

Newtonized Orthogonal Matching Pursuit: Frequency Estimation over the Continuum

Babak Mamandipoor, Dinesh Ramasamy, and Upamanyu Madhow

Abstract—We propose a fast sequential algorithm for the fundamental problem of estimating frequencies and amplitudes of a noisy mixture of sinusoids. The algorithm is a natural generalization of Orthogonal Matching Pursuit (OMP) to the continuum using Newton refinements, and hence is termed Newtonized OMP (NOMP). Each iteration consists of two phases: detection of a new sinusoid, and sequential Newton refinements of the parameters of already detected sinusoids. The refinements play a critical role in two ways: (1) sidestepping the potential *basis mismatch* from discretizing a continuous parameter space, (2) providing *feedback* for locally refining parameters estimated in previous iterations. We characterize convergence, and provide a Constant False Alarm Rate (CFAR) based termination criterion. By benchmarking against the Cramér Rao Bound, we show that NOMP achieves near-optimal performance under a variety of conditions. We compare the performance of NOMP with classical algorithms such as MUSIC and more recent Atomic norm Soft Thresholding (AST) and Lasso algorithms, both in terms of frequency estimation accuracy and run time.

Index Terms—Gridless Compressed Sensing, Orthogonal matching pursuit, Sparse approximation, Line spectral estimation, Frequency estimation, Newton refinement, Decision-feedback.

I. INTRODUCTION

Frequency estimation from a mixture of sinusoids in AWGN is a fundamental problem that arises in a variety of communication and radar applications, including estimation of spatial channels (e.g., for phased arrays), temporal multipath channels (e.g. for equalization), and spatiotemporal channels (e.g., range and direction of arrival estimation for a target). In all of these applications, the spectrum of the measured signal consists of multiple discrete frequencies over a continuous interval. The problem of estimating the frequencies of such a signal is also known as line spectral estimation.

In this paper, we propose an algorithm to estimate frequencies from N equi-spaced noisy samples in time, denoted by $\mathbf{y} \in \mathbb{C}^N$. Defining the unit norm sinusoid of frequency $\omega \in [0, 2\pi)$, by

$$\mathbf{x}(\omega) \triangleq \frac{1}{\sqrt{N}} [1 \ e^{j\omega} \ \dots \ e^{j(N-1)\omega}]^T, \quad (1)$$

Copyright (c) 2015 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubspermissions@ieee.org.

This work was supported in part by Systems on Nanoscale Information fabriCs (SONIC), one of the six SRC STARnet Centers, sponsored by MARCO and DARPA.

B. Mamandipoor and U. Madhow are with the Department of Electrical and Computer Engineering, University of California Santa Barbara, Santa Barbara, CA 93106 USA (e-mail: bmamandi@ece.ucsb.edu; madhow@ece.ucsb.edu).

D. Ramasamy is with Amazon Corporation, Sunnyvale, CA, USA (e-mail: dineshr@ece.ucsb.edu).

the observed signal is a mixture of K sinusoids:

$$\mathbf{y} = \sum_{l=1}^K g_l \mathbf{x}(\omega_l) + \mathbf{z}, \quad \mathbf{z} \sim \mathcal{CN}(\mathbf{0}, \sigma^2 \mathbb{I}_N), \quad (2)$$

where $g_l \in \mathbb{C}$ are the unknown complex gains. The signal to noise ratio for l^{th} sinusoid is given by $\text{SNR}_l = |g_l|^2 / \sigma^2$. The goal of the algorithm is to provide reliable estimates of $\{(g_l, \omega_l) : l = 1, 2, \dots, K\}$ and K , the number of sinusoids in the mixture.

The preceding model and its variants have many applications. For a linear array with N elements with inter-element spacing d , the response corresponding to angle of arrival or departure θ relative to broadside is given by $\mathbf{x}(\omega)$, where $\omega = 2\pi(d/\lambda) \sin(\theta)$ is the spatial frequency corresponding to θ , and λ denotes the carrier wavelength. For estimation of a multipath channel $h(t) = \sum_{l=1}^K g_l \delta(t - \tau_l)$, the channel transfer function is given by $H(f) = \sum_{l=1}^K g_l e^{-j2\pi f \tau_l}$. Sampling uniformly in the frequency domain with spacing Δf , yields a mixture of sinusoids with $\omega_l = -2\pi \Delta f \tau_l$, reducing the problem of estimating delays to that of frequency estimation. This directly models the operation of stepped frequency continuous wave (SFCW) for imaging a collection of point scatterers. When the channel is “seen through” a pulse $p(t)$, as is often the case for channel estimation in communication applications, then the noiseless frequency domain signal is given by $Y(f) = \sum_{l=1}^K g_l P(f) e^{-j2\pi f \tau_l}$. A simple extension of our algorithm to handle weighted sinusoids applies to this setting. We do not provide detailed discussion here due to lack of space, but the code that we have made available [1] does provide the required flexibility.

