Sparse Fourier Transforms on Rank-1 Lattices for the Rapid and Low-Memory Approximation of Functions of Many Variables

Craig Gross[*]   Mark Iwen[†]   Lutz Kämmerer[‡]   Toni Volkmer[§]

December 21, 2020

Abstract

This paper considers fast and provably accurate algorithms for approximating smooth functions on the d-dimensional torus, $f : \mathbb{T}^d \to \mathbb{C}$, that are sparse (or compressible) in the multidimensional Fourier basis. In particular, suppose that the Fourier series coefficients of f, $\{c_{\mathbf{k}}(f)\}_{\mathbf{k}\in\mathbb{Z}^d}$, are concentrated in a given arbitrary finite set $\mathcal{I} \subset \mathbb{Z}^d$ so that

$$\min_{\Omega \subset \mathcal{I} \; s.t. \; |\Omega|=s} \left\| f - \sum_{\mathbf{k}\in\Omega} c_{\mathbf{k}}(f)\, \mathbb{e}^{-2\pi i \mathbf{k}\cdot\circ} \right\|_2 < \epsilon\|f\|_2$$

holds for $s \ll |\mathcal{I}|$ and $\epsilon \in (0,1)$ small. In such cases we aim to both identify a near-minimizing subset $\Omega \subset \mathcal{I}$ and accurately approximate its associated Fourier coefficients $\{c_{\mathbf{k}}(f)\}_{\mathbf{k}\in\Omega}$ as rapidly as possible. In this paper we present both deterministic and explicit as well as randomized algorithms for solving this problem using $\mathcal{O}(s^2 d \log^c(|\mathcal{I}|))$-time/memory and $\mathcal{O}(sd \log^c(|\mathcal{I}|))$-time/memory, respectively. Most crucially, all of the methods proposed herein achieve these runtimes while simultaneously satisfying theoretical best s-term approximation guarantees which guarantee their numerical accuracy and robustness to noise for general functions.

These results are achieved by modifying several different one-dimensional Sparse Fourier Transform (SFT) methods to subsample a function along a reconstructing rank-1 lattice for the given frequency set $\mathcal{I} \subset \mathbb{Z}^d$ in order to rapidly identify a near-minimizing subset $\Omega \subset \mathcal{I}$ as above without having use anything about the lattice beyond its generating vector. This requires the development of new fast and low-memory frequency identification techniques capable of rapidly recovering vector-valued frequencies in $\mathbb{Z}^d$ as opposed to recovering simple integer frequencies as required in the univariate setting. Two different multivariate frequency identification strategies are proposed, analyzed, and shown to lead to their own best s-term approximation methods herein, each with different accuracy versus computational speed and memory tradeoffs.

Keywords: Multivariate Fourier approximation, Approximation algorithms, Fast Fourier transforms, Sparse Fourier transforms, Rank-1 lattices, Fast algorithms

Mathematics Subject Classification (2010): 65T40, 65D15, 42B05, 65Y20, 65T50

1 Introduction

This paper considers methods for efficiently computing sparse Fourier transforms of multivariate periodic functions using rank-1 lattices. In particular, for a function $f : \mathbb{T}^d \to \mathbb{C}$ (where $\mathbb{T} := [0,1]$ with the endpoints identified), our goal is to compute the Fourier coefficients of f,

$$c_{\mathbf{k}}(f) := \int_{\mathbb{T}^d} f(\mathbf{x}) \, \mathrm{e}^{-2\pi\mathrm{i}\mathbf{k}\cdot\mathbf{x}} \, \mathrm{d}\mathbf{x},$$

via samples of f at points in $\mathbb{T}^d$. Here, we assume that f is from the Wiener algebra $\mathcal{W}(\mathbb{T}^d) := \{f \in L^1(\mathbb{T}^d) : \|f\|_{\mathcal{W}(\mathbb{T}^d)} := \sum_{\mathbf{k}\in\mathbb{Z}^d} |c_{\mathbf{k}}(f)| < \infty\}$, and that f is well approximated by just a few of the dominant terms in its Fourier expansion (i.e., has an accurate sparse approximation in the Fourier basis).

One quasi-Monte Carlo approach which is especially popular in the context of Fourier approximations is sampling along rank-1 lattices adapted to frequency spaces of interest [35, 36, 26, 24, 29, 17, 20, 30, 23]. In the standard rank-1 lattice approach, a one-dimensional, length-M discrete Fourier transform (DFT) is applied to samples of f along a rank-1 lattice $\Lambda(\mathbf{z}, M)$ with generating vector $\mathbf{z} \in \mathbb{Z}^d$ over $\mathbb{T}^d$ defined by

$$\Lambda(\mathbf{z}, M) := \left\{ \frac{j}{M}\mathbf{z} \bmod 1 \mid j \in [M] := \{0, 1, \ldots, M-1\} \right\}.$$

Writing any function $f \in \mathcal{W}(\mathbb{T}^d)$ in terms of its Fourier series, $f = \sum_{\mathbf{k}\in\mathbb{Z}^d} c_{\mathbf{k}}(f) \, \mathrm{e}^{2\pi\mathrm{i}\mathbf{k}\cdot\circ}$, the DFT of these samples along the lattice is exactly equivalent to the DFT of the univariate function

$$a(t) := \sum_{\mathbf{k}\in\mathbb{Z}^d} c_{\mathbf{k}}(f) \, \mathrm{e}^{2\pi\mathrm{i}\mathbf{k}\cdot\mathbf{z}\,t} \tag{1}$$

using the equispaced samples $\left(a(j/M)\right)_{j\in[M]} = \left(f(\mathbf{x})\right)_{\mathbf{x}\in\Lambda(\mathbf{z},M)}$. Just as the DFT of equispaced samples of a can be used to approximate its Fourier coefficients, so then can this DFT be used to help approximate the original Fourier coefficients $c_{\mathbf{k}}(f)$ of f. Though the process of matching the discrete coefficients to corresponding coefficients of f is nontrivial (see the following paragraph for further discussion), this multivariate to univariate transformation allows us to carry over many standard one-dimensional DFT results in a straightforward manner. In particular, under our assumption of f being Fourier sparse or compressible, one-dimensional sparse Fourier transform (SFT) techniques [11, 12, 15, 14, 25, 34, 16, 10, 7, 22, 31, 32, 27, 2, 1] become particularly appealing as they can sidestep runtimes which depend polynomially on the bandwidth, in this case M, instead running sublinearly in the magnitude of the underlying frequency space under consideration. Additionally, these techniques often furnish recovery guarantees for Fourier compressible functions in terms of best s-term approximations in the same vein as compressed sensing results [8, 9].

However, in order for our univariate DFT to be properly related to the original multivariate Fourier coefficients $c_{\mathbf{k}}(f)$, any aliasing must not produce extraneous collisions which perturb the multivariate to univariate transformation. Specifically, after applying a length-M DFT to the univariate function a in (1), all one-dimensional frequencies $\mathbf{k} \cdot \mathbf{z}$ are aliased to their residues modulo M. Restricting our attention to some finite multivariate frequency set $\mathcal{I} \subset \mathbb{Z}^d$, we then consider *reconstructing* rank-1 lattices for $\mathcal{I}$, that is, lattices where the mapping $m_{\mathbf{z},M} : \mathcal{I} \to [M]$ given by $\mathbf{k} \mapsto \mathbf{k} \cdot \mathbf{z} \bmod M$ is injective. In this case, each coefficient produced by the DFT of a can be uniquely mapped back to the corresponding multivariate frequency $\mathbf{k}$ of f by inverting $m_{\mathbf{z},M}$.

2

In order to know or store this inverse map we require the calculation of $m_{\mathbf{z},M}(\mathcal{I})$. When we consider a function with a sparse Fourier series however, any benefit in using an SFT to calculate the DFT of samples along the lattice is lost in comparison to the $\mathcal{O}(d\,|\mathcal{I}|)$ size and operation count of the inverse computation. For potentially large search spaces of multivariate frequencies $\mathcal{I}$ such as a full cube $\mathcal{I} = \left(\left(-\left\lceil\frac{N}{2}\right\rceil, \left\lfloor\frac{N}{2}\right\rfloor\right] \cap \mathbb{Z}\right)^d$, both the time and memory complexity of this algorithm suffers from the curse of dimensionality.

The methods given in this paper instead work to use samples along possible larger lattices to produce sparse approximations of the Fourier transform of f without directly inverting $m_{\mathbf{z},M}$. The two algorithms considered below are able to operate on SFTs of manipulations of a in order to relate the univariate coefficients to their multivariate counterparts in $o(|\mathcal{I}|)$-time. This will allow the methods developed herein to run faster and with less memory than it takes to simply enumerate the frequency set $\mathcal{I}$ and/or store $m_{\mathbf{z},M}(\mathcal{I})$ whenever f has a sufficiently accurate sparse approximation.

1.1 Prior work

Much recent work has considered the problem of quickly recovering both exactly sparse multivariate trigonometric polynomials as well as approximating more general functions by sparse trigonometric polynomials using dimension-incremental approaches [37, 33, 5, 6]. These methods recover multivariate frequencies adaptively by searching for energetic frequencies on projections of the potential coefficient space $\mathcal{I} \subset \left(\left(-\left\lceil\frac{N}{2}\right\rceil, \left\lfloor\frac{N}{2}\right\rfloor\right] \cap \mathbb{Z}\right)^d$ into lower dimensional spaces. These lower dimensional candidate sets are then paired together to build up a fully d-dimensional search space smaller than the original one, which is expected to support the most energetic frequencies (see e.g., [19, Section 3] and the references within for a general overview).

In the context of Fourier methods, lattice-based techniques work well to handle support identification on the intermediary, lower-dimensional candidate sets, and especially recently, techniques based on multiple rank-1 lattices have shown success [21, 19]. Though the total complexity in each of these steps is manageable and can be kept linear in the sparsity, these steps must be repeated in general to ensure that no potential frequencies have been left out. In particular, this results in at least $\mathcal{O}(ds^2N)$ operations (up to logarithmic factors) for functions supported on arbitrary frequency sets in order to obtain approximations that are guaranteed to be accurate with high probability. Though from an implementational perspective, this runtime can be mitigated by completing many of the repetitions and initial one-dimensional searches in parallel, once pairing begins, the results of previous iterations must be synchronized and communicated to future steps, necessitating serial interruptions.

Other earlier works include [16] in which previously existing univariate SFT results [15, 34] were refined and adapted to the multivariate setting. Though the resulting complexity on the dimension is well above the dimension-incremental approaches, deterministic guarantees are given for multivariate Fourier approximation in $\mathcal{O}(d^4 s^2)$ (up to logarithmic factors) time and memory, as well as a random variant which dropped to linear scaling in s, leading to a runtime on the order of $\mathcal{O}(d^4 s)$ with respect to s and d. Additionally, the compressed sensing type guarantees in terms of Fourier compressibility of the function under consideration carry over from the univariate SFT analysis. The scheme essentially makes use of a reconstructing rank-1 lattice on a superset of the full integer cube $\mathcal{I} = \left(\left(-\left\lceil\frac{dN}{2}\right\rceil, \left\lfloor\frac{dN}{2}\right\rfloor\right] \cap \mathbb{Z}\right)^d$ with certain number theoretic properties that allow for fast inversion of the resulting one-dimensional coefficients by the Chinese Remainder Theorem.

We note that this necessarily inflated frequency domain accounts for the suboptimal scaling in d above.

In [28], another fully deterministic sampling strategy and reconstruction algorithm is given. Like [16] though, the method can only be applied to Fourier approximations over an ambient frequency space $\mathcal{I}$ that is full d-dimensional cube. Moreover, the vector space structure exploited to construct the sampling sets necessitates that the side length N of this cube is the power of a prime. However, the benefits to this construction are among the best considered so far: the method is entirely deterministic, has noise-robust recovery guarantees in terms of best s-term estimates, the sampling sets used are on the order of $\mathcal{O}(d^3 s^2 N)$, and the reconstruction algorithm's runtime complexity is on the order of $\mathcal{O}(d^2 s^2 N^2)$ both up to logarithmic factors. On the other hand, this algorithm still does not scale linearly in s.

Finally, we discuss [4, 3], a pair of papers detailing high-dimensional Fourier recovery algorithms which offer a simplified (and therefore faster) approach to lattice transforms and dimension-incremental methods. These algorithms make heavy use of a one-dimensional SFT [25, 7] based on a phase modulation approach to discover energetic frequencies in a fashion similar to our Algorithm 1 below. The main idea is to recover entries of multivariate frequencies by using equispaced evaluations of the function along a coordinate axis as well as samples of the function at the same points slightly shifted (the remaining dimensions are generally ignored). This shift in space produces a modulation in frequency from which frequency data can be recovered (cf. Lemma 4 and Algorithm 1). By supplementing this approach with simple reconstructing rank-1 lattice analysis for repetitions of the full integer cube, the runtime and number of samples are given on average as $\mathcal{O}(ds)$ up to logarithmic factors.

However, due to the possibility of collisions of multivariate frequencies under the hashing algorithms employed, these results hold only for random signal models. In particular, theoretical results are only stated for functions with randomly generated Fourier coefficients on the unit circle with randomly chosen frequencies from a given frequency set. Additionally, the analysis of these techniques assumes that the algorithm applied to the randomly generated signal does not encounter certain low probability (with respect to the random signal model considered therein) energetic frequency configurations. Furthermore, the method is restricted in stability, allowing for spatial shifts in sampling bounded by at most the reciprocal of the side length of the multivariate frequency cube under consideration, and only exact recovery is considered (or recovery up to factors related to sample corruption by gaussian noise in [3]). In addition, no results given are proven concerning the approximation of more general periodic functions, e.g., compressible functions.

1.2   Main contributions

We begin with a brief summary of the benefits provided by our approach in comparison to the methods discussed above. Below, we ignore logarithmic factors in our summary of the runtime/sampling complexities.

- All variants, deterministic and random, of both algorithms presented in this paper have runtime and sampling complexities **linear in** d with best s-term estimates for **arbitrary signals**. This is in contrast to the complexities of dimension-incremental approaches [6, 5, 21, 19] and the number theoretic approaches [16, 28] while still achieving similarly strong best s-term guarantees.

4

- Both algorithms proposed herein have randomized variants with runtime and sampling complexities **linear in** s with best s-term estimates on **arbitrary signals** that hold **with high probability**. Thus, the randomized methods proposed in this paper achieve the efficient runtime complexities of [4, 3] while simultaneously exhibiting best s-term approximation guarantees for general periodic functions thereby improving on the non-deterministic dimension incremental approaches [6, 5, 21, 19].

- Both algorithms proposed herein have a deterministic variant with runtime and sampling complexities **quadratic** s with best s-term estimates on **arbitrary signals** that also hold **deterministically**. This is in contrast to all previously discussed methods without deterministic guarantees, [6, 5, 21, 19, 3, 4], as well as improving on prior deterministic results [16, 28] for functions whose energetic frequency support sets $\mathcal{I}$ are smaller than the full cube.

Overview of the methods and related theory

We will build on the structure of the fast and potentially deterministic one-dimensional SFTs from [16] and its discrete variant from [27] by applying those techniques along rank-1 lattices. As previously discussed, the primary difficulty in doing so is determining a way to extract the desired multidimensional frequency information for those most energetic one-dimensional frequencies identified in an efficient and provably accurate way. We propose and analyze two different methods for solving this problem herein.

The first frequency identification approach involves modifications of the phase shifting from [25, 7, 4, 3] in Algorithm 1. By employing the phase shifting process from these works in conjunction with SFTs over an arbitrary reconstructing rank-1 lattice for our multivariate frequency search space $\mathcal{I}$, we achieve a new class of fast method with several new benefits. Notably, we are able to maintain error guarantees for any function (not just random signals) in terms of best Fourier s-term approximations. Additionally, we factor the instability and potential for collisions from [4, 3] into these best s-term approximations, suffering only a linear factor of N from the more typical results produced by our second technique discussed in the next paragraph (cf. Corollaries 1 and 2). Finally, we are able to maintain quadratic in s/deterministic and linear in s/random runtime and sampling complexities while reducing the dependence on the dimension of the function's domain d from [16] to a linear one (cf. Lemma 5).

Our second technique in Algorithm 2 uses a more novel approach to applying SFTs to modifications of the multivariate function along a reconstructing rank-1 lattice. By using a potentially larger rank-1 lattice than one that suffices only for being reconstructing on $\mathcal{I}$, we restrict the function to only $d-1$ dimensions of the lattice at a time, allowing one dimension to remain free. Applying SFTs along the lattice constrained variables, FFTs in the free dimension, and synchronizing based on known Fourier coefficients (for example, from simply applying an SFT on the entire function restricted to the lattice) allows one to reconstruct the full multivariate coefficients with fewer possibilities for numerical instability. In particular this produces more accurate best Fourier s-term approximation guarantees (cf. Corollaries 3 and 4). We again maintain the linear in d, quadratic in s/deterministic, and linear in s/random sampling and runtime complexities, however, we now additionally incur a penalty of a quadratic factor of N (cf. Lemma 7).

We stress here that by compartmentalizing the translation from multivariate analysis to univariate analysis in Algorithms 1 and 2 into the theory of rank-1 lattices, we additionally abstract our interaction with the multivariate frequency domain of interest. As such, our techniques are suitable

for any arbitrary frequency set of interest $\mathcal{I}$ constrained only by our necessity for a reconstructing rank-1 lattice for $\mathcal{I}$ (and potentially projections of $\mathcal{I}$ in the case of Algorithm 2). This flexibility allows our methods to supersede the results from [16], primarily with respect to the polynomial factor of d in our runtime and sampling complexities. We remark that though the existence of the necessary reconstructing rank-1 lattice is a nontrivial requirement, there exist efficient construction algorithms for arbitrary frequency sets via deterministic component by component methods, see e.g., [18, 23, 30].

We also emphasize that the multivariate structures we employ are also entirely modular with respect to their underlying univariate components. More specifically, they can make use of any reasonably fast and theoretically sound SFT approach to produce resulting runtime and sampling estimates which scale well in the complexity of the underlying SFT algorithm (by only a factor of d in Algorithm 1 and a factor of dN in Algorithm 2). Lemmas 5 and 7 are therefore presented in a similarly modular fashion. To provide specific recovery results we therefore use the univariate SFTs from [16, 27]. Notably, we also improve the theoretical approximation guarantees of these univariate SFTs in their own right and in the process include the addition of a robust variant of the discrete method in [27].

Finally, the methods we present are trivially parallelizable so that in particular, a large majority of the additional factors of d or dN respectively in the runtimes of Algorithm 1 or Algorithm 2 discussed above can occur in parallel.

1.3 Organization

The remainder of this paper is presented as follows: in Section 2, we set the notation, the notions of the Fourier transform, and the various types of manipulations we will be using in the sequel. Section 3 reviews and further refines the univariate SFTs from [16, 27] to suit our multivariate analysis. Section 4 presents our main multivariate approximation algorithms and their analysis. In particular, Section 4.1 discusses the phase-shifting approach, while Section 4.2 discusses the two-dimensional SFT/DFT combination approach. Finally, we implement these two algorithms numerically and present the empirical results in Section 5.

2 Notation and assumptions

2.1 Multivariate

We begin by defining a one-dimensional frequency band of length N as $\mathcal{B}_N := \left(-\left\lceil\frac{N}{2}\right\rceil, \left\lfloor\frac{N}{2}\right\rfloor\right] \cap \mathbb{Z}$. For a potentially large but finite multivariate frequency set $\mathcal{I}$, which we think of as containing the most significant frequencies of the function under consideration, we choose $N = \max_{\ell\in[d]}(\max_{\mathbf{k}\in\mathcal{I}} k_\ell - \min_{\tilde{\mathbf{k}}\in\mathcal{I}} \tilde{k}_\ell) + 1$ as the minimal width such that $\mathcal{I} \subset \mathbf{h} + \mathcal{B}_N^d$ for some $\mathbf{h} \in \mathbb{Z}^d$. By appropriately modulating any multivariate function $f : \mathbb{T}^d \to \mathbb{C}$ under consideration, i.e., considering $e^{-2\pi i \mathbf{h}\cdot\circ} f$, we shift the frequencies of Fourier coefficients of f originally in $\mathcal{I}$ to $\mathcal{I} - \mathbf{h} \subset \mathcal{B}_N^d$. Thus, we assume without loss of generality that $\mathcal{I} \subset \mathcal{B}_N^d$ with N as above. Without loss of generality, we will also assume that for a reconstructing rank-1 lattice $\Lambda(\mathbf{z}, M)$, the generating vector satisfies $\mathbf{z} \in [M]^d$.

