [23] A. Joffe. On a set of almost deterministic k -independent random variables. *Annals of Probability*, 2:161–162, 1974.
- [24] H. Karloff and Y. Mansour. On construction of k -wise independent random variables. In *Proc. 26th Annual ACM Symposium on the Theory of Computing*, pages 564–573, 1994.
- [25] R. Karp and A. Wigderson. A fast parallel algorithm for the maximal independent set problem. *Journal of the ACM*, 32(4):762–773, 1985.
- [26] D. Koller and N. Megiddo. Constructing small sample spaces satisfying given constraints. In *Proc. 25th Annual ACM Symposium on the Theory of Computing*, pages 268–277, 1993.
- [27] R. Kumar and R. Rubinfeld. Sublinear time algorithms. *SIGACT News*, 34:57–67, 2003.
- [28] M. Luby. A simple parallel algorithm for the maximal independent set problem. *SIAM Journal on Computing*, 15(4):1036–1053, 1986. Earlier version in STOC’85.
- [29] M. Luby. Removing randomness in parallel computation without a processor penalty. In *Proc. 29th Annual IEEE Symposium on Foundations of Computer Science*, pages 162–173, 1988.

- [30] E. Mossel. Gaussian bounds for noise correlation of functions and tight analysis of long codes. In *Proc. 49th Annual IEEE Symposium on Foundations of Computer Science*, pages 156–165, 2008.
- [31] J. Naor and M. Naor. Small-bias probability spaces: efficient constructions and applications. *SIAM Journal on Computing*, 22(4):838–856, 1993. Earlier version in STOC’90.
- [32] C. P. Neuman and D. I. Schonbach. Discrete (Legendre) orthogonal polynomials - A survey. *International Journal for Numerical Methods in Engineering*, 8:743–770, 1974.
- [33] A. F. Nikiforov, S. K. Suslov, and V. B. Uvarov. *Classical Orthogonal Polynomials of a Discrete Variable*. Springer-Verlag, 1991.
- [34] S. Raskhodnikova, D. Ron, A. Shpilka, and A. Smith. Strong lower bounds for approximating distribution support size and the distinct elements problem. In *Proc. 48th Annual IEEE Symposium on Foundations of Computer Science*, pages 559–569, 2007.
- [35] D. Ron. Property testing (a tutorial). In P.M. Pardalos, S. Rajasekaran, J. Reif, and J.D.P. Rolim, editors, *Handbook of Randomized Computing*, pages 597–649. Kluwer Academic Publishers, 2001.
- [36] R. Rubinfeld and M. Sudan. Robust characterizations of polynomials with applications to program testing. *SIAM Journal on Computing*, 25:252–271, 1996.
- [37] A. Sahai and S. Vadhan. A complete problem for statistical zero knowledge. *Journal of the ACM*, 50(2):1–54, 2003.
- [38] J.R. Sylvester. Determinants of block matrices. *Maths Gazette*, 84:460–467, 2000.
- [39] H. J. S. Smith. On systems of linear indeterminate equations and congruences. *Phil. Trans. Royal Soc. London*, A151:293–326, 1861.
- [40] D. Štefankovič. Fourier transform in computer science. Master’s thesis, University of Chicago, 2000.
- [41] A. Terras. *Fourier Analysis on Finite Groups and Applications*. Cambridge University Press, 1999.
- [42] S. Yekhanin. Towards 3-query locally decodable codes of subexponential length. *Journal of the ACM*, 55(1):1–16, 2008.