In this paper, we are interested in the setting where K is a small integer, so that the underlying signal has a sparse representation in the atomic set of unit norm sinusoids. The goal is to find the best sparse approximation of \mathbf{y} using atoms in the dictionary of sinusoids. A promising algorithm in sparse approximation theory is Orthogonal Matching Pursuit (OMP) [2], [3], a greedy algorithm that iteratively identifies the atom that yields the greatest improvement in approximation quality. If the structure of the atomic set is “simple,” so that identifying the “best” atom in each iteration needs a small amount of computation, then OMP becomes computationally attractive. Unfortunately, for our atomic set, searching over a continuum of atoms is not possible. An approximation that is often employed to overcome this problem is to discretize the set of frequencies, hoping that the signal still admits a sparse representation in this finite set. However, as discussed in [4], no matter how finely we grid the parameter space, the

true underlying atoms of \mathbf{y} need not lie on the grid. This “off-grid effect,” or basis mismatch, degrades the performance of reconstruction algorithms significantly. The authors in [5] propose inserting a gradient-based local search in a Matching Pursuit [6] framework, as a means of alleviating the off-grid effect. In this paper, we go one step further by incorporating a Newton-based Cyclic Refinement step into the OMP framework: this not only sidesteps the off-grid effect, but also enhances performance by refining previously estimated atoms in each iteration.

A. Contributions

Our key contributions are as follows:

- (1) We propose Newtonized OMP (NOMP), which detects the best atom over a discrete grid, but avoids basis mismatch by adding a Newton refinement step, thus emulating pursuit over the continuum. In addition, we go beyond OMP by locally refining all estimated atoms in each iteration, thus re-evaluating estimates of previously detected sinusoids to incorporate the effect of the newly detected sinusoid. This second refinement step can be interpreted as feedback presented to the atoms we have already detected. We show that this feedback mechanism plays a crucial role in handling *interference* among the underlying atoms of \mathbf{y} , yielding estimation accuracies far better than what would be possible by standard greedy pursuit algorithms (e.g., OMP and Matching Pursuit). We do not require explicit estimates of model order, and provide a stopping criterion based on CFAR (constant false alarm rate) based on an estimate of the noise variance.
- (2) We prove convergence of NOMP by providing upper bounds on the number of iterations. Moreover, we derive a bound on the convergence rate, showing that by choosing a fine enough grid for detection, the bound approaches the rate of convergence for OMP on the continuum.
- (3) We show that the algorithm is near-optimal by numerical comparisons against the Cramér Rao Bound (CRB) [7] in a variety of settings. When the frequencies of the sinusoids in the mixture are well-separated, NOMP is able to achieve the CRB for any sinusoid that has an SNR greater than a certain threshold. Moreover, NOMP is able to resolve closely-spaced frequencies with near-optimal accuracy, as long as there is enough disparity in the SNRs of different sinusoids.
- (4) We analyze the computational complexity of NOMP. Our numerical experiments show that its run time is significantly smaller than that of recently proposed state-of-the-art Atomic norm Soft Thresholding (AST) algorithm [8], and is even smaller than that of the classical MUSIC algorithm [9]. We evaluate the run time of the algorithm in a variety of settings. A freely downloadable software package implementing the proposed algorithm can be found in [1].

B. Related work

Line spectral estimation is a fundamental problem in statistical signal processing. Classical (and popular) subspace methods such as MUSIC and ESPRIT [9], [10] exploit the low-rank structure of the autocorrelation matrix. One of the

major advantages of these methods is the capability of resolving multiple closely-spaced frequencies at high SNR. Both MUSIC and ESPRIT have been shown to be asymptotically optimal in the limit of infinite SNR [11], but their performance degrades at medium and low SNRs. Another family of DFT-based classical methods [12], [13], typically have lower computational complexity and estimation accuracy similar to that of subspace methods [13], [14].

More recent techniques using convex optimization cast the frequency estimation problem as that of finding a sparse approximation of the received signal using an infinite-dimensional dictionary of sinusoids. It is shown in [15] that, in the absence of noise, total-variation norm is able to locate frequencies with infinite precision, as long as the minimum frequency separation exceeds $4 \times \Delta_{\text{dft}}$ where $\Delta_{\text{dft}} = 2\pi/N$ is the Discrete Fourier Transform (DFT) grid spacing. The sufficient condition on required minimum separation has been recently improved to $2.52 \times \Delta_{\text{dft}}$ in [16]. An extension to noisy scenarios is provided in [17]. Another approach is Atomic norm Soft Thresholding (AST) [8], [18], which provides theoretical guarantees of noise robustness in terms of mean squared error (MSE). Both total-variation norm and atomic norm are generalizations of the ℓ_1 norm to infinite-dimensional settings. Solving these optimization problems involves solving the Lagrange dual which takes the form of a semi-infinite program (SIP) with finite-dimensional decision variables and infinitely many constraints. For the sparse frequency estimation problem, [17] and [8] reformulate the dual as a semidefinite program (SDP), which enables numerical optimization. Similar reformulation for other problems seems to be difficult [18]. A pragmatic approach is to use Lasso optimization on a highly oversampled grid as an approximation for AST [8], [18]. Both AST and Lasso are benchmarks that we compare our proposed algorithm against in our numerical experiments.

