

A Robust Sparse Fourier Transform in the Continuous Setting

Eric Price and Zhao Song
ecprice@cs.utexas.edu zhaos@utexas.edu
The University of Texas at Austin

January 17, 2017

Abstract

In recent years, a number of works have studied methods for computing the Fourier transform in sublinear time if the output is sparse. Most of these have focused on the discrete setting, even though in many applications the input signal is continuous and naive discretization significantly worsens the sparsity level.

We present an algorithm for robustly computing sparse Fourier transforms in the continuous setting. Let $x(t) = x^*(t) + g(t)$, where x^* has a k -sparse Fourier transform and g is an arbitrary noise term. Given sample access to $x(t)$ for some duration T , we show how to find a k -Fourier-sparse reconstruction $x'(t)$ with

$$\frac{1}{T} \int_0^T |x'(t) - x(t)|^2 dt \lesssim \frac{1}{T} \int_0^T |g(t)|^2 dt.$$

The sample complexity is linear in k and logarithmic in the signal-to-noise ratio and the frequency resolution. Previous results with similar sample complexities could not tolerate an infinitesimal amount of i.i.d. Gaussian noise, and even algorithms with higher sample complexities increased the noise by a polynomial factor. We also give new results for how precisely the individual frequencies of x^* can be recovered.

1 Introduction

The Fourier transform is ubiquitous in digital signal processing of a diverse set of signals, including sound, image, and video. Much of this is enabled by the Fast Fourier Transform (FFT) [CT65], which computes the n -point discrete Fourier transform in $O(n \log n)$ time. But can we do better?

In many situations, much of the reason for using Fourier transforms is because the transformed signal is *sparse*—i.e., the energy is concentrated in a small set of k locations. In such situations, one could hope for a dependency that depends nearly linearly on k rather than n . Moreover, one may be able to find these frequencies while only sampling the signal for some period of time. This idea has led to a number of results on *sparse Fourier transforms*, including [GGI⁺02, GMS05, HIKP12a, IK14], that can achieve $O(k \log(n/k) \log n)$ running time and $O(k \log(n/k))$ sample complexity (although not quite both at the same time) in a robust setting.

These works apply to the discrete Fourier transform, but lots of signals including audio or radio originally come from a continuous domain. The standard way to convert a continuous Fourier transform into a discrete one is to apply a window function then subsample. Unfortunately, doing so “smears out” the frequencies, blowing up the sparsity. Thus, one can hope for significant efficiency gains by directly solving the sparse Fourier transform problem in the continuous setting. This has led researchers to adapt techniques from the discrete setting to the continuous both in theory [BCG⁺14, TBSR13, CF14, DB13] and in practice [SAH⁺13]. However, these results are not robust to noise: if the signal is sampled with a tiny amount of Gaussian noise or decays very slightly over time, no method has been known for computing a sparse Fourier transform in the continuous setting with sample complexity linear in k and logarithmic in other factors. That is what we present in this paper.

Formally, a vector $x^*(t)$ has a k -sparse Fourier transform if it can be written as

$$x^*(t) = \sum_{i=1}^k v_i e^{2\pi i f_i t}$$

for some *tones* $\{(v_i, f_i)\}$. We consider the problem where we can sample some signal

$$x(t) = x^*(t) + g(t)$$

at any t we choose in some interval $[0, T]$, where $x^*(t)$ has a k -sparse Fourier transform and $g(t)$ is arbitrary noise. As long as g is “small enough,” one would like to recover a good approximation to x (or to x^* , or to $\{(v_i, f_i)\}$) using relatively few samples $t \in [0, T]$ and fast running time. Our algorithm achieves several results of this form, but a simple one is an ℓ_2/ℓ_2 guarantee: we reconstruct an $x'(t)$ with k -sparse Fourier transform such that

$$\frac{1}{T} \int_0^T |x'(t) - x(t)|^2 dt \lesssim \frac{1}{T} \int_0^T |g(t)|^2 dt$$

using a number of samples that is k times logarithmic factors¹. To the best of our knowledge, this is the first algorithm achieving such a constant factor approximation with a sample complexity sublinear in T and the signal-to-noise ratio.

Our algorithm also gives fairly precise estimates of the individual tones (v_i, f_i) of the signal x^* . To demonstrate what factors are important, it is helpful to think about a concrete setting. Let us consider sound from a simplified model of a piano.

¹We use $f \lesssim g$ to denote that $f \leq Cg$ for some universal constant C .

Thought experiment: piano tuning In a simplified model of a piano, we have keys corresponding to frequencies over some range $[-F, F]$. The noise $g(t)$ comes from ambient noise and the signals not being pure tones (because, for example, the notes might decay slowly over time). For concrete numbers, a modern piano has 88 keys spaced from about 27.5 Hz to $F = 4200\text{Hz}$. The space between keys ranges from a few Hz to a few hundred Hz, but most chords will have an $\eta = 30\text{Hz}$ or more gap between the frequencies being played. One typically would like to tune the keys to within about $\pm\nu = 1\text{Hz}$. And piano music typically has k around 5.

Now, suppose you would like to build a piano tuner that can listen to a chord and tell you what notes are played and how they are tuned. For such a system, how long must we wait for the tuner to identify the frequencies? How many samples must the tuner take? And how robust is it to the noise?

If you have a constant signal-to-noise ratio, you need to sample for a time T of at least order $1/\nu = 1$ second in order to get 1Hz precision—frequencies within 1Hz of each other will behave very similarly over small fractions of a second, which noise can make indistinguishable. You also need at least $\Omega(k \log \frac{F}{k\nu}) \approx 50$ samples, because the support of the signal contains that many bits of information and you only get a constant number per measurement (at constant SNR). At higher signal-to-noise ratios ρ , these results extend to $\Omega(\frac{1}{\nu\rho})$ duration and $\Omega(k \log_\rho \frac{F}{k\nu})$ samples. But as the signal-to-noise ratio gets very high, there is another constraint on the duration: for $T < \frac{1}{\eta} \approx 33$ milliseconds the different frequencies start becoming hard to distinguish, which causes the robustness to degrade exponentially in k [Moi15] (though the lower bound there only directly applies to a somewhat restricted version of our setting).

This suggests the form of a result: with a duration $T > \frac{1}{\eta}$, one can hope to recover the frequencies to within $\frac{1}{\rho T}$ using $O(k \log_\rho \frac{FT}{k})$ samples. We give an algorithm that is within logarithmic factors of this ideal: with a duration $T > \frac{O(\log(k/\delta))}{\eta}$, we recover the frequencies to within $O(\frac{1}{\rho T})$ using $O(k \log_\rho(FT) \cdot \log(k/\delta) \log k)$ samples, where ρ and $1/\delta$ are (roughly speaking) the minimum and maximum signal-to-noise ratios that you can tolerate, respectively.

Instead of trying to tune the piano by recovering the frequencies precisely, one may simply wish to record the sound for future playback with relatively few samples. Our algorithm works for this as well: the combination $x'(t)$ of our recovered frequencies satisfies

$$\frac{1}{T} \int_0^T |x'(t) - x(t)|^2 dt \lesssim \frac{1}{T} \int_0^T |g(t)|^2 dt.$$

Let us now state our main theorems. The first shows how well we can estimate the frequencies f_i and their weights v_i ; we refer to this (v_i, f_i) pair as a *tone*.

Theorem 1.1 (Tone estimation). *Consider any signal $x(t) : [0, T] \rightarrow \mathbb{C}$ of the form*

$$x(t) = x^*(t) + g(t),$$

for arbitrary “noise” $g(t)$ and an exactly k -sparse $x^ = \sum_{i \in [k]} v_i e^{2\pi i f_i t}$ with frequencies $f_i \in [-F, F]$ and frequency separation $\eta = \min_{i \neq j} |f_i - f_j|$. For some parameter $\delta > 0$, define the “noise level”*

$$\mathcal{N}^2 := \frac{1}{T} \int_0^T |g(t)|^2 dt + \delta \sum_{i=1}^k |v_i|^2.$$

We give an algorithm that takes samples from $x(t)$ over any duration $T > O(\frac{\log(k/\delta)}{\eta})$ and returns a set of k tones $\{(v'_i, f'_i)\}$ that approximates x^ with error proportional to \mathcal{N} . In particular, every*

large tone is recovered: for any v_i with $|v_i| \gtrsim \mathcal{N}$, we have for an appropriate permutation of the indices that

$$|f'_i - f_i| \lesssim \frac{\mathcal{N}}{T|v_i|} \quad \text{and} \quad |v'_i - v_i| \lesssim \mathcal{N}. \quad (1)$$

In fact, we satisfy a stronger guarantee that the total error is bounded:

$$\sum_{i=1}^k \frac{1}{T} \int_0^T |v'_i e^{2\pi i f'_i t} - v_i e^{2\pi i f_i t}|^2 dt \lesssim \mathcal{N}^2. \quad (2)$$

The algorithm takes $O(k \log(FT) \log(\frac{k}{\delta}) \log(k))$ samples and $O(k \log(FT) \log(\frac{FT}{\delta}) \log(k))$ running time, and succeeds with probability at least $1 - 1/k^c$ for an arbitrarily large constant c .

We then show that the above approximation of the individual tones is good enough to estimate the overall signal $x(t)$ to within constant factors:

Theorem 1.2 (Signal estimation). *In the same setting as Theorem 1.1, if the duration is slightly longer at $T > O(\frac{\log(1/\delta) + \log^2 k}{\eta})$, the reconstructed signal $x'(t) = \sum_{i=1}^k v'_i e^{2\pi i f'_i t}$ achieves a constant factor approximation to the complete signal x :*

$$\frac{1}{T} \int_0^T |x'(t) - x(t)|^2 dt \lesssim \mathcal{N}^2. \quad (3)$$

The algorithm takes $O(k \log(FT) \log(\frac{k}{\delta}) \log(k))$ samples and $O(k \log(FT) \log(\frac{FT}{\delta}) \log(k))$ running time, and succeeds with probability at least $1 - 1/k^c$ for an arbitrarily large constant c .

The above theorems give three different error guarantees, which are all in terms of a “noise level” \mathcal{N}^2 that is the variance of the noise $g(t)$ plus δ times the energy of the signal. The algorithm depends logarithmically on δ , so one should think of \mathcal{N}^2 as being the variance of the noise, e.g. σ^2 if samples have error $N(0, \sigma^2)$.

Error guarantees. Our algorithm does a good job of estimating the signal, but how exactly should we quantify this? Because very few previous results have shown robust recovery in the continuous setting, there is no standard error measure to use. We therefore bound the error in three different ways: the maximum error in the estimation of any tone; the weighted total error in the estimation of all tones; and the difference between the reconstructed signal and the true signal over the sampled interval. The first measure has been studied before, while the other two are to the best of our knowledge new but useful to fully explain the robustness we achieve.

The error guarantee (1) says that we achieve good recovery of any tones with magnitude larger than $C\mathcal{N}$ for some constant C . Note that such a requirement is necessary: for tones with $|v_i| \leq \mathcal{N}$, one could have $g(t) = -v_i e^{2\pi i f_i t}$, completely removing the tone (v_i, f_i) from the observed $x(t)$ and making it impossible to find. For the tones of sufficiently large magnitude, we find them to within $\frac{\mathcal{N}}{T|v_i|}$. This is always less than $1/T$, and converges to 0 as the noise level decreases. This is known as *superresolution*—one can achieve very high frequency resolution in sparse, nearly noiseless settings. Moreover, by Lemma 3.15 our “superresolution” precision $|f'_i - f_i| \lesssim \frac{\mathcal{N}}{T|v_i|}$ is optimal.

While the guarantee of (1) is simple and optimal given its form, it is somewhat unsatisfying. It shows that the maximum error over all k tones is \mathcal{N} , while one can hope to bound the *total* error

over all k tones by \mathcal{N} . This is precisely what Equation (2) does. The guarantee (1) is the precision necessary to recover the tone to within $O(\mathcal{N})$ average error in time, that is (1) is equivalent to

$$\frac{1}{T} \int_0^T |v'_i e^{2\pi i f'_i t} - v_i e^{2\pi i f_i t}|^2 dt \lesssim \mathcal{N}^2 \quad \forall i \in [k].$$

In (2), we show that this bound holds even if we sum the left hand side over all $i \in [k]$, so the *average* error is a factor k better than would be implied by (1). It also means that the total mass of all tones that are not recovered to within $\pm 1/T$ is $O(\mathcal{N})$, not just that every tone larger than $O(\mathcal{N})$ is recovered to within $\pm 1/T$.

The stronger bound (2) can be converted to the guarantee (3), which is analogous to the ℓ_2/ℓ_2 recovery guarantee standard in compressive sensing. It bounds the error of our estimate in terms of the input noise, on average over the sampled duration. The standard form of the ℓ_2/ℓ_2 guarantee in the discrete sparse Fourier transform setting [GGI⁺02, GMS05, HIKP12a, IKP14, IK14] compares the energy in frequency domain rather than time domain. This cannot be achieved directly in the continuous setting, since frequencies infinitesimally close to each other are indistinguishable over a bounded time interval T . But if the signal is periodic with period T (the case where a discrete sparse Fourier transform applies directly), then (3) is equivalent to the standard guarantee by Parseval's identity. So (3) seems like the right generalization of ℓ_2/ℓ_2 recovery to our setting.

Other factors. Our algorithm succeeds with high probability in k , which could of course be amplified by repetition (but increasing the sample complexity and running time). Our running time, at $O(k \log(FT) \log(FT/\delta) \log k)$, is (after translating between the discrete and continuous setting) basically a $\log k$ factor larger than the fastest known algorithm for discrete sparse Fourier transforms ([HIKP12a]). But since that result only succeeds with constant probability, and repetition would also increase that by a $\log k$ factor, our running time is actually equivalent to the fastest known results that succeed with high probability in k .

Our sample complexity is $O(\log(k/\delta) \log k)$ worse than the presumptive optimal, which is known to be achievable in the discrete setting ([IK14]). However, the techniques that [IK14] uses to avoid the $\log(k/\delta)$ factor seem hard to adapt to the continuous setting without losing some of the robustness/precision guarantees.

Another useful property of our method, not explicitly stated in the theorem, is that the sampling method is relatively simple: it chooses $O(\log(FT) \log k)$ different random arithmetic sequences of points, where each sequence has $O(k \log(k/\delta))$ points spread out over a constant fraction of the time domain T . Thus one could implement this in hardware using a relatively small number of samplers, each of which performs regular sampling at a rate of $\frac{k \log(k/\delta)}{T}$. This is in contrast to the Nyquist rate for non-sparse signals of $2F$ — in the piano example, each sampler has a rate on the order of 50Hz rather than 8000Hz.

The superresolution frequency is optimal because two signals of magnitude v_i and frequency separation $\nu < 1/T$ will differ by $O(\nu^2 T^2 |v_i|^2)$ over the duration T , so for ν below our threshold the difference is just \mathcal{N}^2 . Hence if the observed signal $x(t)$ looks identical to (v_i, f_i) , it might actually be (v_i, f'_i) with noise equaling the difference between the two.

The only previous result known in the form of (1) was [Moi15], which lost a $\text{poly}(k, \frac{\max |v_i|}{\min |v_i|}, \eta)$ factor in noise tolerance and also did not optimize for sample complexity.

1.1 Comparison to Naive Methods

This section compares our result to some naive ways one could try to solve this problem by applying algorithms not designed for a continuous, sparse Fourier transform setting. The next section will

compare our result to algorithms that are designed for such a setting.

Nyquist Sampling. The traditional theory of band-limited signals from discrete samples says that, from samples taken at the Nyquist rate $2F$, one can reconstruct an F -band-limited signal exactly. The Whittaker-Shannon interpolation formula then says that

$$x^*(t) = \sum_{i=-\infty}^{\infty} x^*(2Fi) \operatorname{sinc}(2Ft - i) \quad (4)$$

where $\operatorname{sinc}(t)$ is the normalized sinc function $\frac{\sin(\pi t)}{\pi t}$. This is for the band-limited “pure” signal x^* , but one could then get a relationship for samples of the actual signal $x(t)$. This has no direct implications for learning the tones (e.g. our (1) or (2)), but for learning the signal (our (3)) there is also an issue. Even in the absence of noise and for $k = 1$, this method will have error polynomially rather than exponentially small in the number of samples.

That is, if there is no noise the method has zero error given infinitely many samples. But we only receive samples over the interval $[0, T]$, leading to error. Consider the trivial setting of $x(t) = 1$. The partial sum of (4) at a given t will be missing terms for $i > 2Ft$ and $i < 0$, which (for a random t in $[0, T/2]$) have magnitude at most $1/(Ft)$. The terms alternate in sign, so the sum has error approximately $1/(Ft)^2$. This means that the error over the first $1/F$ time is a constant, leading to average error of $\frac{1}{FT}$. This is with an algorithm that uses FT samples and time. By contrast, our algorithm in the noiseless setting has error exponentially small in the samples and running time.

Discrete Sparse Fourier Transforms. An option related to the previous would be to discretize very finely, then apply a discrete sparse Fourier transform algorithm to keep the sample complexity and runtime small. The trouble here is that sparse Fourier transforms require sparsity, and this process decreases the sparsity. In particular, this process supposes that the signal is periodic with period T , so one can analyze this process as first converting the signal to one, equivalent over $[0, T]$, but with frequency spectrum only containing integer multiples of $1/T$. This is done by convolving each frequency f_i with a sinc function (corresponding to windowing to $[0, T]$) then restricting to multiples of $1/T$ (corresponding to aliasing). The result is that a one-sparse signal $e^{2\pi i f_i t}$ is viewed as having Fourier spectrum

$$\widehat{x''}[j] = \operatorname{sinc}(f_i T - j)$$

for $j \in \mathbb{Z}$. When f_i is not a multiple of $1/T$, this means the signal is not a perfectly sparse signal. And this is true regardless of the discretization level, which only affects the subsequent error from aliasing $j \in \mathbb{Z}$ down to \mathbb{Z}_n . To have error proportional to $\delta \|\widehat{x}^*\|_2$, one would need to run such methods for a sparsity level of k/δ . Thus, as with Nyquist sampling, the sample and runtime will be polynomial, rather than logarithmic, in δ .

The above discussion refers to methods for learning the signal (our (3)). In terms of learning the tones, one could run the algorithm for sparsity $O(k)$ so that δ is a small constant, which would let one learn roughly where the peaks are and get most of the frequencies to the nearest $1/T$. This would give a similar bound to our (1), but without the superresolution effect as the noise becomes small. On the plus side, the duration could be just $O(1/\eta)$ —which is sufficient for the different peaks to be distinguishable—rather than $O(\frac{\log k}{\eta})$ as our method would require, and the time and sample complexities could save a $\log k$ factor (if one did not want to recover *all* the tones, just most of them).

Essentially, this algorithm tries to round each frequency to the nearest multiple of $1/T$, which introduces noise that is a constant fraction of the signal. If the signal-to-noise ratio is already low, this does not increase the noise level by that much so such an algorithm will work reasonably well. If the signal-to-noise ratio is fairly high, however, then the added noise leads to much worse performance. Getting a constant factor approximation to the whole signal is only nontrivial for high SNR, so such a method does badly in that setting. For approximating the tones, it is comparable to our method in the low SNR setting but does not improve much as SNR increases.

1.2 Previous Work In Similar Settings

There have been several works that recover continuous frequencies from samples in the time domain. Some of these are in our setting where the samples can be taken at arbitrary positions over $[0, T]$ and others are in the discrete-time (DTFT) setting where the samples must be taken at multiples of the Nyquist frequency $\frac{1}{2F}$.

The results of [TBSR13, CF14, YX15, DB13] show that a convex program can solve the problem in the DTFT setting using $O(k \log k \log(FT))$ samples if the duration is $T > O(\frac{1}{\eta})$, in the setting where $g(t) = 0$ and the coefficients of x^* have random phases. The sample complexity can be one log factor better than ours, which one would expect for the noiseless setting. None of these results show robustness to noise, and some additionally require a running time polynomial in FT rather than k .

The result of [BCG⁺14] is in a similar setting to our paper, using techniques of the same lineage. It achieves very similar sample complexity and running time to our algorithm, and a guarantee similar in spirit to (1) with some notion of robustness. However, the robustness is weaker than ours in significant ways. They consider the noise $g(t)$ in frequency space (i.e. $\hat{g}(f)$), then require that $\hat{g}(f)$ is zero at any frequency within η of the signal frequencies f_i , and bound the error in terms of $\mathcal{N}' = \|\hat{g}\|_1/k$ instead of $\|g\|_2$. This fails to cover simple examples of noise, including i.i.d. Gaussian noise $g(t) \sim N(0, \sigma^2)$ and the noise one would get from slow decay of the signal over time (e.g. $x(t) = x^*(t)e^{-\frac{t}{100T}}$). Both types of noise violate both assumptions on the noise: $\hat{g}(f)$ will be nonzero arbitrarily close to each f_i and $\|\hat{g}\|_1$ will be unbounded. Their result also requires a longer duration than our algorithm and has worse precision for any fixed duration.

The result of [Moi15] studies noise tolerance in the DTFT setting, ignoring sample complexity and running time. It shows that the matrix pencil method [HS90], using FT samples, achieves a guarantee of the form (1), except that the bounds are an additional $\text{poly}(FT, k, \delta)$ factor larger. Furthermore, it shows a sharp characterization of the minimal T for which this is possible by any algorithm: $T = (1 \pm o(1))\frac{2}{\eta}$ is necessary and sufficient. It is an interesting question whether the lower bound generalizes to our non-DTFT setting, where the samples are not necessarily taken from an even grid.

Lastly, [SAH⁺13] tries to apply sparse Fourier transforms to a domain with continuous signals. They first apply a discrete sparse Fourier transform then use hill-climbing to optimize their solution into a decent set of continuous frequencies. They have interesting empirical results but no theoretical ones.

2 Algorithm Overview

At a high level, our algorithm is an adaptation of the framework used by [HIKP12a] to the continuous setting. However, getting our result requires a number of subtle changes to the algorithm. This section will describe the most significant ones. We assume some familiarity with previous work in the area [GGI⁺02, CCF02, GMS05, GLPS12, HIKP12a].