To avoid confusion with the hat notation which will be reserved for univariate functions below, we denote the sequence of all Fourier coefficients (i.e., the Fourier transform) of a periodic function $f : \mathbb{T}^d \to \mathbb{C}$ as $c(f) = (c_{\mathbf{k}}(f))_{\mathbf{k}\in\mathbb{Z}^d}$, also writing this as simply c when the function is clear from context. Its restriction to $\mathcal{I}$ is denoted $c(f)|_{\mathcal{I}} = (c_{\mathbf{k}}(f))_{\mathbf{k}\in\mathcal{I}}$, and the best s-term approximation,

that is, its restriction to the support of the s-largest magnitude entries, is denoted c_s^{opt} or $(c|_{\mathcal{I}})_s^{\mathrm{opt}}$ on $\mathbb{Z}^d$ or $\mathcal{I}$, respectively. We denote multiindexed vectors only defined on finite index sets (which are not restrictions of infinitely indexed sequences) in boldface, e.g., $\mathbf{b} = (b_{\mathbf{k}})_{\mathbf{k} \in \mathcal{I}}$, as well as identify this multivariate vector as a one-dimensional vector $\mathbf{b} \in \mathbb{C}^{|\mathcal{I}|}$ via lexicographic ordering when dictated by context. Again dictated by context, we also extend these multiindexed vectors to larger index sets by setting them to zero outside of their original domain. For example, if $\mathbf{b} = (b_{\mathbf{k}})_{\mathbf{k} \in \mathcal{I}}$ and $c = (c_{\mathbf{k}})_{\mathbf{k} \in \mathbb{Z}^d}$,

$$\|\mathbf{b} - c\|_{\ell^1(\mathbb{Z}^d)} = \sum_{\mathbf{k} \in \mathcal{I}} |b_{\mathbf{k}} - c_{\mathbf{k}}| + \sum_{\mathbf{k} \in \mathbb{Z}^d \setminus \mathcal{I}} |c_{\mathbf{k}}|.$$

In the multivariate approaches which follow, we will also make use of the shift operator $S_{\ell,\alpha}$ for dimension $\ell \in [d]$ and shift $\alpha \in \mathbb{R}$ defined by its action on the multivariate periodic function $f : \mathbb{T}^d \to \mathbb{C}$ as

$$S_{\ell,\alpha}(f)(x_1, \ldots, x_d) := f(x_1, \ldots, x_{\ell-1}, (x_\ell + \alpha) \bmod 1, x_{\ell+1}, \ldots, x_d).$$

When necessary, we will separate out coordinate ℓ of a multivariate point $\mathbf{x} \in \mathbb{T}^d$ or frequency $\mathbf{k} \in \mathbb{Z}^d$, denoting the remaining coordinates as $\mathbf{x}'_\ell \in \mathbb{T}^{d-1}$ or $\mathbf{k}'_\ell \in \mathbb{Z}^{d-1}$. With a slight abuse of notation, we can rewrite the original point or frequency as $\mathbf{x} = (x_\ell, \mathbf{x}'_\ell)$ or $\mathbf{k} = (k_\ell, \mathbf{k}'_\ell)$.

2.2 Univariate

For any univariate periodic function $a : \mathbb{T} \to \mathbb{C}$, we define the vector $\mathbf{a} \in \mathbb{C}^M$ as the vector of M equispaced samples of a on $\mathbb{T}$, that is, $\mathbf{a} = (a(j/M))_{j \in [M]}$. As in the multivariate case, we define the Fourier transform of $a : \mathbb{T} \to \mathbb{C}$ as the sequence $\hat{a} = (\hat{a}_\omega)_{\omega \in \mathbb{Z}}$ with

$$\hat{a}_\omega := \int_{\mathbb{T}} a(t)\, e^{-2\pi i \omega x}\, \mathrm{d}x \text{ for all } \omega \in \mathbb{Z}.$$

Additionally, we define the vector $\hat{\mathbf{a}} \in \mathbb{C}^M$ as the restriction of $\hat{a}$ to $\mathcal{B}_M$. If not explicitly stated, the length of the discretized function $\mathbf{a}$ and Fourier transform $\hat{\mathbf{a}}$ will be clear from context. Note that $\hat{\mathbf{a}}$ is not necessarily the discrete Fourier transform of $\mathbf{a}$, which we define as

$$(\mathbf{F}_M\, \mathbf{a})_\omega := \frac{1}{M} \sum_{j \in [M]} a_j\, e^{-2\pi i \omega j / M} = \frac{1}{M} \sum_{j \in [M]} a\left(\frac{j}{M}\right) e^{-2\pi i \omega j / M}, \text{ where}$$

$$\mathbf{F}_M := \left(e^{-2\pi i \omega j / M} / M \right)_{j \in [M], \omega \in \mathcal{B}_M}$$

is the discrete Fourier matrix. Our convention here and in the remainder of the paper is to use zero-based indexing which is always taken implicitly modulo the length of the dimension, e.g., $(\mathbf{F}_M)_{0,-1} = (\mathbf{F}_M)_{0,M-1}$.

For any vector $\mathbf{b} \in \mathbb{C}^M$, we denote its best s-term approximation $\mathbf{b}_s^{\mathrm{opt}}$, where as above, $\mathbf{b}_s^{\mathrm{opt}}$ is the restriction of $\mathbf{b}$ to its s largest magnitude entries. In the sequel, we always assume that our sparsity parameters s are at most half the size of the vectors under consideration so that, e.g., $\mathbf{b}_{2s}^{\mathrm{opt}}$ is well-defined. Additionally as above, vectors can also be compared with other vectors on larger index sets than they are defined by simply setting the smaller vectors to zero outside of their original domain.

As for one-dimensional approximations, we will be considering SFT algorithms which, given sparsity parameter s and bandwidth M, produce an s-sparse approximation of the Fourier transform of a function $a \in C(\mathbb{T})$ restricted to $\mathcal{B}_M$. Note that these are not necessarily discrete algorithms which take in $\mathbf{a}$ as input. We denote these algorithms $\mathcal{A}_{s,M} : C(\mathbb{T}) \to \mathbb{C}^M$, which produce $\mathcal{A}_{s,M}a =: \mathbf{v} \in \mathbb{C}^M$ as approximations to $\hat{\mathbf{a}} \in \mathbb{C}^M$ using some fixed number of samples of a.

3 One-dimensional sparse Fourier transform results

Below, we summarize some of the previous work on one-dimensional sparse Fourier transforms which will be used in our multivariate algorithms. Note that we will consider algorithms which produce $2s$-sparse approximations of the Fourier coefficients of a given signal and satisfy error guarantees in terms of the best $2s$ *and* s-term approximations. We first review the sublinear-time algorithm from [16] which uses fewer than M nonequispaced samples of a function. Below, we will present slightly improved error bounds which necessitate the following lemma.

Lemma 1. *For* $\mathbf{x} \in \mathbb{C}^K$ *and* $\mathcal{S}_\tau := \{k \in [K] \mid |x_k| \geq \tau\}$, *if* $\tau \geq \frac{\|\mathbf{x} - \mathbf{x}_s^{\mathrm{opt}}\|_1}{s}$, *then* $|\mathcal{S}_\tau| \leq 2s$ *and*

$$\|\mathbf{x} - \mathbf{x}|_{\mathcal{S}_\tau}\|_2 \leq \|\mathbf{x} - \mathbf{x}_{2s}^{\mathrm{opt}}\|_2 + \tau\sqrt{2s}.$$

Proof. Ordering the entries of $\mathbf{x}$ in descending order (with ties broken arbitrarily) as $|x_{k_1}| \geq |x_{k_2}| \geq \ldots$, we first note that

$$\|\mathbf{x} - \mathbf{x}_s^{\mathrm{opt}}\|_1 \geq \sum_{j=s+1}^{2s} |x_{k_j}| \geq s|x_{k_{2s}}|.$$

By assumption then, $\tau \geq |x_{k_{2s}}|$, and since $\mathcal{S}_\tau$ contains the $|\mathcal{S}_\tau|$-many largest entries of $\mathbf{x}$, we must have $\mathcal{S}_\tau \subset \mathrm{supp}(\mathbf{x}_{2s}^{\mathrm{opt}})$. Note then that $|\mathcal{S}_\tau| \leq 2s$. Finally, we calculate

$$
\begin{aligned}
\|\mathbf{x} - \mathbf{x}|_{\mathcal{S}_\tau}\|_2 &\leq \|\mathbf{x} - \mathbf{x}_{2s}^{\mathrm{opt}}\|_2 + \|\mathbf{x}_{2s}^{\mathrm{opt}} - \mathbf{x}_{\mathcal{S}_\tau}\|_2 \\
&\leq \|\mathbf{x} - \mathbf{x}_{2s}^{\mathrm{opt}}\|_2 + \sqrt{\sum_{k \in \mathrm{supp}(\mathbf{x}_{2s}^{\mathrm{opt}}) \setminus \mathcal{S}_\tau} x_k^2} \\
&\leq \|\mathbf{x} - \mathbf{x}_{2s}^{\mathrm{opt}}\|_2 + \tau\sqrt{2s}
\end{aligned}
$$

completing the proof. $\square$

Theorem 1 (Robust sublinear-time, nonequispaced SFT: [16], Theorem 7/[27], Lemma 4)**.** *For a signal* $a \in \mathcal{W}(\mathbb{T}) \cap C(\mathbb{T})$ *corrupted by some arbitrary noise* $\mu : \mathbb{T} \to \mathbb{C}$, *Algorithm 3 of [16], denoted* $\mathcal{A}_{2s,M}^{\mathrm{sub}}$, *will output a $2s$-sparse coefficient vector* $\mathbf{v} \in \mathbb{C}^M$ *which*

    1. reconstructs every frequency of $\hat{\mathbf{a}} \in \mathbb{C}^M$, $\omega \in \mathcal{B}_M$, *with corresponding Fourier coefficients meeting the tolerance*

$$|\hat{a}_\omega| > (4 + 2\sqrt{2}) \left(\frac{\|\hat{\mathbf{a}} - \hat{\mathbf{a}}_s^{\mathrm{opt}}\|_1}{s} + \|\hat{a} - \hat{\mathbf{a}}\|_1 + \|\mu\|_\infty \right),$$

    2. satisfies the ℓ^∞ error estimate for recovered coefficients

$$\|(\hat{\mathbf{a}} - \mathbf{v})|_{\mathrm{supp}(\mathbf{v})}\|_\infty \leq \sqrt{2} \left(\frac{\|\hat{\mathbf{a}} - \hat{\mathbf{a}}_s^{\mathrm{opt}}\|_1}{s} + \|\hat{a} - \hat{\mathbf{a}}\|_1 + \|\mu\|_\infty \right),$$

8

3. satisfies the ℓ^2 error estimate

$$\|\hat{\mathbf{a}} - \mathbf{v}\|_2 \leq \|\hat{\mathbf{a}} - \hat{\mathbf{a}}_{2s}^{\text{opt}}\|_2 + \frac{(8\sqrt{2} + 6)\|\hat{\mathbf{a}} - \hat{\mathbf{a}}_s^{\text{opt}}\|_1}{\sqrt{s}} + (8\sqrt{2} + 6)\sqrt{s}(\|\hat{a} - \hat{\mathbf{a}}\|_1 + \|\mu\|_\infty),$$

4. and the number of required samples of a and the operation count for $\mathcal{A}_{2s,M}^{\text{sub}}$ are

$$\mathcal{O}\left(\frac{s^2 \log^4 M}{\log s}\right).$$

The Monte Carlo variant of $\mathcal{A}_{2s,M}^{\text{sub}}$, denoted $\mathcal{A}_{2s,M}^{\text{sub,MC}}$, referred to by Corollary 4 of [16] satisfies all of the conditions (1) – (3) simultaneously with probability $(1 - \sigma) \in [2/3, 1)$ and has number of required samples and operation count

$$\mathcal{O}\left(s \log^3(M) \log\left(\frac{M}{\sigma}\right)\right).$$

The samples required by $\mathcal{A}_{2s,M}^{\text{sub,MC}}$ are a subset of those required by $\mathcal{A}_{2s,M}^{\text{sub}}$.

Proof. We refer to [16, Theorem 7] and its modification for noise robustness in [27, Lemma 4] for the proofs of properties (2) and (4). As for (1), [16, Lemma 6] and its modification in [27, Lemma 4] imply that any $\omega \in \mathcal{B}_M$ with $|\hat{a}_\omega| > 4(\|\hat{\mathbf{a}} - \hat{\mathbf{a}}_s^{\text{opt}}\|_1/s + \|\hat{a} - \hat{\mathbf{a}}\|_1 + \|\mu\|_\infty) =: 4\delta$ will be identified in [16, Algorithm 3]. An approximate Fourier coefficient for these and any other recovered frequencies is stored in the vector $\mathbf{x}$ which satisfies the same estimate in property (2) by the proof of [16, Theorem 7] and [27, Lemma 4]. However, only the $2s$ largest magnitude values of $\mathbf{x}$ will be returned in $\mathbf{v}$. We therefore analyze what happens when some of the potentially large Fourier coefficients corresponding to frequencies in $\mathcal{S}_{4\delta}$ do not have their approximations assigned to $\mathbf{v}$.

For the definition of $\mathcal{S}_\tau$ in Lemma 1 applied to $\hat{\mathbf{a}}$, we must have $|\mathcal{S}_{4\delta}| \leq 2s = |\text{supp}(\mathbf{v})|$. If $\omega \in \mathcal{S}_{4\delta} \setminus \text{supp}(\mathbf{v})$, there must then exist some other $\omega' \in \text{supp}(\mathbf{v}) \setminus \mathcal{S}_{4\delta}$ which was identified and took the place of ω in $\text{supp}(\mathbf{v})$. For this to happen, $|\hat{a}_{\omega'}| \leq 4\delta$ and $|x_{\omega'}| \geq |x_\omega|$. But by property (2) (extended to all coefficients in $\mathbf{x}$), we know

$$4\delta + \sqrt{2}\delta \geq |\hat{a}_{\omega'}| + \sqrt{2}\delta \geq |x_{\omega'}| \geq |x_\omega| \geq |\hat{a}_\omega| - \sqrt{2}\delta.$$

Thus, any frequency in $\mathcal{S}_{4\delta}$ not chosen satisfies $|\hat{a}_\omega| \leq (4 + 2\sqrt{2})\delta$, and so every frequency in $\mathcal{S}_{(4+2\sqrt{2})\delta}$ is in fact identified in $\mathbf{v}$ verifying property (1).

As for property (3), we estimate the ℓ^2 error using property (2), Lemma 1, and the above argument as

$$\begin{aligned}
\|\hat{\mathbf{a}} - \mathbf{v}\|_2 &\leq \|\hat{\mathbf{a}} - \hat{\mathbf{a}}|_{\text{supp}(\mathbf{v})}\|_2 + \|(\hat{\mathbf{a}} - \mathbf{v})|_{\text{supp}(\mathbf{v})}\|_2 \\
&\leq \|\hat{\mathbf{a}} - \hat{\mathbf{a}}|_{\mathcal{S}_{4\delta} \cap \text{supp}(\mathbf{v})}\|_2 + \sqrt{2}\delta\sqrt{2s} \\
&\leq \|\hat{\mathbf{a}} - \hat{\mathbf{a}}|_{\mathcal{S}_{4\delta}}\|_2 + \|\hat{\mathbf{a}}|_{\mathcal{S}_{4\delta} \setminus \text{supp}(\mathbf{v})}\|_2 + 2\delta\sqrt{s} \\
&\leq \|\hat{\mathbf{a}} - \hat{\mathbf{a}}_{2s}^{\text{opt}}\|_2 + 4\delta\sqrt{2s} + (4 + 2\sqrt{2})\delta\sqrt{2s} + 2\delta\sqrt{s} \\
&= \|\hat{\mathbf{a}} - \hat{\mathbf{a}}_{2s}^{\text{opt}}\|_2 + (8\sqrt{2} + 6)\sqrt{s}\delta
\end{aligned}$$

as desired. $\qquad\square$

9

Remark 1. *In the noiseless case, if the univariate function a is Fourier s-sparse, i.e., is a trigono-metric polynomial and M is large enough such that $\mathrm{supp}(\hat{a}) \subset \mathcal{B}_M$, both $\mathcal{A}_{2s,M}^{\mathrm{sub}}$ and $\mathcal{A}_{2s,M}^{\mathrm{sub,MC}}$ will exactly recover $\hat{\mathbf{a}}$ (the latter with probability $1 - \sigma$), and therefore $\hat{a}$. In particular, note that the output of either algorithm will then actually be s-sparse.*

Using the above SFT algorithm with the discretization process outlined in [27] leads to a fully *discrete* sparse Fourier transform, requiring only equispaced samples of a. However, rather than separately accounting for the truncation to the frequency band $\mathcal{B}_M$ as above, the equispaced samples allow us to take advantage of aliasing, which is particularly important when we apply the algorithm along reconstructing rank-1 lattices. Thus, instead of approximating $\hat{\mathbf{a}} \in \mathbb{C}^M$, the restriction of $\hat{a}$ to $\mathcal{B}_M$, as above, we prefer to approximate the discrete Fourier transform of $\mathbf{a}$. We now review how these notions of restrictions versus aliasing interact.

Lemma 2. *Let $a \in C(\mathbb{T})$ be bandlimited with $\mathrm{supp}(\hat{a}) \subset \mathcal{B}_M$. Then $\hat{\mathbf{a}} = \mathbf{F}_M \mathbf{a}$.*

Proof. Writing $a(t) = \sum_{\omega \in \mathcal{B}_M} \hat{a}_\omega \, \mathrm{e}^{2\pi \mathrm{i} \omega t}$, for any $\omega \in \mathcal{B}_M$, we calculate

$$
\begin{aligned}
(\mathbf{F}_M \, \mathbf{a})_\omega &= \frac{1}{M} \sum_{j \in \mathcal{B}_M} a\left(\frac{j}{M}\right) \mathrm{e}^{-2\pi \mathrm{i} \omega j / M} = \frac{1}{M} \sum_{j \in \mathcal{B}_M} \left(\sum_{\tilde{\omega} \in \mathcal{B}_M} \hat{a}_{\tilde{\omega}} \, \mathrm{e}^{2\pi \mathrm{i} \tilde{\omega} j / M} \right) \mathrm{e}^{-2\pi \mathrm{i} \omega j / M} \\
&= \frac{1}{M} \sum_{j \in \mathcal{B}_M} \sum_{\tilde{\omega} \in \mathcal{B}_M} \hat{a}_{\tilde{\omega}} \, \mathrm{e}^{2\pi \mathrm{i} (\tilde{\omega} - \omega) j / M} = \sum_{\tilde{\omega} \in \mathcal{B}_M} \hat{a}_{\tilde{\omega}} \, \delta_{0,(\tilde{\omega} - \omega \bmod M)} \\
&= \hat{a}_\omega,
\end{aligned}
$$

as desired. □

Lemma 3. *For any function $a : \mathbb{T} \to \mathbb{C}$ with Fourier series $a(t) = \sum_{\omega \in \mathbb{Z}} \hat{a}_\omega \, \mathrm{e}^{2\pi \mathrm{i} \omega t}$, define the aliased polynomial*

$$
a_{\mathrm{alias}}(t) = \sum_{\omega \in \mathcal{B}_M} \underbrace{\left(\sum_{\tilde{\omega} \equiv \omega \ (\bmod M)} \hat{a}_{\tilde{\omega}} \right)}_{=: (\hat{\mathbf{a}}_{\mathrm{alias}})_\omega} \mathrm{e}^{2\pi \mathrm{i} \omega t}.
$$

Then the equispaced samples coincide, giving $\mathbf{a} = \mathbf{a}_{\mathrm{alias}} \in \mathbb{C}^M$ and $\hat{\mathbf{a}}_{\mathrm{alias}} = \mathbf{F}_M \mathbf{a}$.

Proof. We group frequencies in the Fourier series of a by their residues in $\mathcal{B}_M$, giving

$$
\begin{aligned}
(\mathbf{a})_j &= \sum_{\tilde{\omega} \in \mathbb{Z}} \hat{a}_{\tilde{\omega}} \, \mathrm{e}^{2\pi \mathrm{i} \tilde{\omega} j / M} = \sum_{\omega \in \mathcal{B}_M} \sum_{n \in \mathbb{Z}} \hat{a}_{\omega + nM} \, \mathrm{e}^{2\pi \mathrm{i} (\omega + nM) j / M} = \sum_{\omega \in \mathcal{B}_M} \left(\sum_{\tilde{\omega} \equiv \omega \ (\bmod M)} \hat{a}_{\tilde{\omega}} \right) \mathrm{e}^{2\pi \mathrm{i} \omega j / M} \\
&= (\mathbf{a}_{\mathrm{alias}})_j \ \text{ for all } j \in [M].
\end{aligned}
$$

Now, since $\mathrm{supp}(\hat{\mathbf{a}}_{\mathrm{alias}}) \subset \mathcal{B}_M$, Lemma 2 implies $\hat{\mathbf{a}}_{\mathrm{alias}} = \mathbf{F}_M \mathbf{a}_{\mathrm{alias}} = \mathbf{F}_M \mathbf{a}$. □

Eventually, we will consider techniques for approximation of arbitrary periodic functions rather than simply polynomials. For this reason, we require noise-robust recovery results for the method in [27]. The necessary modifications to account for this robustness as well as the improved guarantees carried over from the previous algorithm are given below.