NOMP can be viewed as coordinate optimization over a continuum. Coordinate-wise descent has been widely used for sparse approximation; for example, such methods are competitive for solving Lasso type problems [19], [20]. Preliminary results in [21] show that coordinate descent to a relaxation of the AST problem can be a means to speed up implementation.

There is a large body of work on feedback for improving the performance of iterative greedy algorithms for sparse approximation [22], [23], [24], [25], [26], [27]: information from recent iterations is used to remove errors introduced by previous steps. For example, [23] introduces a forward-backward greedy algorithm that allows *aggressive* backward steps (discarding small-magnitude atoms) after a greedy forward step. In contrast, NOMP, by virtue of its continuous parameterization, can employ a *mild* form of feedback by locally refining the set of estimated atoms, thus allowing an atom to be replaced by a nearby, highly correlated atom, in the continuum. The idea of embedding a local refinement step in OMP has been proposed in [28] in a discrete setting, but this algorithm does not address basis mismatch, and assumes that the model order is known *a priori*.

Many theoretical results on OMP [3], [29], [30] are based on assumptions such as incoherence or Restricted Isometry Property (RIP) for the underlying dictionary. Such approaches

do not work for analyzing NOMP, since nearby atoms are highly correlated for continuous frequency estimation. Instead, our convergence analysis borrows tools from analysis of AST-based line spectral estimation [8], along with observations following from our CFAR-based design.

For a single sinusoid, frequency estimation using coarse detection followed by Newton refinement was proposed three decades ago by Abatzoglou [31]. This was recently adapted for estimation of a single delay in [32], and shown to approach estimation-theoretic bounds. In prior work involving some of the authors, we have used sequential algorithms similar to NOMP for millimeter wave spatial channel estimation with compressive measurements [33], [34]. The NOMP algorithm in the present paper is an improvement on these, with a principled CFAR-based stopping criterion. We present NOMP within an application-independent abstraction, recognizing the fundamental nature and widespread utility of the frequency estimation problem. We have reported on a version of NOMP in a recent conference paper [35] (although the algorithm in [35] omits a least squares step included here for establishing convergence rate results, but with little impact on practical performance). The present paper goes beyond [35] in providing a detailed convergence analysis, and a comprehensive comparison with the state of the art.

A closely related algorithm to NOMP is proposed in [36], which employs a Bayesian framework for frequency estimation, using Newton refinements for updating the frequencies. The details are different from our non-Bayesian, CFAR framework, and convergence analysis is not provided, but the benefits of Newtonization are also evident in the numerical results in [36].

Outline: We present our algorithm in Section II. In Section III, we present the CFAR-based stopping criterion, together with a simplified analytical characterization of false alarm and miss probabilities. In Section IV, we show why oversampling is essential for discretization followed by refinement to emulate pursuit over the continuum. We discuss convergence in Section V. In Section VI, we report on numerical experiments, and compare NOMP with other methods in terms of estimation accuracy and computational complexity. Section VII discusses extension of the algorithm to more general settings, with illustrative numerical results for compressive measurements. Section VIII concludes the paper.

We set $N = 256$ throughout in our numerical results (other values of N yield entirely similar trends).

Notation: Complex conjugate transpose of v is denoted by v^H . $\Re\{a\}$ is the real part of complex number a . The Moore-Penrose pseudo-inverse of matrix A is denoted by A^\dagger . The distance between any two frequencies $\{\omega_l, \omega_k\}$ is defined by $\text{dist}(\omega_k, \omega_l) \triangleq \min_{a \in \mathbb{Z}} |\omega_k - \omega_l + 2\pi a|$, i.e. the wrap-around distance when we restrict frequencies to lie in $[0, 2\pi)$. The DFT matrix with unit norm columns and the corresponding grid spacing are denoted by \mathcal{F} and Δ_{dft} , respectively. The inner product between $v, u \in \mathbb{C}^N$ is defined as $\langle v, u \rangle = u^H v$.

II. NOMP ALGORITHM

We first discuss estimation of a single sinusoid, and then build on it to generalize to a mixture of sinusoids.

A. Single Frequency

We have $\mathbf{y} = g\mathbf{x}(\omega) + \mathbf{z}$. The Maximum Likelihood (ML) estimate of the gain and frequency are obtained by minimizing the residual power $\|\mathbf{y} - g\mathbf{x}(\omega)\|^2$, or equivalently, by maximizing the function

$$S(g, \omega) = 2\Re\{\mathbf{y}^H g\mathbf{x}(\omega)\} - |g|^2 \|\mathbf{x}(\omega)\|^2. \quad (3)$$

Directly optimizing $S(g, \omega)$ over all gains and frequencies is difficult. Therefore, we adopt a two stage procedure: (1) Detection stage, where we find a coarse estimate of ω by restricting it to a discrete set, (2) Refinement stage, in which we iteratively refine gain and frequency estimates.