First we describe a high-level overview of the structure. The algorithm proceeds in $\log k$ stages, where each stage attempts to recover each tone with a large constant probability (e.g. 9/10). In each stage, we choose a parameter $\sigma \approx \frac{T}{k \log(k/\delta)}$ that we think of as “hashing” the frequencies into random positions. For this σ , we will choose about $\log(FT)$ different random “start times” t_0 and sample an arithmetic sequence starting at t_0 , i.e. observe

$$x(t_0), x(t_0 + \sigma), x(t_0 + 2\sigma), \dots, x(t_0 + (k \log(k/\delta))\sigma)$$

We then scale these observations by a “window function,” which has specific properties but among other things scales down the values near the ends of the sequence, giving a smoother transition between the time before and after we start/end sampling. We alias this down to $B = O(k)$ terms (i.e. add together terms $1, B + 1, 2B + 1, \dots$ to get a B -dimensional vector) and take the B -dimensional DFT. This gives a set of B values \hat{u}_j . The observation made in previous papers is that \hat{u} is effectively a *hashing* of the tones of \hat{x} into B buckets, where σ defines a permutation on the frequencies that affects whether two different tones land in the same bucket, and \hat{u}_j approximately equals the sum of all the tones that land in bucket j , each scaled by a phase shift depending on t_0 .

Because of this phase shift, for each choice of t_0 the value of \hat{u}_j is effectively a sample from the Fourier transform of a signal that contains only the tones of \hat{x}^* that land in bucket j , with zeros elsewhere. And since there are k tones and $O(k)$ buckets, most tones are alone in their bucket. Therefore this sampling strategy reduces the original problem of k -sparse recovery to one of 1-sparse recovery—we simply choose t_0 according to some strategy that lets us achieve 1-sparse recovery, and recover a tone for each bin.

One-sparse recovery. The algorithm for one-sparse recovery in [HIKP12a] is a good choice for adaptation to the continuous setting. It narrows down to the frequency in a locality-aware way, maintaining an interval of frequencies that decreases in size at each stage (in contrast to the method in [GMS05], which starts from the least significant bit rather than most significant bit).

If a frequency is perturbed slightly in time (e.g., by multiplying by a very slow decay over time) this will blur the frequency slightly into a narrow band. The one-sparse recovery algorithm of [HIKP12a] will proceed normally until it gets to the narrow scale, at which point it will behave semi-arbitrarily and return something near that band. This gives a desired level of robustness—the error in the recovered frequency will be proportional to the perturbation.

Still, to achieve our result we need a few changes to the one-sparse algorithm. One is related to the duration T : in the very last stage of the algorithm, when the interval is stretched at the maximal amount, we can only afford one “fold” rather than the typical $O(\log n)$. The only cost to this is in failure probability, and doing it for one stage is fine—but showing this requires a different proof. Another difference is that we need the final interval to have precision $\frac{1}{T\rho}$ if the signal-to-noise ratio is ρ —the previous analysis showed $\frac{1}{T\sqrt{\rho}}$ and needed to be told ρ , but (as we shall see) to achieve an ℓ_2/ℓ_2 guarantee we need the optimal ρ -dependence and for the algorithm to be oblivious to the value of ρ . Doing so requires a modification to the algorithm and slightly more clever analysis.

k -sparse recovery. The changes to the k -sparse recovery structure are broader. First, to make the algorithm simpler we drop the [GLPS12]-style recursion with smaller k , and just repeat an $O(k)$ -size hashing $O(\log k)$ times. This loses a $\log k$ factor in time and sample complexity, but because of the other changes it is not easy to avoid, and at the same time improves our success probability.

The most significant changes come because we can no longer measure the noise in frequency space or rely on the hash function to randomize the energy that collides with a given heavy hitter. Because we only look at a bounded time window T , Parseval’s identity does not hold and the energy of the noise in frequency space may be unrelated to its observed energy. Moreover, if the noise consists of frequencies infinitesimally close to a true frequency, then because σ is bounded the true frequency will always hash to the same bin as the noise. These two issues are what drive the restrictions on noise in the previous work [BCG⁺14]—assuming the noise is bounded in ℓ_1 norm in frequency domain and is zero in a neighborhood of the true frequencies fixes both issues. But we want a guarantee in terms of the average ℓ_2 noise level \mathcal{N}^2 in time domain over the observed duration. If the noise level is \mathcal{N}^2 , because we cannot hash the noise independently of the signal, we can only hope to guarantee reliable recovery of tones with magnitude larger than \mathcal{N}^2 . This is in contrast to the \mathcal{N}^2/k that is possible in the discrete setting, and would naively lose a factor of k in the ℓ_2/ℓ_2 approximation.

The insight here is that, even though the noise is not distributed randomly across bins, the total amount of noise is still bounded. If a heavy hitter of magnitude v^2 is not recovered due to noise, that requires $\Omega(v^2)$ noise mass in the bin that is not in any other bin. Thus the total amount of signal mass not recovered due to noise is $O(\mathcal{N}^2)$, which allows for ℓ_2/ℓ_2 recovery.

This difference is why our algorithm only gets a constant factor approximation rather than the $1 + \epsilon$ guarantee that hashing techniques for sparse recovery can achieve in other settings. These techniques hash into $B = O(k/\epsilon)$ bins so the average noise per bin is $O(\frac{\epsilon}{k}\mathcal{N}^2)$. In our setting, where the noise is not hashed independently of the signal, this would give no benefit.

Another difference arises in the choice of the parameter σ , which is the separation between samples in the arithmetic sequence used for a single hashing, and gives the permutation during hashing. In the discrete setting, one chooses σ uniformly over n , which in our setting would correspond to a scale of $\sigma \approx \frac{1}{\eta}$. Since the arithmetic sequences have $O(k \log(k/\delta))$ samples, the duration would then become at least $\frac{k \log(k/\delta)}{\eta}$ (which is why [BCG⁺14] has this duration). What we observe is that σ can actually be chosen at the scale of $\frac{1}{k\eta}$, giving the desired $O(\frac{\log(k/\delta)}{\eta})$ duration. This causes frequencies at the minimum separation η to always land in bins that are a constant separation apart. This is sufficient because we use [HIKP12a]-style window functions with strong isolation properties (and, in fact, [HIKP12a] could have chosen $\sigma \approx n/B$); it would be an issue if we were using the window functions of [GMS05, IK14] that have smaller supports but less isolation.

Getting an ℓ_2 bound Lastly, converting the guarantee (2) into (3) is a nontrivial task that is trivial in the discrete setting. In the discrete setting, it follows immediately from the different frequencies being orthogonal to each other. In our setting, we use that the recovered frequencies should themselves have $\Omega(\eta)$ separation, and that well-separated frequencies are nearly orthogonal over long enough time scales $T \gg 1/\eta$.

This bears some similarity to issues that arise in sparse recovery with overcomplete dictionaries. It would be interesting to see whether further connections can be made between the problems.

3 Proof outline

In this section we present the key lemmas along the path to producing the algorithm. The full proof are presented in appendix.

Notation. First we define the notation necessary to understand the lemmas. The full notation as used in the proofs appears in Section A.

The algorithm proceeds in stages, each of which hashes the frequencies to B bins. The hash function depends on two parameters σ and b , and so we define it as $h_{\sigma,b}(f) : [-F, F] \rightarrow [B]$.

A tone with a given frequency f can have two “bad events” $E_{coll}(f)$ or $E_{off}(f)$ hold for a given hashing. These correspond to colliding with another frequency of x^* or landing within an α fraction of the edge, respectively; they each will occur with small constant probability.

For a given hashing, we will choose a number of different offsets a that let us perform recovery of the tones that have neither bad event in this stage.

We use $f \lesssim g$ to denote that there exists a constant C such that $f \leq Cg$, and $f \approx g$ to denote $f \lesssim g \lesssim f$.

Key Lemmas First, we need to be able to compare the distance between two pure tone signals in time domain to their differences in parameters. The relation is as follows:

Lemma 3.1. *Let (v, f) and (v', f') denote any two tones, i.e., (magnitude, frequency) pairs. Then for*

$$\text{dist}((v, f), (v', f'))^2 := \frac{1}{T} \int_0^T \left| v e^{2\pi f t i} - v' e^{2\pi f' t i} \right|^2 dt,$$

we have

$$\text{dist}((v, f), (v', f'))^2 \approx (|v|^2 + |v'|^2) \cdot \min(1, T^2 |f - f'|^2) + |v - v'|^2,$$

and

$$\text{dist}((v, f), (v', f')) \approx |v| \cdot \min(1, T |f - f'|) + |v - v'|.$$

The basic building block for our algorithm is a function `HashToBins`, which is very similar to one of the same name in [HIKP12a].

The key property of `HashToBins` is that, if neither “bad” event holds for a frequency f (i.e. it does not collide or land near the boundary of the bin), then for the bin $j = h_{\sigma,b}(f)$ we have that $|\hat{u}_j| \approx |\hat{x}^*(f)|$ with a phase depending on a .

How good is the approximation? In the discrete setting, one can show that each tone has error about \mathcal{N}^2/B in expectation. Here, because the hash function cannot randomize the noise, we instead show that the total error over all tones is about \mathcal{N}^2 :

Lemma 3.2. *Let $\sigma \in [\frac{1}{B\eta}, \frac{2}{B\eta}]$ uniformly at random, then $b \in [0, \frac{[F/\eta]}{\sigma B}]$, $a \in [0, \frac{cT}{\sigma}]$ be sampled uniformly at random for some constant $c > 0$. Let the other parameters be arbitrary in $\hat{u} = \text{HashToBins}(x, P_{\sigma,a,b}, B, \delta, \alpha)$, and consider $H = \{f \in \text{supp}(\hat{x}^*) \mid \text{neither } E_{coll}(f) \text{ nor } E_{off}(f) \text{ holds}\}$ and $I = [B] \setminus h_{\sigma,b}(\text{supp}(\hat{x}^*))$ to be the bins that have no frequencies hashed to them. Then*

$$\mathbb{E}_{\sigma,b,a} \left[\sum_{f \in H} \left| \hat{u}_{h_{\sigma,b}(f)} - \hat{x}^*(f) e^{a\sigma 2\pi f i} \right|^2 + \sum_{j \in I} \hat{u}_j^2 \right] \lesssim \mathcal{N}^2$$

We prove Lemma 3.2 by considering the cases of $x^* = 0$ and $g = 0$ separately; linearity then gives the result. Both follow from properties of our window functions.

Lemma 3.3. *If $x^*(t) = 0, \forall t \in [0, T]$, then*

$$\mathbb{E}_{\sigma,a,b} \left[\sum_{j=1}^B |\hat{u}_j|^2 \right] \lesssim \frac{1}{T} \int_0^T |g(t)|^2 dt.$$

Lemma 3.4. *If $g(t) = 0, \forall t \in [0, T]$. Let H denote a set of frequencies, $H = \{f \in \text{supp}(\widehat{x}^*) \mid \text{neither } E_{\text{coll}}(f) \text{ nor } E_{\text{off}}(f) \text{ holds}\}$. Then,*

$$\mathbb{E}_{\sigma, a, b} \left[\sum_{f \in H} \left| \widehat{u}_{h_{\sigma, b}(f)} - \widehat{x}^*(f) e^{a\sigma 2\pi f i} \right|^2 \right] \leq \delta \|\widehat{x}^*\|_1^2.$$

Lemma 3.2 is essentially what we need for 1-sparse recovery. We first show a lemma about the inner call, which narrows the frequency from a range of size Δl to one of size $\frac{\Delta l}{\rho s t}$ for some parameters $\rho s t$. This gives improved performance (superresolution) when the signal-to-noise ratio ρ within the bucket is high. The parameter s and t provide a tradeoff between success probability, performance, running time, and duration.

Lemma 3.5. *Consider any B, δ, α . Algorithm HashToBins takes $O(B \log(k/\delta))$ samples and runs in $O(\frac{B}{\alpha} \log(k/\delta) + B \log B)$ time.*

Lemma 3.6. *Given σ and b , consider any frequency f for which neither $E_{\text{coll}}(f)$ nor $E_{\text{off}}(f)$ holds, and let $j = h_{\sigma, b}(f)$. Let $\mu^2(f) = \mathbb{E}_a[|\widehat{u}_j - \widehat{x}^*(f) e^{a\sigma 2\pi f i}|^2]$ and $\rho^2 = |\widehat{x}^*(f)|^2 / \mu^2(f)$. For sufficiently large ρ , and $\forall 0 < s < 1, t \geq 4$, consider any run of LocateInner with $f \in [l_j - \frac{\Delta l}{2}, l_j + \frac{\Delta l}{2}]$. It takes $O(R_{\text{loc}})$ random $(\gamma, \beta) \in [\frac{1}{2}, 1] \times [\frac{st}{4\sigma\Delta l}, \frac{st}{2\sigma\Delta l}]$ samples over duration $\beta\sigma = \Theta(\frac{st}{\Delta l})$, runs in $O(stR_{\text{loc}})$ time, to learn f within a region that has length $\Theta(\frac{\Delta l}{t})$ with failure probability at most $(\frac{4}{s\rho})^{R_{\text{loc}}} + t \cdot (60s)^{R_{\text{loc}}/2}$.*

By repeating this inner loop, we can recover the tones in almost every bin that does not have the “bad” events happen, so we recover a large fraction of the heavy hitters in each stage.

Lemma 3.7. *Algorithm LocateKSignal takes $O(k \log_C(FT) \log(k/\delta))$ samples over $O(\frac{\log(k/\delta)}{\eta})$ duration, runs in $O(k \log_C(FT) \log(FT/\delta))$ time, and outputs a set $L \subset [-F, F]$ of $O(k)$ frequencies with minimum separation $\Omega(\eta)$.*

Given σ and b , consider any frequency f for which neither of $E_{\text{coll}}(f)$ or $E_{\text{off}}(f)$ hold. Let $j = h_{\sigma, b}(f)$, $\mu^2(f) = \mathbb{E}_a[|\widehat{u}_j - \widehat{x}^(f) e^{a\sigma 2\pi f i}|^2]$, and $\rho^2 = |\widehat{x}^*(f)|^2 / \mu^2(f)$. If $\rho > C$, then with an arbitrarily large constant probability there exists an $f' \in L$ with*

$$|f - f'| \lesssim \frac{1}{T\rho}.$$

Combining this with estimation of the magnitudes of recovered frequencies, we can show that the total error over all bins without “bad” events—that is, bins with either one well placed frequency or zero frequencies—is small. At this point we give no guarantee for the (relatively few) bins with bad events; the recovered values there may be arbitrarily large.

Lemma 3.8. *Algorithm OneStage takes $O(k \log_C(FT) \log(k/\delta))$ samples over $O(\frac{\log(k/\delta)}{\eta})$ duration, runs in $O(k(\log_C(FT) \log(FT/\delta)))$ time, and outputs a set of $\{(v'_i, f'_i)\}$ of size $O(k)$ with $\min_{i \neq j} |f'_i - f'_j| \gtrsim \eta$. Moreover, one can imagine a subset $S \subseteq [k]$ of “successful” recoveries, where $\Pr[i \in S] \geq \frac{9}{10} \forall i \in [k]$ and for which there exists an injective function $\pi : [k] \rightarrow [O(k)]$ so that*

$$\mathbb{E}_{\sigma, b} \left[\sum_{i \in S} \frac{1}{T} \int_0^T \left| v'_i e^{2\pi f'_i t i} - v_{\pi(i)} e^{2\pi f_{\pi(i)} t i} \right|^2 dt \right] \lesssim C^2 \mathcal{N}^2.$$

with $1 - 1/k^c$ probability for an arbitrarily large constant c .

We can repeat the procedure for $O(\log k)$ stages and merge the results, getting a list of $O(k)$ tones that includes k tones that match up well to the true tones. However, we give no guarantee for the rest of the recovered tones at this point—as far as the analysis is concerned, mistakes from bins with collisions may cause arbitrarily large spurious tones.

Lemma 3.9. *Repeating algorithm OneStage $O(\log k)$ times, MergedStages returns a set $\{(v'_i, f'_i)\}$ of size $O(k)$ with $\min_{i \neq j} |f'_i - f'_j| \gtrsim \eta$ that can be indexed by π such that*

$$\sum_{i=1}^k \frac{1}{T} \int_0^T \left| v'_i e^{2\pi f'_i t i} - v_{\pi(i)} e^{2\pi f_{\pi(i)} t i} \right|^2 dt \lesssim C^2 \mathcal{N}^2.$$

with probability $1 - 1/k^c$ for an arbitrarily large constant c .

To address the issue of spurious tones, we run the above algorithm twice and only take the tones that are recovered in both stages. We show that the resulting $O(k)$ tones are together a good approximation to the vector.

Lemma 3.10. *If we run MergedStages twice and take the tones $\{(v'_i, f'_i)\}$ from the first result that have f'_i within $c\eta$ for small c of some frequency in the second result, we get a set of $k'' = O(k)$ tones that can be indexed by some permutation π such that*

$$\sum_{i=1}^k \frac{1}{T} \int_0^T \left| v'_i e^{2\pi f'_i t i} - v_{\pi(i)} e^{2\pi f_{\pi(i)} t i} \right|^2 dt + \sum_{i=k+1}^{k''} |v'_i|^2 \lesssim C^2 \mathcal{N}^2. \quad (5)$$

Simply picking out the largest k recovered tones then gives the result (2).

Theorem 3.11. *Algorithm ContinuousFourierSparseRecovery returns a set $\{(v'_i, f'_i)\}$ of size k with $\min_{i \neq j} |f'_i - f'_j| \gtrsim \eta$ for which*

$$\sum_{i=1}^k \frac{1}{T} \int_0^T \left| v'_i e^{2\pi f'_i t i} - v_{\pi(i)} e^{2\pi f_{\pi(i)} t i} \right|^2 dt \lesssim C^2 \mathcal{N}^2$$

with probability $1 - 1/k^c$ for an arbitrarily large constant c .

By only considering the term in the sum corresponding to tone i and applying Lemma 3.1, we get result (1):

Corollary 3.12. *With probability $1 - 1/k^c$ for an arbitrarily large constant c , we recover a set of tones $\{(v'_i, f'_i)\}$ such that, for any v_i with $|v_i| \gtrsim \mathcal{N}$, we have for an appropriate permutation of the indices that*

$$|f'_i - f_i| \lesssim \frac{\mathcal{N}}{T|v_i|} \quad \text{and} \quad |v'_i - v_i| \lesssim \mathcal{N}. \quad (6)$$

We then show that (2) implies (3) for sufficiently long durations T . A long duration helps because it decreases the correlation between η -separated frequencies.

Lemma 3.13. *Let $\{(v_i, f_i)\}$ and $\{(v'_i, f'_i)\}$ be two sets of k tones for which $\min_{i \neq j} |f_i - f_j| \geq \eta$ and $\min_{i \neq j} |f'_i - f'_j| \gtrsim \eta$ for some $\eta > 0$. Suppose that $T > O(\frac{\log^2 k}{\eta})$. Then these sets can be indexed such that*

$$\frac{1}{T} \int_0^T \left| \sum_{i=1}^k (v'_i e^{2\pi i f'_i t} - v_i e^{2\pi i f_i t}) \right|^2 dt \lesssim \sum_{i=1}^k \frac{1}{T} \int_0^T |v'_i e^{2\pi i f'_i t} - v_i e^{2\pi i f_i t}|^2 dt. \quad (7)$$

Combining Theorem 3.11 and Lemma 3.13 immediately implies

Theorem 3.14. *Suppose we sample for a duration $T > O(\frac{\log(1/\delta) + \log^2 k}{\eta})$. Then the reconstructed signal $x'(t) = \sum_{i=1}^k v_i' e^{2\pi i f_i' t}$ achieves a constant factor approximation to the complete signal x :*

$$\frac{1}{T} \int_0^T |x'(t) - x(t)|^2 dt \lesssim C^2 \mathcal{N}^2. \quad (8)$$

The algorithm takes $O(k \log \frac{F}{\eta} \log(\frac{k}{\delta}) \log(k))$ samples, runs in $O(k \log \frac{F}{\eta} \log(\frac{FT}{\delta}) \log(k))$ time, and succeeds with probability at least $1 - 1/k^c$ for an arbitrarily large constant c .

That finishes the proof of our main theorem. We also show that our “superresolution” precision from (1) is optimal, which is a simple corollary of Lemma 3.1.

Lemma 3.15. *There exists a constant $c > 0$ such that, for a given sample duration T , one cannot recover the frequency f to within*

$$c \frac{\mathcal{N}}{T |\widehat{x}^*(f)|}$$

with 3/4 probability, for all $\delta > 0$, even if $k = 1$.