Theorem 2 (Robust discrete sublinear-time SFT: see [27], Theorem 5). *For a signal $a \in \mathcal{W}(\mathbb{T}) \cap C(\mathbb{T})$ corrupted by some arbitrary noise $\mu : \mathbb{T} \to \mathbb{C}$, and $1 \leq r \leq \frac{M}{36}$ Algorithm 1 of [27], denoted $\mathcal{A}_{2s,M}^{\mathrm{disc}}$, will output a 2s-sparse coefficient vector $\mathbf{v} \in \mathbb{C}^M$ which*

1. *reconstructs every frequency of $\mathbf{F}_M\,\mathbf{a} \in \mathbb{C}^M$, $\omega \in \mathcal{B}_M$, with corresponding aliased Fourier coefficient meeting the tolerance*

$$|(\mathbf{F}_M\,\mathbf{a})_\omega| > 12(1 + \sqrt{2})\left(\frac{\|\mathbf{F}_M\,\mathbf{a} - (\mathbf{F}_M\,\mathbf{a})_s^{\mathrm{opt}}\|_1}{2s} + 2(\|\mathbf{a}\|_\infty M^{-r} + \|\boldsymbol{\mu}\|_\infty)\right),$$

2. *satisfies the ℓ^∞ error estimate for recovered coefficients*

$$\|(\mathbf{F}_M\,\mathbf{a} - \mathbf{v})|_{\mathrm{supp}(\mathbf{v})}\|_\infty \leq 3\sqrt{2}\left(\frac{\|\mathbf{F}_M\,\mathbf{a} - (\mathbf{F}_M\,\mathbf{a})_s^{\mathrm{opt}}\|_1}{2s} + 2(\|\mathbf{a}\|_\infty M^{-r} + \|\boldsymbol{\mu}\|_\infty)\right),$$

3. *satisfies the ℓ^2 error estimate*

$$\|\mathbf{F}_M\,\mathbf{a} - \mathbf{v}\|_2 \leq \|\mathbf{F}_M\,\mathbf{a} - (\mathbf{F}_M\,\mathbf{a})_{2s}^{\mathrm{opt}}\|_2 + 38\frac{\|\mathbf{F}_M\,\mathbf{a} - (\mathbf{F}_M\,\mathbf{a})_s^{\mathrm{opt}}\|_1}{\sqrt{s}} + 152\sqrt{s}(\|\mathbf{a}\|_\infty M^{-r} + \|\boldsymbol{\mu}\|_\infty),$$

4. *and the number of required samples of $\mathbf{a}$ and the operation count for $\mathcal{A}_{2s,M}^{\mathrm{disc}}$ are*

$$\mathcal{O}\left(\frac{s^2 r^{3/2} \log^{11/2} M}{\log s}\right).$$

The Monte Carlo variant of $\mathcal{A}_{2s,M}^{\mathrm{disc}}$, denoted $\mathcal{A}_{2s,M}^{\mathrm{disc,MC}}$, satisfies the all of the conditions (1) – (3) simultaneously with probability $(1 - \sigma) \in [2/3, 1)$ and has number of required samples and operation count

$$\mathcal{O}\left(s r^{3/2} \log^{9/2}(M) \log\left(\frac{M}{\sigma}\right)\right).$$

Proof. All notation in this proof matches that in [27] (in particular, we use f to denote the one-dimensional function in place of a in the theorem statement and $N = 2M + 1$). We begin by substituting the 2π-periodic gaussian filter given in (3) on page 756 with the 1-periodic gaussian and associated Fourier transform

$$g(x) = \frac{1}{c_1} \sum_{n=-\infty}^{\infty} e^{-\frac{(2\pi)^2(x-n)^2}{2c_1^2}}, \quad \hat{g}_\omega = \frac{1}{\sqrt{2\pi}} e^{-\frac{c_1^2 \omega^2}{2}}.$$

Note then that all results regarding the Fourier transform remain unchanged, and since this 1-periodic gaussian is a just a rescaling of the 2π-periodic one used in [27], the bound in [27, Lemma 1] holds with a similarly compressed gaussian, that is, for all $x \in \left[-\frac{1}{2}, \frac{1}{2}\right]$

$$g(x) \leq \left(\frac{3}{c_1} + \frac{1}{\sqrt{2\pi}}\right) e^{-\frac{(2\pi x)^2}{2c_1^2}}. \tag{2}$$

11

Analogous results up to and including [27, Lemma 10] for 1-periodic functions then hold straightforwardly.

Assuming that our signal measurements $\mathbf{f} = (f(y_j))_{j=0}^{2M} = (f(\frac{j}{N}))_{j=0}^{2M}$ are corrupted by some discrete noise $\boldsymbol{\mu} = (\mu_j)_{j=0}^{2M}$, we consider for any $x \in \mathbb{T}$ a similar bound to [27, Lemma 10]. Here, $j' := \arg\min_j |x - y_j|$ and $\kappa := \lceil \gamma \ln N \rceil + 1$ for some $\gamma \in \mathbb{R}^+$ to be determined. Then,

$$
\left| \frac{1}{N} \sum_{j=0}^{2M} f(y_j) g(x - y_j) - \frac{1}{N} \sum_{j=j'-\kappa}^{j'+\kappa} (f(y_j) + \mu_j) g(x - y_j) \right|
$$

$$
\leq \frac{1}{N} \left| \sum_{j=0}^{2M} f(y_j) g(x - y_j) - \sum_{j=j'-\kappa}^{j'+\kappa} f(y_j) g(x - y_j) \right| + \frac{1}{N} \left| \sum_{j=j'-\kappa}^{j'+\kappa} \mu_j g(x - y_j) \right|
$$

$$
\leq \frac{1}{N} \left| \sum_{j=0}^{2M} f(y_j) g(x - y_j) - \sum_{j=j'-\kappa}^{j'+\kappa} f(y_j) g(x - y_j) \right| + \frac{1}{N} \|\boldsymbol{\mu}\|_\infty \sum_{k=-\kappa}^{\kappa} g(x - y_{j'+k})
$$

We bound the first term in this sum by a direct application of [27, Lemma 10]; however, we take this opportunity to reduce the constant in the bound given there. In particular, bounding this term by the final expression in the proof of [27, Lemma 10] and using our implicit assumption that $36 \leq N$, we have

$$
\left| \frac{1}{N} \sum_{j=0}^{2M} f(y_j) g(x - y_j) - \frac{1}{N} \sum_{j=j'-\kappa}^{j'+\kappa} (f(y_j) + \mu_j) g(x - y_j) \right|
$$

$$
\leq \left(\frac{3}{\sqrt{2\pi}} + \frac{1}{2\pi} \sqrt{\frac{\ln 36}{36}} \right) \|\mathbf{f}\|_\infty N^{-r} + \frac{1}{N} \|\boldsymbol{\mu}\|_\infty \sum_{k=-\kappa}^{\kappa} g(x - y_{j'+k}).
$$

(3)

We now work on bounding the second term. First note that for all $k \in [-\kappa, \kappa] \cap \mathbb{Z}$,

$$
g(x - y_{j'\pm k}) = g\left(x - y_{j'} \pm \frac{k}{N} \right).
$$

Assuming without loss of generality that $0 \leq x - y_{j'}$, we can bound the nonnegatively indexed summands by (2) as

$$
g\left(x - y_{j'} + \frac{k}{N} \right) \leq \left(\frac{3}{c_1} + \frac{2}{\sqrt{2\pi}} \right) \mathrm{e}^{-\frac{(2\pi)^2 k^2}{2c_1^2 N^2}}.
$$

(4)

For the negatively indexed summands, the definition of $j' = \arg\min_j |x - y_j|$ implies that $x - y_{j'} \leq \frac{1}{2N}$. In particular,

$$
x - y_{j'} - \frac{k}{N} \leq \frac{1 - 2k}{2N} < 0 \implies \left(x - y_{j'} - \frac{k}{N} \right)^2 \geq \frac{1 - 2k}{2N} \left(x - y_{j'} - \frac{k}{N} \right) \geq \frac{2k - 1}{2N} \cdot \frac{k}{N},
$$

giving

$$
g\left(x - y_{j'} - \frac{k}{N} \right) \leq \left(\frac{3}{c_1} + \frac{2}{\sqrt{2\pi}} \right) \mathrm{e}^{-\frac{(2\pi)^2 k^2}{2c_1^2 N^2}} \mathrm{e}^{\frac{(2\pi)^2 k}{4c_1^2 N^2}}.
$$

(5)

12

We now bound the final exponential. We first recall from [27] the choices of parameters

$$c_1 = \frac{\beta\sqrt{\ln N}}{N}, \quad \kappa = \lceil \gamma \ln N \rceil + 1, \quad \gamma = \frac{6r}{\sqrt{2\pi}} = \frac{\beta\sqrt{r}}{2\sqrt{\pi}}, \quad \beta = 6\sqrt{r}, \quad \text{where } 1 \le r \le \frac{N}{36}.$$

For $k \in [1, \kappa] \cap \mathbb{Z}$ then,

$$\begin{aligned}
\exp\left(\frac{(2\pi)^2 k}{4c_1^2 N^2}\right) &\le \exp\left(\frac{(2\pi)^2 \kappa}{4c_1^2 N^2}\right) \\
&\le \exp\left(\frac{\pi^2\left(\frac{6r\ln N}{\sqrt{2\pi}} + 2\right)}{36r\ln N}\right) \\
&\le \exp\left(\frac{\pi}{6\sqrt{2}} + \frac{\pi^2}{18r\ln N}\right) \\
&\le \exp\left(\frac{\pi}{6\sqrt{2}} + \frac{\pi^2}{18\ln 36}\right) =: A.
\end{aligned}$$

Combining this with our bounds for the nonnegatively indexed summands (4) and the negatively indexed summands (5), we have

$$\frac{1}{N} \sum_{k=-\kappa}^{\kappa} g(x - y_{j'+k}) \le \left(\frac{3}{\beta\sqrt{\ln N}} + \frac{1}{N\sqrt{2\pi}}\right)\left(1 + (1 + A)\sum_{k=1}^{\kappa} \mathbb{e}^{-\frac{(2\pi)^2 k^2}{2\beta^2 \ln N}}\right)$$

Expressing the final sum as a truncated lower Riemann sum and applying a change of variables on the resulting integral, we have

$$\sum_{k=1}^{\kappa} \mathbb{e}^{-\frac{(2\pi)^2 k^2}{2\beta^2 \ln N}} \le \frac{\beta\sqrt{\ln N}}{\sqrt{2\pi}} \int_0^\infty e^{-x^2}\, dx = \frac{\beta\sqrt{\ln N}}{2\sqrt{2\pi}}.$$

Making use of our parameter values from [27], and the fact that $1 \le r \le \frac{N}{36}$,

$$\begin{aligned}
\frac{1}{N} \sum_{k=-\kappa}^{\kappa} g(x - y_{j'+k}) &\le \left(\frac{3}{\beta\sqrt{\ln N}} + \frac{1}{N\sqrt{2\pi}}\right)\left(1 + \frac{1+A}{2\sqrt{2\pi}}\beta\sqrt{\ln N}\right) \\
&\le \frac{3}{6\sqrt{\ln 36}} + \frac{3(1+A)}{2\sqrt{2\pi}} + \frac{1}{36\sqrt{2\pi}} + \frac{1+A}{4\pi}\sqrt{\frac{\ln 36}{36}} \quad (6) \\
&< 2.
\end{aligned}$$

With our revised bound for (3) above, we reprove [27, Theorem 4] to estimate $g * f$ by the truncated discrete convolution with noisy samples. In particular, we apply [27, Theorem 3], (3),

(2), and finally our same assumption that $1 \leq r \leq \frac{N}{36}$ to obtain

$$
\left| (g * f)(x) - \frac{1}{N} \sum_{j=j'-\left\lceil \frac{6r}{\sqrt{2}\pi} \ln N \right\rceil - 1}^{j'+\left\lceil \frac{6r}{\sqrt{2}\pi} \ln N \right\rceil + 1} (f(y_j) + \mu_j) g(x - y_j) \right|
$$

$$
\leq \frac{N^{1-r}}{6\sqrt{r}\sqrt{\ln N}} \|\mathbf{f}\|_\infty N^{-r} + \left(\frac{3}{\sqrt{2\pi}} + \frac{1}{2\pi} \sqrt{\frac{\ln 36}{36}} \right) \|\mathbf{f}\|_\infty N^{-r} + 2\|\boldsymbol{\mu}\|_\infty
$$

$$
\leq \left(\frac{1}{6\sqrt{\ln 36}} + \frac{3}{\sqrt{2\pi}} + \frac{1}{2\pi} \sqrt{\frac{\ln 36}{36}} \right) \frac{\|\mathbf{f}\|_\infty}{N^r} + 2\|\boldsymbol{\mu}\|_\infty < 2 \left(\frac{\|\mathbf{f}\|_\infty}{N^r} + \|\boldsymbol{\mu}\|_\infty \right).
$$

Replacing all references of $3\|\mathbf{f}\|_\infty N^{-r}$ by $2(\|\mathbf{f}\|_\infty N^{-r} + \|\boldsymbol{\mu}\|_\infty)$ in the remainder of the steps up to proving [27, Theorem 5] gives the desired noise robustness (with a slightly improved constant).

Using the revised error estimates of the nonequispaced algorithm from Theorem 1 and redefining $\delta = 3(\|\hat{\mathbf{f}} - \hat{\mathbf{f}}_s^{\text{opt}}\|_1 / 2s + 2(\|\mathbf{f}\|_\infty N^{-r} + \|\boldsymbol{\mu}\|_\infty))$ as in the proof of [27, Theorem 5] (which also contains the proof of property (2)), the discretization algorithm [27, Algorithm 1] will produce candidate Fourier coefficient approximations in lines 9 and 12 corresponding to every $|\hat{f}_\omega| \geq (4 + 2\sqrt{2})\delta$ in place of 4δ in Theorem 1. The exact same argument as in the proof of Theorem 1 then applies to the selection of the $2s$-largest entries of this approximation with the revised threshold values and error bounds to give properties (1) and (3).

In detail, [27, Lemma 13] and the discussion right after its statement gives that property (2) holds for any approximate coefficient with frequency recovered throughout the algorithm (which, for the purposes of the following discussion, we will store in $\mathbf{x}$ rather than $\hat{R}$ defined in [27, Algorithm 1]), not just those in the final output $\mathbf{v} := \mathbf{x}_s^{\text{opt}}$. Additionally, by the same lemma and our revised bounds from Theorem 1, any frequency $\omega \in [N]$ satisfying $|f_\omega| > (4 + 2\sqrt{2})\delta$ will have an associated coefficient estimate in $\mathbf{x}$.

By Lemma 1, $|\mathcal{S}_{(4+2\sqrt{2})\delta}| \leq 2s = |\text{supp}(\mathbf{v})|$, and so if $\omega \in \mathcal{S}_{(4+2\sqrt{2})\delta} \setminus \text{supp}(\mathbf{v})$, there exists some $\omega' \in \text{supp}(\mathbf{v}) \setminus \mathcal{S}_{(4+2\sqrt{2})\delta}$ such that $v_{\omega'}$ took the place of v_ω in $\mathcal{S}$. In particular, this means that $|x_{\omega'}| \geq |x_\omega|$, $|\hat{f}_{\omega'}| \leq (4 + 2\sqrt{2})\delta$, and $|\hat{f}_\omega| > (4 + 2\sqrt{2}\delta)$. Thus,

$$
(4 + 2\sqrt{2})\delta + \sqrt{2}\delta > |\hat{f}_{\omega'}| + \sqrt{2}\delta \geq |x_{\omega'}| \geq |x_\omega| \geq |\hat{f}_\omega| - \sqrt{2}\delta,
$$

implying that $|\hat{f}_\omega| \leq 4(1 + \sqrt{2})\delta$ and therefore proving (1).

Finally, to prove (3), we use Lemma 1, and consider

$$
\begin{aligned}
\|\hat{\mathbf{f}} - \mathbf{v}\|_2 &\leq \|\hat{\mathbf{f}} - \hat{\mathbf{f}}|_{\text{supp}(\mathbf{v})}\|_2 - \|(\hat{\mathbf{f}} - \mathbf{v})|_{\text{supp}(v)}\|_2 \\
&\leq \|\hat{\mathbf{f}} - \hat{\mathbf{f}}|_{\mathcal{S}_{(4+2\sqrt{2})\delta} \cap \text{supp}(\mathbf{v})}\|_2 + \sqrt{2}\delta\sqrt{2s} \\
&\leq \|\hat{\mathbf{f}} - \hat{\mathbf{f}}|_{\mathcal{S}_{(4+2\sqrt{2})\delta}}\|_2 + \|\hat{\mathbf{f}}|_{\mathcal{S}_{(4+2\sqrt{2})\delta} \setminus \text{supp}(\mathbf{v})}\|_2 + 2\delta\sqrt{s} \\
&\leq \|\hat{\mathbf{f}} - \hat{\mathbf{f}}_{2s}^{\text{opt}}\|_2 + (4 + 2\sqrt{2})\delta\sqrt{2s} + 4(1 + \sqrt{2})\delta\sqrt{2s} + 2\delta\sqrt{s} \\
&\leq \|\hat{\mathbf{f}} - \hat{\mathbf{f}}_{2s}^{\text{opt}}\|_2 + (14 + 8\sqrt{2})\delta\sqrt{s}
\end{aligned}
$$

which finishes the proof. $\qquad\square$

4 Fast multivariate sparse Fourier transforms

Having detailed two sublinear-time, one-dimensional SFT algorithms, we are now prepared to extend these to the multivariate setting. The general approach will be to apply the one-dimensional methods to transformations of our multivariate function of interest with samples taken along rank-1 lattices. The particular approaches for transforming our multivariate function will then allow for the efficient extraction of multidimensional frequency information for the most energetic coefficients identified by univarate SFTs. In particular, our first approach considered in Section 4.1 successively shifts the function in each dimension, whereas our second approach considered in Section 4.2 successively collapses all but one dimension along a rank-1 lattice and samples the resulting two-dimensional function.

   Before continuing, it is important to stress that these two approaches given in Algorithms 1 and 2 below can make use of *any* univariate SFT algorithm $\mathcal{A}_{s,M}$. Thus, the analysis of each algorithm is presented in a similarly modular fashion. Each algorithm is followed by a lemma (Lemma 5 and Lemma 7 respectively) which provides associated error guarantees when any sufficiently accurate univariate SFT $\mathcal{A}_{s,M}$ is employed. The lemmas are then each followed by two corollaries (Corollaries 1 and 2 and Corollaries 3 and 4 respectively) where we apply the lemma to the two example univariate SFTs reviewed in Section 3 specified by Theorems 2 and 1.

4.1 Phase encoding

We begin with a review of how the shift operator $S_{\ell,\alpha}$ interacts with the Fourier transform.

Lemma 4. *For any dimension $\ell \in [d]$, shift $\alpha \in \mathbb{R}$, and function $f : \mathbb{T}^d \to \mathbb{C}$, shifting f with the operator $S_{\ell,\alpha}$ modulates the Fourier coefficients as $c_{\mathbf{k}}(S_{\ell,\alpha}f) = e^{2\pi i k_\ell \alpha} c_{\mathbf{k}}(f)$ for all $\mathbf{k} \in \mathbb{Z}^d$.*

Proof. Using the Fourier series of f in the computation of $c_{\mathbf{k}}(S_{\ell,\alpha}f)$ gives

$$
\begin{aligned}
c_{\mathbf{k}}(S_{\ell,\alpha}f) &= \int_{\mathbb{T}^d} S_{\ell,\alpha}f(\mathbf{x})\, e^{-2\pi i \mathbf{k}\cdot\mathbf{x}}\, \mathrm{d}\mathbf{x} \\
&= \int_{\mathbb{T}^d} \sum_{\mathbf{h}\in\mathbb{Z}^d} c_{\mathbf{h}}(f)\, e^{2\pi i (h_\ell(x_\ell+\alpha)+\mathbf{h}'_\ell\cdot\mathbf{x}'_\ell)}\, e^{-2\pi i \mathbf{k}\cdot\mathbf{x}}\, \mathrm{d}\mathbf{x} \\
&= \sum_{\mathbf{h}\in\mathbb{Z}^d} e^{2\pi i h_\ell \alpha}\, c_{\mathbf{h}}(f) \int_{\mathbb{T}^d} e^{2\pi i (\mathbf{h}-\mathbf{k})\cdot\mathbf{x}}\, \mathrm{d}\mathbf{x} \\
&= e^{2\pi i k_\ell \alpha}\, c_{\mathbf{k}}(f),
\end{aligned}
$$

as desired. ∎

   By performing additional lattice SFTs on shifted versions of our original function, we can then separate out the components of recovered frequencies in modulations of the function's Fourier coefficients. Using the common one-dimensional recovered frequencies between each transform in the form $\mathbf{k}\cdot\mathbf{z} \bmod M$, we can then reconstruct the multivariate frequencies. This approach of encoding frequencies in the phase is summarized in Algorithm 1.