For any given ω , the gain that maximizes $S(g, \omega)$ is given by $\hat{g} = (\mathbf{x}(\omega)^H \mathbf{y}) / \|\mathbf{x}(\omega)\|^2$. Substituting \hat{g} in $S(g, \omega)$ yields that the generalized likelihood ratio test (GLRT) estimate of ω (treating g as a nuisance parameter) is the solution to the following optimization problem:

$$\hat{\omega} = \arg \max_{\omega} G_{\mathbf{y}}(\omega), \quad (4)$$

where

$$G_{\mathbf{y}}(\omega) = |\mathbf{x}(\omega)^H \mathbf{y}|^2 / \|\mathbf{x}(\omega)\|^2 \quad (5)$$

is the GLRT cost function. We use this observation to find a coarse estimate of (g, ω) in the Detection stage.

Detection: We obtain a coarse estimate of ω by restricting it to a finite discrete set denoted by $\Omega \triangleq \{k(2\pi/\gamma N) : k = 0, 1, \dots, (\gamma N - 1)\}$, where γ is the over-sampling factor relative to the DFT grid. For our simulation results, we set $\gamma = 4$. The outputs of this stage are $\omega_c \in \Omega$ that maximizes the GLRT cost function (4), and the corresponding gain $(\mathbf{x}(\omega_c)^H \mathbf{y}) / \|\mathbf{x}(\omega_c)\|^2$.

Refinement: Since ω can take any value in interval $[0, 2\pi)$, we add a Newton-based refinement stage for estimation on the continuum. Let $(\hat{g}, \hat{\omega})$ denote the current estimate. The Newton step for frequency refinement is given by

$$\hat{\omega}' = \hat{\omega} - \dot{S}(\hat{g}, \hat{\omega}) / \ddot{S}(\hat{g}, \hat{\omega}) \quad (6)$$

where

$$\dot{S}(g, \omega) = \Re\{(\mathbf{y} - g\mathbf{x}(\omega))^H g(d\mathbf{x}(\omega)/d\omega)\} \quad (7)$$

$$\ddot{S}(g, \omega) = \Re\{(\mathbf{y} - g\mathbf{x}(\omega))^H g(d^2\mathbf{x}(\omega)/d\omega^2)\} - |g|^2 \|d\mathbf{x}(\omega)/d\omega\|^2. \quad (8)$$

As we want to *maximize* $S(g, \omega)$, we only apply the update rule (6) when the function is locally concave (i.e. $\ddot{S}(\hat{g}, \hat{\omega}) < 0$). The gain parameter is then updated to maximize $S(g, \hat{\omega}')$: $\hat{g}' = (\mathbf{x}(\hat{\omega}')^H \mathbf{y}) / \|\mathbf{x}(\hat{\omega}')\|^2$.

Refinement Acceptance Condition (RAC): We *accept* a refinement only if it leads to a strict improvement in $G_{\mathbf{y}}(\omega)$; that is, if $G_{\mathbf{y}}(\hat{\omega}') > G_{\mathbf{y}}(\hat{\omega})$. This ensures that an accepted refinement can only decrease the overall residual energy, and that the residual energy is non-increasing throughout the course of the algorithm, which ensures convergence, as shown in Section V.

B. Multiple Frequencies

Let $\mathcal{P} = \{(g_l, \omega_l), l = 1, \dots, k\}$ denote a set of estimates of the parameters of the sinusoids in the mixture. Let

$$\mathbf{y}_r(\mathcal{P}) = \mathbf{y} - \sum_{l=1}^{l=k} g_l \mathbf{x}(\omega_l) \quad (9)$$

denote the *residual* measurement corresponding to this estimate. The following procedure is a direct generalization of the single sinusoid refinement algorithm to multiple frequencies. It proceeds by employing the single sinusoid algorithm to perform *Newtonized coordinate descent* on the overall residual energy $\|\mathbf{y}_r(\mathcal{P})\|^2$. One step of this coordinate descent involves cycling through all sinusoids in \mathcal{P} in a predetermined order. In this process, suppose that we wish to refine the l -th sinusoid: we treat $\mathbf{y}_r(\mathcal{P} \setminus \{(g_l, \omega_l)\})$ as our measurement \mathbf{y} and employ the single frequency update step to refine (g_l, ω_l) .