References

- [BCG⁺14] Petros Boufounos, Volkan Cevher, Anna C Gilbert, Yi Li, and Martin J Strauss. What’s the frequency, Kenneth?: Sublinear Fourier sampling off the grid. In *Algorithmica(A preliminary version of this paper appeared in the Proceedings of RANDOM/APPROX 2012, LNCS 7408, pp. 61-72)*, pages 1–28. Springer, 2014.
- [CCF02] Moses Charikar, Kevin Chen, and Martin Farach-Colton. Finding frequent items in data streams. In *Automata, Languages and Programming*, pages 693–703. Springer, 2002.
- [CF14] Emmanuel J Candès and Carlos Fernandez-Granda. Towards a mathematical theory of super-resolution. *Communications on Pure and Applied Mathematics*, 67(6):906–956, 2014.
- [CT65] James W Cooley and John W Tukey. An algorithm for the machine calculation of complex Fourier series. *Mathematics of computation*, 19(90):297–301, 1965.
- [DB13] Marco F Duarte and Richard G Baraniuk. Spectral compressive sensing. *Applied and Computational Harmonic Analysis*, 35(1):111–129, 2013.
- [GGI⁺02] Anna C Gilbert, Sudipto Guha, Piotr Indyk, S Muthukrishnan, and Martin Strauss. Near-optimal sparse Fourier representations via sampling. In *Proceedings of the thirty-fourth annual ACM symposium on Theory of computing*, pages 152–161. ACM, 2002.
- [GLPS12] Anna C Gilbert, Yi Li, Ely Porat, and Martin J Strauss. Approximate sparse recovery: optimizing time and measurements. *SIAM Journal on Computing*, 41(2):436–453, 2012.
- [GMS05] Anna C Gilbert, S Muthukrishnan, and Martin Strauss. Improved time bounds for near-optimal sparse Fourier representations. In *Optics & Photonics 2005*, pages 59141A–59141A. International Society for Optics and Photonics, 2005.
- [HIKP12a] Haitham Hassanieh, Piotr Indyk, Dina Katabi, and Eric Price. Nearly optimal sparse Fourier transform. In *Proceedings of the forty-fourth annual ACM symposium on Theory of computing*. ACM, 2012.
- [HIKP12b] Haitham Hassanieh, Piotr Indyk, Dina Katabi, and Eric 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. SIAM, 2012.
- [HS90] Yingbo Hua and Tapan K Sarkar. Matrix pencil method for estimating parameters of exponentially damped/undamped sinusoids in noise. *Acoustics, Speech and Signal Processing, IEEE Transactions on*, 38(5):814–824, 1990.
- [IK14] Piotr Indyk and Michael Kapralov. Sample-optimal Fourier sampling in any constant dimension. In *Foundations of Computer Science (FOCS), 2014 IEEE 55th Annual Symposium on*, pages 514–523. IEEE, 2014.
- [IKP14] Piotr Indyk, Michael Kapralov, and Eric Price. (Nearly) Sample-optimal sparse Fourier transform. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 480–499. SIAM, 2014.

- [Moi15] Ankur Moitra. The threshold for super-resolution via extremal functions. In *STOC*, 2015.
- [SAH⁺13] Lixin Shi, O Andronesi, Haitham Hassanieh, Badih Ghazi, Dina Katabi, and Elfar Adalsteinsson. Mrs sparse-fft: Reducing acquisition time and artifacts for in vivo 2d correlation spectroscopy. In *ISMRM13, Int. Society for Magnetic Resonance in Medicine Annual Meeting and Exhibition*, 2013.
- [TBSR13] Gongguo Tang, Badri Narayan Bhaskar, Parikshit Shah, and Benjamin Recht. Compressed sensing off the grid. *Information Theory, IEEE Transactions on*, 59(11):7465–7490, 2013.
- [YX15] Zai Yang and Lihua Xie. Achieving high resolution for super-resolution via reweighted atomic norm minimization. In *Acoustics, Speech and Signal Processing (ICASSP), 2015 IEEE International Conference on*, pages 3646–3650. IEEE, 2015.

A Notation and Definitions: Permutation, Hashing, Filters

This section gives definitions about the permutation, hashing, and filters that are used throughout the proofs. Let $[n]$ denote the set $\{1, 2, \dots, n-1, n\}$. \mathbb{R} denotes the real numbers, \mathbb{Z} denotes the integer numbers and \mathbb{C} denotes the complex numbers. The convolution of two continuous functions f and g is written as $f * g$,

$$(f * g)(t) := \int_{-\infty}^{+\infty} f(\tau)g(t - \tau)d\tau$$

and the discrete convolution of f and g is given by,

$$(f * g)[n] := \sum_{m=-\infty}^{+\infty} f[m]g[n - m]$$

Let \mathbf{i} denote $\sqrt{-1}$, and $e^{i\theta} = \cos(\theta) + \mathbf{i}\sin(\theta)$. For any complex number $z \in \mathbb{C}$, we have $z = a + \mathbf{i}b$, where $a, b \in \mathbb{R}$. Define $\bar{z} = a - \mathbf{i}b$, $|z|^2 = z\bar{z} = a^2 + b^2$ and let $\phi(z)$ be the phase of z . Let $\text{supp}(f)$ denote the support of function/vector f , and $\|f\|_0 = |\text{supp}(f)|$. For any $p \in [1, \infty]$, the ℓ_p norm of a vector of x is $\|x\|_p = (\sum_i |x_i|^p)^{\frac{1}{p}}$, defined to be $\max_i |x_i|$ for $p = \infty$. Let k denote the sparsity of frequency domain. All the frequencies $\{f_1, f_2, \dots, f_k\}$ are from $[-F, F]$. Let $B = O(k)$ denote the number of hash bins in our algorithm.

We translate the ‘‘permutation’’ $P_{\sigma, a, b}$ of [HIKP12a] from the DFT setting to the DTFT setting.

Definition A.1. $(P_{\sigma, a, b}x^*)(t) = x^*(\sigma(t - a))e^{-2\pi i \sigma b t}$.

Lemma A.2. $\widehat{P_{\sigma, a, b}x^*}(\sigma(f - b)) = e^{-2\pi f \sigma a i} \widehat{x^*}(f)$

Proof. The time domain representation of the given Fourier definition would be

$$\begin{aligned} x'(t) &= \sum_{f \in \text{supp}(\widehat{x^*})} e^{-2\pi i \sigma a f} \widehat{x^*}(f) e^{2\pi i \sigma (f - b)t} \\ &= \sum_{f \in \text{supp}(\widehat{x^*})} e^{-2\pi i \sigma b t} \widehat{x^*}(f) e^{2\pi i \sigma (t - a)f} \\ &= e^{-2\pi i \sigma b t} x^*(\sigma(t - a)) \end{aligned}$$

which matches, so the formula is right. \square

We also extend the flat window function for the DFT setting [HIKP12a], [HIKP12b] to the DTFT setting:

Definition A.3. Let $M = O(B \log \frac{k}{\delta})$. We say that $(G, \widehat{G}') = (G_{B, \delta, \alpha}, \widehat{G}'_{B, \delta, \alpha}) \in \mathbb{R}^M \times \mathbb{R}^{[-F, F]}$ is a flat window function with parameters $B \geq 1$, $\delta > 0$, and $\alpha > 0$. For simplicity, let's say B is a function of α . Define $|\text{supp}(G)| = M$ and \widehat{G}' satisfies

- $G_i = \frac{\sin(\frac{i}{B})}{i} \cdot e^{-\frac{i^2}{2\sigma^2}}$, where $\sigma = \Theta(B \sqrt{\log(k/\delta)})$.
- $\widehat{G}(f) = \sum_{i=1}^M G_i e^{f \cdot \frac{i}{M} 2\pi i}$.
- $\text{supp}(\widehat{G}') \subset [-\frac{2\pi}{2B}, \frac{2\pi}{2B}]$.
- $\widehat{G}'(f) = 1$ for all $f \in [-\frac{(1-\alpha)2\pi}{2B}, \frac{(1-\alpha)2\pi}{2B}]$.

- $\widehat{G}'(f) = 0$ for all $|f| \geq \frac{2\pi}{2B}$.
- $\widehat{G}'(f) \in [0, 1]$ for all f .
- $\|\widehat{G}' - \widehat{G}\|_\infty^2 < \delta/k$.

Claim A.4. $\sum_{i=1}^M G_i^2 \approx \frac{1}{B}$, where $M = O(B \log k/\delta)$.

Proof. By definition of G_i , we have

$$\sum_{i=1}^M G_i^2 = 2 \sum_{i=1}^{M/2} \frac{\sin^2(i\frac{1}{B})}{(i)^2} (e^{-\frac{(i)^2}{2\sigma^2}})^2.$$

There exists some constant $c \in [0, 2\pi)$, such that $\frac{\sin(i/B)}{i} \approx \frac{1}{B}$ if $i/B < c\pi$.

$$\begin{aligned} \sum_{i=1}^M G_i^2 &= 2 \sum_{i=1}^{\lfloor Bc\pi \rfloor} \frac{\sin^2(i\frac{1}{B})}{(i)^2} (e^{-\frac{(i)^2}{2\sigma^2}})^2 + 2 \sum_{i=\lceil Bc\pi \rceil}^{M/2} \frac{\sin^2(i\frac{1}{B})}{(i)^2} (e^{-\frac{(i)^2}{2\sigma^2}})^2 \\ &\leq 2 \sum_{i=1}^{\lfloor Bc\pi \rfloor} \frac{\sin^2(i\frac{1}{B})}{(i)^2} \cdot 1 + 2 \sum_{i=\lceil Bc\pi \rceil}^{M/2} \frac{\sin^2(i\frac{1}{B})}{(i)^2} \cdot 1 \\ &\lesssim \sum_{i=1}^{\lfloor Bc\pi \rfloor} \frac{1}{B^2} + \sum_{i=\lceil Bc\pi \rceil}^{M/2} \frac{1}{i^2} \\ &\lesssim \frac{1}{B}. \end{aligned}$$

Thus, we show an upper bound. It remains to prove the lower bound.

$$\begin{aligned} \sum_{i=1}^M G_i^2 &\geq 2 \sum_{i=1}^{\lfloor Bc\pi \rfloor} \frac{\sin^2(i\frac{1}{B})}{(i)^2} (e^{-\frac{(i)^2}{2\sigma^2}})^2 \\ &\geq 2 \sum_{i=1}^{\lfloor Bc\pi \rfloor} \frac{\sin^2(i\frac{1}{B})}{(i)^2} (e^{-\frac{(Bc\pi)^2}{2\sigma^2}})^2 \\ &\gtrsim 2 \sum_{i=1}^{\lfloor Bc\pi \rfloor} \frac{1}{B^2} (e^{-\frac{(Bc\pi)^2}{2\sigma^2}})^2 \\ &\gtrsim \frac{1}{B} (e^{-\frac{(Bc\pi)^2}{2\sigma^2}})^2 \\ &\gtrsim \frac{1}{B} (e^{-c_0})^2. \end{aligned}$$

The last inequality follows by there exists some universal constant $c_0 > 0$ such that $-\frac{1}{\log(k/\delta)} \gtrsim -c_0$. Thus, we show $\sum_{i=1}^M G_i^2 \gtrsim \frac{1}{B}$. \square

To analyze the details of our algorithm, we explain some lower-level definitions and claims first. Here we give the definition of three notations that related to hash function.

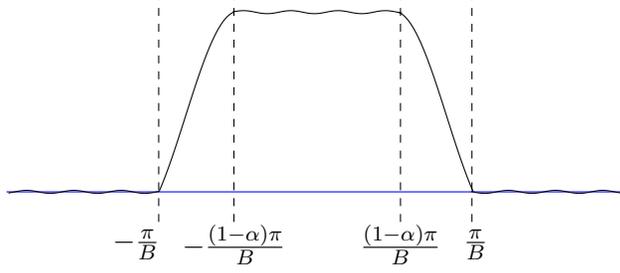


Figure 1: Filter $\widehat{G}'(f)$

Definition A.5. $\pi_{\sigma,b}(f) = 2\pi\sigma(f - b) \pmod{2\pi}$. We maps frequency to a circle $[0, 2\pi)$, since our observation of sample is the phase of some complex number, which also belongs to $[0, 2\pi)$.

Definition A.6. $h_{\sigma,b}(f) = \text{round}(\pi_{\sigma,b}(f) \cdot \frac{B}{2\pi})$. $h_{\sigma,b}(f)$ is a “hash function” that hashes frequency f into one of the B bins. The motivation is, it is very likely that each bin only has 1 heavy hitters if we choose large enough B . Then, for each bin, we can run a 1-sparse algorithm to recover the frequency.

Definition A.7. $o_{\sigma,b}(f) = \pi_{\sigma,b}(f) - \frac{2\pi}{B} \cdot h_{\sigma,b}(f)$. Offset $o_{\sigma,b}(f)$ denotes the distance from $\pi_{\sigma,b}(f)$ to the center of the corresponding bin that frequency f was hashed into.

Then we define some events that might happen after applying hash function to the entire frequency domain.

Definition A.8. “Collision” event $E_{\text{coll}}(f)$: holds iff $h_{\sigma,b}(f) \in h_{\sigma,b}(\text{supp}(x^*) \setminus \{f\})$. The “collision” event happening means there exists some other frequency f' such that both f and f' are hashed into the same bin. Once two frequencies are colliding in one bin, the algorithm will not be able to recover them.

Definition A.9. “Large offset” event $E_{\text{off}}(f)$: holds iff $|o_{\sigma,b}(f)| \geq (1 - \alpha)\frac{2\pi}{B}$. The event holds if frequency f is not within factor $1 - \alpha$ of the radius close to the center of that hash bin. It causes the frequency to be in the intermediate regime of filter and not recoverable, see Figure 1.

Definition A.10. We sample σ uniformly at random from $[\frac{1}{B\eta}, \frac{2}{B\eta}]$. Conditioning on σ is chosen first, we sample b uniformly at random from $[0, \frac{[F/\eta]}{B\sigma}]$. Then we sample γ uniformly at random from $[\frac{1}{2}, 1]$ and β uniformly at random from $[\widehat{\beta}, 2\widehat{\beta}]$, where $\widehat{\beta}$ is dynamically changing during our algorithm(The details of setting $\widehat{\beta}$ are explained in Lemma 3.6). For $P_{\sigma,\gamma,b}$ and $P_{\sigma,\gamma+\beta,b}$, we take the following two sets of samples over time domain,

$$\begin{aligned} & x(\sigma(1 - \gamma)), x(\sigma(2 - \gamma)), x(\sigma(3 - \gamma)), \dots, x(\sigma(B \log(k/\delta) - \gamma)) \\ & x(\sigma(1 - \gamma - \beta)), x(\sigma(2 - \gamma - \beta)), x(\sigma(3 - \gamma - \beta)), \dots, x(\sigma(B \log(k/\delta) - \gamma - \beta)) \end{aligned}$$

Conditioning on drawing σ, b from some distribution, we are able to show that the probability of “Collision” and “Large offset” event holding are small.

Lemma A.11. For any \widetilde{T} , and $0 \leq \widetilde{\epsilon}, \widetilde{\delta} \leq \widetilde{T}$, if we sample $\widetilde{\sigma}$ uniformly at random from $[A, 2A]$, then

$$\frac{2\widetilde{\epsilon}}{\widetilde{T}} - \frac{2\widetilde{\epsilon}}{A} \leq \Pr \left[\widetilde{\sigma} \pmod{\widetilde{T}} \in [\widetilde{\delta} - \widetilde{\epsilon}, \widetilde{\delta} + \widetilde{\epsilon}] \right] \leq \frac{2\widetilde{\epsilon}}{\widetilde{T}} + \frac{4\widetilde{\epsilon}}{A}. \quad (9)$$

Proof. Since we sample $\tilde{\sigma}$ uniformly at random from $[A, 2A]$, let I denote a set of candidate integers, then the smallest one is $\lfloor A \rfloor$ and the largest one is $\lceil 2A \rceil$. Thus, the original probability equation is equivalent to

$$\Pr \left[\tilde{\sigma} \in [s \cdot \tilde{T} + \tilde{\delta} - \tilde{\epsilon}, s \cdot \tilde{T} + \tilde{\delta} + \tilde{\epsilon}] \exists s \in I \right], \quad (10)$$

where $I = \{\lfloor A/\tilde{T} \rfloor, \dots, \lceil 2A/\tilde{T} \rceil\}$.

Consider any $s \in I$, the probability of $\tilde{\sigma}$ belonging to the interval $[s \cdot \tilde{T} + \tilde{\delta} - \tilde{\epsilon}, s \cdot \tilde{T} + \tilde{\delta} + \tilde{\epsilon}]$ is

$$\Pr \left[\tilde{\sigma} \in [s \cdot \tilde{T} + \tilde{\delta} - \tilde{\epsilon}, s \cdot \tilde{T} + \tilde{\delta} + \tilde{\epsilon}] \right] = \frac{2\tilde{\epsilon}}{A}.$$

Taking the summation over all $s \in I$, we obtain

$$\begin{aligned} & \Pr \left[\tilde{\sigma} \in [s \cdot \tilde{T} + \tilde{\delta} - \tilde{\epsilon}, s \cdot \tilde{T} + \tilde{\delta} + \tilde{\epsilon}] \exists s \in I \right] \\ &= \sum_{s \in I} \Pr \left[\tilde{\sigma} \in [s \cdot \tilde{T} + \tilde{\delta} - \tilde{\epsilon}, s \cdot \tilde{T} + \tilde{\delta} + \tilde{\epsilon}] \right] \\ &= \sum_{s \in I} \frac{2\tilde{\epsilon}}{A} \\ &= \frac{2\tilde{\epsilon}|I|}{A}. \end{aligned}$$

It remains to bound $\frac{2\tilde{\epsilon}|I|}{A}$. Since $|I| = \lceil 2A/\tilde{T} \rceil - \lfloor A/\tilde{T} \rfloor + 1$, then we have an upper bound for I ,

$$|I| \leq A/\tilde{T} + 2.$$

On the other side, we have an lower bound,

$$|I| \geq A/\tilde{T} - 1.$$

Plugging upper bound of I into $\frac{2\tilde{\epsilon}|I|}{A}$,

$$\frac{2\tilde{\epsilon}|I|}{A} \leq \frac{2\tilde{\epsilon}}{A}(A/\tilde{T} + 2) = \frac{2\tilde{\epsilon}}{\tilde{T}} + \frac{4\tilde{\epsilon}}{A}.$$

Using the lower bound of $|I|$, we have

$$\frac{2\tilde{\epsilon}|I|}{A} \geq \frac{2\tilde{\epsilon}}{A}(A/\tilde{T} - 1) = \frac{2\tilde{\epsilon}}{\tilde{T}} - \frac{2\tilde{\epsilon}}{A}.$$

Thus, we complete the proof. \square

The following corollary will be used many times in this paper. The proof directly follows by Lemma A.11.

Corollary A.12. *For any \tilde{T} , Δf , and $0 \leq \tilde{\epsilon}, \tilde{\delta} \leq \tilde{T}$, if we sample $\tilde{\sigma}$ uniformly at random from $[A, 2A]$, then*

$$\frac{2\tilde{\epsilon}}{\tilde{T}} - \frac{2\tilde{\epsilon}}{A\Delta f} \leq \Pr \left[\tilde{\sigma}\Delta f \pmod{\tilde{T}} \in [\tilde{\delta} - \tilde{\epsilon}, \tilde{\delta} + \tilde{\epsilon}] \right] \leq \frac{2\tilde{\epsilon}}{\tilde{T}} + \frac{4\tilde{\epsilon}}{A\Delta f}. \quad (11)$$

Proof. Since $\tilde{\sigma}$ is sampled uniformly at random from $[A, 2A]$, then $\tilde{\sigma}\Delta f$ is sampled uniformly at random from $[A\Delta f, 2A\Delta f]$. Now applying Lemma A.11 by only replacing $A\Delta f$ by A . \square

Claim A.13. Let σ be sampled uniformly at random from $[\frac{1}{B\eta}, \frac{2}{B\eta}]$ and $\min_{i \neq j} |f_i - f_j| > \eta$. $\forall i, j \in [k]$, if $i \neq j$, then $\Pr[h_{\sigma,b}(f_i) = h_{\sigma,b}(f_j)] \lesssim \frac{1}{B}$.

Proof. To simplify the proof, define $\Delta f = |f_i - f_j|$. We consider two cases: (I) $\eta \leq |f_i - f_j| < \frac{(B-1)\eta}{2}$, (II) $\frac{(B-1)\eta}{2} \leq |f_i - f_j|$.

(I) If $\Delta f = \eta$, then $2\pi\sigma\Delta f$ is at least $\frac{2\pi}{\eta B} \cdot \eta = \frac{2\pi}{B}$, which means two frequencies have to go to different bins after hashing. If $\Delta f = \frac{(B-1)\eta}{2}$, then $2\pi\sigma\Delta f$ is at most $\frac{4\pi}{B\eta} \cdot \frac{(B-1)\eta}{2} = (1 - 1/B)2\pi$. In order to make two frequencies collide, $2\pi\sigma\Delta f$ should belong to $[(1 - 1/B)2\pi, (1 + 1/B)2\pi)$. Since for any $\Delta f \in [\eta, \frac{(B-1)\eta}{2}]$, we have $2\pi\sigma\Delta f \in [\frac{1}{B}2\pi, (1 - 1/B)2\pi)$, which does not intersect interval $[(1 - 1/B)2\pi, (1 + 1/B)2\pi)$. Thus,

$$\Pr_{\sigma,b}[h_{\sigma,b}(f_i) = h_{\sigma,b}(f_j)] = 0.$$

(II) We apply Corollary A.12 by setting $\tilde{T} = 2\pi$, $\tilde{\sigma} = 2\pi\sigma$, $\tilde{\delta} = 0$, $\tilde{\epsilon} = \frac{2\pi}{2B}$, $A = 2\pi\frac{1}{B\eta}$. Then we have

$$\Pr_{\sigma,b}[h_{\sigma,b}(f_i) = h_{\sigma,b}(f_j)] = \Pr_{\sigma,b} \left[2\pi\sigma\Delta f \in [s \cdot 2\pi - \frac{2\pi}{2B}, s \cdot 2\pi + \frac{2\pi}{2B}] \exists s \in I \right], \quad (12)$$

where

$$I = \{ \lfloor \frac{1}{B\eta}\Delta f \rfloor, \dots, \lceil \frac{2}{B\eta}\Delta f \rceil \}.$$

By upper bound of Corollary A.12, Equation (12) is at most

$$\frac{1}{\frac{2\pi}{B\eta}\Delta f} \cdot \frac{2\pi}{B} \cdot (\frac{1}{B\eta}\Delta f + 2) = \frac{1}{B} + \frac{2\eta}{\Delta f} \leq \frac{1}{B} + \frac{4}{B-1} \lesssim \frac{1}{B},$$

where the first inequality follows by $\frac{(B-1)\eta}{2} \leq |f_i - f_j|$ which is the assumption of part (II). \square

Claim A.14. $\forall f, \forall 0 < \alpha < 1$, $\Pr_{\sigma,b} [|o_{\sigma,b}(f)| \leq (1 - \alpha)\frac{2\pi}{2B}] \geq 1 - O(\alpha)$.