Algorithm 1 Simple Frequency Index Recovery by Phase Encoding

---

Input: A multivariate periodic function $f \in \mathcal{W}(\mathbb{T}^d \cap C(\mathbb{T}^d)$ (from which we are able to obtain potentially noisy samples), a multivariate frequency set $\mathcal{I} \subset \mathcal{B}_N^d$, a reconstructing rank-1 lattice $\Lambda(\mathbf{z}, M)$ for $\mathcal{I}$, and an SFT algorithm $\mathcal{A}_{s,M}$.

Output: Sparse coefficient vector $\mathbf{b} = (b_{\mathbf{k}})_{\mathbf{k} \in \mathcal{B}_N^d}$ (optionally supported on $\mathcal{I}$, see Line (X)), an approximation to $(c|_{\mathcal{I}})_s^{\mathrm{opt}}$.

(I) Apply $\mathcal{A}_{s,M}$ to the univariate restriction of f to the lattice, $a(t) = f(t\mathbf{z})$, to produce $\mathbf{v} = \mathcal{A}_{s,M}a$, a sparse approximation of $\mathbf{F}_M \mathbf{a} \in \mathbb{C}^M$.

(II) **for all** $\ell \in [d]$ **do**

(III)   Apply $\mathcal{A}_{s,M}$ to $a^{\ell}(t) = S_{\ell,1/N}f(t\mathbf{z})$ to produce $\mathbf{v}^{\ell} = \mathcal{A}_{s,M}a^{\ell}$, a sparse approximation of $\mathbf{F}_M \mathbf{a}^{\ell} \in \mathbb{C}^M$.

(IV) **end for**

(V) $\mathbf{b} \leftarrow \mathbf{0}$

(VI) **for all** $\omega \in \mathrm{supp}(\mathbf{v}) \subset \mathcal{B}_M$ **do**

(VII)   **for all** $\ell \in [d]$ **do**

(VIII) $(k_{\omega})_{\ell} \leftarrow \mathrm{round}(N \arg(v_{\omega}^{\ell}/v_{\omega})/2\pi)$

(IX)   **end for**

(X)   **if** $\mathbf{k}_{\omega} \cdot \mathbf{z} \equiv \omega \pmod{M}$ (and optionally $\mathbf{k}_{\omega} \in \mathcal{I}$; see Remark 2) **then**

(XI) $b_{\mathbf{k}_{\omega}} \leftarrow b_{\mathbf{k}_{\omega}} + v_{\omega}$

(XII)   **end if**

(XIII) **end for**

4.1.1 Analysis of Algorithm 1

Lemma 5 (General recovery result for Algorithm 1)**.** *Let $\mathcal{A}_{s,M}$ in the input to Algorithm 1 be a noise-robust SFT algorithm which, for a function $a \in \mathcal{W}(\mathbb{T}) \cap C(\mathbb{T})$ corrupted by some arbitrary noise $\mu : \mathbb{T} \to \mathbb{C}$, constructs an s-sparse Fourier approximation $\mathcal{A}_{s,M}(a + \mu) =: \mathbf{v} \in \mathbb{C}^M$ which*

1. reconstructs every frequency (up to s many) of $\mathbf{F}_M \mathbf{a} \in \mathbb{C}^M$, $\omega \in \mathcal{B}_M$, with corresponding Fourier coefficient meeting the tolerance $|(\mathbf{F}_M \mathbf{a})_{\omega}| > \tau$,

2. satisfies the ℓ^{∞} error estimate for recovered coefficients

$$\left\| (\mathbf{F}_M \mathbf{a} - \mathbf{v})|_{\mathrm{supp}(\mathbf{v})} \right\|_{\infty} \leq \eta_{\infty} < \tau,$$

3. satisfies the ℓ^2 error estimate

$$\|\mathbf{F}_M \mathbf{a} - \mathbf{v}\|_2 \leq \eta_2,$$

4. and requires $\mathcal{O}(P(s, M))$ total evaluations of a, operating with computational complexity $\mathcal{O}(R(s, M))$.

Additionally, assume that the parameters τ and η_{∞} hold uniformly for each SFT performed in Algorithm 1.

Let f, $\mathcal{I}$, and $\Lambda(\mathbf{z}, M)$ be as specified in the input to Algorithm 1. Collecting the τ-significant frequencies of f into the set $\mathcal{S}_{\tau} := \{\mathbf{k} \in \mathcal{I} \mid |c_{\mathbf{k}}(f)| > \tau\}$, assume that $|\mathrm{supp}(c) \cap \mathcal{S}_{\tau}| \leq s$, and set

$$\beta = \max\left(\tau, \eta_{\infty}\left(1 + \frac{2}{\sin\left(\frac{\pi}{N}\right)}\right)\right).$$

16

Then Algorithm 1 (ignoring the optional check on Line (X)) will produce an s-sparse approximation **b** *of the Fourier coefficients of f satisfying the error estimate*

$$\|\mathbf{b} - c(f)\|_{\ell^2(\mathbb{Z}^d)} \leq \eta_2 + (\beta + \eta_\infty)\sqrt{\max(s - |\mathcal{S}_\beta|, 0)}$$
$$+ \|c(f)|_\mathcal{I} - c(f)|_{\mathcal{S}_\beta}\|_{\ell^2(\mathbb{Z}^d)} + \|c(f) - c(f)|_\mathcal{I}\|_{\ell^2(\mathbb{Z}^d)}$$

requiring $\mathcal{O}(d \cdot P(s, M))$ *total evaluations of f, in* $\mathcal{O}(d \cdot (R(s, M) + s))$ *total operations.*

Proof. We begin by assuming that f is a trigonometric polynomial with $\mathrm{supp}(c(f)) \subset \mathcal{I}$. Since $\Lambda(\mathbf{z}, M)$ is a reconstructing rank-1 lattice for $\mathcal{I}$, Lemmas 2 and 3 ensure that for each $\mathbf{k} \in \mathcal{I}$, $c_\mathbf{k}(f) = \hat{a}_{\mathbf{k} \cdot \mathbf{z} \bmod M} = (\mathbf{F}_M \mathbf{a})_{\mathbf{k} \cdot \mathbf{z} \bmod M}$. Thus, Lines (I) and (III) of Algorithm 1 will produce coefficient estimates in the lattice DFT for every $\mathbf{k} \in \mathcal{S}_\tau$. We then write these SFT approximations as $v_{\mathbf{k} \cdot \mathbf{z} \bmod M} = c_\mathbf{k}(f) + \eta_\mathbf{k}$ and $v^\ell_{\mathbf{k} \cdot \mathbf{z} \bmod M} = \mathbb{e}^{2\pi \mathrm{i} k_\ell / N}(c_\mathbf{k}(f) + \eta^\ell_\mathbf{k})$ respectively, where we have made use of Lemma 4. Note that $|\eta_\mathbf{k}|, |\eta^\ell_\mathbf{k}| \leq \eta_\infty$. Now, considering the estimate for k_ℓ, we have

$$\frac{N}{2\pi} \arg\left(\frac{v^\ell_{\mathbf{k} \cdot \mathbf{z} \bmod M}}{v_{\mathbf{k} \cdot \mathbf{z} \bmod M}}\right) = \frac{N}{2\pi} \arg\left(\mathbb{e}^{2\pi \mathrm{i} k_\ell / N} \frac{c_\mathbf{k}(f) + \eta^\ell_\mathbf{k}}{v_{\mathbf{k} \cdot \mathbf{z} \bmod M}}\right)$$
$$= k_\ell + \frac{N}{2\pi} \arg\left(\frac{c_\mathbf{k}(f) + \eta^\ell_\mathbf{k}}{v_{\mathbf{k} \cdot \mathbf{z} \bmod M}}\right)$$
$$= k_\ell + \frac{N}{2\pi} \arg\left(1 + \frac{\eta^\ell_\mathbf{k} - \eta_\mathbf{k}}{v_{\mathbf{k} \cdot \mathbf{z} \bmod M}}\right).$$

We now only consider $|c_\mathbf{k}| > \beta \geq \max(\tau, 3\eta_\infty)$, that is $\mathbf{k} \in \mathcal{S}_\beta \subset \mathcal{S}_\tau$, and therefore, the corresponding approximate coefficient satisfies $|v_{\mathbf{k} \cdot \mathbf{z} \bmod M}| > \beta - \eta_\infty$. Thus, the magnitude of the fraction in the argument must be strictly less than $\frac{2\eta_\infty}{\beta - \eta_\infty} \leq 1$. Therefore, we consider the argument of a point lying in the right half of the complex plane, in the open disc of radius $\frac{2\eta_\infty}{\beta - \eta_\infty}$ centered at 1. The maximal absolute argument of a point in this disc will be that of a point lying on a tangent line passing through the origin. This point, the origin, and 1 then form a right triangle from which we deduce that

$$\left|\arg\left(1 + \frac{\eta^\ell_\mathbf{k} - \eta_\mathbf{k}}{v_{\mathbf{k} \cdot z \bmod M}}\right)\right| < \arcsin\left(\frac{2\eta_\infty}{\beta - \eta_\infty}\right) \leq \frac{\pi}{N}.$$

Our choice of $\beta \geq \eta_\infty(1 + 2/\sin(\pi/N))$ then implies that

$$\left|\frac{N}{2\pi} \arg\left(\frac{v^\ell_{\mathbf{k} \cdot \mathbf{z} \bmod M}}{v_{\mathbf{k} \cdot \mathbf{z} \bmod M}}\right) - k_\ell\right| < \frac{1}{2},$$

and so after rounding to the nearest integer, Algorithm 1 will recover k_ℓ for all $\ell \in [d]$ and $\mathbf{k} \in \mathcal{S}_\beta$.

By this mapping constructed in the final loop of Algorithm 1, we set $b_\mathbf{k} := v_{\mathbf{k} \cdot \mathbf{z} \bmod M}$ for each $\mathbf{k} \in \mathcal{S}_\beta$. Additionally, the $\max(s - |\mathcal{S}_\beta|, 0)$ many coefficients v_ω for which $\omega \neq \mathbf{k} \cdot \mathbf{z} \bmod M$ for any $\mathbf{k} \in \mathcal{S}_\beta$ are still available for potential assignment. If any multivariate frequency $\mathbf{k}_\omega \in \mathcal{I}$ is reconstructed and passes the mandatory check in Line (X) then the approximate Fourier coefficient v_ω properly corresponds to $(\mathbf{F}_M \mathbf{a})_{\mathbf{k}_\omega \cdot \mathbf{z} \bmod M} = c_{\mathbf{k}_\omega}(f)$.

On the other hand, if some error introduced in the SFTs reconstructs a multivariate frequency $\mathbf{k}_\omega \notin \mathcal{I}$, the reconstructing property does not allow us to conclude anything about a (k_ω, ω) pair passing the check in Line (X). Thus, it is possible that v_ω will contribute to some component of **b** not corresponding to any frequency in $\mathcal{I}$. At the least however, since we know that all entries of **v**

corresponding to frequencies in $\mathcal{S}_\beta$ are correctly assigned, the remaining ones satisfy $|v_\omega| \le \beta + \eta_\infty$. Using these facts allows us to estimate the error as

$$
\begin{aligned}
\|\mathbf{b} - c\|_{\ell^2(\mathbb{Z}^d)} &\le \|\mathbf{b}|_{\mathbb{Z}^d \setminus \mathcal{I}}\|_{\ell^2(\mathbb{Z}^d)} + \|\mathbf{b}|_\mathcal{I} - c|_{\mathrm{supp}(\mathbf{b}) \cap \mathcal{I}}\|_{\ell^2(\mathbb{Z}^d)} + \|c - c|_{\mathcal{S}_\beta}\|_{\ell^2(\mathbb{Z}^d)} \\
&\le (\beta + \eta_\infty)\sqrt{\max(s - |\mathcal{S}_\beta|, 0)} + \eta_2 + \|c - c|_{\mathcal{S}_\beta}\|_{\ell^2(\mathcal{I})}
\end{aligned}
\tag{7}
$$

where we have additionally used the accuracy of the initial one-dimensional SFT and the assumption that c is supported on $\mathcal{I}$.

We now handle the case when f is not necessarily a polynomial with Fourier support contained in $\mathcal{I}$. Rather than aiming to approximate $c_\mathbf{k}(f)$ for every $\mathbf{k} \in \mathbb{Z}^d$, we restrict attention to only frequencies in $\mathcal{I}$, instead attempting to approximate the Fourier coefficients of $f_\mathcal{I} = \sum_{\mathbf{k} \in \mathcal{I}} c_\mathbf{k}(f) e^{2\pi i \mathbf{k} \cdot \circ}$. We then have that $f =: f_\mathcal{I} + f_{\mathbb{Z}^d \setminus \mathcal{I}}$ and view potentially noisy input $f + \mu$ to our algorithm as

$$
f + \mu = f_\mathcal{I} + \underbrace{f_{\mathbb{Z}^d \setminus \mathcal{I}} + \mu}_{\mu'}.
$$

Algorithm 1 applied to $f + \mu$ is then equivalent to applying it to $f_\mathcal{I} + \mu'$, where now τ, η_∞, and η_2 depend on μ', and the output is an approximation of $c|_\mathcal{I}$. Since μ' represents noise on the input to $\mathcal{A}_{s,M}$ in its applications to $f_\mathcal{I}(t\mathbf{z})$ and $S_{\ell, 1/N} f_\mathcal{I}(t\mathbf{z})$ we remark here that

$$
\|\mu'\|_\infty \le \|f_{\mathbb{Z}^d \setminus \mathcal{I}}\|_\infty + \|\mu\|_\infty \le \|c(f) - c(f)|_\mathcal{I}\|_{\ell^1(\mathbb{Z}^d)} + \|\mu\|_\infty
\tag{8}
$$

so as to help us estimate τ, η_∞, and η_2 in future applications of the lemma. Accounting for the truncation to $\mathcal{I}$ in the ℓ^2 error bound and using (7) applied to $c|_\mathcal{I}$, we estimate

$$
\begin{aligned}
\|\mathbf{b} - c\|_{\ell^2(\mathbb{Z}^d)} &\le \|\mathbf{b} - c|_\mathcal{I}\|_{\ell^2(\mathbb{Z}^d)} + \|c - c|_\mathcal{I}\|_{\ell^2(\mathbb{Z}^d)} \\
&\le (\beta + \eta_\infty)\sqrt{\max(s - |\mathcal{S}_\beta|, 0)} + \eta_2 + \|c|_\mathcal{I} - c|_{\mathcal{S}_\beta}\|_{\ell^2(\mathbb{Z}^d)} + \|c - c|_\mathcal{I}\|_{\ell^2(\mathbb{Z}^d)}.
\end{aligned}
$$

Since $1 + d$ SFTs are required, the number of f evaluations is $\mathcal{O}(d \cdot P(s, M))$ and the associated computational complexity is $\mathcal{O}(d \cdot R(s, M))$. The complexity of Lines (VI)–(XIII) is $\mathcal{O}(sd)$. $\qquad\square$

Remark 2. *Since the only possible misassigned values of v_ω contribute to coefficients in $\mathbf{b}$ outside the chosen frequency set $\mathcal{I}$ for which $\Lambda(\mathbf{z}, M)$ is reconstructing, if it is possible to quickly (e.g., in $\mathcal{O}(d)$ time) check a multivariate frequency's inclusion in $\mathcal{I}$ (e.g., a hyperbolic cross), entries outside of $\mathcal{I}$ in $\mathbf{b}$ can be identified in the optional check on Line (X) and remain (correctly) unassigned. This has the effect of removing the $(\beta + \eta_\infty)\sqrt{\max(s - |\mathcal{S}_\beta|, 0)}$ term in the error bound while not increasing the computational complexity. Additionally, this outputs an approximation to $(c|_\mathcal{I})_s^{\mathrm{opt}}$ which is supported only on our supplied frequency set $\mathcal{I}$ as we may expect or prefer.*

We now apply Lemma 5 with the discrete sublinear-time SFT from Theorem 2 to give specific error bounds in terms of best s-term approximation errors as well as detailed runtime and sampling complexities.

Corollary 1 (Algorithm 1 with discrete sublinear-time SFT). *Let $N \ge 9$. For $\mathcal{I} \subset \mathcal{B}_N^d$ with reconstructing rank-1 lattice $\Lambda(\mathbf{z}, M)$ and the function $f \in \mathcal{W}(\mathbb{T}^d) \cap C(\mathbb{T}^d)$, we consider applying Algorithm 1 where each function sample may be corrupted by noise at most $e_\infty \ge 0$ in absolute magnitude. Using the discrete sublinear-time SFT algorithm $\mathcal{A}_{2s,M}^{\mathrm{disc}}$ or $\mathcal{A}_{2s,M}^{\mathrm{disc,MC}}$ with parameter*

$1 \le r \le \frac{M}{36}$, *Algorithm 1 will produce* $\mathbf{b} = (b_{\mathbf{k}})_{\mathbf{k} \in \mathcal{B}_N^d}$ *a 2s-sparse approximation of c satisfying the error estimate*

$$\|\mathbf{b} - c\|_2$$
$$\le \|c - c|_{\mathcal{I}}\|_2 + (48 + 4N) \frac{\|c|_{\mathcal{I}} - (c|_{\mathcal{I}})_s^{\text{opt}}\|_1}{\sqrt{s}} + (188 + 16N) \sqrt{s}(\|f\|_\infty M^{-r} + \|c - c|_{\mathcal{I}}\|_1 + e_\infty)$$

albeit with probability $1 - \sigma \in [0, 1)$ *for the Monte Carlo version. The total number of evaluations of f and computational complexity will be*

$$\mathcal{O}\left(\frac{ds^2 r^{3/2} \log^{11/2} M}{\log s} \right) \quad \text{or} \quad \mathcal{O}\left(dsr^{3/2} \log^{9/2} M \log\left(\frac{dM}{\sigma} \right) \right)$$

for $\mathcal{A}_{2s,M}^{\text{disc}}$ *or* $\mathcal{A}_{2s,M}^{\text{disc,MC}}$ *respectively.*

Proof. For the definitions of τ and β in Lemma 5 with associated values given by Theorem 2, Lemma 1 applied with $\mathbf{x} = c|_{\mathcal{I}}$ implies that $\mathcal{S}_\beta$ can contain at most $2s$ elements, and we have the bound

$$\|c|_{\mathcal{I}} - c|_{\mathcal{S}_\beta}\|_{\ell^2(\mathbb{Z}^d)} \le \|c|_{\mathcal{I}} - (c|_{\mathcal{I}})_{2s}^{\text{opt}}\|_{\ell^2(\mathbb{Z}^d)} + \beta\sqrt{2s} \le \frac{\|c|_{\mathcal{I}} - (c|_{\mathcal{I}})_s^{\text{opt}}\|_{\ell^1(\mathbb{Z}^d)}}{2\sqrt{s}} + \beta\sqrt{2s}, \quad (9)$$

where the last inequality follows from [9, Theorem 2.5] applied to $c|_{\mathcal{I}} - (c|_{\mathcal{I}})_s^{\text{opt}}$. Lemma 5 then holds with s replaced by $2s$ for the $2s$-sparse approximations given by $\mathcal{A}_{2s,M}^{\text{disc}}$ or $\mathcal{A}_{2s,M}^{\text{disc}}$ in Algorithm 1.

Assuming $N \ge 9$, the specific values of τ and η_∞ from Theorem 2 give

$$\beta = \max\left(\tau, \eta_\infty \left(1 + \frac{2}{\sin\left(\frac{\pi}{N}\right)} \right) \right) = \eta_\infty \left(1 + \frac{2}{\sin\left(\frac{\pi}{N}\right)} \right) \le \eta_\infty \left(1 + \frac{2}{9\sin\left(\frac{\pi}{9}\right)} N \right).$$

Using our bound (8) from treating the truncation error as measurement noise additionally accounting for any noise in our input bounded by e_∞ we obtain

$$\beta \le \eta_\infty \left(1 + \frac{2}{9\sin\left(\frac{\pi}{9}\right)} N \right)$$
$$\le 3\sqrt{2} \left(\frac{\|c|_{\mathcal{I}} - (c|_{\mathcal{I}})_s^{\text{opt}}\|_1}{2s} + 2(\|f\|_\infty M^{-r} + \|c - c|_{\mathcal{I}}\|_1 + e_\infty) \right) \left(1 + \frac{2}{9\sin\left(\frac{\pi}{9}\right)} N \right).$$

Inserting the estimate for $\|c|_{\mathcal{I}} - c|_{\mathcal{S}_\beta}\|_2$ from (9), our bound for β above, and the value for η_2 from Theorem 2 (where again we use [9, Theorem 2.5]) into the recovery bound in Lemma 5 gives the final error estimate.