Algorithm 1 Newtonized Orthogonal Matching Pursuit

- 1: **Procedure** EXTRACTSPECTRUM(\mathbf{y}, τ):
 - 2: $m \leftarrow 0, \mathcal{P}_0 = \{\}$
 - 3: **while** $\max_{\omega \in \text{DFT}} G_{\mathbf{y}_r(\mathcal{P}_m)}(\omega) > \tau$ **do**
 - 4: $m \leftarrow m + 1$
 - 5: **IDENTIFY**
 $\hat{\omega} = \arg \max_{\omega \in \Omega} G_{\mathbf{y}_r(\mathcal{P}_{m-1})}(\omega)$
 and its corresponding gain
 $\hat{g} \leftarrow (\mathbf{x}(\hat{\omega})^H \mathbf{y}_r(\mathcal{P}_{m-1})) / \|\mathbf{x}(\hat{\omega})\|^2$
 - 6: $\mathcal{P}'_m \leftarrow \mathcal{P}_{m-1} \cup \{(\hat{g}, \hat{\omega})\}$
 - 7: **SINGLE REFINEMENT:** Refine $(\hat{g}, \hat{\omega})$ using single frequency Newton update algorithm (R_s Newton steps) to obtain improved estimates $(\hat{g}', \hat{\omega}')$.
 - 8: $\mathcal{P}''_m \leftarrow \mathcal{P}_{m-1} \cup \{(\hat{g}', \hat{\omega}')\}$
 - 9: **CYCLIC REFINEMENT:** Refine parameters in \mathcal{P}''_m one at a time: For each $(g, \omega) \in \mathcal{P}''_m$ we treat $\mathbf{y}_r(\mathcal{P}''_m \setminus \{(g, \omega)\})$ as the measurement \mathbf{y} , and apply single frequency Newton update algorithm. We perform R_c rounds of cyclic refinements. Let \mathcal{P}'''_m denote the new set of parameters.
 - 10: **UPDATE** all gains in \mathcal{P}'''_m by least squares:
 $X \triangleq [\mathbf{x}(\omega_1) \dots \mathbf{x}(\omega_m)], \{\omega_l\}$ are the frequencies in \mathcal{P}'''_m
 $[g_1 \dots g_m]^T = X^\dagger \mathbf{y}$
 Let \mathcal{P}_m denote the new set of parameters.
 - 11: **return** \mathcal{P}_m
-

The NOMP procedure is summarized in Algorithm 1. We now briefly discuss the role of its main components:

- **SINGLE REFINEMENT** (Step 7): Ideally, we want to identify $\hat{\omega}$ which maximizes the GLRT cost function over the continuum. The **SINGLE REFINEMENT** step emulates search over the continuum by locally refining the estimate of $\hat{\omega}$ obtained by picking the maximum over the discrete set Ω .
- **CYCLIC REFINEMENT** (Step 9) is where NOMP diverges from *forward greedy methods* [37] (in particular OMP) by

providing *feedback* for local refinements of previously detected sinusoids. This gives them an opportunity to better explain the received signal in light of new information regarding the presence of another sinusoid. This feedback is presented in the form of an updated residue. As we see in Sections V-C and VI, this step is crucial for fast convergence and high estimation accuracy.

- **UPDATE** by least squares (Step 10): Here we update gains by projecting the received signal \mathbf{y} onto the subspace spanned by the estimated frequencies. This ensures that the residual energy is the minimum possible for the current set of estimated frequencies. We see in Section V that performing this projection step just prior to detecting a new sinusoid enables us to lower bound the convergence rate of NOMP by mirroring arguments used to establish bounds on OMP convergence [37].

We have left the number of refinement steps unspecified so far. For the simulations in this paper, we set $R_s = 1$ for every newly detected sinusoid in the **SINGLE REFINEMENT** step, and $R_c \in \{1, 3, 5\}$ refinement rounds in the **CYCLIC REFINEMENT** step, depending on the difficulty of the estimation problem.

C. Complexity analysis

We analyze the computational complexity of the main steps of NOMP assuming that the algorithm runs for exactly K iterations (i.e., perfect stopping). Checking whether the stopping criterion is satisfied is efficiently implemented using Fast Fourier Transform (FFT), with complexity $\mathcal{O}(KN \log(N))$. The **IDENTIFY** Step involves computing the GLRT cost function over the set of frequencies defined by Ω . This can also be computed using FFTs in $\mathcal{O}(\gamma KN \log(\gamma N))$ time. The **SINGLE REFINEMENT** Step takes only $\mathcal{O}(R_s N)$ operations per sinusoid, hence the total cost for **SINGLE REFINEMENT** is $\mathcal{O}(R_s KN)$. The **CYCLIC REFINEMENT** involves refining all frequencies that have been estimated so far, and has overall complexity $\mathcal{O}(R_c R_s K^2 N)$. If we directly compute the pseudo-inverse and apply it to the vector of observation (i.e., $X^\dagger \mathbf{y}$) in the **UPDATE** Step, the complexity is $\mathcal{O}(NK^2 + K^3)$ per iteration, and the overall cost is $\mathcal{O}(NK^3 + K^4)$. However, we note that iterative methods (such as Richardson's iteration or conjugate gradient [22]) are extremely efficient in computing the least squares solution and can be used for speeding up the **UPDATE** Step. Empirical observations suggest that the **CYCLIC REFINEMENT** step dominates the overall computational cost of NOMP.

Minimum frequency separation: When two frequencies, say ω_1 and ω_2 , are “very close”, intuitively, the mixture $g_1 \mathbf{x}(\omega_1) + g_2 \mathbf{x}(\omega_2)$ is explained “very well” by a single frequency, say as $(g_1 + g_2) \mathbf{x}(\omega_1)$. Thus, a natural metric to characterize regimes for testing algorithms for mixture frequency estimation is the *minimum frequency separation* between any two sinusoids. We denote this by $\Delta\omega_{\min} = \min_{k \neq l} \text{dist}(\omega_k, \omega_l)$ and we would like our algorithms to work well even for small values of $\Delta\omega_{\min}$.