Proof. Since we draw σ uniformly at random from $[\frac{1}{B\eta}, \frac{2}{B\eta}]$, then $2\pi\sigma(f-b) \in [\frac{2\pi}{B\eta}(f-b), \frac{4\pi}{B\eta}(f-b)]$ uniformly at random. The probability is equal to

$$\Pr_{\sigma,b} \left[2\pi\sigma(f-b) \in [\cdot \frac{2\pi}{B} - (1-\alpha)\frac{2\pi}{2B}, s \cdot \frac{2\pi}{B} + (1-\alpha)\frac{2\pi}{2B}] \exists s \in I \right],$$

where

$$I = \{ \lfloor \frac{2\pi}{B\eta}(f-b)\frac{B}{2\pi} - \frac{1-\alpha}{2} \rfloor, \dots, \lceil \frac{4\pi}{B\eta}(f-b)\frac{B}{2\pi} + \frac{1-\alpha}{2} \rceil \}.$$

We apply Corollary A.12 by setting $\tilde{T} = \frac{2\pi}{B}$, $\tilde{\sigma} = 2\pi\sigma$, $\tilde{\delta} = 0$, $\tilde{\epsilon} = (1-\alpha)\frac{2\pi}{2B}$, $A = 2\pi\frac{1}{B\eta}$, $\Delta f = |f-b|$.

By lower bound of Corollary A.12, we have

$$\Pr \left[|o_{\sigma,b}(f)| \leq (1-\alpha)\frac{2\pi}{2B} \right] \geq (1-\alpha) - (1-\alpha) \cdot \frac{\eta}{|f-b|}.$$

Since $\alpha \cdot \frac{\eta}{|f-b|} > 0$, then

$$\Pr \left[|o_{\sigma,b}(f)| \leq (1-\alpha)\frac{2\pi}{2B} \right] \geq (1-\alpha) - \frac{\eta}{|f-b|}.$$

Recall that we sample σ uniformly at random from $[\frac{1}{B\eta}, \frac{2}{B\eta}]$ and sample b uniformly at random from $[0, \frac{[F/\eta]}{B\sigma}]$. Since $f \in [-F, F]$ and b is uniformly chosen from range $(0, \frac{[F/\eta]}{B\sigma}]$, thus for any $C > 0$, $\Pr_b[|f - b| \leq C\eta] \lesssim \frac{C\eta}{F}$. Replacing C by $1/\alpha$, we have $\Pr_b[\frac{\eta}{|f-b|} \leq \alpha] \geq 1 - \Theta(\frac{\eta}{\alpha F})$. Compared to $\frac{\eta}{F}$, α is just a constant. Thus, we finish the proof. \square

B Proofs of basic hashing-related lemmas

Lemma 3.1. *Let (v, f) and (v', f') denote any two tones, i.e., (magnitude, frequency) pairs. Then for*

$$\text{dist}((v, f), (v', f'))^2 := \frac{1}{T} \int_0^T |ve^{2\pi fti} - v'e^{2\pi f'ti}|^2 dt,$$

we have

$$\text{dist}((v, f), (v', f'))^2 \approx (|v|^2 + |v'|^2) \cdot \min(1, T^2|f - f'|^2) + |v - v'|^2,$$

and

$$\text{dist}((v, f), (v', f')) \approx |v| \cdot \min(1, T|f - f'|) + |v - v'|.$$

Proof. Define $\nu = |f - f'|$. First, let's show the first upper bound. We have that

$$\begin{aligned} \text{LHS} &= \frac{1}{T} \int_0^T |ve^{2\pi fti} - v'e^{2\pi f'ti}|^2 dt \\ &\leq 2 \cdot \frac{1}{T} \int_0^T |ve^{2\pi fti} - ve^{2\pi f'ti}|^2 + |ve^{2\pi f'ti} - v'e^{2\pi f'ti}|^2 dt \\ &= 2|v|^2 \cdot \frac{1}{T} \int_0^T |e^{2\pi\nu ti} - 1|^2 dt + 2|v - v'|^2 \\ &\leq 2|v|^2 \cdot \frac{1}{T} \int_0^T \min(2, 2\pi\nu t)^2 dt + 2|v - v'|^2 \\ &\leq 2|v|^2 \cdot \min(4, \frac{4\pi^2}{3}\nu^2 T^2) + 2|v - v'|^2, \end{aligned}$$

as desired.

Now consider the lower bound. First we show this in the setting where $|v| = |v'|$. Suppose $v' = ve^{-\theta i}$. Then we want to bound

$$\begin{aligned} \text{LHS} &= \frac{1}{T} \int_0^T |ve^{2\pi fti} - ve^{(2\pi f't - \theta)i}|^2 dt \\ &= |v|^2 \frac{1}{T} \int_0^T |e^{(2\pi\nu t + \theta)i} - 1|^2 dt, \end{aligned}$$

as being at least $\Omega(|v|^2(\min(1, \nu^2 T^2) + \theta^2))$. In the case that $\nu T < 1/10$, then

$$|e^{(2\pi\nu t + \theta)i} - 1| \gtrsim |2\pi\nu t + \theta|,$$

and

$$\mathbb{E}_t[(2\pi\nu t + \theta)^2] \geq (\theta - \frac{2\pi\nu T}{2})^2 + \mathbb{E}[(2\pi\nu(t - T/2))^2] \gtrsim \nu^2 T^2 + (\theta - \frac{2\pi\nu T}{2})^2 \gtrsim \nu^2 T^2 + \theta^2.$$

On the other hand, if $\nu T > 1/10$, then $2\pi\nu t - \theta$ is $\Omega(1)$ for at least a constant fraction of the t , giving that

$$\left| e^{(2\pi\nu t + \theta)\mathbf{i}} - 1 \right| \gtrsim 1.$$

Hence the lower bound holds whenever $|v| = |v'|$.

Finally, consider the lower bound for $|v| \neq |v'|$. Without loss of generality assume $|v'| \geq |v|$, and define $v^* = \frac{|v|}{|v'|} v''$. For any two angles θ, θ' we have that

$$|ve^{\theta\mathbf{i}} - v'e^{\theta'\mathbf{i}}|^2 \geq |ve^{\theta\mathbf{i}} - v^*e^{\theta'\mathbf{i}}|^2 + |v^* - v'|^2,$$

because the angle $\angle vv^*v'$ is obtuse. Therefore

$$\begin{aligned} \text{LHS} &\geq \frac{1}{T} \int_0^T \left| ve^{2\pi f t \mathbf{i}} - v^* e^{2\pi f' t \mathbf{i}} \right|^2 + |v^* - v'|^2 dt \\ &\gtrsim |v|^2 \min(1, \nu^2 T^2) + |v - v^*|^2 + |v^* - v'|^2 \\ &\gtrsim |v|^2 \min(1, \nu^2 T^2) + |v - v'|^2. \end{aligned}$$

Now, if $|v'|^2 \leq 2|v|^2$, this gives the desired bound. But otherwise, it also gives the desired bound because $|v - v'|^2 \gtrsim |v'|^2$. So we get the bound in all settings.

The second equation follows from the first, using that $(a+b)^2 \approx a^2 + b^2$ for positive a, b to show

$$\text{dist}((v, f), (v', f')) \approx (|v| + |v'|) \cdot \min(1, T|f - f'|) + |v - v'|.$$

We can then replace $|v| + |v'|$ with $|v|$ because either they are equivalent up to constants or $|v - v'|$ is within a constant factor of $|v| + |v'|$. \square

Lemma 3.3. *If $x^*(t) = 0, \forall t \in [0, T]$, then*

$$\mathbb{E}_{\sigma, a, b} \left[\sum_{j=1}^B |\hat{u}_j|^2 \right] \lesssim \frac{1}{T} \int_0^T |g(t)|^2 dt.$$

Proof. Since $\hat{u}_j = \text{FFT}(u_j)$, then $\sum_{j=1}^B |\hat{u}_j|^2 = B \sum_{j=1}^B |u_j|^2$. Recall that $(P_{\sigma, a, b} x)(t) = x(\sigma(t - a))e^{2\pi\sigma b t \mathbf{i}}$, $u_j = \sum_{i=1}^{\log(k/\delta)} y_{j+Bi}$ and $y_j = G_j \cdot (P_{\sigma, a, b} g)_j = G_j \cdot g(\sigma(j - a))e^{2\pi i \sigma b j}$. Then

$$\begin{aligned} \mathbb{E}_{\sigma, a, b} \sum_{j=1}^B |u_j|^2 &= \mathbb{E}_{\sigma, a, b} \sum_{j=1}^B \left| \sum_{i=1}^{\log(k/\delta)} y_{j+Bi} \right|^2 \\ &= \mathbb{E}_{\sigma, a} \sum_{j=1}^B \mathbb{E}_b \left| \sum_{i=1}^{\log(k/\delta)} y_{j+Bi} \right|^2 \\ &= \mathbb{E}_{\sigma, a} \sum_{j=1}^B \mathbb{E}_b \left(\sum_{i=1}^{\log(k/\delta)} y_{j+Bi} \right) \left(\sum_{i'=1}^{\log(k/\delta)} \overline{y_{j+Bi'}} \right) \\ &= \mathbb{E}_{\sigma, a} \sum_{j=1}^B \mathbb{E}_b \left(\sum_{i=1}^{\log(k/\delta)} y_{j+Bi} \overline{y_{j+Bi}} + \sum_{i \neq i'}^{\log(k/\delta)} y_{j+Bi} \overline{y_{j+Bi'}} \right). \end{aligned}$$

For any $(i, j) \in [\log(k/\delta)] \times [B]$, let $S_{i,j} = G_{j+Bi}g(\sigma(j + Bi - a)) = y_{j+Bi}e^{-2\pi i\sigma b(j+Bi)}$, then

$$\mathbb{E}_{\sigma,a,b} \sum_{j=1}^B |u_j|^2 = \mathbb{E}_{\sigma,a} \sum_{j=1}^B \mathbb{E}_b \left(\underbrace{\sum_{i=1}^{\log(k/\delta)} |S_{i,j}|^2}_{C_1} + \underbrace{\sum_{i \neq i'}^{\log(k/\delta)} S_{i,j} \overline{S_{i',j}} e^{2\pi i\sigma b B(i-i')}}_{C_2} \right).$$

Consider the expectation of C_2 :

$$\begin{aligned} \mathbb{E}_b C_2 &= \mathbb{E}_b \sum_{i \neq i'}^{\log(k/\delta)} S_{i,j} \overline{S_{i',j}} e^{2\pi i\sigma b B(i-i')} \\ &= \sum_{i \neq i'}^{\log(k/\delta)} S_{i,j} \overline{S_{i',j}} \mathbb{E}_b e^{2\pi i\sigma b B(i-i')} \\ &= 0 \quad \text{by Definition A.10} \end{aligned} \tag{13}$$

Note that term C_1 is independent of b which means $\mathbb{E}_b C_1 = C_1$. Thus, we can remove the expectation over b . Then,

$$\begin{aligned} \mathbb{E}_{\sigma,a,b} \sum_{j=1}^B |u_j|^2 &= \mathbb{E}_{\sigma,a} \sum_{j=1}^B \sum_{i=1}^{\log(k/\delta)} |G_{j+Bi}|^2 \cdot |g(\sigma(j + Bi - a))|^2 \\ &= \mathbb{E}_{\sigma,a} \sum_{i=1}^{B \log(k/\delta)} |G_i|^2 \cdot |g(\sigma(i - a))|^2. \end{aligned}$$

Now, the idea is to replace the expectation term \mathbb{E}_a by an integral term $\int_{a \in A} (\star) da$. Then, replace it by another integral term $\int_0^T (\star) dt$. Let A denote a set of intervals that we will sample a from. It is easy to verify that $|A| \lesssim T/\sigma$, since $(\sigma(i - a))$ is sampled from $[0, T]$. If we choose T to be a constant factor larger than $\sigma |\text{supp}(G)|$, then we also have $|A| \gtrsim T/\sigma$.

$$\begin{aligned} \mathbb{E}_{\sigma} \mathbb{E}_a \sum_{i=1}^{B \log(k/\delta)} |G_i|^2 \cdot |g(\sigma(i - a))|^2 &= \mathbb{E}_{\sigma} \frac{1}{|A|} \int_{a \in A} \sum_{i=1}^{B \log(k/\delta)} |G_i|^2 \cdot |g(\sigma(i - a))|^2 da \\ &= \mathbb{E}_{\sigma} \sum_{i=1}^{B \log(k/\delta)} |G_i|^2 \cdot \frac{1}{|A|} \int_{a \in A} |g(\sigma(i - a))|^2 da \\ &= \mathbb{E}_{\sigma} \sum_{i=1}^{B \log(k/\delta)} |G_i|^2 \cdot \frac{1}{\sigma |A|} \int_{a \in A} |g(\sigma(i - a))|^2 d\sigma a \\ &\lesssim \mathbb{E}_{\sigma} \|G\|_2^2 \cdot \frac{1}{T} \int_0^T |g(t)|^2 dt. \end{aligned}$$

By Claim A.4, we know that $\|G\|_2^2 \approx \frac{1}{B}$. Combining $\sum_{j=1}^B |\hat{u}_j|^2 = B \sum_{j=1}^B |u_j|^2$ and $\|G\|_2^2 \approx \frac{1}{B}$ gives the desired result. \square

Lemma 3.4. *If $g(t) = 0, \forall t \in [0, T]$. Let H denote a set of frequencies, $H = \{f \in \text{supp}(\widehat{x}^*) \mid \text{neither } E_{\text{coll}}(f) \text{ nor } E_{\text{off}}(f) \text{ holds}\}$. Then,*

$$\mathbb{E}_{\sigma, a, b} \left[\sum_{f \in H} \left| \widehat{u}_{h_{\sigma, b}(f)} - \widehat{x}^*(f) e^{a\sigma 2\pi f i} \right|^2 \right] \leq \delta \|\widehat{x}^*\|_1^2.$$

Proof. For simplicity, let $G = G_{B, \delta, \alpha}$ and $\widehat{G}' = \widehat{G}'_{B, \delta, \alpha}$. we have

$$\begin{aligned} \widehat{y} &= G \cdot \widehat{P_{\sigma, a, b} x} \\ &= \widehat{G} * \widehat{P_{\sigma, a, b} x} \\ &= \widehat{G}' * \widehat{P_{\sigma, a, b} x} + (\widehat{G} - \widehat{G}') * \widehat{P_{\sigma, a, b} x}. \end{aligned}$$

The ℓ_∞ norm of second term can be bounded :

$$\|(\widehat{G} - \widehat{G}') * \widehat{P_{\sigma, a, b} x}\|_\infty \leq \|\widehat{G} - \widehat{G}'\|_\infty \|\widehat{P_{\sigma, a, b} x}\|_1 \leq \sqrt{\delta/k} \|\widehat{x}^*\|_1.$$

Thus, consider the j th term of \widehat{u} ,

$$\begin{aligned} \widehat{u}_j &= \widehat{y}_{jF/B} \\ &= \sum_{|l| < F/(2B)} \widehat{G}'_{-l} (\widehat{P_{\sigma, a, b} x})_{jF/B+l} \pm \sqrt{\delta/k} \|\widehat{x}^*\|_1 \\ &= \sum_{|\pi_{\sigma, b}(f) - jF/B| < F/(2B)} \widehat{G}'_{jF/B - \pi_{\sigma, b}(f)} \widehat{P_{\sigma, a, b} x}_{\pi_{\sigma, b}(f)} \pm \sqrt{\delta/k} \|\widehat{x}^*\|_1 \\ &= \sum_{h_{\sigma, b}(f)=j} \widehat{G}'_{-o_{\sigma, b}(f)} \widehat{x}^*(f) e^{2\pi f \sigma a i} \pm \sqrt{\delta/k} \|\widehat{x}^*\|_1. \end{aligned}$$

If neither $E_{\text{coll}}(f)$ nor $E_{\text{off}}(f)$ happens, then we know that frequency f is the only heavy hitter hashed into bin j and $\widehat{G}'_{-o_{\sigma, b}(f)} = 1$ for frequency f . Thus,

$$\mathbb{E}_{\sigma, a, b} \left[\left| \widehat{u}_{h_{\sigma, b}(f)} - \widehat{x}^*(f) e^{a\sigma 2\pi f i} \right|^2 \right] \leq \delta/k \|\widehat{x}^*\|_1^2.$$

Since the above equation holds for all $f \in H$, we get

$$\sum_{f \in H} \mathbb{E}_{\sigma, a, b} \left[\left| \widehat{u}_{h_{\sigma, b}(f)} - \widehat{x}^*(f) e^{a\sigma 2\pi f i} \right|^2 \right] \leq k\delta/k \|\widehat{x}^*\|_1^2 = \delta \|\widehat{x}^*\|_1^2.$$

□

C Proofs for one stage of recovery

Binary search of one-sparse algorithm We first explain a simple, clean, but not optimal one-sparse algorithm, then we try to optimize the algorithm step by step. Let “heavy” frequency $f \in [-F, F]$, we can split the frequency interval into two regions: left region $[-F, 0)$ and right region $[0, F]$. We can observe $\theta \equiv 2\pi\beta f \pmod{2\pi}$ by checking the phase difference between using $P_{1, a, b}$ and $P_{1, a+\beta, b}$, where β is uniformly at random sampled from some suitable range $[\widehat{\beta}, 2\widehat{\beta}]$. For each observation θ , we can guess m different possibilities for f , say $\theta_1, \theta_2, \dots, \theta_m$, if θ_i belong to

the left region, we add a vote to that, otherwise we add a vote to the right region. After taking enough samples, we choose the region that has the largest vote. This decision will let us narrow down the searching range of the true frequency by half with some good probability. Suppose we decide to choose the right region $[0, F]$, then we can just repeat the above binary search over $[0, F]$ again to get into a region that has size $F/2$. Repeating it D times, we can learn the frequency with a region that has size at most $2F/2^D$. But the binary search is not the best approach, actually, we can do much better with using t -ary search.

k -sparse To locate those k heavy signals in the frequency domain, we need to consider the “bins” computed by `HashToBins` with $P_{\sigma,a,b}$. One of the main difference from previous work [HIKP12a] is, instead of permuting the discrete coordinates according to $P_{\sigma,a,b}$ and partitioning the coordinates into $B = O(k)$ bins, we permute the continuous frequency domain and partition the frequency domain into $B = O(k)$ bins. With a large constant probability, we can obtain for each heavy signal f that neither E_{coll} nor E_{off} happens. After splitting those k frequencies into different bins, then we can run the one-sparse algorithm for all the bins simultaneously.

t -ary search To argue the final succeed probability of our algorithm, we need to take the union bound for each array/region in each round and also take the union bound over each round. There are several benefits of changing binary search to t -ary search, (1) to reach the same accuracy, the number of rounds D for t -ary search is smaller than the number of rounds for binary search; (2) our searching procedure is a “noisy” searching problem, for the noiseless version of the searching problem, we do not need to take care of the union bound argument. Having a parameter for the number of arrays/regions is important to optimize the entire procedure.

Recall the traditional binary search problem(noiseless version), given a list of sorted numbers $a[1, 2, \dots, n]$ in increasing order. We want to determine if some number x belongs to $a[1, 2, \dots, n]$. When we compare some $a[i]$ with x , we will know the true answer.

But the noisy binary search problem is slightly harder. $a[1, 2, \dots, n]$ is still a list of sorted numbers in increasing order, we want to determine if some number x belongs to $a[1, 2, \dots, n]$. But when we compare the $a[i]$ with x , we will know the true answer with 9/10 probability, and get the false answer with 1/10 probability. In this case, we can not finish the task by following the procedure of traditional binary search algorithm, e.g. making the decision by just comparing $a[i]$ with x once. The reason is after taking the union bound of $\log n$ rounds, the failure probability can be arbitrarily large. One idea to fix this issue is independently comparing $a[i]$ with x multiple times, e.g. $\log \log n$ times. Then we can amplify the succeed probability of each round from 9/10 to $1 - \frac{1}{\text{poly}(\log n)}$, after taking the union bound over $\log n$ rounds, we still have $1 - \frac{1}{\text{poly}(\log n)}$ succeed probability.

Our problem is still more complicated than the above noisy binary search problem. In each round of algorithm `LocateInner`, we do a t -array search by splitting the candidate frequency region(that has length Δl) into t consecutive regions, Q_1, Q_2, \dots, Q_t , each of them has the equal size $\frac{\Delta l}{t}$. By using the hash values of $P_{\sigma,a,b}$ and $P_{\sigma,a+\beta,b}$ (Line 26-27 in Algorithm 2), we can have an observation over $[0, 2\pi)$. For each such observation over $[0, 2\pi)$, it was in fact scaled by $2\pi\sigma\beta$ and rounded over $[0, 2\pi)$. Thus the corresponding frequency location of this observation might belong to $m = \Theta(\sigma\beta\Delta l)$ different possible regions. Since we do not know which one, we just add a vote to all of the possible regions. To understand t -ary search, let’s consider this example. Suppose we split the frequency into 20 regions, the true frequency belongs to region 9 and each observation is correct with probability 4/5. For each observation, we will add a vote to a batch of roughly evenly spaced regions.

1. For the observation 1, we add a vote to region 1, 5, 9, 13, 17.
2. For the observation 2, we add a vote to region 3, 9, 15.
3. For the observation 3, we add a vote to region 2, 9, 16.
4. For the observation 4, we add a vote to region 4, 9, 14, 19.
5. For the observation 5, we add a vote to region 2, 6, 10, 14, 20.

where the first four observations are the correct observations and the last one is wrong. Then the true region will have more than half of the $R_{loc} = 5$ votes with some good probability.

Details of adding vote In previous description, to let people have a better understanding of t -ary search, we simplify the step of adding vote. In fact, adding a vote to each of the m candidate region is not enough. The right thing to do is, not only adding a vote to those m regions, but also adding a vote to a constant number c_n of neighbors of each possible region (e.g. two neighbors nearby that region, line 33 in Algorithm 2). The reason for doing this is, if the frequency f is located very close to the boundary of two regions, then neither of them will get more than $R_{loc}/2$ votes. After we already have R_{loc} independent observations, we just choose any region that contains more than $R_{loc}/2$ votes and enlarge it by the same constant factor c_n to be the next candidate frequency region for t -array search, and plugging the new parameters into algorithm `LocateInner` and running it again.