In detail, let

$$A = \frac{\|c|_{\mathcal{I}} - (c|_{\mathcal{I}})_s^{\text{opt}}\|_1}{\sqrt{s}}, \quad \delta = \left(\frac{A}{2\sqrt{s}} + 2(\|f\|_\infty M^{-r} + \|c - c|_{\mathcal{I}}\|_1 + e_\infty) \right).$$

19

Then

$$\beta \le 3\sqrt{2}\left(1 + \frac{2}{9\sin\left(\frac{\pi}{9}\right)}N\right)\delta, \quad \eta_2 \le \frac{A}{2} + 76\sqrt{s}\delta, \quad \|c|_{\mathcal{I}} - c|_{\mathcal{S}_\beta}\|_2 \le \frac{A}{2} + \sqrt{2}\beta\sqrt{s}.$$

Our error bound is then

$$
\begin{aligned}
\|\mathbf{b} - c\|_2 &\le \eta_2 + (\beta + \eta_\infty)\sqrt{2s} + \|c|_{\mathcal{I}} - c|_{\mathcal{S}_\beta}\|_2 + \|c - c|_{\mathcal{I}}\|_2 \\
&\le A + 76\sqrt{s}\delta + 6\left(3 + \frac{4}{9\sin\left(\frac{\pi}{9}\right)}N\right)\sqrt{s}\delta + \|c - c|_{\mathcal{I}}\|_2 \\
&= A + \left(94 + \frac{8}{3\sin\left(\frac{\pi}{9}\right)}N\right)\sqrt{s}\delta + \|c - c|_{\mathcal{I}}\|_2 \\
&= \left(48 + \frac{4}{3\sin\left(\frac{\pi}{9}\right)}N\right)A + \left(188 + \frac{16}{3\sin\left(\frac{\pi}{9}\right)}N\right)\sqrt{s}(\|f\|_\infty M^{-r} + \|c - c|_{\mathcal{I}}\|_1 + e_\infty) \\
&\quad + \|c - c|_{\mathcal{I}}\|_2 \\
&\le (48 + 4N)\frac{\|c|_{\mathcal{I}} - (c|_{\mathcal{I}})_s^{\mathrm{opt}}\|_1}{\sqrt{s}} + (188 + 16N)\sqrt{s}(\|f\|_\infty M^{-r} + \|c - c|_{\mathcal{I}}\|_1 + e_\infty) \\
&\quad + \|c - c|_{\mathcal{I}}\|_2.
\end{aligned}
$$

The change to the complexity of the random algorithm arises from distributing the probability of failure σ over the $d+1$ SFTs in a union bound. $\qquad\square$

Though the nonequispaced SFTs discussed in Theorem 1 do not approximate the discrete Fourier transform and therefore do not alias the one-dimensional frequencies $\mathbf{k} \cdot \mathbf{z}$ into frequencies in $\mathcal{B}_M$, slightly modifying Algorithm 1 to use SFTs with a larger bandwidth allows for the following recovery result.

Corollary 2 (Algorithm 1 with nonequispaced sublinear-time SFT). *For $\mathcal{I} \subset \mathcal{B}_N^d$ with $N \ge 6$, fix the new bandwidth parameter $\tilde{M} := 2\max_{\mathbf{k}\in\mathcal{I}}|\mathbf{k}\cdot\mathbf{z}| + 1$. For $\Lambda(\mathbf{z}, M)$, a reconstructing rank-1 lattice for $\mathcal{I}$ with $M \le \tilde{M}$, and the function $f \in \mathcal{W}(\mathbb{T}^d) \cap C(\mathbb{T}^d)$, we consider applying Algorithm 1 where each function sample may be corrupted by noise at most $e_\infty \ge 0$ in absolute magnitude with the following modifications:*

1. *use the sublinear-time SFT algorithm $\mathcal{A}_{2s,\tilde{M}}^{\mathrm{sub}}$ or $\mathcal{A}_{2s,\tilde{M}}^{\mathrm{sub,MC}}$*

2. *and only check equality against ω in Line (X) (rather than equivalence modulo M),*

to produce $\mathbf{b} = (b_{\mathbf{k}})_{\mathbf{k}\in\mathcal{B}_N^d}$ a $2s$-sparse approximation of c satisfying the error estimate

$$\|\mathbf{b} - c\|_{\ell^2(\mathbb{Z}^d)} \le (24 + 3N)\left[\frac{\|c|_{\mathcal{I}} - (c|_{\mathcal{I}})_s^{\mathrm{opt}}\|_1}{\sqrt{s}} + \sqrt{s}\|c - c|_{\mathcal{I}}\|_1 + \sqrt{s}e_\infty\right] + \|c - c|_{\mathcal{I}}\|_2.$$

albeit with probability $1 - \sigma \in [0, 1)$ for the Monte Carlo version. For $\mathcal{A}_{2s,\tilde{M}}^{\mathrm{sub}}$ and $\mathcal{A}_{2s,\tilde{M}}^{\mathrm{sub,MC}}$ respectively, the total number of evaluations of f and computational complexity will be

$$\mathcal{O}\left(\frac{ds^2\log^4\tilde{M}}{\log s}\right) \quad or \quad \mathcal{O}\left(ds\log^3(\tilde{M})\log\left(\frac{d\tilde{M}}{\sigma}\right)\right).$$

Proof. The bandwidth specified ensures that $\mathcal{B}_{\tilde{M}} \supset \{\mathbf{k} \cdot \mathbf{z} \mid \mathbf{k} \in \mathcal{I}\}$. In the case where f is a trigonometric polynomial with $\operatorname{supp}(c(f)) \subset \mathcal{I}$, so long as there exists some M such that $\Lambda(\mathbf{z}, M)$ is reconstructing for $\mathcal{I}$, the one-dimensional Fourier transforms truncated to $\mathcal{B}_{\tilde{M}}$ coincide with length $\tilde{M}$ DFTs. Thus, we can view an approximation from the algorithm in Theorem 1 as one of a length $\tilde{M}$ DFT. The reasoning in the proofs of Lemma 5 and Corollary 1 then holds with the SFT algorithms, parameters, numbers of samples, and complexities of Theorem 1.

In detail, we first note that for $N \geq 6$, we have

$$\beta = \eta_\infty \left(1 + \frac{2}{\sin\left(\frac{\pi}{N}\right)}\right) \leq \eta_\infty \left(1 + \frac{2}{6\left(\sin\frac{\pi}{6}\right)}N\right) = \eta_\infty \left(1 + \frac{2}{3}N\right).$$

Now let

$$A = \frac{\|c|_{\mathcal{I}} - (c|_{\mathcal{I}})_s^{\mathrm{opt}}\|_1}{\sqrt{s}}, \quad \delta = \left(\frac{A}{\sqrt{s}} + \|c - c|_{\mathcal{I}}\|_1 + e_\infty\right).$$

Then

$$\beta \leq \sqrt{2}\left(1 + \frac{2}{3}N\right)\delta, \quad \eta_2 \leq \frac{A}{2} + (8\sqrt{2} + 6)\sqrt{s}\delta, \quad \|c|_{\mathcal{I}} - c|_{\mathcal{S}_\beta}\|_2 \leq \frac{A}{2} + \sqrt{2}\beta\sqrt{s}.$$

Our error bound is then

$$\begin{aligned}
\|\mathbf{b} - c\|_2 &\leq \eta_2 + (\beta + \eta_\infty)\sqrt{2s} + \|c|_{\mathcal{I}} - c|_{\mathcal{S}_\beta}\|_2 + \|c - c|_{\mathcal{I}}\|_2 \\
&\leq A + (8\sqrt{2} + 6)\sqrt{s}\delta + \left(6 + \frac{8}{3}N\right)\sqrt{s}\delta + \|c - c|_{\mathcal{I}}\|_2 \\
&= A + \left(8\sqrt{2} + 12 + \frac{8}{3}N\right)\sqrt{s}\delta + \|c - c|_{\mathcal{I}}\|_2 \\
&= \left(8\sqrt{2} + 13 + \frac{8}{3}N\right)A + \left(8\sqrt{2} + 12 + \frac{8}{3}N\right)\sqrt{s}(\|c - c|_{\mathcal{I}}\|_1 + e_\infty) + \|c - c|_{\mathcal{I}}\|_2 \\
&\leq (24 + 3N)\left(\frac{\|c|_{\mathcal{I}} - (c|_{\mathcal{I}})_s^{\mathrm{opt}}\|_1}{\sqrt{s}} + \sqrt{s}(\|c - c|_{\mathcal{I}}\|_1 + e_\infty)\right) + \|c - c|_{\mathcal{I}}\|_2.
\end{aligned}$$

$\square$

Remark 3. *As in [13], we can estimate $\tilde{M}$ above with two different techniques:*

$$\tilde{M} = 1 + 2\max_{\mathbf{k} \in \mathcal{I}}\left|\sum_{\ell \in [d]} k_\ell z_\ell\right| \leq 1 + 2\sum_{\ell \in [d]} |z_\ell| \max_{\mathbf{k} \in \mathcal{I}} |k_\ell| = \mathcal{O}(dNM),$$

$$\tilde{M} = 1 + 2\max_{\mathbf{k} \in \mathcal{I}}\left|\sum_{\ell \in [d]} k_\ell z_\ell\right| \leq 1 + 2\|\mathbf{z}\|_\infty \max_{\mathbf{k} \in \mathcal{I}} \|\mathbf{k}\|_1 = \mathcal{O}\left(M \max_{\mathbf{k} \in \mathcal{I}} \|\mathbf{k}\|_1\right).$$

The latter case is especially useful when $\mathcal{I}$ is a subset of a known ℓ^1 ball as it will provide a dimension independent upper bound on $\tilde{M}$. Either of these upper bounds may then be used in practice to avoid having to estimate $\tilde{M}$.

That being said however, if one is willing to perform the one-time search through the frequency set $\mathcal{I}$ to more accurately calculate $\tilde{M}$, one can go even further to use the minimal bandwidth $\tilde{M}' = \max_{\mathbf{k} \in \mathcal{I}}(\mathbf{k} \cdot \mathbf{z}) - \min_{\mathbf{k} \in \mathcal{I}}(\mathbf{k} \cdot \mathbf{z}) + 1$ so long as the function samples are properly modulated to shift the one-dimensional frequencies into $\mathcal{B}_{\tilde{M}'}$. For example, running $\mathcal{A}^{\mathrm{sub}}_{2s,\tilde{M}'}$ or $\mathcal{A}^{\mathrm{sub,MC}}_{2s,\tilde{M}'}$ on $a(t) = e^{2\pi i \phi t} f(t\mathbf{z})$ and $a^\ell(t) = e^{2\pi i \phi t} S_{\ell,1/N} f(t\mathbf{z})$ with $\phi = \left\lfloor \frac{\tilde{M}'}{2} \right\rfloor - \max_{\mathbf{k} \in \mathcal{I}}(\mathbf{k} \cdot \mathbf{z})$ is acceptable so long as this shift is accounted for in the frequency check on Line (X). Of course, these improvements will only have the effect of reducing the logarithmic factors in the computational complexity.

4.2 Two-dimensional DFT technique

Below, we will consider a method for recovering frequencies which, rather than shifting one dimension of the multivariate periodic function f at a time, leaves one dimension of f out at a time. We will fix one dimension $\ell \in [d]$ of f at equispaced nodes over $\mathbb{T}$ and apply a lattice SFT to the other $d - 1$ components. Applying a standard FFT to the results will produce a two-dimensional DFT. The indices corresponding to the standard FFT will represent frequency components in dimension ℓ while the indices corresponding to the lattice SFT will be used to synchronize with known one-dimensional frequencies $\mathbf{k} \cdot \mathbf{z} \bmod M$. The approach is summarized in the following lemma and in Algorithm 2.

Lemma 6. *Fix some finite multivariate frequency set $\mathcal{I} \subset \mathcal{B}^d_N$, let $\Lambda(\mathbf{z}, M)$ be a reconstructing rank-1 lattice for $\{\mathbf{k} - k_\ell \mathbf{e}_\ell \mid \mathbf{k} \in \mathcal{I}\}$, and assume that f has Fourier support $\mathrm{supp}(c) \subset \mathcal{I}$. Fixing one dimension $\ell \in [d]$, and writing the generating vector as $\mathbf{z} = (z_\ell, \mathbf{z}'_\ell) \in \mathbb{Z}^d$, define the polynomials*

$$a^\ell_j(t) := f\left(\frac{j}{N}, t\mathbf{z}'_\ell\right) \text{ for all } j \in [N],$$

that is, fix coordinate ℓ at j/N and restrict the remaining coordinates to dimensions $[d] \setminus \{\ell\}$ of the rank-1 lattice. Then for all one-dimensional frequencies $\omega \in [M]$,

$$\left(\mathbf{F}_M \, \mathbf{a}^\ell_j\right)_\omega = \begin{cases} \displaystyle\sum_{\substack{h_\ell \in \mathcal{B}_N \text{ s.t.} \\ (h_\ell, \mathbf{k}'_\ell) \in \mathcal{I}}} e^{2\pi i j h_\ell / N} c_{(h_\ell, \mathbf{k}'_\ell)}(f) & \text{if there exists } \mathbf{k} \in \mathcal{I} \text{ with } \omega \equiv \mathbf{k}'_\ell \cdot \mathbf{z}'_\ell \pmod{M}, \\ 0 & \text{otherwise.} \end{cases}$$

Moreover, defining the matrix $\mathbf{A}^\ell = \left(\left(\mathbf{F}_M \, \mathbf{a}^\ell_j\right)_\omega\right)_{j \in [N], \omega \in [M]}$, we have

$$\left(\mathbf{F}_N \, \mathbf{A}^\ell\right)_{k_\ell \bmod N, \, \mathbf{k}'_\ell \cdot \mathbf{z}'_\ell \bmod M} = c_{\mathbf{k}}(f) \text{ for all } \mathbf{k} \in \mathcal{I},$$

and the remaining entries of the matrix $\mathbf{F}_N \, \mathbf{A}^\ell \in \mathbb{C}^{N \times M}$ are zero.

Proof. Using the Fourier series representation of f, we have

$$a^\ell_j(t) := \sum_{\mathbf{k} \in \mathcal{I}} c_{\mathbf{k}}(f) \, e^{2\pi i \left(\frac{j k_\ell}{N} + \mathbf{k}'_\ell \cdot \mathbf{z}'_\ell t\right)}.$$

We calculate for $\omega \in [M]$

$$
\begin{aligned}
\left(\mathbf{F}_M\,\mathbf{a}_j^\ell\right)_\omega &= \frac{1}{M} \sum_{i\in[M]} \sum_{\mathbf{h}\in\mathcal{I}} \mathrm{e}^{\frac{2\pi\mathrm{i}jh_\ell}{N}} c_{\mathbf{h}}(f)\, \mathrm{e}^{\frac{2\pi\mathrm{i}(\mathbf{h}_\ell'\cdot\mathbf{z}_\ell'-\omega)i}{M}} \\
&= \sum_{\mathbf{h}\in\mathcal{I}} \mathrm{e}^{\frac{2\pi\mathrm{i}jh_\ell}{N}} c_{\mathbf{h}}(f)\, \delta_{0,(\mathbf{h}_\ell'\cdot\mathbf{z}_\ell'-\omega \bmod M)} \\
&= \sum_{\substack{h_\ell\in\mathcal{B}_N \text{ s.t.} \\ (h_\ell,\mathbf{k}_\ell')\in\mathcal{I}}} \mathrm{e}^{\frac{2\pi\mathrm{i}jh_\ell}{N}} c_{(h_\ell,\mathbf{k}_\ell')}(f),
\end{aligned}
$$

when $\mathbf{k}\in\mathcal{I}$ is such that $\mathbf{k}_\ell'\cdot\mathbf{z}_\ell' \equiv \omega \pmod{M}$, and clearly is zero when no such $\mathbf{k}\in\mathcal{I}$ exists. Note that the final equality uses that $\Lambda(\mathbf{z},M)$ is a reconstructing rank-1 lattice for $\{\mathbf{k}-k_\ell\mathbf{e}_\ell \mid \mathbf{k}\in\mathcal{I}\}$. Applying $\mathbf{F}_N$ to $\mathbf{A}^\ell$ then gives

$$
\left(\mathbf{F}_N\,\mathbf{A}^\ell\right)_{k_\ell \bmod N,\,\mathbf{k}_\ell'\cdot\mathbf{z}_\ell' \bmod M} = \frac{1}{N}\sum_{j\in[N]}\sum_{h_\ell\in\mathcal{B}_N} c_{(h_\ell,\mathbf{k}_\ell')}(f)\, \mathrm{e}^{\frac{2\pi\mathrm{i}(h_\ell-k_\ell \bmod N)j}{N}} = c_{\mathbf{k}}(f).
$$

$\square$

Remark 4. *We bring special attention to the fact that Algorithm 2 requires as input a rank-1 lattice $\Lambda(\mathbf{z},M)$ which is reconstructing for not only $\mathcal{I}$, but also the projections of $\mathcal{I}$ of the form $\{\mathbf{k}-k_\ell \mid \mathbf{k}\in\mathcal{I}\}$ for any $\ell\in[d]$. For frequency sets $\mathcal{I}$ which are downward closed, that is, if $\mathcal{I}$ is such that for any $\mathbf{k}\in\mathcal{I}$ and $\mathbf{h}\in\mathbb{Z}^d$, $|\mathbf{h}|\leq|\mathbf{k}|$ component-wise implies that $\mathbf{h}\in\mathcal{I}$, any reconstructing rank-1 lattice for $\mathcal{I}$ is necessarily one for the considered projections as well. Thus, for many frequency spaces of interest, e.g., hyperbolic crosses (cf. Remarks 2 and 3 as well as Section 5 below), any reconstructing rank-1 lattice for $\mathcal{I}$ will suffice as input to Algorithm 2.*

4.2.1 Analysis of Algorithm 2

Lemma 7 (General recovery result for Algorithm 2.). *Let f, $\mathcal{I}$, and $\Lambda(\mathbf{z},M)$ be as specified in the input to Algorithm 2. Additionally, let $\mathcal{A}_{s,M}$ be a noise-robust SFT algorithm satisfying the same constraints as in Lemma 5 with parameters τ and η_∞ holding uniformly for each SFT performed in Algorithm 2.*

Collect the τ-significant frequencies of f into the set $\mathcal{S}_\tau := \{\mathbf{k}\in\mathcal{I} \mid |c_{\mathbf{k}}(f)|>\tau\}$ and assume that $|\operatorname{supp}(c)\cap\mathcal{S}_\tau|\leq s$. Then Algorithm 2 (ignoring the optional check on Line (XVI)) will produce an s-sparse approximation of the Fourier coefficients of f satisfying the error estimate

$$
\|\mathbf{b}-c\|_{\ell^2(\mathbb{Z}^d)} \leq \eta_2 + (4\tau+\eta_\infty)\sqrt{\max(s-|\mathcal{S}_{4\tau}|,0)} + \|c|_\mathcal{I}-c|_{\mathcal{S}_{4\tau}}\|_{\ell^2(\mathbb{Z}^d)} + \|c-c|_\mathcal{I}\|_{\ell^2(\mathbb{Z}^d)},
$$

requiring $\mathcal{O}\left(dN\cdot P(s,M)\right)$ total evaluations of f, in $\mathcal{O}\left(dN(R(s,M)+sN\log N)\right)$ total operations.

Proof. We begin by assuming that f is a trigonometric polynomial with $\operatorname{supp}(c(f))\subset\mathcal{I}$. Since $\Lambda(\mathbf{z},M)$ is a reconstructing rank-1 lattice for $\mathcal{I}$, the DFT-aliasing ensures that Line (I) of Algorithm 2 will return approximate coefficients uniquely corresponding to all τ-significant frequencies

---

Algorithm 2 Frequency Index Recovery by Two Dimensional DFT

Input: A multivariate periodic function $f \in \mathcal{W}(\mathbb{T}^d) \cap C(\mathbb{T}^d)$ (from which we are able to obtain potentially noisy samples), a multivariate frequency set $\mathcal{I} \subset \mathcal{B}_N^d$, a rank-1 lattice $\Lambda(\mathbf{z}, M)$ which is reconstructing for $\mathcal{I}$ and $\{\mathbf{k} - k_\ell \mathbf{e}_\ell \mid \mathbf{k} \in \mathcal{I}\}$ for all $\ell \in [d]$, and an SFT algorithm $\mathcal{A}_{s,M}$.