Without a minimum frequency separation condition, the estimation problem can be hopelessly ill-posed. This has been studied in detail in [15] using Slepian's work on prolate

spheroidal sequences [38]. It has been shown that if the frequencies are clustered together, it becomes impossible for *any* method to recover the information from the *noisy* observations. It is important to note that in the limit of infinite SNR, however, one can still estimate a sparse clustered set of frequencies regardless of their separation (e.g., using Prony’s method of polynomial interpolation [39]).

Estimation Theoretic Bounds: Estimation theoretic quantities such as Cramér Rao Bound (CRB) and Ziv-Zakai Bound (ZZB) [40] provide lower bounds on the variance of estimators. For single frequency estimation, these bounds are given by [41]

$$\text{CRB}(\text{SNR}) = \frac{6}{\text{SNR} \times (N^2 - 1)}, \quad (10)$$

and

$$\text{ZZB}(\text{SNR}) = \int_0^\pi Q \left(\sqrt{\text{SNR} \left(1 - \left| \frac{\sin(Nh/2)}{N \sin(h/2)} \right| \right)} \right) h \, dh. \quad (11)$$

In Figure 1, we plot estimation error bounds as a function of $\text{SNR} = \|\mathbf{x}(\omega)\|^2 / \sigma^2 = 1/\sigma^2$. The ZZB has a distinct threshold behavior: large frequency estimation errors are inevitable in a “low SNR” regime below a threshold, while the ZZB converges to the CRB when the SNR is large enough compared to this threshold. In Section VI, we compare algorithms in terms of high SNR behavior relative to the CRB, and low SNR behavior relative to the ZZB threshold. Note that even if per sinusoid SNRs are higher than the ZZB threshold, the joint estimation problem for multiple sinusoids may still be ill-posed (e.g., if the frequency separation is too small).

Remark 1: We have defined “integrated SNR” obtained by dividing the total power of a sinusoid by the noise power per complex dimension, i.e., $\text{SNR} = \frac{\mathbb{E}[\|\mathbf{x}(\omega)\|^2]}{\sigma^2}$. An alternative definition of signal to noise ratio is the “per-sample SNR”, given by $\text{SNR}_{\text{sample}} = \frac{\mathbb{E}[\|\mathbf{x}_i(\omega)\|^2]}{\mathbb{E}[\|\mathbf{z}_i\|^2]} = \frac{1}{N} \text{SNR}$. For example, $\text{SNR} = 25$ dB (which is the nominal SNR value in our simulations), corresponds to $\text{SNR}_{\text{sample}} \approx 1$ dB.




Fig. 1: CRB and ZZB for estimating the frequency of a single sinusoid.

III. CFAR-BASED STOPPING CRITERION

Detection problems face a tension between false alarm and missed detection (or miss, for short). In many detection

problems, a model for the signal can be elusive. Thus, a common strategy is based on the Constant False Alarm Rate (CFAR) criterion [42], [43], which only requires a model for noise. Here, we use the CFAR criterion to estimate model order (i.e. number of sinusoids in the mixture K): if the residual signal can be explained “well enough” by noise, up to a target false alarm rate, then we stop. We show by simulation that the actual false alarm rate is close to the nominal being designed for.

We also estimate the probability of miss, taking into account the effect of noise but ignoring “interference” from other sinusoids. The resulting receiver operating characteristic (ROC) turns out to be in remarkable agreement with simulations. This shows that, when the sinusoids are separated beyond a minimum separation, and when their SNRs exceed the ZZB thresholds for individual sinusoids, then the probabilities of false alarm and miss are dominated by noise rather than by inter-sinusoid interference.

A. Stopping Criterion

The algorithm terminates when

$$G_{\mathbf{y}_r(\mathcal{P})}(\omega) = |\langle \mathbf{y}_r(\mathcal{P}), \mathbf{x}(\omega) \rangle|^2 < \tau$$

for all DFT frequencies $\{2\pi k/N : k = 0, \dots, N-1\}$. In other words, we stop when $\|\mathcal{F}\mathbf{y}_r(\mathcal{P})\|_\infty^2 < \tau$, where $\mathcal{F}\mathbf{a}$ is the Discrete Fourier Transform of \mathbf{a} , and report \mathcal{P} as our estimate of the sinusoids in the mixture.