The slightly different last round Recall that, in the previous description of each round of t -ary search, we're able to decrease the searching range of frequency geometrically. To achieve this progress, we have to increase the sample duration of β is geometrically. At some point, the sample duration will reach T . Suppose the sample duration of β is geometrically increasing during the first $D - 1$ rounds (Line 13-15 in Algorithm 2) and it becomes T after $D - 1$ rounds. If we want to use the same duration T to perform one more round again, can we get some benefit for learning the frequency by doing some tricks? (1) Suppose in all the first $D - 1$ rounds, we use $R_{loc}(D - 1)$ samples. Then we can use $R_{loc}(D - 1)$ for the last round, since it does not increase the sample complexity. This way allows us to learn frequency within $\frac{1}{cT}$. But can we do better than that? (2) Using more samples is a nice observation, another right thing to do is changing the algorithm from reporting region to reporting frequency. In the previous $D - 1$ rounds, as the description of t -ary search, we just report the region that has more than $R_{loc}/2$ votes. But, we can take the median over all the values that were assigned to any region that has votes more than $R_{loc}/2$. This way can actually allow us to learn frequency location more accurately ($\approx \frac{1}{\rho T}$) without increasing the resolution for β ! Another question people might ask is, if we take median in the last round, why not just do it in every round? The answer is, in the first $D - 1$ rounds, our goal is just narrowing down the search range of frequency location and we do not need to report a frequency location. Another reason is reporting a candidate region needs less samples and has higher succeed probability, but taking the median is more expensive than reporting a candidate region. We cannot pay for taking the median in every round.

To prove main Lemma 3.6 for one stage recovery, we introduce Lemma C.1. Before explaining the proof, we give some definitions. Consider a frequency f , and define $j = h_{\sigma,b}(f)$ to be the bin that frequency f was hashed into. Define $\theta = f - b \pmod{F}$. Define

$$\hat{u} = \text{HashToBins}(x, P_{\sigma,\gamma,b}, B, \delta, \alpha) \quad \text{and} \quad \hat{u}' = \text{HashToBins}(x, P_{\sigma,\gamma+\beta,b}, B, \delta, \alpha).$$

Note that “bin” and “region” are representing different things in this paper. “bin” is related to hash function $h_{\sigma,b}(f)$. “region” is only used in the algorithm of one stage recovery in this section. Define “true” region to be the region that contains frequency f . Define “wrong” region to be the region that is not within a constant number c_n of neighbors of the “true” region. Part (I) of Lemma C.1 shows that for each observation generated by a batch of samples drawn from time domain, the counter $v_{j,q'}$ corresponding to the true region will increase by one, with some “good” probability. On the other side, Part (II) of Lemma C.1 shows that the counter $v_{j,q}$ corresponding to wrong region will not increase by one, with some “good” probability. Note that, [HIKP12a] proved the case when $c_n = 6$ under discrete setting, we translate it into continuous setting, here.

Lemma C.1. *Given σ and b . Assume $f \in \text{region}(j, q')$ and $\mathbb{E}_\gamma[|\hat{u}_j - e^{2\pi\gamma\sigma\theta\mathbf{i}}\hat{x}(\theta)|^2] \leq \frac{1}{\rho^2}|\hat{x}(\theta)|^2$.*

$\forall 0 < s < 1$, for each two samples $(\gamma, 0)$ and (γ, β) , where γ is sampled from $[\frac{1}{2}, 1]$ uniformly at random and β is sampled from $[\frac{st}{4\sigma\Delta l}, \frac{st}{2\sigma\Delta l}]$ uniformly at random, we have

(I) for the q' , with probability at least $1 - (\frac{2}{\rho s})^2$, $v_{j,q'}$ will increase by one.

(II) for any q such that $|q - q'| > 3$, with probability at least $1 - 15s$, $v_{j,q}$ will not increase.

Proof. Part (I) of Lemma C.1. We have,

$$\mathbb{E}_\gamma\left[|\hat{u}_j - e^{2\pi\gamma\sigma\theta\mathbf{i}}\hat{x}(\theta)|^2\right] \leq \frac{1}{\rho^2}|\hat{x}(\theta)|^2.$$

By Chebyshev’s Inequality, we have $\forall g > 0$, with probability $1 - g$ we have

$$\begin{aligned} \left|\hat{u}_j - e^{2\pi\gamma\sigma\theta\mathbf{i}}\hat{x}(\theta)\right| &\leq \frac{1}{\rho}\sqrt{\frac{1}{g}}|\hat{x}(\theta)| \\ \|\phi(\hat{u}_j) - (\phi(\hat{x}(\theta)) - 2\pi\gamma\sigma\theta)\|_{\circ} &\leq \sin^{-1}\left(\frac{1}{\rho}\sqrt{\frac{1}{g}}\right), \end{aligned}$$

where $\|x - y\|_{\circ} = \min_{z \in \mathbb{Z}} |x - y + 2\pi z|$ denote the “circular distance” between x and y . Similarly, replacing γ by $\gamma + \beta$, with probability $1 - g$, we also have that

$$\|\phi(\hat{u}'_j) - (\phi(\hat{x}(\theta)) - 2\pi(\gamma + \beta)\sigma\theta)\|_{\circ} \leq \sin^{-1}\left(\frac{1}{\rho}\sqrt{\frac{1}{g}}\right).$$

Define $c_j = \phi(\hat{u}_j/\hat{u}'_j)$. Combining the above two results, with probability $1 - 2g$ we have

$$\begin{aligned} \|c_j - 2\pi\beta\sigma\theta\|_{\circ} &= \|\phi(\hat{u}_j) - \phi(\hat{u}'_j) - 2\pi\beta\sigma\theta\|_{\circ} \\ &= \|(\phi(\hat{u}_j) - (\phi(\hat{x}(\theta)) - 2\pi\gamma\sigma\theta)) - (\phi(\hat{u}'_j) - (\phi(\hat{x}(\theta)) - 2\pi(\gamma + \beta)\sigma\theta))\|_{\circ} \\ &\leq \|\phi(\hat{u}_j) - (\phi(\hat{x}(\theta)) - 2\pi\gamma\sigma\theta)\|_{\circ} + \|\phi(\hat{u}'_j) - (\phi(\hat{x}(\theta)) - 2\pi(\gamma + \beta)\sigma\theta)\|_{\circ} \\ &\leq 2\sin^{-1}\left(\frac{1}{\rho}\sqrt{\frac{1}{g}}\right). \end{aligned}$$

Here, we want to set $g = (\frac{4}{s\pi\rho})^2$, thus, with probability at least $1 - (\frac{2}{s\rho})^2$

$$\|c_j - 2\pi\beta\sigma\theta\|_{\circ} < s\pi/2. \tag{14}$$

The above equation shows that c_j is a good estimate for $2\pi\beta\sigma\theta$ with good probability. We will now show that this means the true region $Q_{q'}$ gets a vote with large probability.

For each q' with $f \in [l_j - \frac{\Delta l}{2} + \frac{q'-1}{t}\Delta l, l_j - \frac{\Delta l}{2} + \frac{q'}{t}\Delta l] \subset [-F, F]$, we have that $m_{j,q'} = l_j - \frac{\Delta l}{2} + \frac{q'-0.5}{t}\Delta l$ and $\theta_{j,q'} = m_{j,q'} - b \pmod{F}$ satisfies

$$|f - m_{j,q'}| \leq \frac{\Delta l}{2t} \text{ and } |\theta - \theta_{j,q'}| \leq \frac{\Delta l}{2t}.$$

Since we sample β uniformly at random from $[\frac{st}{4\sigma\Delta l}, \frac{st}{2\sigma\Delta l}]$, then $\beta \leq \frac{st}{2\sigma\Delta l}$, which implies that $2\pi\beta\sigma\frac{\Delta l}{2t} \leq \frac{s\pi}{2}$. Thus, we can show the observation c_j is close to the true region in the following sense,

$$\begin{aligned} & \|c_j - 2\pi\beta\sigma\theta_{j,q'}\|_{\circ} \\ & \leq \|c_j - 2\pi\beta\sigma\theta\|_{\circ} + \|2\pi\beta\sigma\theta - 2\pi\beta\sigma\theta_{j,q'}\|_{\circ} \text{ by triangle inequality} \\ & < \frac{s\pi}{2} + 2\pi\|\beta\sigma\theta - \beta\sigma\theta_{j,q'}\|_{\circ} \text{ by Equation (14)} \\ & \leq \frac{s\pi}{2} + 2\pi\beta\sigma\frac{\Delta l}{2t} \\ & = \frac{s\pi}{2} + \frac{s\pi}{2} \\ & \leq s\pi. \end{aligned}$$

Thus, $v_{j,q'}$ will increase in each round with probability at least $1 - (\frac{2}{\rho s})^2$. \square

Proof. Part (II) of Lemma C.1.

Consider q with $|q - q'| > 3$. Then $|f - m_{j,q}| \geq \frac{7\Delta l}{2t}$, and (assuming $\beta \geq \frac{st}{4\sigma\Delta l}$) we have

$$2\pi\beta\sigma|f - m_{j,q}| \geq 2\pi\frac{st}{4\sigma\Delta l}\sigma|f - m_{j,q}| = \frac{s\pi t}{2\Delta l}|f - m_{j,q}| \geq \frac{7s\pi}{4} > \frac{3s\pi}{2}. \quad (15)$$

There are two cases: $|f - m_{j,q}| \leq \frac{\Delta l}{st}$ and $|f - m_{j,q}| > \frac{\Delta l}{st}$.

First, if $|f - m_{j,q}| \leq \frac{\Delta l}{st}$. In this case, from the definition of β it follows that

$$2\pi\beta\sigma|f - m_{j,q}| \leq \frac{s\pi t}{\sigma\Delta l}\sigma|f - m_{j,q}| \leq \pi. \quad (16)$$

Combining equations (15) and (16) implies that

$$\Pr[2\pi\beta\sigma(f - m_{j,q}) \pmod{2\pi} \in [-\frac{3s}{4}2\pi, \frac{3s}{4}2\pi]] = 0.$$

Second, if $|f - m_{j,q}| > \frac{\Delta l}{st}$. We show this claim is true: $\Pr[2\pi\beta\sigma(f - m_{j,q}) \pmod{2\pi} \in [-\frac{3s}{4}2\pi, \frac{3s}{4}2\pi]] \lesssim s$. To prove it, we apply Corollary A.12 by setting $\tilde{T} = 2\pi$, $\tilde{\sigma} = 2\pi\sigma\beta$, $\tilde{\delta} = 0$, $\tilde{\epsilon} = \frac{3s}{4}2\pi$, $A = 2\pi\sigma\hat{\beta}$, $\Delta f = |f - m_{j,q}|$. By upper bound of Corollary A.12, the probability is at most

$$\frac{2\tilde{\epsilon}}{\tilde{T}} + \frac{4\tilde{\epsilon}}{A\Delta f} = \frac{3s}{2} + \frac{3s}{\sigma\hat{\beta}\Delta f} \leq \frac{3s}{2} + \frac{3s}{\sigma\frac{st}{4\sigma\Delta l}\frac{\Delta l}{st}} < 15s.$$

Then in either case, with probability at least $1 - 15s$, we have

$$\|2\pi\beta\sigma m_{j,q} - 2\pi\beta\sigma f\|_{\circ} > \frac{3s}{4}2\pi.$$

which implies that $v_{j,q}$ will not increase. \square

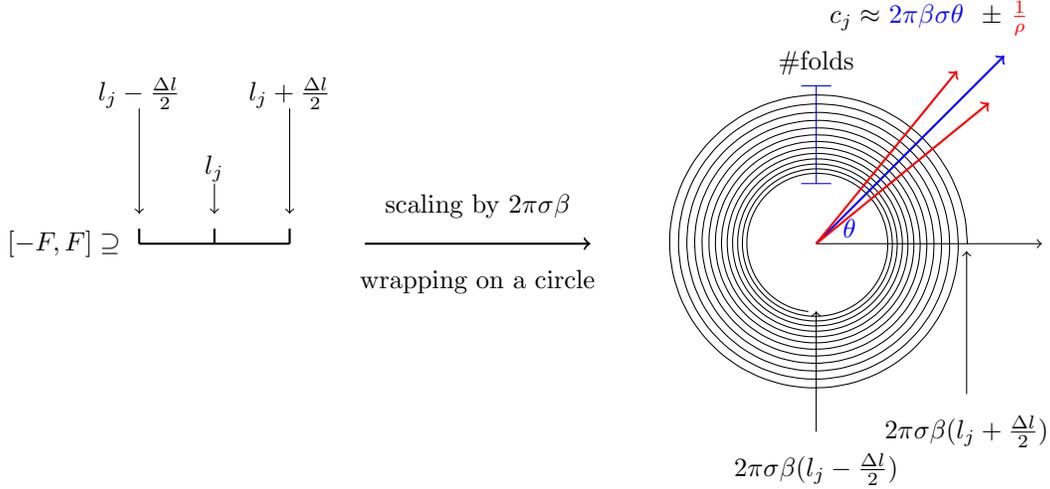


Figure 2: For an arbitrary frequency interval $[l_j - \frac{\Delta l}{2}, l_j + \frac{\Delta l}{2}]$, we scale it by $2\pi\sigma\beta$ to get a longer interval $[2\pi\sigma\beta(l_j + \frac{\Delta l}{2}), 2\pi\sigma\beta(l_j - \frac{\Delta l}{2})]$. Then, we wrap the longer interval on a circle $[0, 2\pi)$. The number of folds after wrapping is $\lceil \sigma\beta\Delta l \rceil$. For any random sample, the observation c_j is close to the true answer within $1/\rho$ with some “good” probability.

Lemma 3.6. *Given σ and b , consider any frequency f for which neither $E_{coll}(f)$ nor $E_{off}(f)$ holds, and let $j = h_{\sigma,b}(f)$. Let $\mu^2(f) = \mathbb{E}_a[|\hat{u}_j - \widehat{x}^*(f)e^{a\sigma 2\pi f i}|^2]$ and $\rho^2 = |\widehat{x}^*(f)|^2/\mu^2(f)$. For sufficiently large ρ , and $\forall 0 < s < 1, t \geq 4$, consider any run of `LocateInner` with $f \in [l_j - \frac{\Delta l}{2}, l_j + \frac{\Delta l}{2}]$. It takes $O(R_{loc})$ random $(\gamma, \beta) \in [\frac{1}{2}, 1] \times [\frac{st}{4\sigma\Delta l}, \frac{st}{2\sigma\Delta l}]$ samples over duration $\beta\sigma = \Theta(\frac{st}{\Delta l})$, runs in $O(stR_{loc})$ time, to learn f within a region that has length $\Theta(\frac{\Delta l}{t})$ with failure probability at most $(\frac{4}{s\rho})^{R_{loc}} + t \cdot (60s)^{R_{loc}/2}$.*

Proof. Let t denote the number of regions, and $[l_j - \frac{\Delta l}{2}, l_j + \frac{\Delta l}{2}]$ be the interval that contains frequency f . Let Q_q denote a region that is $[l_j - \frac{\Delta l}{2} + (q-1)\frac{\Delta l}{t}, l_j - \frac{\Delta l}{2} + q\frac{\Delta l}{t}]$. Let $\theta = f - b \pmod{F}$. Recall that we sample σ uniformly at random from $[\frac{1}{B\eta}, \frac{2}{B\eta}]$. Then we sample γ uniformly at random from $[\frac{1}{2}, 1]$ and sample β uniformly at random from $[\frac{st}{4\sigma\Delta l}, \frac{st}{2\sigma\Delta l}]$. Define $c_j = \phi(\hat{u}_j/\hat{u}'_j)$. Let m denote the number of folds, which is equal to $\lceil \sigma\beta\Delta l \rceil$. Let $v_{j,q}$ denote the vote of region (j, q) .

We hope to show that in any round r , each observed c_j is close to $2\pi\sigma\beta\theta$ with good probability. On the other hand, for each observed c_j , we need to assign it to some regions and increase the vote of the corresponding region. The straightforward way is just checking all possible t regions, which takes $O(t)$ time. In fact, there are only $\Theta(m)$ regions close enough to the observation c_j , where $m = \Theta(\frac{2\pi\sigma\beta\Delta l}{2\pi}) = \Theta(st)$. The reason is $2\pi\sigma\beta$ will scale the original length Δl interval to a new interval that has length $2\pi\sigma\beta\Delta l$. This new interval can only wrap around circle $[0, 2\pi)$ at most $\lceil \sigma\beta\Delta l \rceil$ times. The running time is $O(stR_{loc})$, since we take R_{loc} independent observations.

For each observed c_j : Part (I) of Lemma C.1 says, we assign it to the true region with some good probability; Part (II) of Lemma C.1 says, we do not assign it to the wrong region with some good probability. Thus taking R_{loc} independent c_j , we can analyze the failure probability of this algorithm based on these three cases.

(I) What's the probability of true fold fails?

$$\begin{aligned}
& \Pr[\text{True fold fails}] \\
&= \Pr[\text{True region fails} \geq R_{loc}/2 \text{ times}] \\
&\leq 2 \cdot \binom{R_{loc}}{R_{loc}/2} \cdot (\Pr[\text{True region fails once}])^{R_{loc}/2} \\
&\leq 2 \cdot \binom{R_{loc}}{R_{loc}/2} \cdot \left(\frac{2}{s\rho}\right)^{2R_{loc}/2} \quad \text{by Lemma C.1} \\
&\leq \left(\frac{4}{s\rho}\right)^{R_{loc}}.
\end{aligned}$$

(II) What if the region that is “near” (within c_n neighbors) true region becomes true?

Any of those region gets a vote only if true region also gets a vote. Since our algorithm choosing any region that has more than $R_{loc}/2$ votes and enlarging the region size by containing c_n nearby neighbors of that chosen region, then the new larger region must contain the “real” true region.

(III) What if the region that is “far away” (not within c_n neighbors) from true region becomes true?

By Part (II) of Lemma C.1, the probability of one such wrong region gets a vote is at most $15s$. Thus, one of the wrong region gets more than $R_{loc}/2$ votes is at most $(60s)^{R_{loc}/2}$. By taking the union bound over all t regions, we have the probability of existing one wrong region getting more than $R_{loc}/2$ vote is at most $t \cdot (60s)^{R_{loc}/2}$.

Thus, if we first find any region Q_q that has more than $R_{loc}/2$ votes, and report a slightly larger region $[l_j - \frac{\Delta l}{2} + (q-1)\frac{\Delta l}{t} - \frac{c_n}{2}\frac{\Delta l}{t}, l_j - \frac{\Delta l}{2} + q\frac{\Delta l}{t} + \frac{c_n}{2}\frac{\Delta l}{t}]$, it is very likely this large region contains the frequency f . Finally, the failure probability of this algorithm is at most $\Theta\left(\left(\frac{4}{s\rho}\right)^{R_{loc}} + t \cdot (60s)^{R_{loc}/2}\right)$. \square

Lemma C.2. *Taking the median of values belong to any region getting at least $\frac{1}{2}R_{loc}$ votes, then we can learn frequency f within $\Theta\left(\frac{\Delta l}{\rho st}\right)$ with probability $1 - \exp(-\Omega(R_{loc}))$.*

Proof. Let $\text{region}(j, q')$ be the region that getting at least $\frac{1}{2}R_{loc}$ votes. Let $R = |\text{region}(j, q')|$ denote the number of observations/votes assigned to $\text{region}(j, q')$. Since this region getting at least $\frac{1}{2}R_{loc}$ votes, then $\frac{1}{2}R_{loc} \leq R \leq R_{loc}$. Using Equation (14) in Lemma C.1, $\forall g > 0$, we have

$$\|c_j - 2\pi\beta\sigma\theta\|_{\circ} \leq 2 \left(\sin^{-1}\left(\frac{1}{\rho}\sqrt{\frac{1}{g}}\right) \right),$$

holds with probability $1 - 2g$. Choosing $g = \Theta(1)$, we have with constant success probability $p > \frac{1}{2}$,

$$\|c_j - 2\pi\beta\sigma\theta\|_{\circ} \lesssim \frac{1}{\rho},$$

holds.