Output: Sparse coefficient vector $\mathbf{b} = (b_{\mathbf{k}})_{\mathbf{k} \in \mathcal{B}_N^d}$ (optionally supported on $\mathcal{I}$, see Line (XVI)), an approximation to $(c|_{\mathcal{I}})_s^{\mathrm{opt}}$.

(I) Apply $\mathcal{A}_{s,M}$ to the univariate restriction of f to the lattice, $a(t) := f(t\mathbf{z})$, to produce $\mathbf{v} := \mathcal{A}_{s,M}a$, a sparse approximation of $\mathbf{F}_M \mathbf{a} \in \mathbb{C}^M$.

(II) **for all** $\ell \in [d]$ **do**

(III)   **for all** $j \in [N]$ **do**

(IV)   Apply $\mathcal{A}_{s,M}$ to $a_j^\ell(t) := f(\frac{j}{N}, t\mathbf{z}_\ell')$ to produce $\mathbf{v}_j^\ell := \mathcal{A}_{s,M}a_j^\ell$, a sparse approximation of $\mathbf{F}_M \mathbf{a}_j^\ell$.

(V)   Row j of $\mathbf{V}^\ell \leftarrow \mathbf{v}_j^\ell$.

(VI)   **end for**

(VII)   **for all** nonzero columns ω of $\mathbf{V}^\ell$ **do**

(VIII)   Apply $\mathbf{F}_N$ to column ω of $\mathbf{V}^\ell$ to produce $\mathbf{F}_N \mathbf{V}^\ell$.

(IX)   **end for**

(X) **end for**

(XI) $\mathbf{b} \leftarrow \mathbf{0}$

(XII) **for all** $\omega \in \mathrm{supp}(\mathbf{v})$ **do**

(XIII)   **for all** $\ell \in [d]$ **do**

(XIV) $((k_\omega)_\ell, \sim) \leftarrow \arg\min\{|v_\omega - (\mathbf{F}_N \mathbf{V}^\ell)_{h,\omega'}| \mid (h, \omega') \in \mathcal{B}_N \times [M] \text{ with } hz_\ell + \omega' \equiv \omega \pmod{M}\}$

(XV)   **end for**

(XVI)   **if** $\mathbf{k}_\omega \cdot \mathbf{z} \equiv \omega \pmod{M}$ (and optionally $\mathbf{k}_\omega \in \mathcal{I}$) **then**

(XVII) $b_{\mathbf{k}_\omega} \leftarrow b_{\mathbf{k}_\omega} + v_\omega$

(XVIII)   **end if**

(XIX) **end for**

$\mathbf{k} \in \mathcal{S}_\tau$ which we can label $v_{\mathbf{k} \cdot \mathbf{z} \bmod M}$. Additionally, Line (IV) recovers approximations to all τ-significant frequencies of $\mathbf{F}_M \mathbf{a}_j^\ell$ which have the form given in Lemma 6. In particular, if $\mathbf{k} \in \mathcal{S}_\tau$, we have

$$
\tau < |c_{\mathbf{k}}(f)| = \left| \left(\mathbf{F}_N \mathbf{A}^\ell \right)_{k_\ell \bmod N, \mathbf{k}_\ell' \cdot \mathbf{z}_\ell' \bmod M} \right|
$$

$$
= \left| \frac{1}{N} \sum_{j \in [N]} \left(\mathbf{F}_M \mathbf{a}_j^\ell \right)_{\mathbf{k}_\ell' \cdot \mathbf{z}_\ell' \bmod M} \mathrm{e}^{\frac{-2\pi i j k_\ell \bmod N}{N}} \right|
$$

$$
\leq \frac{1}{N} \sum_{j \in [N]} \left| \left(\mathbf{F}_M \mathbf{a}_j^\ell \right)_{\mathbf{k}_\ell' \cdot \mathbf{z}_\ell' \bmod M} \right|
$$

$$
\leq \max_{j \in [N]} \left| (\mathbf{F}_M \mathbf{a}_j^\ell)_{\mathbf{k}_\ell' \cdot \mathbf{z}_\ell' \bmod M} \right| .
$$

Thus, there exists at least one $\mathbf{F}_M \mathbf{a}_j^\ell$ with $\mathbf{k}_\ell' \cdot \mathbf{z}_\ell' \bmod M$ recovered as a τ-significant frequency in the SFT of Line (IV), and $\mathbf{k}_\ell' \cdot \mathbf{z}_\ell' \bmod M$ will be a nonzero column in $\mathbf{V}^\ell$ for all $\mathbf{k} \in \mathcal{S}_\tau$.

Analyzing these SFTs in more detail for any $\mathbf{k} \in \mathcal{I}$ such that $\mathbf{k}_\ell' \cdot \mathbf{z}_\ell' \bmod M$ is a nonzero column of $\mathbf{V}^\ell$, we write

$$
\left(\mathbf{v}_j^\ell \right)_{\mathbf{k}_\ell' \cdot \mathbf{z}_\ell' \bmod M} = \left(\mathbf{F}_M \mathbf{a}_j^\ell \right)_{\mathbf{k}_\ell' \cdot \mathbf{z}_\ell' \bmod M} + \left(\eta_j^\ell \right)_{\mathbf{k}_\ell' \cdot \mathbf{z}_\ell' \bmod M}
$$

where, by the ℓ^∞ and recovery guarantees for $\mathcal{A}_{s,M}$, the error satisfies

$$
\left| \left(\eta_j^\ell \right)_{\mathbf{k}_\ell' \cdot \mathbf{z}_\ell' \bmod M} \right| \leq \begin{cases} \eta_\infty & \text{if } \left(\mathbf{v}_j^\ell \right)_{\mathbf{k}_\ell' \cdot \mathbf{z}_\ell' \bmod M} \neq 0 \\ \tau & \text{if } \left(\mathbf{v}_j^\ell \right)_{\mathbf{k}_\ell' \cdot \mathbf{z}_\ell' \bmod M} = 0 \end{cases} \leq \tau.
$$

Thus, in the application of $\mathbf{F}_N$ to column $\mathbf{k}_\ell' \cdot \mathbf{z}_\ell' \bmod M$ of $\mathbf{V}^\ell$, we have

$$
\left(\mathbf{F}_N \mathbf{V}^\ell \right)_{k_\ell \bmod N, \mathbf{k}_\ell' \cdot \mathbf{z}_\ell' \bmod M}
$$

$$
= \left(\mathbf{F}_N \mathbf{A}^\ell \right)_{k_\ell \bmod N, \mathbf{k}_\ell' \cdot \mathbf{z}_\ell' \bmod M} + \left(\mathbf{F}_N \left(\left(\eta_j^\ell \right)_{\mathbf{k}_\ell' \cdot \mathbf{z}_\ell' \bmod M} \right)_{j \in [N]} \right)_{k_\ell \bmod N}
$$

$$
=: c_{\mathbf{k}}(f) + \eta_{\mathbf{k}}^\ell
$$

with

$$
|\eta_{\mathbf{k}}^\ell| = \left| \frac{1}{N} \sum_{j \in [N]} \left(\eta_j^\ell \right)_{\mathbf{k}_\ell' \cdot \mathbf{z}_\ell' \bmod M} \mathrm{e}^{\frac{-2\pi i j k_\ell \bmod N}{N}} \right| \leq \max_{j \in [N]} \left| \left(\eta_j^\ell \right)_{\mathbf{k}_\ell' \cdot \mathbf{z}_\ell' \bmod M} \right| \leq \tau.
$$

These same calculations apply to the computed columns of $\mathbf{F}_N \mathbf{V}^\ell$ which do not correspond to values of $\mathbf{k}_\ell' \cdot \mathbf{z}_\ell' \bmod M$ for some $\mathbf{k} \in \mathcal{I}$ since we assume $\mathrm{supp}(c(f)) \subset \mathcal{I}$, and so at worst, these columns are filled with noise bounded in magnitude by τ.

Restricting our attention to $\mathbf{k} \in \mathcal{S}_{4\tau} \subset \mathcal{S}_\tau$, we know that Line (XIV) will be run with $\omega = \mathbf{k} \cdot \mathbf{z} \bmod M$ and $(k_\ell \bmod N, \mathbf{k}_\ell' \cdot \mathbf{z}_\ell' \bmod M)$ as an admissible index in the minimization. By the reconstructing property of $\Lambda(\mathbf{z}, M)$, no other $\mathbf{h} \in \mathcal{I}$ will correspond to an admissible index $(h_\ell \bmod N, \mathbf{h}_\ell' \cdot \mathbf{z}_\ell' \bmod M)$, and so the only remaining values of $(\mathbf{F}_N \mathbf{V}^\ell)_{h, \omega'}$ in the minimization correspond

to pure noise η bounded in magnitude by τ. Analyzing the objective at $(k_\ell \bmod N, \mathbf{k}'_\ell \cdot \mathbf{z}'_\ell \bmod M)$, we find

$$|v_{\mathbf{k}\cdot\mathbf{z} \bmod M} - (\mathbf{F}_N \mathbf{V}^\ell)_{k_\ell \bmod N, \mathbf{k}'_\ell \cdot \mathbf{z}'_\ell \bmod M}| \le 2\tau < |c_{\mathbf{k}}(f)| - 2\tau \le |v_{\mathbf{k}\cdot\mathbf{z} \bmod M} - \eta|,$$

and so the value for $(k_\omega)_\ell$ will in fact be assigned k_ℓ. Thus, after all d components of $\mathbf{k}_\omega = \mathbf{k}$ have been recovered, $b_{\mathbf{k}}$ will be assigned $v_{\mathbf{k}\cdot\mathbf{z} \bmod M}$.

The remaining $\max(s - |\mathcal{S}_{4\tau}|, 0)$ nonzero entries of $\mathbf{v}$ can be distributed to entries of $\mathbf{b}$ possibly correctly but with no guarantee; at the very least however, these values must be at most $4\tau + \eta_\infty$ in magnitude. We split $\mathbf{b}$ as $\mathbf{b} = \mathbf{b}^{\mathrm{correct}} + \mathbf{b}^{\mathrm{incorrect}}$ to account for the values of $\mathbf{v}$ respectively assigned correctly and incorrectly and note that $\mathrm{supp}(\mathbf{b}^{\mathrm{correct}}) \supset \mathcal{S}_{4\tau}$. We then estimate the error as

$$\|\mathbf{b} - c\|_{\ell^2(\mathbb{Z}^d)} \le \|\mathbf{b}^{\mathrm{correct}} - c|_{\mathrm{supp}(\mathbf{b}^{\mathrm{correct}})}\|_{\ell^2(\mathbb{Z}^d)} + \|\mathbf{b}^{\mathrm{incorrect}}\|_{\ell^2(\mathbb{Z}^d)} + \|c - c|_{\mathrm{supp}(\mathbf{b}^{\mathrm{correct}})}\|_{\ell^2(\mathbb{Z}^d)}$$
$$\le \eta_2 + (4\tau + \eta_\infty)\sqrt{\max(s - |\mathcal{S}_{4\tau}|, 0)} + \|c - c|_{\mathcal{S}_{4\tau}}\|_{\ell^2(\mathbb{Z}^d)}.$$

As in the proof of Lemma 5, we note that the mandatory check in Line (XVI) helps ensure that all misassigned values v_ω which contribute to $\mathbf{b}^{\mathrm{incorrect}}$ correspond to reconstructed $\mathbf{k}_\omega$ *outside* of $\mathcal{I}$, with the optional check in this line (see Remark 2) eliminating $\mathbf{b}^{\mathrm{incorrect}}$ and the corresponding term in the error estimate entirely.

Now, supposing that the Fourier support of f is not limited to only $\mathcal{I}$, just as in the analysis for Algorithm 1, we treat f as a perturbation of $f_{\mathcal{I}}$, and use the robust SFT algorithm and the previous argument to approximate $c|_{\mathcal{I}}$. Note again that in each SFT, the noise added when using measurements of f as proxies for those of $f_{\mathcal{I}}$ is compounded by $\|f_{\mathbb{Z}^d \setminus \mathcal{I}}\|_\infty$ and is bounded by $\|c - c|_{\mathcal{I}}\|_{\ell^1(\mathbb{Z}^d)}$. Applying the guarantees above gives

$$\|\mathbf{b} - c\|_{\ell^2(\mathbb{Z}^d)} \le \|\mathbf{b} - c|_{\mathcal{I}}\|_{\ell^2(\mathbb{Z}^d)} + \|c - c|_{\mathcal{I}}\|_{\ell^2(\mathbb{Z}^d)}$$
$$\le \eta_2 + (4\tau + \eta_\infty)\sqrt{\max(s - |\mathcal{S}_{4\tau}|, 0)} + \|c|_{\mathcal{I}} - c|_{\mathcal{S}_{4\tau}}\|_{\ell^2(\mathbb{Z}^d)} + \|c - c|_{\mathcal{I}}\|_{\ell^2(\mathbb{Z}^d)}.$$

Employing fast Fourier transforms for the at most dsN DFTs, the computational complexity of Lines (II)–(X) is $\mathcal{O}\left(d(N \cdot R(s, M) + sN^2 \log N)\right)$ (which dominates the complexity of the remainder of the algorithm). Since $1 + dN$ SFTs are required, the number of f evaluations is $\mathcal{O}(dN \cdot P(s, M))$. □

Remark 5. *Though the number of nonzero columns of $\mathbf{V}^\ell$ can be theoretically at most sN, in practice with a high quality algorithm, each of the N SFTs should recover nearly the same frequencies, meaning that there are actually $\mathcal{O}(s)$ columns. This would remove a power of N in the second term of the runtime estimate.*

Note however, that even with near exact SFT algorithms, recovering exactly s total frequencies is not a certainty. There can be cancellations for certain terms in $\mathbf{F}_M \mathbf{a}_j^\ell$ depending interactions between the coefficients sharing the same values on their $[d] \setminus \{\ell\}$ entries, which makes it possible that an SFT on $\mathbf{F}_M \mathbf{a}_j^\ell$ will miss coefficients. If required to output s-entries, an SFT algorithm could favor some noisy value corresponding to a frequency outside the support.

Remark 6. *Though we perform an exact FFT of the nonzero columns of $\mathbf{V}^\ell$ in Line (VIII) of Algorithm 2, Lemma 6 implies that the resulting matrix will be as sparse as the original function's Fourier transform. Thus, for a truly compressible function, an SFT down the columns of $\mathbf{V}^\ell$*

26

would be feasible as well. However, in especially higher dimensions, even small N can support large frequency spaces $\mathcal{I}$. In these large frequency spaces, what is perceived as relatively sparse can therefore quickly surpass N, rendering an s-sparse, length N SFT useless.

Applying the discrete sublinear-time SFT from Theorem 2 to Lemma 7 analogously to the derivation of Corollary 1 from Lemma 5 allows for the following recovery bound for Algorithm 2. In particular, we observe asymptotically improved error guarantees over Corollary 1 at the cost of a slight increase in runtime.

Corollary 3 (Algorithm 2 with discrete sublinear-time SFT). *For $\mathcal{I} \subset \mathbb{Z}^d$ with reconstructing rank-1 lattice $\Lambda(\mathbf{z}, M)$ and the function $f \in \mathcal{W}(\mathbb{T}^d) \cap C(\mathbb{T}^d)$, we consider applying Algorithm 2 where each function sample may be corrupted by noise at most $e_\infty \geq 0$ in absolute magnitude. Using the discrete sublinear-time SFT algorithm $\mathcal{A}_{2s,M}^{\mathrm{disc}}$ or $\mathcal{A}_{2s,M}^{\mathrm{disc,MC}}$ with parameter $1 \leq r \leq \frac{M}{36}$ will produce $\mathbf{b} = (b_\mathbf{k})_{\mathbf{k} \in \mathcal{B}_N^d}$ a $2s$-sparse approximation of c satisfying the error estimate*

$$\|\mathbf{b} - c\|_{\ell^2(\mathbb{Z}^d)} \leq 206 \frac{\|c|_{\mathcal{I}} - (c|_{\mathcal{I}})_s^{\mathrm{opt}}\|_1}{\sqrt{s}} + 820\sqrt{s}(\|f\|_\infty M^{-r} + \|c - c|_{\mathcal{I}}\|_1 + e_\infty)$$
$$+ \|c - c|_{\mathcal{I}}\|_2$$

albeit with probability $1 - \sigma \in [0, 1)$ for the Monte Carlo version.

The total number of evaluations of f and the computational complexity will be

$$\mathcal{O}\left(dsN\left(\frac{sr^{3/2}\log^{11/2} M}{\log s} + N\log N\right)\right) \quad or \quad \mathcal{O}\left(dsN\left(r^{3/2}\log^{9/2}(M)\log\left(\frac{dNM}{\sigma}\right) + N\log N\right)\right)$$

for $\mathcal{A}_{2s,M}^{\mathrm{disc}}$ or $\mathcal{A}_{2s,M}^{\mathrm{disc,MC}}$ respectively.

Proof. In detail, let

$$A = \frac{\|c|_{\mathcal{I}} - (c|_{\mathcal{I}})_s^{\mathrm{opt}}\|_1}{\sqrt{s}}, \quad B = \left(274 + 96\sqrt{2}\right)$$
$$\delta = \left(\frac{A}{2\sqrt{s}} + 2(\|f\|_\infty M^{-r} + \|c - c|_{\mathcal{I}}\|_1 + e_\infty)\right).$$

Then

$$\tau \leq 12(1 + \sqrt{2})\delta, \quad \eta_\infty \leq 3\sqrt{2}\delta$$
$$\eta_2 \leq \frac{A}{2} + 76\sqrt{s}\delta,$$
$$\|c|_{\mathcal{I}} - c|_{\mathcal{S}_{4\tau}}\|_2 \leq \frac{A}{2} + 4\sqrt{2}\tau\sqrt{s} = \frac{A}{2} + 48(2 + \sqrt{2})\delta\sqrt{s}.$$

Our error bound is then

$$\|\mathbf{b} - c\|_{\ell^2(\mathbb{Z}^d)} \leq \eta_2 + (4\tau + \eta_\infty)\sqrt{2s} + \|c|_{\mathcal{I}} - c|_{\mathcal{S}_{4\tau}}\|_{\ell^2(\mathbb{Z}^d)} + \|c - c|_{\mathcal{I}}\|_{\ell^2(\mathbb{Z}^d)}$$
$$\leq A + B\sqrt{s}\delta + \|c - c|_{\mathcal{I}}\|_2$$
$$= \left(1 + \frac{B}{2}\right)A + 2B\sqrt{s}(\|f\|_\infty M^{-r} + \|c - c|_{\mathcal{I}}\|_1 + e_\infty) + \|c - c|_{\mathcal{I}}\|_2$$
$$\leq 206\frac{\|c|_{\mathcal{I}} - (c|_{\mathcal{I}})_s^{\mathrm{opt}}\|_1}{\sqrt{s}} + 820\sqrt{s}(\|f\|_\infty M^{-r} + \|c - c|_{\mathcal{I}}\|_1 + e_\infty)$$
$$+ \|c - c|_{\mathcal{I}}\|_2.$$

$\square$

Again, the same strategy from Corollary 2 of widening the frequency band and shifting the one-dimensional transforms accordingly allows us to use the nonequispaced SFT algorithm from Theorem 1 in Algorithm 2. Note here that the widening and shifting occurs on a dimension by dimension basis so as to account for the differing one-dimensional frequencies of the form $\mathbf{k}'_\ell \cdot \mathbf{z}'_\ell$ for $\mathbf{k} \in \mathcal{I}$.

Corollary 4 (Algorithm 2 with nonequispaced sublinear-time SFT). *For $\mathcal{I} \subset \mathcal{B}^d_N$, let $\tilde{M}$ be the larger one-dimensional bandwidth parameter from Corollary 2, and additionally define $\tilde{M}^\ell :=$ $2 \max_{\mathbf{k} \in \mathcal{I}} |\mathbf{k}'_\ell \cdot \mathbf{z}'_\ell| + 1$. For $\Lambda(\mathbf{z}, M)$, a reconstructing rank-1 lattice for $\mathcal{I}$ with $M \leq \min\{\tilde{M}, \min_{\ell \in [d]} \tilde{M}^\ell\}$, and the function $f \in \mathcal{W}(\mathbb{T}^d) \cap C(\mathbb{T}^d)$, we consider applying Algorithm 2 where each function sample may be corrupted by noise at most $e_\infty \geq 0$ in absolute magnitude with the following modifications:*

1. *use the sublinear-time SFT algorithm $\mathcal{A}^{\mathrm{sub}}_{2s,\tilde{M}}$ or $\mathcal{A}^{\mathrm{sub,MC}}_{2s,\tilde{M}}$ in Line (I) and $\mathcal{A}^{\mathrm{sub}}_{2s,\tilde{M}^\ell}$ or $\mathcal{A}^{\mathrm{sub,MC}}_{2s,\tilde{M}^\ell}$ in Line (IV)*

2. *and only check equality against ω in Line (XIV) (rather than equivalence modulo M),*

to produce $\mathbf{b} = (b_{\mathbf{k}})_{\mathbf{k} \in \mathcal{B}^d_N}$ a $2s$-sparse approximation of c satisfying the error estimate

$$\|\mathbf{b} - c\|_{\ell^2(\mathbb{Z}^d)} = 98 \left(\frac{\|c|_{\mathcal{I}} - (c|_{\mathcal{I}})^{\mathrm{opt}}_s\|_1}{\sqrt{s}} + \sqrt{s}\|c - c|_{\mathcal{I}}\|_1 + \sqrt{s}e_\infty \right) + \|c - c|_{\mathcal{I}}\|_2$$

albeit with probability $1 - \sigma \in [0, 1)$ for the Monte Carlo version.