Suppose that we have already correctly detected all sinusoids in the mixture. In this case, the residual is $\mathbf{y}_r(\mathcal{P}) \approx \mathbf{z}$, where $\mathbf{z} \sim \mathcal{CN}(\mathbf{0}, \sigma^2 \mathbb{I}_N)$ (since \mathcal{F} is a projection matrix, the statistics of WGN are unchanged by it). It is easy to show that

$$\Pr \left\{ \|\mathcal{F}\mathbf{y}_r(\mathcal{P})\|_\infty^2 > \tau \right\} = 1 - (1 - \exp(-\tau/\sigma^2))^N. \quad (12)$$

We choose our stopping criterion threshold τ so that $\Pr \left\{ \|\mathcal{F}\mathbf{y}_r(\mathcal{P})\|_\infty^2 > \tau \right\} = P_{\text{fa}}$, where P_{fa} is a nominal false alarm rate. Using (12), we can explicitly compute this threshold as

$$\tau = -\sigma^2 \log \left(1 - (1 - P_{\text{fa}})^{1/N} \right).$$

A more easily interpreted expression can be obtained via asymptotics for large N [44] (which provide excellent approximations for the moderate values of N used in our numerical results). Let $M_N \triangleq \|\mathcal{F}\mathbf{y}_r(\mathcal{P})\|_\infty^2$. If $\mathbf{y}_r(\mathcal{P}) \approx \mathbf{z}$, we have $\mathbb{E}[M_N] = \sigma^2 \sum_{k=1}^N \frac{1}{k} \approx \sigma^2 \log N$, and the asymptotic distribution of $E \triangleq M_N - \sigma^2 \log N$ is given by $\Pr\{E \leq x\} = \exp(-\exp(-x/\sigma^2))$. We set $\tau = \sigma^2 \log(N) + x$, for x so that $\Pr\{E \geq x\}$ is equal to the nominal false alarm rate P_{fa} . This is given by $x = -\sigma^2 \log \log(1/(1 - P_{\text{fa}}))$ and the resulting expression for the threshold τ is

$$\tau = \sigma^2 \log(N) - \sigma^2 \log \log(1/(1 - P_{\text{fa}})). \quad (13)$$

The “measured” versus “nominal” false alarm rates for different values of nominal P_{fa} , are shown in Figure 2. Each point in the plot is generated by 300 runs of NOMP algorithm for estimating frequencies in a mixture of $K = 16$ sinusoids of fixed nominal SNR. The minimum frequency separation $\Delta\omega_{\min} = 2.5\Delta_{\text{dft}}$. We declare a false alarm whenever NOMP

overestimates the model order K . As shown in Figure 2, the empirical false alarm rate closely follows the nominal value at various SNRs.

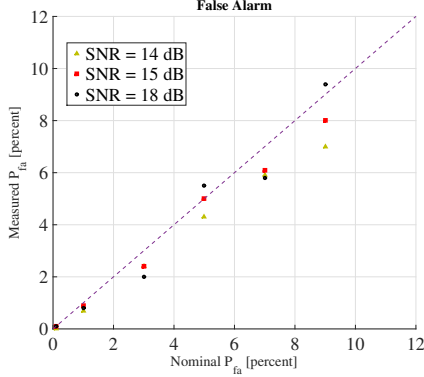


Fig. 2: Nominal v.s. measured probability of false alarm.

B. Probability of Miss

Define neighborhood N_{ω_i} around each true frequency ω_i by $N_{\omega_i} \triangleq \{\omega : \text{dist}(\omega, \omega_i) < 0.25 \times \Delta_{\text{dft}}\}$. We declare successful detection of ω_i if at least one of the estimated frequencies lies in N_{ω_i} , otherwise we declare a miss for ω_i . A miss results from noise and inter-sinusoid interference, but we only model noise here. For the minimum separation considered here, we show using simulations that the empirical probability of miss is only a little higher than the analytical estimate that we derive below.

A sinusoid of amplitude A leads to a maximum FFT value of αA , where $\alpha \in [0, 1]$ captures the amplitude reduction due to the grid mismatch. The magnitude of the maximum FFT coefficient, denoted by M_{fft} , is Rician with $\Pr\{M_{\text{fft}} < x\} = 1 - Q_1\left(\frac{\sqrt{2}\alpha A}{\sigma}, \frac{\sqrt{2}x}{\sigma}\right)$, where Q_1 is Marcum Q-function. The sinusoid is not detected by the algorithm if $M_{\text{fft}} < \sqrt{\tau}$, hence

$$P_{\text{miss}} = 1 - Q_1\left(\alpha\sqrt{2SNR}, \sqrt{2\tau/\sigma^2}\right). \quad (14)$$

Assuming uniform distribution for a frequency within a DFT grid interval gives us $\mathbb{E}[\alpha] = \mathbb{E}[\sin(N\omega/2)/(N \sin(\omega/2))] = 0.88$, where $\omega \sim \text{Uniform}[-\pi/N, \pi/N]$. Therefore,

$$P_{\text{miss}} \approx 1 - Q_1\left(0.88\sqrt{2SNR}, \sqrt{2\tau/\sigma^2}\right). \quad (15)$$

Putting equations (13) and (15) together, we can characterize the ROC at various SNRs as shown in Figure 3. The simulation parameters are the same as Figure 2. We see that for SNR=18 dB the probability of miss is negligible. When SNR goes below the ZZB threshold e.g. SNR=14 dB, the probability of miss is bounded away from zero. This behavior is predicted by the ZZB threshold and our ROC analysis for single frequency estimation. However, as we see in Figure 3, they serve as excellent approximations for multiple frequency estimation.