Taking the median over all the observations that belong to $\text{region}(j, q')$ gives

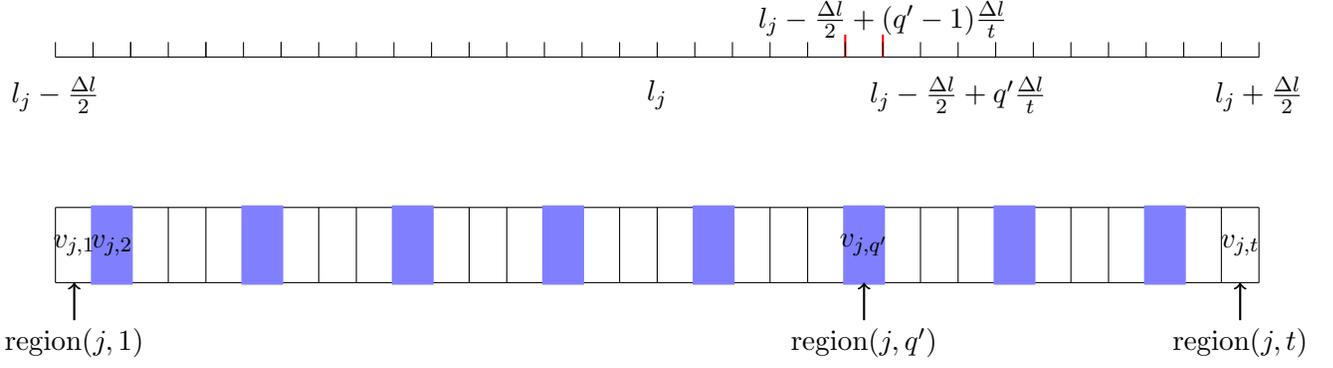


Figure 3: The number of “blue” regions is equal to the number of folds m . The number of total regions is t . For each observed c_j , instead of checking all the t regions, we only assign vote to the these “blue” regions. Since only these m “blue” regions can be the candidate region that contains frequency f .

$$\begin{aligned}
\Pr\left(\left|\text{median}_{r \in \text{region}(j, q')} c_j^r - 2\pi\beta^r \sigma\theta\right| \gtrsim \frac{1}{\rho}\right) &< \sum_{i=R/2}^{R_{loc}} \binom{R_{loc}}{i} (1-p)^i p^{R_{loc}-i} \\
&= \sum_{i=R_{loc}/2}^{R_{loc}} \binom{R_{loc}}{i} (1-p)^i p^{R_{loc}-i} + \sum_{i=R/2}^{R_{loc}/2} \binom{R_{loc}}{i} (1-p)^i p^{R_{loc}-i} \\
&< \sum_{i=R_{loc}/2}^{R_{loc}} \binom{R_{loc}}{R_{loc}/2} (1-p)^i + \sum_{i=R/2}^{R_{loc}/2} \binom{R_{loc}}{R_{loc}/2} (1-p)^i \\
&< 2 \binom{R_{loc}}{R_{loc}/2} (1-p)^{R/2} \\
&\leq 2(2e)^{R_{loc}/2} (1-p)^{R_{loc}/4} \\
&\leq e^{-cR_{loc}}, \tag{17}
\end{aligned}$$

where the the second inequality follows by $\binom{R}{i} \leq \binom{R}{R/2}$ and $p^{R_{loc}-i} < 1, \forall i$; the fourth inequality follows by $\binom{n}{k} \leq (ne/k)^k$; the last inequality follows by choosing some p such that $\frac{1}{2} \log_{2e} \frac{1}{1-p} - 1 > 2c$ where $c > 0$ is some constant. Equation (17) implies that

$$\Pr\left(\left|\text{median}_{r \in \text{region}(j, q')} \theta^r - \theta\right| < \frac{1}{\rho 2\pi\sigma\hat{\beta}}\right) > 1 - \exp(-\Omega(R_{loc})),$$

where $\forall r \in [R_{loc}]$, β^r is sampled uniformly at random from $[\hat{\beta}, 2\hat{\beta}] = [\frac{st}{4\sigma\Delta l}, \frac{st}{2\sigma\Delta l}]$. Thus, we can learn f within $\Theta(\frac{1}{\rho 2\pi\sigma\hat{\beta}}) = \Theta(\frac{\Delta l}{\rho st})$. □

Lemma 3.7. *Algorithm LocateKSignal takes $O(k \log_C(FT) \log(k/\delta))$ samples over $O(\frac{\log(k/\delta)}{\eta})$ duration, runs in $O(k \log_C(FT) \log(FT/\delta))$ time, and outputs a set $L \subset [-F, F]$ of $O(k)$ frequencies with minimum separation $\Omega(\eta)$.*

Given σ and b , consider any frequency f for which neither of $E_{coll}(f)$ or $E_{off}(f)$ hold. Let $j = h_{\sigma,b}(f)$, $\mu^2(f) = \mathbb{E}_a[|\widehat{u}_j - \widehat{x}^*(f)e^{a\sigma 2\pi f i}|^2]$, and $\rho^2 = |\widehat{x}^*(f)|^2/\mu^2(f)$. If $\rho > C$, then with an arbitrarily large constant probability there exists an $f' \in L$ with

$$|f - f'| \lesssim \frac{1}{T\rho}.$$

Proof. Algorithm `LocateKSignal` rerun procedure `LocateInner` D times. For the first $D - 1$ rounds, the sampling range for β is increased by t every time. For the last round, the sampling range for β is not increasing any more. On the other hand, the sampling range for β for $D - 1$ round and the last round are the same.

Recall that, we sample σ uniformly at random from $[\frac{1}{B\eta}, \frac{2}{B\eta}]$. Then we sample γ uniformly at random from $[\frac{1}{2}, 1]$ and β uniformly at random from $[\frac{st}{4\sigma\Delta l}, \frac{st}{2\sigma\Delta l}]$. c_n is some constant for the number of neighbor regions nearby “true” region. In the first $D - 1$ rounds, we set $s = 1/\sqrt{C}$, $\Delta l = F/(t')^{i-1}, \forall i \in [D - 1]$, $t \approx \log(FT)$, and $t' = \frac{t}{c_n+1}$. For the last round, we set $s \approx 1/C$, $\Delta l \approx st/T$ and $t \approx \log(FT)/s$. C is known as “approximation” factor. Finally, we need to choose depth $D = \log_{t'}(FT/st) \approx \log_{t'}(FT)$ and set $R_{loc} = O(\log_C(tC))$ for all the rounds. Define $DR = \sum_{i=1}^D R_{loc}^i$, which is $(D - 1)R_{loc} + R_{loc}^D = O(\log_C(FT))$.

After setting all the parameters for the first $D - 1$ rounds and the last round, we explain some intuitions and motivations for setting the last round in a different way of the first $D - 1$ rounds. For the first $D - 1$ rounds, it is not acceptable to have constant failure probability for each round, since we need to take the union bound over $D - 1$ rounds. But, for the last round, it is acceptable to allow just constant failure probability, since it is just a single round. That’s the reason for setting C in a different way.

We have the following reason for choosing the number of regions(= t) in the last round larger than that of first $D - 1$ rounds. For the first $D - 1$ rounds, we do not need to learn frequency within $\approx \frac{1}{T\rho}$. It is enough to know which region does frequency belong to, although the diameter of the region is large at the beginning. The algorithm is making progress round by round, since the diameter of each region is geometrically decreasing while $\widehat{\beta}$ is geometrically increasing. For the last round, by Lemma C.2, we can learn f within $\Theta(\frac{\Delta l}{\rho st})$. Since after the last round, we hope to learn frequency within $\frac{1}{T\rho}$, thus we need to choose some s, t and Δl such that $\frac{1}{T} \approx \frac{\Delta l}{st}$. To get more accuracy result in the last round, we’d like to choose a larger t . But there is no reason to increase $\widehat{\beta}$ again, since the $\widehat{\beta}$ of the $(D - 1)$ th rounds can tolerance the t we choose at the last round.

To show the constant succeed probability of this Lemma, we separately consider about the failure probability of the first $D - 1$ rounds and the last round. By the union bound, the probability of existing one of the first $D - 1$ rounds is failing is,

$$\begin{aligned} & (D - 1) \left(\left(\frac{4}{s\rho}\right)^{R_{loc}} + t \cdot (60s)^{R_{loc}/2} \right) \\ & \leq D \left(\left(\frac{4}{s\rho}\right)^{R_{loc}} + t \cdot (60s)^{R_{loc}/2} \right) \\ & \leq D \left(\left(\frac{4}{sC}\right)^{R_{loc}} + t \cdot (60s)^{R_{loc}/2} \right) \text{ by } \rho > C > 1 \\ & = D \left(\left(\frac{4}{\sqrt{C}}\right)^{R_{loc}} + t \cdot \left(\frac{60}{\sqrt{C}}\right)^{R_{loc}/2} \right) \text{ by setting } s \approx 1/\sqrt{C} \\ & \leq D \cdot \frac{1}{(Ct)^c} \text{ by setting } R_{loc} = O(\log_C(tC)), \end{aligned}$$

where c is some arbitrarily large constant. Using $t > D$, we can show that failure happening in any of the first $D - 1$ rounds is small. Then, we still need to show that the probability of the last round is failing is also small,

$$\begin{aligned}
& \left(\frac{4}{s\rho}\right)^{R_{loc}} + t \cdot (60s)^{R_{loc}/2} + e^{-\Theta(R_{loc})} \\
& \leq \left(\Theta\left(\frac{C}{\rho}\right)\right)^{R_{loc}} + t \cdot \left(\Theta\left(\frac{1}{C}\right)\right)^{R_{loc}/2} + e^{-\Theta(R_{loc})} \text{ by setting } s \approx 1/C \\
& \leq \frac{1}{c_1} + \frac{1}{(tC)^{c_2}} + e^{-\Theta(R_{loc})} \text{ by setting } R_{loc} = O(\log_C(tC)) \\
& \leq \frac{1}{c_1} + \frac{1}{(tC)^{c_2}} + \frac{1}{c_3},
\end{aligned}$$

where in the first line, the first two terms are from Lemma 3.6 and the third term comes from Lemma C.2; in the last line c_1 , c_2 and c_3 are some arbitrarily large constants.

The expected running time includes the following part: Running **HashtoBins** algorithm $O(DR)$ times, each run takes $O\left(\frac{B}{\alpha} \log \frac{k}{\delta} + B \log B\right)$. Updating the counter v , which takes $O(DR \cdot Bt)$ time. The total running time should be

$$\begin{aligned}
& O\left(DR\left(\frac{B}{\alpha} \log \frac{k}{\delta} + B \log B\right) + (DR_{loc}Bt)\right) \\
& = O(DRB \log(k/\delta \cdot B \cdot FT)) \\
& = O(B \log_C(FT) \log\left(\frac{k}{\delta} BFT\right)) \\
& = O(B \log_C(FT) \log(FT/\delta)) \quad \text{by } FT \gg F \frac{1}{\eta} \gg k.
\end{aligned}$$

The total number of samples is

$$O\left(DR \cdot B \log\left(\frac{k}{\delta}\right)\right) = O(B \log_C(FT) \log(k/\delta)).$$

The sample duration of Algorithm **LocateKSignal** is $O\left(\frac{\log \frac{k}{\delta}}{\eta}\right)$.

In conclusion, we can show that for frequency where neither E_{coll} nor E_{off} holds, we recover an f' with $|f - f'| \lesssim \frac{1}{T\rho}$ as long as $\rho > C$, with an arbitrarily large constant probability. \square

Lemma 3.8. *Algorithm **OneStage** takes $O(k \log_C(FT) \log(k/\delta))$ samples over $O\left(\frac{\log(k/\delta)}{\eta}\right)$ duration, runs in $O(k(\log_C(FT) \log(FT/\delta)))$ time, and outputs a set of $\{(v'_i, f'_i)\}$ of size $O(k)$ with $\min_{i \neq j} |f'_i - f'_j| \gtrsim \eta$. Moreover, one can imagine a subset $S \subseteq [k]$ of “successful” recoveries, where $\Pr[i \in S] \geq \frac{9}{10} \forall i \in [k]$ and for which there exists an injective function $\pi : [k] \rightarrow [O(k)]$ so that*

$$\mathbb{E}_{\sigma, b} \left[\sum_{i \in S} \frac{1}{T} \int_0^T \left| v'_i e^{2\pi f'_i t i} - v_{\pi(i)} e^{2\pi f_{\pi(i)} t i} \right|^2 dt \right] \lesssim C^2 \mathcal{N}^2.$$

with $1 - 1/k^c$ probability for an arbitrarily large constant c .

Proof. Let H denote the set of frequencies f for which neither $E_{coll}(f)$ nor $E_{off}(f)$ holds. For each such f , let $v = \widehat{x^*}(f)$ and denote

$$\mu^2(f) = \mathbb{E}_a [|\widehat{u}_j - v e^{a\sigma 2\pi f i}|^2] \tag{18}$$

and $\rho^2(f) = |v|^2/\mu^2(f)$, as in Lemma 3.7. We have that, if $\rho^2(f) > C^2$, then with an arbitrarily large constant probability one of the recovered $f' \in L$ has $|f' - f| \lesssim \frac{1}{T\rho}$. If this happens, then OneStage will estimate v using $v' = \hat{u}_j e^{-a\sigma 2\pi f' i}$. By triangle inequality,

$$|v' - v|^2 \lesssim |v|^2 |e^{a2\pi\sigma(f'-f)i} - 1|^2 + |\hat{u}_j - v e^{a\sigma 2\pi f i}|^2. \quad (19)$$

Since $a\sigma \leq T$ and $|f' - f| \lesssim \frac{1}{T\rho}$, then the first term of RHS of Equation (19) have

$$|v|^2 |e^{a2\pi\sigma(f'-f)i} - 1|^2 \lesssim |v|^2 |a\sigma(f' - f)|^2.$$

For the second term of RHS of Equation (19). Using Equation (18), we have

$$|\hat{u}_j - v e^{a\sigma 2\pi f i}|^2 \lesssim \mu^2(f),$$

with arbitrarily large constant probability. Combining the bounds for those two terms gives

$$|v' - v|^2 \lesssim |v|^2 |a\sigma(f' - f)|^2 + \mu^2(f),$$

with arbitrarily large constant probability. Since $a\sigma \leq T$, the first term is $|v|^2/\rho^2 = \mu^2$, for

$$|v' - v|^2 \lesssim \mu^2(f).$$

On the other hand, if $\rho^2(f) < C^2$, then $|v| = \rho(f)\mu(f) \lesssim C\mu(f)$ so regardless of the frequency f' recovered, the estimate v' will have

$$|v' - v|^2 \lesssim C^2 \mu^2(f).$$

with arbitrarily large constant probability.

Combining with Lemma 3.1, we get for any $f \in H$ that the recovered f', v' will have

$$\frac{1}{T} \int_0^T |v' e^{2\pi f' t i} - v e^{2\pi f t i}|^2 dt \lesssim C^2 \mu^2(f).$$

with arbitrarily large constant probability. Let $S \subset H$ be the set of frequencies for which this happens. We can choose our permutation π to match frequencies in S to their nearest approximation. By Lemma 3.2, this means that

$$\mathbb{E}_{\sigma, b} \left[\sum_{i \in S} \frac{1}{T} \int_0^T |v'_i e^{2\pi f'_i t i} - v_{\pi(i)} e^{2\pi f_{\pi(i)} t i}|^2 dt \right] \lesssim C^2 \mathcal{N}^2.$$

as desired. □

D Proofs for combining multiple stages

We first prove that the median of a bunch of estimates of a frequency has small error if most of the estimates have small error.

Lemma D.1. *Let (v_i, f_i) be a set of tones for $i \in S$. Define v' and f' to be the (coordinate-wise) median of the (v_i, f_i) . Then for any (v^*, f^*) we have*

$$\frac{1}{T} \int_0^T |v^* e^{2\pi f^* t i} - v' e^{2\pi f' t i}|^2 dt \lesssim \text{median}_i \frac{1}{T} \int_0^T |v^* e^{2\pi f^* t i} - v_i e^{2\pi f_i t i}|^2 dt.$$

Proof. By Lemma 3.1 we have that

$$\frac{1}{T} \int_0^T \left| v^* e^{2\pi f^* t i} - v' e^{2\pi f' t i} \right|^2 dt \approx |v^*|^2 \min(1, T^2 |f^* - f'|^2) + |v^* - v'|^2.$$

Using that v' is taken as a two dimensional median, it suffices to show: if $x^{(1)}, x^{(2)}, \dots \in \mathbb{R}^3$ then $x' = \underset{i}{\text{median}} x^{(i)}$ has

$$\|x'\|_2^2 \lesssim \underset{i}{\text{median}} \|x^{(i)}\|_2^2. \quad (20)$$

This follows because in each of the three coordinates j , we have

$$(x'_j)^2 = (\underset{i}{\text{median}} x_j^{(i)})^2 \leq \underset{i}{\text{median}} (x_j^{(i)})^2 \leq \underset{i}{\text{median}} \|x^{(i)}\|_2^2.$$

so summing over the three coordinates gives (20), as desired.

Therefore, for two dimensional median and one dimension median, we have

$$|v^* - v'|^2 \lesssim \underset{i}{\text{median}} |v^* - v_i|^2 \quad (21)$$

and

$$|f^* - f'|^2 \lesssim \underset{i}{\text{median}} |f^* - f_i|^2.$$

Moreover,

$$\begin{aligned} |v^*|^2 \cdot \min(1, T^2 |f^* - f'|^2) &\lesssim |v^*|^2 \cdot \min(1, T^2 \underset{i}{\text{median}} |f^* - f_i|^2) \\ &= |v^*|^2 \cdot \min(1, \underset{i}{\text{median}} T^2 |f^* - f_i|^2) \\ &= |v^*|^2 \cdot \underset{i}{\text{median}} \min(1, T^2 |f^* - f_i|^2) \end{aligned} \quad (22)$$

Combining Equation (21) and (22), we have

$$\begin{aligned} |v^* - v'|^2 + |v^*|^2 \cdot \min(1, T^2 |f^* - f'|^2) &\lesssim \underset{i}{\text{median}} |v^* - v_i|^2 + \underset{i}{\text{median}} |v^*|^2 \cdot \min(1, T^2 |f^* - f_i|^2) \\ &= \underset{i}{\text{median}} |v^* - v_i|^2 + |v^*|^2 \cdot \min(1, T^2 |f^* - f_i|^2). \end{aligned}$$

which completes the proof. \square

Lemma 3.9. *Repeating algorithm OneStage $O(\log k)$ times, MergedStages returns a set $\{(v'_i, f'_i)\}$ of size $O(k)$ with $\min_{i \neq j} |f'_i - f'_j| \gtrsim \eta$ that can be indexed by π such that*

$$\sum_{i=1}^k \frac{1}{T} \int_0^T \left| v'_i e^{2\pi f'_i t i} - v_{\pi(i)} e^{2\pi f_{\pi(i)} t i} \right|^2 dt \lesssim C^2 \mathcal{N}^2.$$

with probability $1 - 1/k^c$ for an arbitrarily large constant c .

Proof. Our only goal here is to recover the actual tones well, not to worry about spurious tones.

Suppose the number of stages we perform is $R = O(\log k)$. The algorithm for merging the various stages is to scan over a $c\eta$ size region for small c , and take the median (in both frequency and magnitude) over $3c\eta$ region around that $c\eta$ if there are at least $\frac{6}{10}R$ results in that $c\eta$ region.

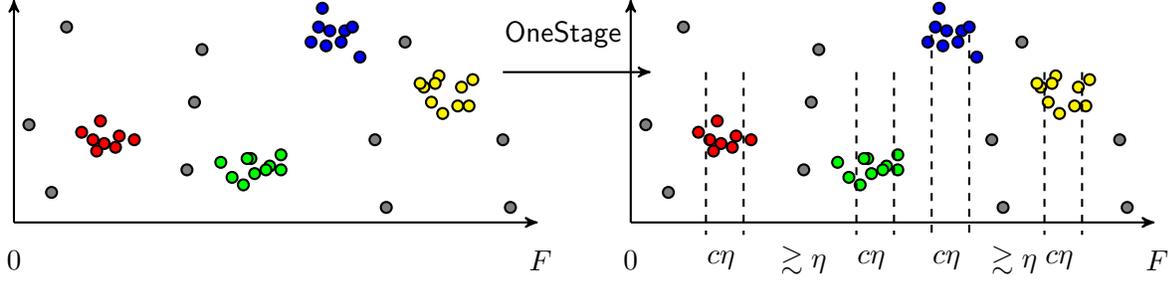


Figure 4: We demonstrate the algorithm for merging various stages ($R = O(\log k)$) on 2-dimensional data. Note that the true data should be 3-dimensional, since for each tone (v_i, f_i) , $v_i \in \mathbb{C}$ and $f_i \in \mathbb{R}$. The x -axis represents the frequency and the y -axis represents the real part of magnitude.

If so, the algorithm will jump to the first right point that is at least η far away from current region and look for the next $c\eta$ region. Because the minimum separation between frequencies for a given stage is $\Omega(\eta)$, this will have minimum separation η in the output, and because there are $O(k)$ tones output at each stage so will this method. What remains is to show that the total error is small.

We say a stage is “good” if the term inside the expectation of Lemma 3.8 is less than 10 times its expectation, as happens with 9/10 probability:

$$\Pr \left(\sum_{i \in S} \frac{1}{T} \int_0^T \left| v'_i e^{2\pi f'_i t i} - v_{\pi(i)} e^{2\pi f_{\pi(i)} t i} \right|^2 dt \leq 10 \mathbb{E}_{\sigma, b} \left[\sum_{i \in S} \frac{1}{T} \int_0^T \left| v'_i e^{2\pi f'_i t i} - v_{\pi(i)} e^{2\pi f_{\pi(i)} t i} \right|^2 dt \right] \right) \geq \frac{9}{10}.$$

We say a frequency f_i is “successful” in a given stage if the stage is good and i lies in S for that stage. Therefore a frequency is successful in each stage with at least 8/10 probability. For sufficiently large $R = O(\log k)$, this will cause all k frequencies to be successful more than $\frac{7}{10}R$ times with high probability. Suppose this happens.

Let $\mu^2(f_i)$ denote the error of (v_i, f_i) . By Lemma 3.8, the total error over all good stages and every successful recovery of a tone in a good stage is $O(C^2 \mathcal{N}^2 R)$. We define $\mu^2(f)$ to be the 6/10 R worst amount of error in the recovery of f over all stages. Because there are $R/10$ worse successful recoveries for each f , we have that $\sum_i \mu^2(f_i) \lesssim C^2 \mathcal{N}^2$.

We will match each f_i with a recovered frequency f'_i with cost at most $\mu^2(f_i)$. If $\mu^2(f_i) \gtrsim |v_i|^2$, then we can set an arbitrary f'_i with $v'_i = 0$. Otherwise, more than $\frac{6}{10}R$ successful recoveries of f_i also yield f' that are within $O(\frac{1}{T}) \ll c\eta$ of f_i . Thus the algorithm for merging tones will find enough tones to report something. What it reports will be the median of at most R values, 6/10 of which have less error than $\mu^2(f)$. Therefore by Lemma D.1 the reported frequency and magnitude will have error $O(\mu^2(f))$. This suffices to get the result. \square

Lemma 3.10. *If we run MergedStages twice and take the tones $\{(v'_i, f'_i)\}$ from the first result that have f'_i within $c\eta$ for small c of some frequency in the second result, we get a set of $k'' = O(k)$ tones that can be indexed by some permutation π such that*

$$\sum_{i=1}^k \frac{1}{T} \int_0^T \left| v'_i e^{2\pi f'_i t i} - v_{\pi(i)} e^{2\pi f_{\pi(i)} t i} \right|^2 dt + \sum_{i=k+1}^{k''} |v'_i|^2 \lesssim C^2 \mathcal{N}^2. \quad (5)$$

Proof. By Lemma 3.9 and Lemma 3.1, the first run gives us a set of $k' = O(k)$ pairs $\{(v'_i, f'_i)\}$ such that they may be indexed by using permutation π such that the first k are a good approximation

to $\{(v_i, f_i)\}$ in the sense that

$$\sum_{i=1}^k (|v_{\pi(i)}|^2 + |v'_i|^2) \cdot \min(1, T^2 |f_{\pi(i)} - f'_i|^2) + |v_{\pi(i)} - v'_i|^2 \lesssim C^2 \mathcal{N}^2.$$

We may as well index to match tones to their nearest match in frequency (because the separation for both f'_i and f_i is at least $\Omega(\eta) > 1/T$). That is, we may index such that $|f'_j - f_i| \gtrsim \eta$ for any $j \neq i$.