Letting $\bar{M} = \max(\tilde{M}, \max_{\ell \in [d]} \tilde{M}^\ell)$, the total number of evaluations of f will be

$$\mathcal{O}\left(\frac{dNs^2 \log^4 \bar{M}}{\log s} \right) \quad \text{or} \quad \mathcal{O}\left(dNs \log^3 \bar{M} \log\left(\frac{dN\bar{M}}{\sigma} \right) \right)$$

with associated computational complexities

$$\mathcal{O}\left(dNs \left(\frac{s \log^4 \bar{M}}{\log s} + N \log N \right) \right) \quad \text{or} \quad \mathcal{O}\left(dNs \left(\log^3 \bar{M} \log\left(\frac{dN\bar{M}}{\sigma} \right) + N \log N \right) \right)$$

for $\mathcal{A}^{\mathrm{sub}}_{2s,\cdot}$ and $\mathcal{A}^{\mathrm{sub,MC}}_{2s,\cdot}$ respectively.

Proof. In detail, let

$$A = \frac{\|c|_{\mathcal{I}} - (c|_{\mathcal{I}})^{\mathrm{opt}}_s\|_1}{\sqrt{s}}, \quad B = 40(1 + \sqrt{2})$$

$$\delta = \left(\frac{A}{\sqrt{s}} + \|c - c|_{\mathcal{I}}\|_1 + e_\infty \right).$$

Then

$$\tau \leq (4 + 2\sqrt{2})\delta, \quad \eta_\infty \leq \sqrt{2}\delta$$

$$\eta_2 \leq \frac{A}{2} + (8\sqrt{2} + 6)\sqrt{s}\delta,$$

$$\|c|_{\mathcal{I}} - c|_{\mathcal{S}_{4\tau}}\|_2 \leq \frac{A}{2} + 4\sqrt{2}\tau\sqrt{s} = \frac{A}{2} + 16(1 + \sqrt{2})\delta\sqrt{s}.$$

Our error bound is then

$$\|\mathbf{b} - c\|_2 \le \eta_2 + (4\tau + \eta_\infty)\sqrt{2s} + \||c|_{\mathcal{I}} - c|_{\mathcal{S}_{4\tau}}\|_{\ell^2(\mathbb{Z}^d)} + \|c - c|_{\mathcal{I}}\|_{\ell^2(\mathbb{Z}^d)}$$

$$= A + B\sqrt{s}\delta + \|c - c|_{\mathcal{I}}\|_2$$

$$= (1 + B)\frac{\||c|_{\mathcal{I}} - (c|_{\mathcal{I}})_s^{\mathrm{opt}}\|_1}{\sqrt{s}} + B\sqrt{s}\|c - c|_{\mathcal{I}}\|_1 + B\sqrt{s}e_\infty + \|c - c|_{\mathcal{I}}\|_2$$

$$= 98\left(\frac{\||c|_{\mathcal{I}} - (c|_{\mathcal{I}})_s^{\mathrm{opt}}\|_1}{\sqrt{s}} + \sqrt{s}\|c - c|_{\mathcal{I}}\|_1 + \sqrt{s}e_\infty\right) + \|c - c|_{\mathcal{I}}\|_2$$

$\square$

Remark 7. *The bounds in Remark 3 will still hold for $\tilde{M}^\ell$ as well; thus one of these upper bounds can be used as the effective bandwidth parameter for every SFT without having to calculate the $d+1$ bandwidths by scanning $\mathcal{I}$. Again however, if this scan is tolerable, one can reduce the overall complexity by using analogous minimal bandwidths discussed in Remark 3 along with corresponding frequency shifts.*

5 Numerics

We now demonstrate the effectiveness of our phase encoding and two-dimensional DFT algorithms for computing Fourier coefficients of multivariate functions in a series of empirical tests. The two techniques are implemented in MATLAB, with the code for the algorithms and tests in this section publicly available[1]. The results below use a MATLAB implementation[2] of the randomized univariate sublinear-time nonequispaced algorithm $\mathcal{A}_{2s,M}^{\mathrm{sub,MC}}$ (cf. Theorem 1) as the underlying SFT for both multivariate approaches as this allows for the fastest runtime and most sample efficient implementations.

In the univariate code, all parameters but one are qualitatively tuned below theoretical upper bounds to increase efficiency while maintaining accuracy and are kept constant between tests below. In particular, we fix the values $\texttt{C} := 1$, $\texttt{sigma} := 2/3$, and $\texttt{primeShift} := 0$ (see the documentation and the original paper [16] for more detail). The only parameter we vary is $\texttt{randomScale}$ which affects the rate at which the deterministic algorithm $\mathcal{A}_{2s,M}^{\mathrm{sub}}$ is randomly sampled to produce the Monte Carlo version $\mathcal{A}_{2s,M}^{\mathrm{sub,MC}}$. This parameter represents a multiplicative scaling on logarithmic factors of the bandwidth which determines how many prime numbers are randomly selected from those used in the deterministic SFT implementation. Therefore, lower values of $\texttt{randomScale}$ will result in using fewer prime numbers, decreasing the number of function samples and overall runtime at the risk of a higher probability of failure. We consider values well below the code default and theoretical upper bound of 21 given in [16].

5.1 Exactly sparse case

In the beginning, we consider the case of multivariate trigonometric polynomials with frequencies supported within hyperbolic cross index sets. We define the d-dimensional hyperbolic cross

---
[1]available at https://gitlab.com/grosscra/Rank1LatticeSparseFourier
[2]available at https://gitlab.com/grosscra/SublinearSparseFourierMATLAB

frequency set

$$\mathcal{H}_N^d := \left\{ \mathbf{k} \in \mathbb{Z}^d \colon \prod_{\ell=1}^d \max(1, |k_\ell|) \leq \frac{N}{2} \quad \text{and} \quad \max_{\ell=1,\ldots,d} k_\ell < \frac{N}{2} \right\} \subset \mathcal{B}_N^d$$

of expansion $N \in \mathbb{N}$. For a given sparsity s, we choose s many frequencies uniformly at random from $\mathcal{H}_N^d$, and we randomly draw corresponding Fourier coefficients $c_{\mathbf{k}}$ from $[-1, 1] + \mathrm{i}\,[-1, 1]$, $|c_{\mathbf{k}}| \geq 10^{-3}$. For each parameter setting, we perform the tests 100 times and determine the success rate, i.e., the relative number of cases (out of the 100) where all frequencies were correctly detected, as well as the average number of samples.

5.1.1 Random frequency sets within 10-dimensional hyperbolic cross and high-dimensional full cuboids

We set the spatial dimension $d := 10$, the expansion $N := 33$, and use $\mathcal{I} := \mathcal{H}_{33}^{10}$ as set of possible frequencies with cardinality $|\mathcal{I}| = 45\,548\,649$. Then, the rank-1 lattice with generating vector

$$\mathbf{z} := (1, 33, 579, 3\,628, 21\,944, 169\,230, 1\,105\,193, 7\,798\,320, 49\,768\,670, 320\,144\,128)^\top \qquad (10)$$

and lattice size $M := 2\,040\,484\,044$ is a reconstructing one. We apply Algorithm 1 and Algorithm 2 with the SFT algorithm $\mathcal{A}_{2s,\tilde{M}}^{\text{sub,MC}}$.

In Figure 1a, the average numbers of samples (over 100 test runs) are plotted against the used sparsities $s \in \{10, 20, 50, 100, 200, 500, 1000\}$ for Algorithm 1 and $s \in \{10, 20, 50, 100\}$ for Algorithm 2. The magenta line with circles corresponds to Algorithm 1 with bandwidth parameter $\tilde{M} = dNM \approx 6.7 \cdot 10^{11}$ and `randomScale` := 0.3. We observe that the number of samples grow nearly linearly with respect to the sparsity s. Moreover, the success rate is at least 0.99 (99 out of 100 test runs), where we define success such that the support of output (sparse coefficient vector) contains the true frequencies. Next, we reduce the bandwidth $\tilde{M}$ to $1 + 2\|\mathbf{z}\|_\infty \max_{\mathbf{k}\in\mathcal{I}} \|\mathbf{k}\|_1 \approx 1.6 \cdot 10^{10}$, see also Remark 3, and visualize this as solid blue line with squares. This smaller bandwidth causes a decrease in the number of samples of up to 50 percent while only mildly decreasing the success rates to values not below 0.90. Increasing the `randomScale` parameter to 0.5, denoted by dashed blue line with squares, raises the success rate to 1.00 while achieving still fewer samples than bandwidth parameter $\tilde{M} = dNM \approx 6.7 \cdot 10^{11}$ and `randomScale` = 0.3 (solid magenta line with circles). The numbers of samples for Algorithm 2 are plotted as solid and dashed red lines with triangles for `randomScale` := 0.3 and 0.5, respectively, choosing the bandwidth $\tilde{M} := 1 + 2\|\mathbf{z}\|_\infty \max_{\mathbf{k}\in\mathcal{I}} \|\mathbf{k}\|_1 \approx 1.6 \cdot 10^{10}$. We observe that Algorithm 2 requires much more samples, more than one order of magnitude, compared to Algorithm 1, while achieving similar success rates. For comparison, in case of sparsity $s = 100$ and `randomScale` = 0.5, Algorithm 2 takes almost $M = 2\,040\,484\,044$ samples, where the latter would be required by a non-SFT approach which uses all rank-1 lattice nodes.

In Figure 1b, we investigate the dependence of the required number of samples of Algorithm 1 and 2 on the spatial dimension d, where we consider the values $d \in \{10, 11, \ldots, 20\}$. For this, we use a slightly different setting, where we choose $s = 100$ random frequencies from a full cuboid of cardinality $\approx 10^{12}$. For instance, we utilize the cuboid $\mathcal{I} := \{-8, -7, \ldots, 7\}^9 \times \{-7, -6, \ldots, 7\}$, $|\mathcal{I}| \approx 1.03 \cdot 10^{12}$, in the case $d = 10$ and $\mathcal{I} := \{-2, -1, \ldots, 2\} \times \{-2, -1, 0, 1\}^{18} \times \{-1, 0, 1\}$, $|\mathcal{I}| \approx 1.03 \cdot 10^{12}$, for $d = 20$. The rank-1 lattice size and the bandwidth parameter are chosen
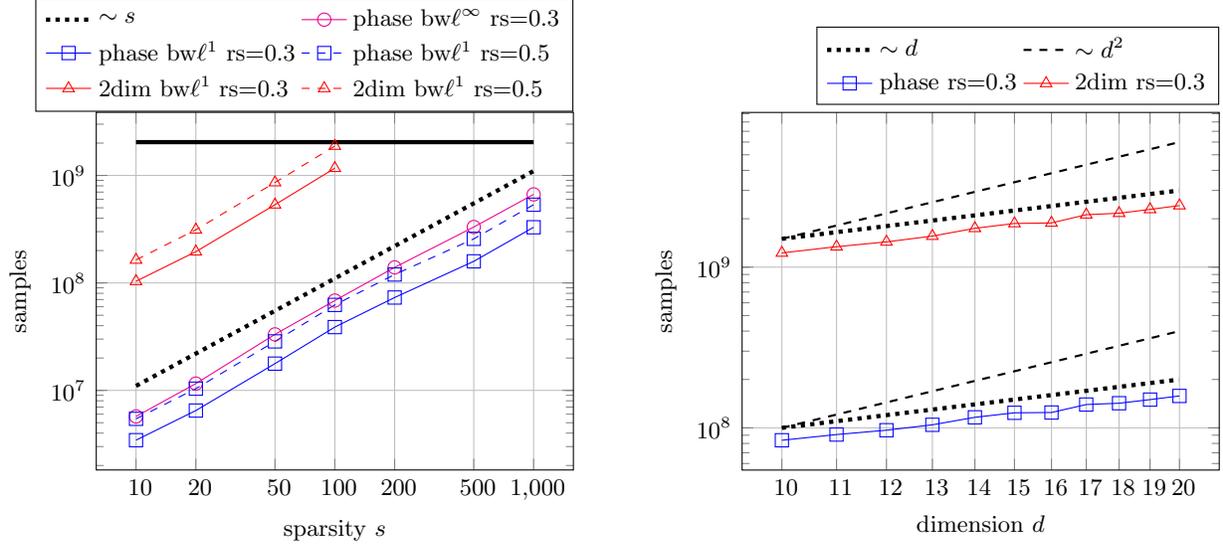
(a) Samples vs. sparsity s. Random frequencies are chosen from hyperbolic cross $\mathcal{I} := \mathcal{H}^{10}_{33}$. "bw$\ell^\infty$" and "bw$\ell^1$" correspond to bandwidth parameters $\tilde{M} = dNM \approx 6.7 \cdot 10^{11}$ and $\tilde{M} = 1 + 2\|\mathbf{z}\|_\infty \max_{\mathbf{k}\in\mathcal{I}} \|\mathbf{k}\|_1 \approx 1.6 \cdot 10^{10}$, respectively.

(b) Samples vs. spatial dimension d. Random frequencies are chosen from **full cuboid** $\mathcal{I}$ of cardinality $|\mathcal{I}| \approx 10^{12}$ with corresponding lattice of size $M = |\mathcal{I}|$ and bandwidth parameter $\tilde{M} = M$.

Figure 1: Average number of samples over 100 test runs for Algorithm 1 with $\mathcal{A}^{\mathrm{sub,MC}}_{2s,\tilde{M}}$, denoted by "phase", and Algorithm 2 with $\mathcal{A}^{\mathrm{sub,MC}}_{2s,\tilde{M}}$, denoted by "2dim", on random multivariate trigonometric polynomials, setting `randomScale := rs`.

to be $M = \tilde{M} = |\mathcal{I}|$. The generating vector $\mathbf{z}$ follows from the edge lengths of the cube, e.g., $\mathbf{z} := (1, 16, 16 \cdot 15, 16 \cdot 15^2, 16 \cdot 15^3, \ldots, 16 \cdot 15^8)^\top$ for $d = 10$ and $\mathbf{z} := (1, 5, 5 \cdot 4, 5 \cdot 4^2, 5 \cdot 4^3, \ldots, 5 \cdot 4^{18})^\top$ for $d = 20$. Since the expansion N is a factor in the number of samples of Algorithm 2, cf. Corollary 4, and we want to concentrate on the dependence on the spatial dimension d, we now fix this parameter to $N := 16$ independent of d. Moreover, the `randomScale` parameter is set to 0.3. The plots indicate that the numbers of samples grow approximately linearly with respect to the dimension d as stated by Corollaries 2 and 4 for Algorithms 1 and 2, respectively. The success rates are slightly better compared to the tests from Figure 1a.

5.1.2 Random frequency sets within 10-dimensional hyperbolic cross and noisy samples

In this section, we again consider random multivariate trigonometric polynomials with frequencies supported within the hyperbolic cross index set $\mathcal{H}^{10}_{33}$ of expansion $N = 33$ and use the reconstructing rank-1 lattice with generating vector $\mathbf{z}$ as stated in (10) and size $M := 2\,040\,484\,044$. Similarly as in [21, Section 5.2], we perturb the samples of the trigonometric polynomial by additive complex (white) Gaussian noise $\varepsilon_j \in \mathbb{C}$ with zero mean and standard deviation σ. The noise is generated by $\varepsilon_j := \sigma/\sqrt{2}\,(\varepsilon_{1,j} + \mathrm{i}\varepsilon_{2,j})$ where $\varepsilon_{1,j}, \varepsilon_{2,j}$ are independent standard normal distributed. Since the

signal-to-noise ratio (SNR) can be approximately computed by

$$\text{SNR} \approx \frac{\sum_{j=0}^{M-1} |f(\mathbf{x}_j)|^2/M}{\sum_{j=0}^{M-1} |\varepsilon_j|^2/M} \approx \frac{\sum_{\mathbf{k}\in\text{supp}(c)} |c_\mathbf{k}(f)|^2}{\sigma^2},$$

this leads to the choice $\sigma := \sqrt{\sum_{\mathbf{k}\in\text{supp}(c)} |c_\mathbf{k}(f)|^2}/\sqrt{\text{SNR}}$ for a targeted SNR value. The SNR is often expressed in the logarithmic decibel scale (dB), $\text{SNR}_{\text{dB}} = 10 \log_{10} \text{SNR}$ and $\text{SNR} = 10^{\text{SNR}_{\text{dB}}/10}$, i.e., a linear $\text{SNR} = 10^2$ corresponds to a logarithmic $\text{SNR}_{\text{dB}} = 20$ and $\text{SNR} = 10^3$ corresponds to $\text{SNR}_{\text{dB}} = 30$. Here, we perform tests with sparsity $s = 100$ and signal-to-noise ratios $\text{SNR}_{\text{dB}} \in \{0, 5, 10, 15, 20, 25, 30\}$. Moreover, we only use the bandwidth parameter $\tilde{M} = 1 + 2\|\mathbf{z}\|_\infty \max_{\mathbf{k}\in\mathcal{I}} \|\mathbf{k}\|_1 \approx 1.6 \cdot 10^{10}$. Besides that, we choose the algorithm parameters as in Figure 1a.
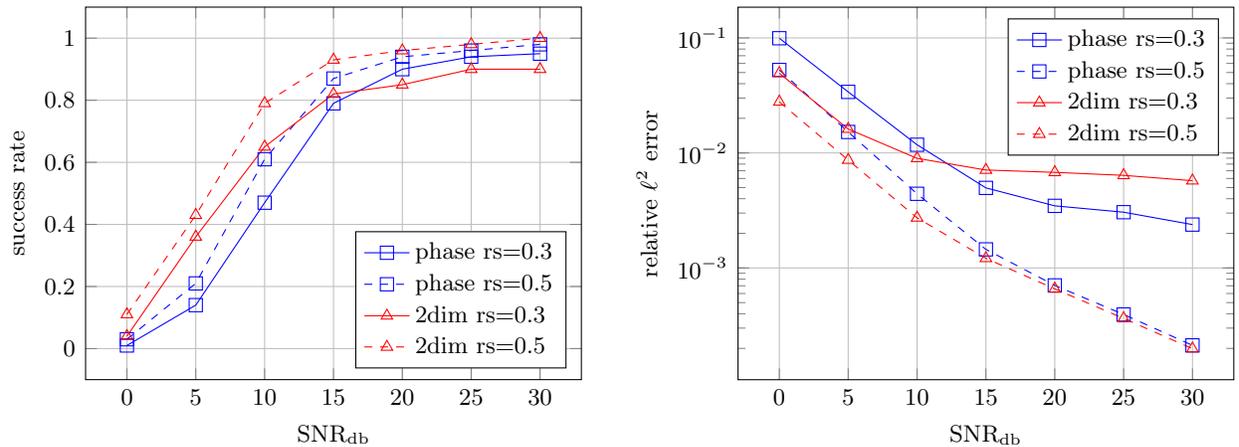
(a) success rates vs. noise level for $s = 100$ (b) relative ℓ^2 errors vs. noise level for $s = 100$

Figure 2: Average success rates (all frequencies detected) and relative ℓ^2 errors over 100 test runs for Algorithm 1 with $\mathcal{A}_{2s,\tilde{M}}^{\text{sub,MC}}$, denoted by "phase", and Algorithm 2 with $\mathcal{A}_{2s,\tilde{M}}^{\text{sub,MC}}$, denoted by "2dim", on random multivariate trigonometric polynomials within hyperbolic cross $\mathcal{I} := \mathcal{H}_{33}^{10}$, setting `randomScale` := rs $\in \{0.3, 0.5\}$ and bandwidth parameter $\tilde{M} = 1 + 2\|\mathbf{z}\|_\infty \max_{\mathbf{k}\in\mathcal{I}} \|\mathbf{k}\|_1 \approx 1.6 \cdot 10^{10}$.

In Figure 2a, we visualize the success rates in dependence on the noise level. For `randomScale` $\in \{0.3, 0.5\}$ and both algorithms, the success rates start at less than 0.12 for $\text{SNR}_{\text{dB}} = 0$ and grow for increasing signal-to-noise ratios until at least 0.90 for $\text{SNR}_{\text{dB}} = 30$. The success rates of Algorithm 2 with $\mathcal{A}_{2s,\tilde{M}}^{\text{sub,MC}}$ ("2dim") are often higher than for Algorithm 1 with $\mathcal{A}_{2s,\tilde{M}}^{\text{sub,MC}}$ ("phase"), which may be caused by the larger numbers of samples for Algorithm 2 and the noise model used. Note that the numbers of samples correspond to those in Figure 1a for $s = 100$ independent of the noise level. For Algorithm 2 with `randomScale` = 0.3, the increase of the success rate seems to stagnate at $\text{SNR}_{\text{dB}} = 20$, while this does not seem to be the case for `randomScale` = 0.5 or Algorithm 1. In particular, this behavior can also be observed in Figure 2b, where we plot the average relative ℓ^2 error of the Fourier coefficients against the signal-to-noise ratio. Here, we observe that for `randomScale` = 0.3, the decrease of the errors for increasing SNR_{dB} values almost

stops once reaching $SNR_{dB} = 20$ for both algorithms. Initially, the average error of Algorithm 2 is smaller, but at $SNR_{dB} = 15$ and higher, the average error of Algorithm 1 is smaller. In case of `randomScale` $= 0.5$, we observe a distinct decrease for growing signal-to-noise ratios for both algorithms.

5.1.3 Deterministic frequency set within 10-dimensional hyperbolic cross and noisy samples

Next, instead of randomly chosen frequencies, we consider frequencies on a d-dimensional weighted hyperbolic cross

$$\mathcal{H}_N^{d,\alpha} := \left\{ \mathbf{k} \in \mathbb{Z}^d \colon \prod_{\ell=1}^d \max(1, \ell^\alpha |k_\ell|) \leq \frac{N}{2} \quad \text{and} \quad \max_{\ell=1}^d k_\ell < \frac{N}{2} \right\}.$$

Here, we use $d = 10$, $N = 33$, $\mathcal{I} := \mathcal{H}_{33}^{10}$, and $\alpha = 1.7$, which yields $s = |\mathcal{H}_{33}^{10,1.7}| = 101$. As before, the Fourier coefficients $c_{\mathbf{k}}$ are randomly chosen from $[-1,1] + \mathfrak{i}\,[-1,1]$, $|c_{\mathbf{k}}| \geq 10^{-3}$. We use the same lattice and bandwidth parameter as in the last subsection as well as the same noise model and parameters.
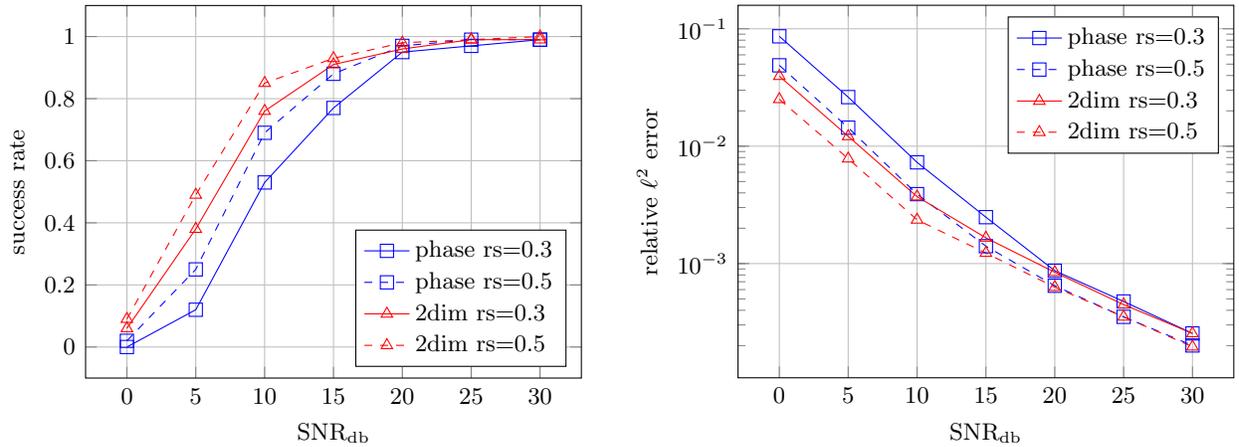
(a) success rates vs. noise level for $s = 100$ (b) relative ℓ^2 errors vs. noise level for $s = 100$

Figure 3: Average success rates (all frequencies detected) and relative ℓ^2 errors over 100 test runs for Algorithm 1 with $\mathcal{A}_{2s,\tilde{M}}^{\mathrm{sub,MC}}$, denoted by "phase", and Algorithm 2 with $\mathcal{A}_{2s,\tilde{M}}^{\mathrm{sub,MC}}$, denoted by "2dim", on multivariate trigonometric polynomials with (deterministic) frequencies on weighted hyperbolic cross within hyperbolic cross $\mathcal{I} := \mathcal{H}_{33}^{10}$, setting `randomScale` $:=$ rs $\in \{0.3, 0.5\}$ and bandwidth parameter $\tilde{M} = 1 + 2\|\mathbf{z}\|_\infty \max_{\mathbf{k} \in \mathcal{I}} \|\mathbf{k}\|_1 \approx 1.6 \cdot 10^{10}$.

In Figure 3, we depict the obtained results. In particular, the results in Figure 3a are very similar to the ones for randomly chosen frequencies in Figure 2a. For the case of deterministic frequencies in Figure 3a, the success rates are slightly better. Moreover, we do not observe the "stagnation" of the success rates for Algorithm 2 with `randomScale` $= 0.3$. Correspondingly, the relative ℓ^2 errors, as shown in Figure 3b, decrease distinctly for growing signal-to-noise ratios. Algorithm 2 performs slightly better than Algorithm 1, but also requires more than one order of magnitude more samples, similar to the results shown in Figure 1a for $s = 100$.

5.2 Compressible case in 10 dimensions

In this section, we apply the methods on a test function which is not exactly sparse but compressible. In addition, we also consider noisy samples as in Section 5.1.2. We use the 10-variate periodic test function $f : \mathbb{T}^{10} \to \mathbb{R}$,

$$f(\mathbf{x}) := \prod_{\ell \in \{0,2,7\}} N_2(x_\ell) + \prod_{\ell \in \{1,4,5,9\}} N_4(x_\ell) + \prod_{\ell \in \{3,6,8\}} N_6(x_\ell), \tag{11}$$

from [33, Section 3.3] and [21, Section 5.3] with infinitely many non-zero Fourier coefficients $c_{\mathbf{k}}(f)$, where $N_m : \mathbb{T} \to \mathbb{R}$ is the B-Spline of order $m \in \mathbb{N}$,

$$N_m(x) := C_m \sum_{k \in \mathbb{Z}} \mathrm{sinc}\left(\frac{\pi}{m} k\right)^m (-1)^k \, \mathrm{e}^{2\pi \mathrm{i} k x},$$

with a constant $C_m > 0$ such that $\|N_m\|_{L^2(\mathbb{T})} = 1$. We remark that each B-Spline N_m of order $m \in \mathbb{N}$ is a piece-wise polynomial of degree $m - 1$. We apply Algorithm 1 with $\mathcal{A}_{2s,\tilde{M}}^{\mathrm{sub,MC}}$ and use the sparsity parameters $s \in \{50, 100, 250, 500, 1000, 2000\}$, which corresponds to $2s \in \{100, 200, 500, 1000, 2000, 4000\}$ frequencies and Fourier coefficients for the output of Algorithm 1. We use the frequency set $\mathcal{I} := \mathcal{H}_{33}^{10}$ and $\mathtt{randomScale} := \mathrm{rs} \in \{0.05, 0.1\}$. Moreover, we work with the same rank-1 lattice as in Section 5.1.2.

The obtained basis index sets $\mathrm{supp}(\mathbf{b})$ should "consist of" the union of three lower dimensional manifolds, a three-dimensional hyperbolic cross in the dimensions $1, 3, 8$; a four-dimensional hyperbolic cross in the dimensions $2, 5, 6, 10$; and a three-dimensional hyperbolic cross in the dimensions $4, 7, 9$. All tests are performed 100 times and the relative L^2 approximation error

$$\frac{\|f - \sum_{\mathbf{k} \in \mathrm{supp}(\mathbf{b})} b_{\mathbf{k}} \, \mathrm{e}^{2\pi \mathrm{i} \mathbf{k} \cdot \circ}\|_{L^2}}{\|f\|_{L^2}} = \frac{\sqrt{\|f\|_{L^2}^2 - \sum_{\mathbf{k} \in \mathrm{supp}(\mathbf{b})} |c_{\mathbf{k}}(f)|^2 + \sum_{\mathbf{k} \in \mathrm{supp}(\mathbf{b})} |b_{\mathbf{k}} - c_{\mathbf{k}}(f)|^2}}{\|f\|_{L^2}}$$

is computed each time.

In Figure 4a, we visualize the average number of samples against the sparsity $2s$ of the approximation. We observe an almost linear increase with respect to $2s$. In Figure 4b, we show the average relative errors for $\mathtt{randomScale} \in \{0.05, 0.1\}$ in the noiseless case as well as $\mathtt{randomScale} = 0.1$ for $\mathrm{SNR}_{\mathrm{db}} \in \{10, 20, 30\}$. In general, for increasing sparsity, the errors become smaller. For $\mathtt{randomScale} = 0.05$ in the noiseless case and $\mathtt{randomScale} = 0.1$ with $\mathrm{SNR}_{\mathrm{db}} = 10$, the average error are similar and stay above $3 \cdot 10^{-2}$ even for sparsity $2s = 4000$. For higher signal-to-noise ratio, the error decreases further. For $\mathrm{SNR}_{\mathrm{db}} = 30$, the obtained average error is $6.1 \cdot 10^{-3}$ for $2s = 4000$, which is only approximately twice as high as the best possible error when using the $2s$ largest (by magnitude) Fourier coefficients $c_{\mathbf{k}}(f)$ with the restriction $\mathbf{k} \in \mathcal{I} := \mathcal{H}_{33}^{10}$. The latter is plotted in Figure 4b as dashed line without markers.

References

[1] S. Bittens and G. Plonka. Real sparse fast DCT for vectors with short support. *Linear Algebra and its Applications*, 582:359–390, 2019.

[2] S. Bittens and G. Plonka. Sparse fast DCT for vectors with one-block support. *Numerical Algorithms*, 82(2):663–697, 2019.

(a) samples vs. sparsity $2s$ (b) relative L^2 errors vs. sparsity $2s$
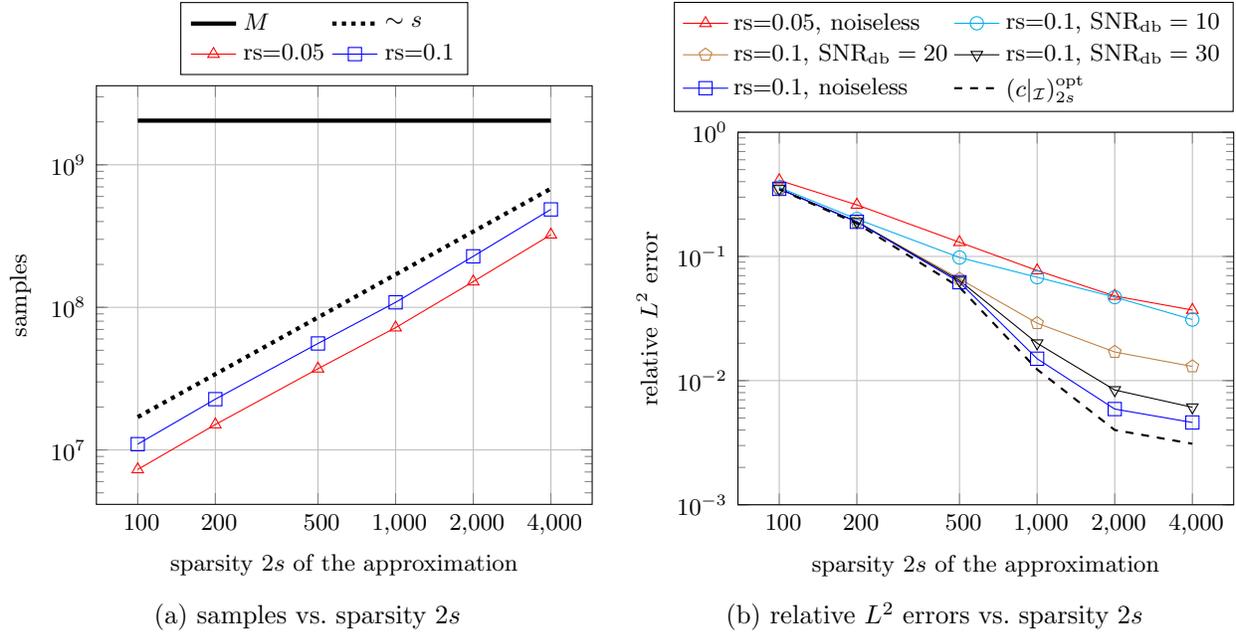
Figure 4: Average number of samples and relative L^2 errors over 100 test runs for Algorithm 1 with $\mathcal{A}^{\text{sub,MC}}_{2s,\tilde{M}}$ on 10-dimensional test function (11) consisting of tensor products of B-Splines of different order. Search space is unweighted hyperbolic cross $\mathcal{I} := \mathcal{H}^{10}_{33}$ with SFT parameters $\texttt{randomScale} := \text{rs} \in \{0.3, 0.5\}$ and $\tilde{M} = 1 + 2\|\mathbf{z}\|_\infty \max_{\mathbf{k} \in \mathcal{I}} \|\mathbf{k}\|_1 \approx 1.6 \cdot 10^{10}$.

[3] B. Choi, A. Christlieb, and Y. Wang. Multiscale High-Dimensional Sparse Fourier Algorithms for Noisy Data. *ArXiv e-prints*, 2019. arXiv:1907.03692.

[4] B. Choi, A. Christlieb, and Y. Wang. High-dimensional sparse Fourier algorithms. *Numerical Algorithms*, 2020.

[5] B. Choi, M. A. Iwen, and F. Krahmer. Sparse harmonic transforms: A new class of sublinear-time algorithms for learning functions of many variables. *Foundations of Computational Mathematics*, 2020.

[6] B. Choi, M. A. Iwen, and T. Volkmer. Sparse Harmonic Transforms II: Best s-Term Approximation Guarantees for Bounded Orthonormal Product Bases in Sublinear-Time. *ArXiv e-prints*, 2020. arXiv:1909.09564.

[7] A. Christlieb, D. Lawlor, and Y. Wang. A multiscale sub-linear time Fourier algorithm for noisy data. *Applied and Computational Harmonic Analysis*, 40(3):553 – 574, 2016.

[8] A. Cohen, W. Dahmen, and R. DeVore. Compressed sensing and best k-term approximation. *Journal of the American mathematical society*, 22(1):211–231, 2009.

[9] S. Foucart and H. Rauhut. *A mathematical introduction to compressive sensing*. Springer, 2013.

[10] A. C. Gilbert, P. Indyk, M. Iwen, and L. Schmidt. Recent developments in the sparse Fourier transform: A compressed Fourier transform for big data. *IEEE Signal Processing Magazine*, 31(5):91–100, 2014.

[11] A. C. Gilbert, S. Muthukrishnan, and M. Strauss. Improved time bounds for near-optimal sparse Fourier representations. In M. Papadakis, A. F. Laine, and M. A. Unser, editors, *Wavelets XI*, volume 5914, pages 398 – 412. International Society for Optics and Photonics, SPIE, 2005.

[12] A. C. Gilbert, M. J. Strauss, and J. A. Tropp. A tutorial on fast fourier sampling. *IEEE Signal Processing Magazine*, 25(2):57–66, 2008.

[13] C. Gross, M. A. Iwen, L. Kämmerer, and T. Volkmer. A deterministic algorithm for constructing multiple rank-1 lattices of near-optimal size. *ArXiv e-prints*, 2020. arXiv:2003.09753.

[14] H. Hassanieh, P. Indyk, D. Katabi, and E. Price. Simple and practical algorithm for sparse Fourier transform. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1183–1194. ACM, New York, 2012.

[15] M. A. Iwen. Combinatorial sublinear-time Fourier algorithms. *Foundations of Computational Mathematics*, 10(3):303–338, 2010.

[16] M. A. Iwen. Improved approximation guarantees for sublinear-time Fourier algorithms. *Applied and Computational Harmonic Analysis*, 34:57–82, 2013.

[17] L. Kämmerer. *High Dimensional Fast Fourier Transform Based on Rank-1 Lattice Sampling*. Dissertation. Universitätsverlag Chemnitz, 2014.

[18] L. Kämmerer. Reconstructing multivariate trigonometric polynomials from samples along rank-1 lattices. In G. E. Fasshauer and L. L. Schumaker, editors, *Approximation Theory XIV: San Antonio 2013*, pages 255–271. Springer International Publishing, 2014.

[19] L. Kämmerer, F. Krahmer, and T. Volkmer. A sample efficient sparse FFT for arbitrary frequency candidate sets in high dimensions. *ArXiv e-prints*, 2020. arXiv:2006.13053.

[20] L. Kämmerer, D. Potts, and T. Volkmer. Approximation of multivariate periodic functions by trigonometric polynomials based on rank-1 lattice sampling. *Journal of Complexity*, 31(4):543–576, 2015.

[21] L. Kämmerer, D. Potts, and T. Volkmer. High-dimensional sparse FFT based on sampling along multiple rank-1 lattices. *Applied and Computational Harmonic Analysis*, 51:225–257, 2021.

[22] M. Kapralov. *Sparse Fourier Transform in Any Constant Dimension with Nearly-Optimal Sample Complexity in Sublinear Time*, page 264–277. Association for Computing Machinery, New York, NY, USA, 2016.

[23] F. Y. Kuo, G. Migliorati, F. Nobile, and D. Nuyens. Function integration, reconstruction and approximation using rank-1 lattices. *ArXiv e-prints*, 2020. arXiv:1908.01178.

[24] F. Y. Kuo, I. H. Sloan, and H. Woźniakowski. Lattice rules for multivariate approximation in the worst case setting. In *Monte Carlo and quasi-Monte Carlo methods 2004*, pages 289–330. Springer, Berlin, 2006.

[25] D. Lawlor, Y. Wang, and A. Christlieb. Adaptive sub-linear time Fourier algorithms. *Advances in Adaptive Data Analysis*, 05(01):1350003, 2013.

[26] D. Li and F. J. Hickernell. Trigonometric spectral collocation methods on lattices. In *Recent advances in scientific computing and partial differential equations (Hong Kong, 2002)*, volume 330 of *Contemp. Math.*, pages 121–132. Amer. Math. Soc., Providence, RI, 2003.

[27] S. Merhi, R. Zhang, M. A. Iwen, and A. Christlieb. A new class of fully discrete sparse Fourier transforms: Faster stable implementations with guarantees. *Journal of Fourier Analysis and Applications*, 25(3):751–784, 2019.

[28] L. Morotti. Explicit universal sampling sets in finite vector spaces. *Applied and Computational Harmonic Analysis*, 43(2):354–369, 2017.

[29] H. Munthe-Kaas and T. Sørevik. Multidimensional pseudo-spectral methods on lattice grids. *Applied Numerical Mathematics. An IMACS Journal*, 62(3):155–165, 2012.

[30] G. Plonka, D. Potts, G. Steidl, and M. Tasche. *Numerical Fourier Analysis*. Springer, 2018.

[31] G. Plonka and K. Wannenwetsch. A sparse fast Fourier algorithm for real non-negative vectors. *Journal of Computational and Applied Mathematics*, 321:532–539, 2017.

[32] G. Plonka, K. Wannenwetsch, A. Cuyt, and W.-s. Lee. Deterministic sparse FFT for M-sparse vectors. *Numerical Algorithms*, 78(1):133–159, 2018.

[33] D. Potts and T. Volkmer. Sparse high-dimensional FFT based on rank-1 lattice sampling. *Applied and Computational Harmonic Analysis*, 41(3):713–748, 2016.

[34] B. Segal and M. Iwen. Improved sparse Fourier approximation results: faster implementations and stronger guarantees. *Numerical Algorithms*, 63(2):239–263, 2013.

[35] V. N. Temlyakov. Reconstruction of periodic functions of several variables from the values at the nodes of number-theoretic nets. *Analysis Mathematica*, 12(4):287–305, 1986.

[36] V. N. Temlyakov. *Approximation of periodic functions*. Computational Mathematics and Analysis Series. Nova Science Publishers, Inc., Commack, NY, 1993.

[37] T. Volkmer. *Multivariate Approximation and High-Dimensional Sparse FFT Based on Rank-1 Lattice Sampling*. Dissertation. Universitätsverlag Chemnitz, 2017.