Now, for indices where $|f'_i - f_{\pi(i)}| > 1/T$, one would do better by setting the corresponding v_i to zero. In particular, suppose you knew the true frequencies $f_{\pi(i)}$ and only took the (v'_i, f'_i) where f'_i is close to $f_{\pi(i)}$. That is, there exists some permutation π , for the subset $S \subseteq [k]$ containing $\{i : |f_{\pi(i)} - f'_i| \leq c/T\}$ for any $c \gtrsim 1$, we have

$$\sum_{i \in S} (|v_{\pi(i)}|^2 + |v'_i|^2) \cdot \min(1, T^2 |f_{\pi(i)} - f'_i|^2) + |v_{\pi(i)} - v'_i|^2 + \sum_{i \in [k] \setminus S} (|v_{\pi(i)}|^2 + |v'_i|^2) \lesssim C^2 \mathcal{N}^2. \quad (23)$$

The problem is that we do not know the set S , so we can not throw out the other frequencies. However, we can fake it by running the algorithm again on the signal. In particular, we apply Lemma 3.9 again to sparse recovery of the signal defined by

$$\{(v_i, f_i) \mid i \in [k]\} \cup \{(0, f'_i) \mid i \in \{k+1, \dots, k'\}\}.$$

That is, we pretend the “signal” has terms at the other f'_i for $k < i \leq k'$, but with magnitude zero. This is an identical signal in time domain, so it’s really just an analytical tool; for the analysis, it has $k' = O(k)$ sparsity and $\Omega(\eta)$ separation, so Lemma 3.9 applies again and gets a set of $k'' = O(k)$ pairs $\{(v''_j, f''_j)\}$ such that, for the subset $S^* \subset [k']$ containing the indices i where $|f''_j - f'_i| \leq c/T$ for some $j \in [k'']$. In other words, there exists some permutation τ such that for the subset $S^* \subset [k']$ containing $\{i : |f''_{\tau(i)} - f'_i| \leq c/T\}$. We have by analogy to (23) that

$$\begin{aligned} C^2 \mathcal{N}^2 &\gtrsim \sum_{i \in S^* \cap [k]} (|v_{\pi(i)}|^2 + |v''_{\tau(i)}|^2) \cdot \min(1, T^2 |f_{\pi(i)} - f''_{\tau(i)}|^2) + |v_{\pi(i)} - v''_{\tau(i)}|^2 \\ &+ \sum_{i \in S^* \setminus [k]} (|0|^2 + |v''_{\tau(i)}|^2) \cdot \min(1, T^2 |f'_i - f''_{\tau(i)}|^2) + |v''_{\tau(i)} - 0|^2 \\ &+ \sum_{i \in [k] \setminus S^*} (|v_{\pi(i)}|^2 + |v''_{\tau(i)}|^2) + \sum_{i \in [k'] \setminus S^* \setminus [k]} (|v''_{\tau(i)}|^2) \\ &\gtrsim \sum_{i \in S^* \cap [k]} (|v_{\pi(i)}|^2 + |v''_{\tau(i)}|^2) \cdot \min(1, T^2 |f_{\pi(i)} - f''_{\tau(i)}|^2) + |v_{\pi(i)} - v''_{\tau(i)}|^2 \\ &+ \sum_{i \in S^* \setminus [k]} |v''_{\tau(i)}|^2 + \sum_{i \in [k] \setminus S^*} |v_{\pi(i)}|^2, \end{aligned}$$

where π and τ are two permutations and $|S^*| = k^* = O(k)$. This last term is precisely the desired total error for the set of tones $\{(v''_{\tau(i)}, f''_{\tau(i)}) : i \in S^*\}$, giving the result. \square

To prove Theorem 3.11, we still need the following “local” Lemma.

Lemma D.2. For any three tones $(v_{\pi(i)}, f_{\pi(i)})$, (v_i^*, f_i^*) and (v_i', f_i') , if $|v_i'| \geq |v_i^*|$ then,

$$\begin{aligned} & (|v_i'|^2 + |v_{\pi(i)}|^2) \cdot \min(1, T^2|f_i' - f_{\pi(i)}|^2) + |v_i' - v_{\pi(i)}|^2 \\ \lesssim & (|v_i^*|^2 + |v_{\pi(i)}|^2) \cdot \min(1, T^2|f_i^* - f_{\pi(i)}|^2) + |v_i^* - v_{\pi(i)}|^2 + |v_i'|^2. \end{aligned}$$

Proof. First, we can show an upper bound for LHS. Using inequality $\min(1, T^2|f_i' - f_{\pi(i)}|^2) \leq 1$,

$$\text{LHS} \leq |v_i'|^2 + |v_{\pi(i)}|^2 + |v_i' - v_{\pi(i)}|^2.$$

By triangle inequality,

$$|v_i' - v_{\pi(i)}|^2 \leq 2|v_i'|^2 + 2|v_{\pi(i)}|^2.$$

Thus, we obtain,

$$\text{LHS} \lesssim |v_i'|^2 + |v_{\pi(i)}|^2.$$

Second, we can show a lower bound for RHS. Since first term of RHS is nonnegative, then

$$\text{RHS} \geq |v_i^* - v_{\pi(i)}|^2 + |v_i'|^2.$$

Using $|v_i'| \geq |v_i^*|$, we have

$$\text{RHS} \gtrsim |v_i'|^2 + 2|v_i^*|^2 + 2|v_i^* - v_{\pi(i)}|^2.$$

By triangle inequality,

$$2|v_i^*|^2 + 2|v_i^* - v_{\pi(i)}|^2 \geq |v_{\pi(i)}|^2.$$

Then, we prove the lower bound for RHS,

$$\text{RHS} \gtrsim |v_i'|^2 + |v_{\pi(i)}|^2.$$

Combining the lower bound of RHS and the upper bound of LHS completes the proof. \square

By plugging Lemma 3.1 into Lemma D.2, we have

Corollary D.3. For any three tones $(v_{\pi(i)}, f_{\pi(i)})$, (v_i^*, f_i^*) and (v_i', f_i') , if $|v_i'| \geq |v_i^*|$, then

$$\frac{1}{T} \int_0^T \left| v_i' e^{2\pi f_i' t i} - v_{\pi(i)} e^{2\pi f_{\pi(i)} t i} \right|^2 dt \lesssim \frac{1}{T} \int_0^T \left| v_i^* e^{2\pi f_i^* t i} - v_{\pi(i)} e^{2\pi f_{\pi(i)} t i} \right|^2 dt + |v_i'|^2.$$

We use Corollary D.3 to present the proof of Theorem 3.11,

Theorem 3.11. Algorithm `ContinuousFourierSparseRecovery` returns a set $\{(v_i', f_i')\}$ of size k with $\min_{i \neq j} |f_i' - f_j'| \gtrsim \eta$ for which

$$\sum_{i=1}^k \frac{1}{T} \int_0^T \left| v_i' e^{2\pi f_i' t i} - v_{\pi(i)} e^{2\pi f_{\pi(i)} t i} \right|^2 dt \lesssim C^2 \mathcal{N}^2$$

with probability $1 - 1/k^c$ for an arbitrarily large constant c .

Proof. Let $\{(v_i^*, f_i^*)\}_{i=1,2,\dots,k''}$ denote the set of tones returned by Lemma 3.10, where $k'' = O(k)$ and $k'' > k$. For any $i \in [k]$, tone (v_i^*, f_i^*) was mapped to tone $(v_{\pi(i)}, f_{\pi(i)})$ in Lemma 3.10. Let $\{(v'_i, f'_i)\}_{i=1,2,\dots,k}$ denote a subset of $\{(v_i^*, f_i^*)\}_{i=1,2,\dots,k''}$ that satisfies the following two conditions (1) for any $i \in [k]$, $|v'_i|$ is one of the top- k largest magnitude tones; (2) for any $i \in [k]$, $|v'_i| \geq |v_i^*|$. By Lemma 3.10, we also know that $\min_{i \neq j} |f'_i - f'_j| \gtrsim \eta$.

For any $i \in [k]$, we consider these three tones $(v_{\pi(i)}, f_{\pi(i)})$, (v_i^*, f_i^*) , (v'_i, f'_i) . If $(v_i^*, f_i^*) \neq (v'_i, f'_i)$, then applying Corollary D.3 we have

$$\frac{1}{T} \int_0^T \left| v'_i e^{2\pi f'_i t i} - v_{\pi(i)} e^{2\pi f_{\pi(i)} t i} \right|^2 dt \lesssim \frac{1}{T} \int_0^T \left| v_i^* e^{2\pi f_i^* t i} - v_{\pi(i)} e^{2\pi f_{\pi(i)} t i} \right|^2 dt + |v'_i|^2.$$

Otherwise $(v_i^*, f_i^*) = (v'_i, f'_i)$, we also have

$$\frac{1}{T} \int_0^T \left| v'_i e^{2\pi f'_i t i} - v_{\pi(i)} e^{2\pi f_{\pi(i)} t i} \right|^2 dt = \frac{1}{T} \int_0^T \left| v_i^* e^{2\pi f_i^* t i} - v_{\pi(i)} e^{2\pi f_{\pi(i)} t i} \right|^2 dt,$$

which means there exists some universal constant α such that $\forall i \in [k]$, if $(v_i^*, f_i^*) = (v'_i, f'_i)$ then

$$\frac{1}{T} \int_0^T \left| v'_i e^{2\pi f'_i t i} - v_{\pi(i)} e^{2\pi f_{\pi(i)} t i} \right|^2 dt \leq \alpha \cdot \left(\frac{1}{T} \int_0^T \left| v_i^* e^{2\pi f_i^* t i} - v_{\pi(i)} e^{2\pi f_{\pi(i)} t i} \right|^2 dt + |v'_i|^2 \right)$$

holds. Otherwise $(v_i^*, f_i^*) \neq (v'_i, f'_i)$, then

$$\frac{1}{T} \int_0^T \left| v'_i e^{2\pi f'_i t i} - v_{\pi(i)} e^{2\pi f_{\pi(i)} t i} \right|^2 dt = \alpha \frac{1}{T} \int_0^T \left| v_i^* e^{2\pi f_i^* t i} - v_{\pi(i)} e^{2\pi f_{\pi(i)} t i} \right|^2 dt$$

holds. Let S denote the set of indices i such that $(v_i^*, f_i^*) \neq (v'_i, f'_i)$. Taking the summation from $i = 1$ to $i = k$,

$$\begin{aligned} \sum_{i=1}^k \frac{1}{T} \int_0^T \left| v'_i e^{2\pi f'_i t i} - v_{\pi(i)} e^{2\pi f_{\pi(i)} t i} \right|^2 dt &\leq \sum_{i \in S} \alpha \cdot \left(\frac{1}{T} \int_0^T \left| v_i^* e^{2\pi f_i^* t i} - v_{\pi(i)} e^{2\pi f_{\pi(i)} t i} \right|^2 dt + |v'_i|^2 \right) \\ &\quad + \sum_{i \in [k] \setminus S} \alpha \cdot \left(\frac{1}{T} \int_0^T \left| v_i^* e^{2\pi f_i^* t i} - v_{\pi(i)} e^{2\pi f_{\pi(i)} t i} \right|^2 dt \right), \end{aligned}$$

furthermore, we have

$$\sum_{i=1}^k \frac{1}{T} \int_0^T \left| v'_i e^{2\pi f'_i t i} - v_{\pi(i)} e^{2\pi f_{\pi(i)} t i} \right|^2 dt \leq \alpha \sum_{i=1}^k \left(\frac{1}{T} \int_0^T \left| v_i^* e^{2\pi f_i^* t i} - v_{\pi(i)} e^{2\pi f_{\pi(i)} t i} \right|^2 dt \right) + \alpha \sum_{i \in S} |v'_i|^2.$$

To finish the proof, we need to show that $\sum_{i \in S} |v'_i|^2 \leq \sum_{i=k+1}^{k''} |v_i^*|^2$. The point is, for any $i \in S$, we know that $(v'_i, f'_i) \neq (v_i^*, f_i^*)$ which implies that $(v'_i, f'_i) \notin \{(v_i^*, f_i^*)\}_{i=1}^k$. Thus,

$$\begin{aligned} &\sum_{i=1}^k \frac{1}{T} \int_0^T \left| v'_i e^{2\pi f'_i t i} - v_{\pi(i)} e^{2\pi f_{\pi(i)} t i} \right|^2 dt \\ &\leq \alpha \sum_{i=1}^k \left(\frac{1}{T} \int_0^T \left| v_i^* e^{2\pi f_i^* t i} - v_{\pi(i)} e^{2\pi f_{\pi(i)} t i} \right|^2 dt \right) + \alpha \sum_{i=k+1}^{k''} |v_i^*|^2 \\ &\lesssim \sum_{i=1}^k \left(\frac{1}{T} \int_0^T \left| v_i^* e^{2\pi f_i^* t i} - v_{\pi(i)} e^{2\pi f_{\pi(i)} t i} \right|^2 dt \right) + \sum_{i=k+1}^{k''} |v_i^*|^2 \\ &\lesssim C^2 \mathcal{N}^2, \end{aligned}$$

where $k'' = O(k)$ is defined in Lemma 3.10 and the last inequality follows by using Equation (5) in Lemma 3.10. □

E Proofs for converting (2) into (3)

In this section, we show that as long as the sample duration T is sufficiently large, it is possible to convert Equation (2) to Equation (3). First, we show an auxiliary lemma, Lemma E.3, which bounds an integral that will appear in the analysis.

We will show that

$$\int_{-\infty}^{+\infty} \min(T, \frac{1}{|f_i - \theta|}) \cdot \min(T, \frac{1}{|f_j - \theta|}) d\theta \lesssim \frac{\log(T|f_i - f_j|)}{|f_i - f_j|}.$$

for $f_j - f_i \geq 2/T$. We split this into two pieces.

Claim E.1. *Given two frequencies f_i, f_j and $f_j - f_i \geq \frac{2}{T}$, we have*

$$\int_{f_i - \frac{1}{T}}^{f_i} \min(T, \frac{1}{|f_i - \theta|}) \cdot \frac{1}{|f_j - \theta|} d\theta \lesssim \frac{1}{f_j - f_i}.$$

Proof. By $f_i - \frac{1}{T} < \theta < f_i$, we have

$$\text{LHS} = \int_{f_i - \frac{1}{T}}^{f_i} T \cdot \frac{1}{f_j - \theta} d\theta.$$

Since $\frac{1}{f_j - \theta} \approx \frac{1}{f_j - f_i}$ for all $\theta \in [f_i - \frac{1}{T}, f_i]$,

$$\text{LHS} \lesssim \int_{f_i - \frac{1}{T}}^{f_i} \frac{T}{f_j - f_i} d\theta = \frac{1}{f_j - f_i}.$$

□

Claim E.2. *Given two frequencies f_i, f_j and $f_j - f_i \geq \frac{2}{T}$, we have*

$$\int_{-\infty}^{f_i - \frac{1}{T}} \min(T, \frac{1}{|f_i - \theta|}) \cdot \frac{1}{|f_j - \theta|} d\theta \lesssim \frac{\log(T|f_j - f_i|)}{f_j - f_i}.$$

Proof. By $\theta < f_i - \frac{1}{T} < f_j$, we have that

$$\begin{aligned}
\text{LHS} &= \int_{-\infty}^{f_i - \frac{1}{T}} \frac{1}{f_i - \theta} \cdot \frac{1}{f_j - \theta} d\theta \\
&= \frac{1}{f_j - f_i} \int_{-\infty}^{f_i - \frac{1}{T}} \frac{f_j - f_i}{(f_i - \theta)(f_j - \theta)} d\theta \\
&= \frac{1}{f_j - f_i} \int_{-\infty}^{f_i - \frac{1}{T}} \frac{1}{f_i - \theta} - \frac{1}{f_j - \theta} d\theta \\
&= -\frac{1}{f_j - f_i} \log \frac{f_i - f_i + \frac{1}{T}}{f_j - f_i + \frac{1}{T}} \\
&= -\frac{1}{f_j - f_i} \log \frac{1}{T(f_j - f_i) + 1} \\
&\lesssim \frac{\log(T(f_j - f_i))}{f_j - f_i}.
\end{aligned}$$

□

Lemma E.3. *Given two frequencies f_i, f_j and $f_j - f_i \geq \frac{2}{T}$, we have*

$$\int_{-\infty}^{+\infty} \min(T, \frac{1}{|f_i - \theta|}) \cdot \min(T, \frac{1}{|f_j - \theta|}) d\theta \lesssim \frac{\log(T|f_i - f_j|)}{|f_i - f_j|}.$$

Proof. By symmetry, we have

$$\text{LHS} = 2 \int_{-\infty}^{\frac{f_i + f_j}{2}} \min(T, \frac{1}{|f_i - \theta|}) \cdot \min(T, \frac{1}{|f_j - \theta|}) d\theta.$$

Since $T > \frac{1}{f_j - \theta}$ when $\theta < \frac{f_i + f_j}{2}$,

$$\text{LHS} \leq 2 \int_{-\infty}^{\frac{f_i + f_j}{2}} \min(T, \frac{1}{|f_i - \theta|}) \cdot \frac{1}{|f_j - \theta|} d\theta.$$

We also observe that $\frac{1}{|f_j - \theta|} \approx \frac{1}{|f_j - f_i|}$ for all $\theta \in [f_i - \frac{f_j - f_i}{2}, f_i + \frac{f_j - f_i}{2}]$,

$$\int_{f_i - \frac{f_j - f_i}{2}}^{\frac{f_i + f_j}{2}} \min(T, \frac{1}{|f_i - \theta|}) \cdot \frac{1}{|f_j - \theta|} d\theta \approx \int_{f_i - \frac{f_j - f_i}{2}}^{f_i} \min(T, \frac{1}{|f_i - \theta|}) \cdot \frac{1}{|f_j - \theta|} d\theta.$$

Thus, we get

$$\text{LHS} \lesssim \int_{-\infty}^{f_i} \min(T, \frac{1}{|f_i - \theta|}) \cdot \frac{1}{|f_j - \theta|} d\theta.$$

Plugging Claim E.1 and E.2 into the above formula completes the proof. □

Lemma E.4. *For any i , let $a_i(t) = v_i e^{2\pi f_i t i} - v'_i e^{2\pi f'_i t i}$, then for $i \neq j$,*

$$\frac{1}{T} \int_0^T a_i(t) \overline{a_j(t)} dt \lesssim \frac{\log(\Delta f_{i,j} T)}{\Delta f_{i,j} T} \cdot \left(\frac{1}{T} \int_0^T |a_i(t)|^2 dt \cdot \frac{1}{T} \int_0^T |a_j(t)|^2 dt \right)^{\frac{1}{2}}, \quad (24)$$

where $\Delta f_{i,j} = \min(|f_i - f_j|, |f_i - f'_j|, |f'_i - f_j|, |f_i - f'_j|)$.

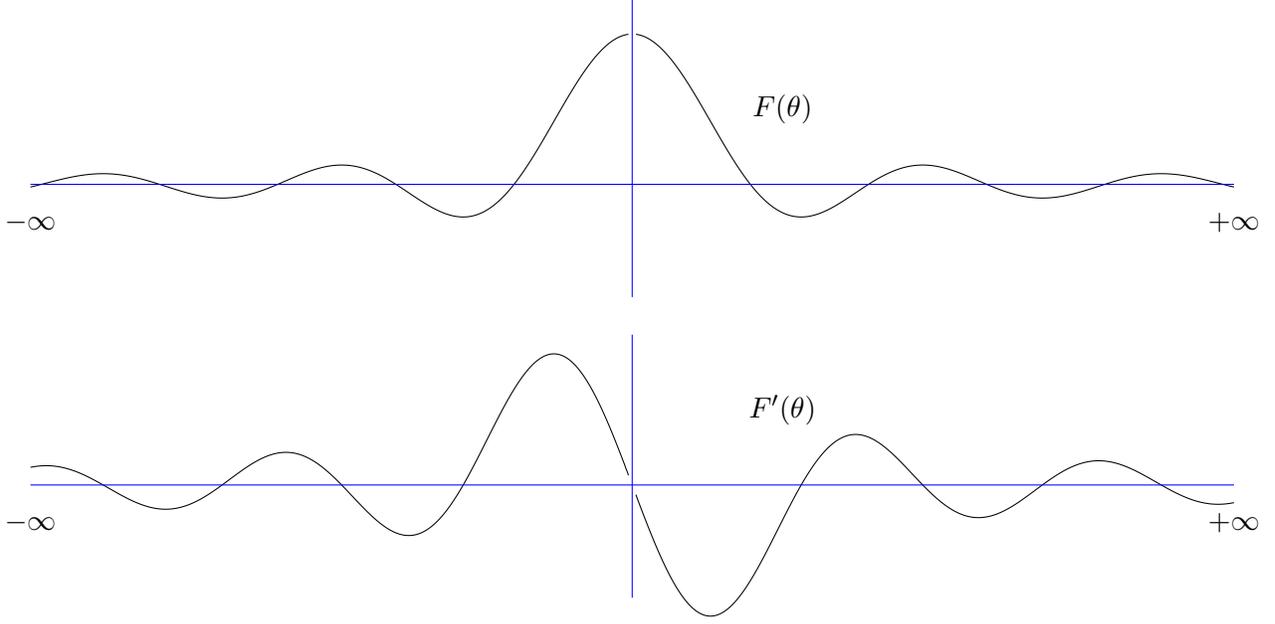


Figure 5: $F(\theta)$ is a sinc function and has derivate function $F'(\theta)$.

Proof. Let $\nu_i = |f_i - f'_i|$ and $\nu_j = |f_j - f'_j|$. Define $\|a_i\| = \sqrt{\frac{1}{T} \int_0^T |a_i(t)|^2 dt}$. We define $f(t)$ and $F(\theta)$ to be a rectangle and sinc function respectively:

$$f(t) = \begin{cases} 1 & \text{if } 0 \leq t \leq T \\ 0 & \text{otherwise} \end{cases}$$

$$F(\theta) = \frac{\sin(2\pi\theta T)}{2\pi\theta}$$

where $F = \widehat{f}$.

$F(\theta)$ has the derivative,

$$F'(\theta) = \frac{2\pi\theta T \cos(2\pi\theta T) - \sin(2\pi\theta T)}{2\pi\theta^2},$$

which means that

$$|F'(\theta)| \lesssim \begin{cases} T^2 & \text{if } \theta \leq 1/T \\ T/|\theta| & \text{otherwise} \end{cases}$$

Let $y_i(t) = a_i(t) \cdot f(t)$, then

$$\begin{aligned} \widehat{y}_i(\theta) &= \widehat{a}_i(\theta) * \widehat{f}(\theta) \\ &= \widehat{a}_i(\theta) * F(\theta) \text{ by } F = \widehat{f} \\ &= v_i F(f_i - \theta) - v'_i F(f'_i - \theta) \\ &= (v_i - v'_i) F(f_i - \theta) + v'_i (f_i - f'_i) \cdot F'(x - \theta) \text{ some } x \in [f_i, f'_i]. \end{aligned}$$

We split into two cases. First, if $\nu_i \leq \frac{1}{T}$, then

$$\begin{aligned} |\widehat{y}_i(\theta)| &\lesssim (|v_i - v'_i| + \nu_i T |v'_i|) \cdot \min\left(T, \frac{1}{|f_i - \theta|}\right) \\ &\lesssim \|a_i\| \cdot \min\left(T, \frac{1}{|f_i - \theta|}\right) \text{ by Lemma 3.1,} \end{aligned} \tag{25}$$

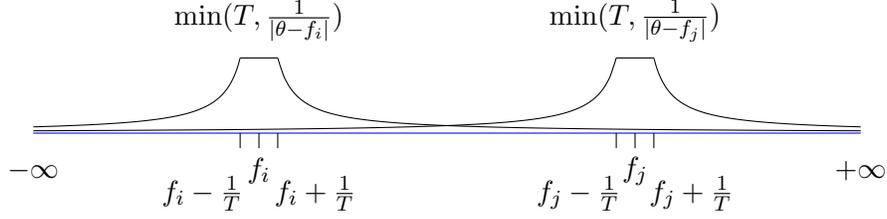


Figure 6: $\int_{-\infty}^{+\infty} \min(T, \frac{1}{|\theta - f_i|}) \cdot \min(T, \frac{1}{|\theta - f_j|}) d\theta$

where the first line holds for both $f_i > f'_i$ and $f_i \leq f'_i$ since the triangle inequality. Therefore

$$\begin{aligned}
& \frac{1}{T} \int_0^T a_i(t) \overline{a_j(t)} dt \\
&= \frac{1}{T} \int_{-\infty}^{\infty} y_i(t) \overline{y_j(t)} dt \text{ by } y_i(t) = a_i(t) \cdot f(t) \text{ and } f(t) = 1 \text{ if } t \in [0, T] \\
&= \frac{1}{T} \int_{-\infty}^{\infty} \widehat{y}_i(\theta) \overline{\widehat{y}_j(\theta)} d\theta \text{ by the property of Fourier Transform} \\
&\lesssim \frac{1}{T} \|a_i\| \|a_j\| \underbrace{\int_{-\infty}^{+\infty} \min(T, \frac{1}{|f_i - \theta|}) \cdot \min(T, \frac{1}{|f_j - \theta|}) d\theta}_C \text{ by Equation (25)}. \tag{26}
\end{aligned}$$

Using Lemma E.3, we have following bound for term C ,

$$\int_{-\infty}^{+\infty} \min(T, \frac{1}{|f_i - \theta|}) \cdot \min(T, \frac{1}{|f_j - \theta|}) d\theta \lesssim \frac{\log T |f_j - f_i|}{|f_j - f_i|}.$$

This gives the result for $\nu_i \leq \frac{1}{T}$. In the alternate case, we have $\nu_i > \frac{1}{T}$, then

$$\begin{aligned}
|\widehat{y}_i(\theta)| &\lesssim \nu_i \cdot \min(T, \frac{1}{|f_i - \theta|}) + \nu'_i \cdot \min(T, \frac{1}{|f'_i - \theta|}) \\
&\lesssim \|a_i\| \cdot \left(\min(T, \frac{1}{|f_i - \theta|}) + \min(T, \frac{1}{|f'_i - \theta|}) \right) \text{ by Lemma 3.1.}
\end{aligned}$$

By similar reason for Equation (26), we have

$$\begin{aligned}
& \frac{1}{T} \int_0^T a_i(t) \overline{a_j(t)} dt \\
\lesssim & \frac{1}{T} \|a_i\| \|a_j\| \underbrace{\int_{-\infty}^{+\infty} \min(T, \frac{1}{|f_i - \theta|}) \cdot \min(T, \frac{1}{|f_j - \theta|}) d\theta}_{C_1} \\
& + \frac{1}{T} \|a_i\| \|a_j\| \underbrace{\int_{-\infty}^{+\infty} \min(T, \frac{1}{|f_i - \theta|}) \cdot \min(T, \frac{1}{|f'_j - \theta|}) d\theta}_{C_2} \\
& + \frac{1}{T} \|a_i\| \|a_j\| \underbrace{\int_{-\infty}^{+\infty} \min(T, \frac{1}{|f'_i - \theta|}) \cdot \min(T, \frac{1}{|f_j - \theta|}) d\theta}_{C_3} \\
& + \frac{1}{T} \|a_i\| \|a_j\| \underbrace{\int_{-\infty}^{+\infty} \min(T, \frac{1}{|f'_i - \theta|}) \cdot \min(T, \frac{1}{|f'_j - \theta|}) d\theta}_{C_4}.
\end{aligned}$$

Applying Lemma E.3 on the term C_1 , C_2 , C_3 and C_4 respectively,

$$\frac{1}{T} \int_0^T a_i(t) \overline{a_j(t)} \lesssim \|a_i\| \|a_j\| \frac{\log(T \Delta f_{i,j})}{\Delta f_{i,j}},$$

where $\Delta f_{i,j} = \min(|f_i - f_j|, |f_i - f'_j|, |f'_i - f_j|, |f'_i - f'_j|)$. □

Lemma E.5. *Let $\{(v_i, f_i)\}$ and $\{(v'_i, f'_i)\}$ be two sets of k tones for which $\min_{i \neq j} |f_i - f_j| \geq \eta$ and $\min_{i \neq j} |f'_i - f'_j| \geq \eta$ for some $\eta > 0$. Suppose that $T > C/\eta$ for a sufficiently large constant C . Then these sets can be indexed such that*

$$\frac{1}{T} \int_0^T \left| \sum_{i=1}^k (v'_i e^{2\pi i f'_i t} - v_i e^{2\pi i f_i t}) \right|^2 dt \leq (1 + O(\frac{\log(k\eta T) \log(k)}{\eta T})) \sum_{i=1}^k \frac{1}{T} \int_0^T |v'_i e^{2\pi i f'_i t} - v_i e^{2\pi i f_i t}|^2 dt. \tag{27}$$

Proof. For simplicity, let $a_i(t) = v_i e^{2\pi i f_i t} - v'_i e^{2\pi i f'_i t}$. Let's express the square of summations by diagonal term and off-diagonal term, and then bound them separately.

$$\begin{aligned}
& \int_0^T \left| \sum_{i=1}^k a_i(t) \right|^2 dt \\
&= \int_0^T \left(\sum_{i=1}^k a_i(t) \right) \left(\sum_{i=1}^k \overline{a_i(t)} \right) dt \\
&= \int_0^T \underbrace{\sum_{i=1}^k a_i(t) \overline{a_i(t)}}_{\text{diagonal}} + \underbrace{\sum_{i \neq j} a_i(t) \overline{a_j(t)}}_{\text{off-diagonal}} dt. \tag{28}
\end{aligned}$$

Using the result of Lemma E.4 and $a^2 + b^2 \geq 2ab$, we can upper bound the off-diagonal term,

$$\begin{aligned}
& \int_0^T a_i(t) \overline{a_j(t)} dt \\
& \lesssim \frac{\log(\Delta f_{i,j} T)}{\Delta f_{i,j} T} \cdot \sqrt{\int_0^T |a_i(t)|^2 dt \int_0^T |a_j(t)|^2 dt} \quad \text{by Lemma E.4} \\
& \lesssim \frac{\log(\Delta f_{i,j} T)}{\Delta f_{i,j} T} \cdot \left(\int_0^T |a_i(t)|^2 dt + \int_0^T |a_j(t)|^2 dt \right) \quad \text{by } 2ab \leq a^2 + b^2.
\end{aligned}$$

where $\Delta f_{i,j} = \min(|f_i - f_j|, |f_i - f'_j|, |f'_i - f_j|, |f'_i - f'_j|)$. If we index such that any f_i is matched with any f'_j with $|f_i - f'_j| < \eta/3$ – which is possible, since at most one such f'_j will exist by the separation among the f'_j , and that f'_j will be within $\eta/3$ of at most one f_i – then we have $\Delta f_{i,j} \gtrsim |f_i - f_j|$. If we order the f_i in increasing order, then in fact $\Delta f_{i,j} \gtrsim \eta|i - j|$.

If $T > C/\eta$ for a sufficiently large constant C , this means that $\Delta f_{i,j} T \gtrsim |i - j|\eta T \geq e$. Since $\frac{\log x}{x}$ is decreasing on the region, this implies

$$\frac{\log(\Delta f_{i,j} T)}{\Delta f_{i,j} T} \lesssim \frac{\log(|i - j|\eta T)}{|i - j|\eta T}.$$

Thus, we have

$$\int_0^T a_i(t) \overline{a_j(t)} dt \lesssim \frac{\log(|i - j|\eta T)}{|i - j|\eta T} \cdot \left(\int_0^T |a_i(t)|^2 dt + \int_0^T |a_j(t)|^2 dt \right). \quad (29)$$

Finally, we have

$$\begin{aligned}
& \int_0^T \left| \sum_{i=1}^k a_i(t) \right|^2 dt - \sum_{i=1}^k \int_0^T |a_i(t)|^2 dt \\
& = \sum_{i \neq j}^k \int_0^T a_i(t) \overline{a_j(t)} dt \quad \text{by Equation (28)} \\
& \lesssim \sum_{i \neq j}^k \frac{\log(|i - j|\eta T)}{|i - j|\eta T} \cdot \left(\int_0^T |a_i(t)|^2 dt + \int_0^T |a_j(t)|^2 dt \right) \quad \text{by Equation (29)} \\
& \leq \frac{\log(k\eta T)}{\eta T} \sum_{i=1}^k \sum_{j \neq i}^k \frac{1}{|i - j|} \cdot \left(\int_0^T |a_i(t)|^2 dt + \int_0^T |a_j(t)|^2 dt \right) \\
& = 2 \frac{\log(k\eta T)}{\eta T} \sum_{i=1}^k \sum_{j \neq i}^k \frac{1}{|i - j|} \int_0^T |a_i(t)|^2 dt \quad \text{by symmetry} \\
& \lesssim \frac{\log(k\eta T)}{\eta T} \sum_{i=1}^k \int_0^T |a_i(t)|^2 dt \sum_{j \neq i}^k \frac{1}{|i - j|} \\
& \lesssim \frac{\log(k\eta T) \log(k)}{\eta T} \sum_{i=1}^k \int_0^T |a_i(t)|^2 dt \quad \text{by } \sum_{i=1}^k \frac{1}{i} \approx \log(k).
\end{aligned}$$

Thus, we complete the proof. \square

Lemma 3.13. *Let $\{(v_i, f_i)\}$ and $\{(v'_i, f'_i)\}$ be two sets of k tones for which $\min_{i \neq j} |f_i - f_j| \geq \eta$ and $\min_{i \neq j} |f'_i - f'_j| \gtrsim \eta$ for some $\eta > 0$. Suppose that $T > O(\frac{\log^2 k}{\eta})$. Then these sets can be indexed such that*

$$\frac{1}{T} \int_0^T \left| \sum_{i=1}^k (v'_i e^{2\pi i f'_i t} - v_i e^{2\pi i f_i t}) \right|^2 dt \lesssim \sum_{i=1}^k \frac{1}{T} \int_0^T |v'_i e^{2\pi i f'_i t} - v_i e^{2\pi i f_i t}|^2 dt. \quad (30)$$

Proof. Directly follows by Lemma E.5. □

F Lower Bound

Lemma 3.15. *There exists a constant $c > 0$ such that, for a given sample duration T , one cannot recover the frequency f to within*

$$\frac{\mathcal{N}}{c T |\hat{x}^*(f)|}$$

with 3/4 probability, for all $\delta > 0$, even if $k = 1$.

Proof. Suppose this were possible, and consider two one-sparse signals y and y' containing tones (v, f) and (v, f') , respectively. By Lemma 3.1,

$$\int_0^T |y(t) - y'(t)|^2 dt \lesssim |v|^2 T^2 |f - f'|^2.$$

Consider recovery of the signal $x(t) = y(t)$, and suppose it outputs some frequency f^* . This must simultaneously be a good recovery for the decomposition $(x^*, g) = (y, 0)$ and $(x^*, g) = (y', y - y')$. These have noise levels \mathcal{N}^2 bounded by $\delta |v|^2$ and $\delta |v|^2 + O(|v|^2 T^2 |f - f'|^2)$, respectively. By the assumption of good recovery, and the triangle inequality, we require

$$c \frac{2\sqrt{\delta |v|^2} + \sqrt{O(|v|^2 T^2 |f - f'|^2)}}{T v} \gtrsim |f - f'|$$

or

$$c \cdot O\left(\sqrt{\frac{\delta}{T |f - f'|}} + 1\right) \geq 1.$$

Because δ may be chosen arbitrarily small, we can choose a small constant c such that this is a contradiction. □

Algorithm 1 Continuous Fourier Sparse Recovery

```
1: procedure ContinuousFourierSparseRecovery( $x, k, \delta, \alpha, C, F, T, \eta$ ) — The Main Algorithm
2:    $F$  is the upper bound of frequency,  $T$  is the sample duration,  $C$  is the approximation factor.
3:    $\delta, \alpha$  are the parameters associated with Hash function.
4:    $c = 1/10$  and  $b = 8/10$ 
5:    $R_1 \leftarrow$  NoisyKSparseCFFT( $x, k, \delta, \alpha, C, F$ )
6:    $S_1 \leftarrow$  MergedStages( $R_1, O(k \log k), \eta, c, b$ )
7:    $R_2 \leftarrow$  NoisyKSparseCFFT( $x, O(k), \delta, \alpha, C, F, T$ )
8:    $S_2 \leftarrow$  MergedStages( $R_2, O(k \log k), \Omega(\eta), c, b$ )
9:    $S \leftarrow S_1 \cap S_2$ , which means only keeping the tones that  $S_1$  agrees with  $S_2$  by Lemma 3.10.
10:   $S^* \leftarrow$  Prune( $S, k$ ), which means only keeping the top- $k$  largest magnitude tones.
11:  return  $S^*$ 
12: end procedure
13: procedure NoisyKSparseCFFT( $x, k, \delta, \alpha, C, F, T$ )
14:   Let  $B = k/\epsilon$ .
15:   for  $c = 1 \rightarrow \log(k)$  do
16:     Choose  $\sigma$  uniformly at random from  $[\frac{1}{B\eta}, \frac{2}{B\eta}]$ .
17:     Choose  $b$  uniformly at random from  $[0, \frac{2\pi[F/\eta]}{\sigma B}]$ .
18:      $R_c \leftarrow$  OneStage( $x, B, \delta, \alpha, \sigma, b, C, F, T$ )
19:   end for
20:   return ( $R_1, R_2, \dots, R_{\log(k)}$ ).
21: end procedure
22: procedure MergedStages( $R, m, \eta, c, b$ )
23:    $R$  is a list of  $m$  tones ( $v'_i, f'_i$ )
24:    $c$  is some constant  $< 1$ .
25:    $b$  is some constant  $< 1$ .
26:   Sort list  $R$  based on  $f'_i$ .
27:   Building the 1D range search Tree based on  $m$  points by regarding each frequency  $f'_i$  as a
   1D point on a line where  $x_i = f'_i$ .
28:    $S \leftarrow \emptyset, i \leftarrow 0$ 
29:   while  $i < m$  do
30:     if Tree.Count( $f'_i, f'_i + c\eta$ )  $\geq b \log k$  then
31:        $f \leftarrow$  median  $\{ f'_j \mid f'_j \in [f'_i - c\eta, f'_i + 2c\eta] \}$ 
32:        $v \leftarrow$  median  $\{ v'_j \mid f'_j \in [f'_i - c\eta, f'_i + 2c\eta] \}$ 
33:        $S \leftarrow S \cup (f, v)$ 
34:        $i \leftarrow$  Tree.Search( $f'_i + 2c\eta + \eta/2$ ), which means walk to the first point that is on the
       right of  $f'_i + 2c\eta + \eta/2$ 
35:     else
36:        $i \leftarrow i + 1$ 
37:     end if
38:   end while
39:   return  $S$ 
40: end procedure
```

Algorithm 2 Continuous Fourier Sparse Recovery

```

1: procedure HashToBins( $x, P_{\sigma,a,b}, B, \delta, \alpha$ )
2:   Compute  $\widehat{y}_{jF/B}$  for  $j \in [B]$ , where  $y = G_{B,\alpha,\delta} \cdot (P_{\sigma,a,b}x)$ 
3:   return  $\widehat{u}$  given by  $\widehat{u}_j = \widehat{y}_{jF/B}$ 
4: end procedure
5: procedure OneStage( $x, B, \delta, \alpha, \sigma, b, C, F, T$ )
6:    $L \leftarrow \text{LocateKSignal}(x, B, \delta, \alpha, \sigma, b, C, F, T)$ 
7:   Choose  $a \in [0, 1]$  uniformly at random.
8:    $\widehat{u} \leftarrow \text{HashToBins}(x, P_{\sigma,a,b}, B, \delta, \alpha)$ 
9:   return  $\{(\widehat{u}_{h_{\sigma,b}(f')})e^{-2\pi\sigma a f' \mathbf{i}}, f'\}$  for  $f' \in L$  if not  $E_{\text{off}}(f')$ .
10: end procedure
11: procedure LocateKSignal( $x, B, \delta, \alpha, \sigma, b, C, F, T$ )
12:   Set  $t \approx \log(FT)$ ,  $t' = t/(c_n + 1)$ ,  $D \approx \log_{t'}(FT)$ ,  $R_{\text{loc}} \approx \log_C(tC)$ ,  $l^{(1)} = F/2$ .
13:   for  $i \in [D - 1]$  do
14:      $\Delta l \approx F/(t')^{i-1}$ ,  $s = \frac{1}{\sqrt{C}}$ ,  $\widehat{\beta} = \frac{ts}{2\sigma\Delta l}$ 
15:      $l^{(i+1)} \leftarrow \text{LocateInner}(x, B, \delta, \alpha, \sigma, b, \widehat{\beta}, l^{(i)}, \Delta l, t, R_{\text{loc}}, \text{false})$ .
16:   end for
17:   Set  $s = 1/C$ ,  $t \approx \log(FT)/s$ ,  $\Delta l \approx st/T$ ,  $\widehat{\beta} = \frac{ts}{2\sigma\Delta l}$ ,  $R_{\text{loc}} \approx \log_C(tC)$ 
18:    $l^{(*)} \leftarrow \text{LocateInner}(x, B, \delta, \alpha, \sigma, b, \widehat{\beta}, l^{(D)}, \Delta l, t, R_{\text{loc}}, \text{true})$ .
19:   return  $l^{(*)}$ .
20: end procedure
21: procedure LocateInner( $x, B, \delta, \sigma, b, \widehat{\beta}, l, \Delta l, t, R_{\text{loc}}, \text{last}$ )
22:   Let  $v_{j,q} = 0$  for  $(j, q) \in [B] \times [t]$ .
23:   for  $r \in [R_{\text{loc}}]$  do
24:     Choose  $\gamma \in [\frac{1}{2}, 1]$  uniformly at random.
25:     Choose  $\beta \in [\frac{1}{2}\widehat{\beta}, 1\widehat{\beta}]$  uniformly at random.
26:      $\widehat{u} \leftarrow \text{HashToBins}(x, P_{\sigma,\gamma,b}, B, \delta, \alpha)$ .
27:      $\widehat{u}' \leftarrow \text{HashToBins}(x, P_{\sigma,\gamma+\beta,b}, B, \delta, \alpha)$ .
28:     for  $j \in [B]$  do
29:       for  $i \in [m]$  do
30:          $\theta_{j,i}^r = \frac{1}{2\pi\sigma\beta}(\phi(\widehat{u}_j/\widehat{u}'_j) + 2\pi s_i)$ ,  $s_i \in [\sigma\beta(l_j - \Delta l/2), \sigma\beta(l_j + \Delta l/2)] \cap \mathbb{Z}_+$ 
31:          $f_{j,i}^r = \theta_{j,i}^r + b \pmod{F}$ 
32:         suppose  $f_{j,i}^r$  belongs to region( $j, q$ ),
33:         add a vote to both region( $j, q$ ) and two neighbors nearby that region, e.g.
           region( $j, q - 1$ ) and region( $j, q + 1$ )
34:       end for
35:     end for
36:   end for
37:   for  $j \in [B]$  do
38:      $q_j^* \leftarrow \{q | v_{j,q} > \frac{R_{\text{loc}}}{2}\}$ 
39:     if  $\text{last} = \text{true}$  then
40:        $l_j^* \leftarrow \text{median}\{f_{j,i}^r | f_{j,i}^r \in \text{region}(j, q_j^*), i \in [f], r \in [R_{\text{loc}}]\}$ 
41:     else
42:        $l_j^* \leftarrow \text{center of region}(j, q_j^*)$ 
43:     end if
44:   end for
45:   return  $l^*$ 
46: end procedure

```