Remark 2: For the simulations in this paper, we have used the CFAR-based stopping criteria for the NOMP algorithm. However, a variety of other stopping rules, e.g., Bayesian Information Criteria (BIC) [45] and Akaike Information Criteria

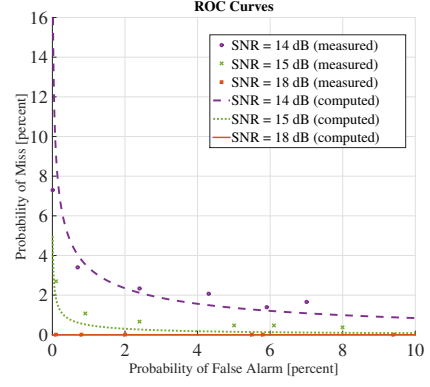


Fig. 3: Measured v.s. computed ROCs

(AIC) [46], can be easily adapted for use with the NOMP algorithm. In Section VI-D, we investigate the performance of the BIC stopping rule (see Appendix I for a quick overview) as well as CFAR-based stopping criteria for NOMP.

IV. THE NEED TO OVERSAMPLE

In this section, we show that oversampling is indeed required at the detection stage, in order for Newton refinements to converge to the maximum of the GLRT cost function. We ignore noise in these discussions. The GLRT cost function is given by $G_{\mathbf{y}}(\omega) = |\sum_{l=1}^K g_l h(\omega - \omega_l)|^2$, where $h(\omega) = \frac{\sin(N\omega/2)}{N \sin(\omega/2)}$ is the Dirichlet kernel. Characterizing the minimum required oversampling factor for this general setting is difficult because of the highly non-convex nature of $G_{\mathbf{y}}(\omega)$. Hence, we focus on a single frequency setting ($K = 1$).

We wish to arrive at ω_1 by optimizing $G_{\mathbf{y}}(\omega)$ using Newton refinements starting on a coarse grid. We note that $\arg \max_{\omega} G_{\mathbf{y}}(\omega) = \arg \max_{\omega} |g_1 h(\omega - \omega_1)|^2 = \arg \max_{\omega} |h(\omega - \omega_1)|^2$. We would like to characterize the minimum oversampling factor such that if we start off from the best guess of the maximum of $G_{\mathbf{y}_r}(\omega)$ on the grid, the Newton refinement stage will take us to ω_1 . That is, we must characterize how close to ω_1 must the nearest grid point lie, so that Newton refinements will always take us to ω_1 from this grid point. Without loss of generality, we set $\omega_1 = 0$ (since we have shifted our frequency axis such that $\omega_1 = 0$, no grid point may lie on 0).

We start by normalizing frequencies by the DFT spacing. In this scaled frequency axis, the Dirichlet kernel is given by $h(x) = \frac{\sin(\pi x)}{N \sin(\pi x/N)}$, where $x = \omega/\Delta_{\text{dft}}$ is sometimes referred to as the normalized frequency. As shown in [47], the Newton method converges to the solution of $h'(x) = 0$ quadratically, if the initial guess x_0 lies in an interval I around the true solution where the following conditions are met:

- $h''(x) \neq 0 \quad \forall x \in I$.
- $h'''(x)$ is finite $\forall x \in I$.
- $|x_0| < 1/M$ where $M \triangleq \sup_{x \in I} 0.5 \left| \frac{h'''(x)}{h''(x)} \right|$.

Figure 4 shows the function $h(x)$ and its derivatives in a window around the origin. The first two conditions are met for any interval $I \subset (-1, 1)$. Therefore, it only remains to satisfy the third condition. By simple algebra one can see that




Fig. 4: Dirichlet kernel $h(x)$ and its derivatives.

if we set $I = (-0.45, +0.45)$, then we have $|x_0| < 1/M$ for any $x_0 \in I$. Therefore the maximum acceptable grid spacing is about 0.9 which is equivalent to minimum oversampling factor ≈ 1.12 . This simple analysis shows that, even for a single sinusoid and no noise, we must sample beyond the DFT grid to ensure that the two-stage detection/refinement procedure successfully identifies the maximum of the GLRT cost function, thereby imitating pursuit over the continuum. Our simulations show that setting the oversampling factor at $\gamma = 4$ or more works very well, independent of the number of the sinusoids in the observations.

V. CONVERGENCE

In this section, we first characterize convergence by providing upper bounds on the number of iterations of NOMP. We then provide a bound on the rate of convergence of NOMP as a function of the “atomic norm” of \mathbf{y} , and the oversampling factor γ . Our convergence results are pessimistic, in that they do not account for the effect of the refinement steps. We show by simulations the dramatic improvement due to refinements, by comparing NOMP to the following variants of OMP:

Discretized OMP (DOMP): This is standard OMP applied to the oversampled grid of sinusoids Ω . We use our NOMP implementation with number of refinement steps set to zero, increasing the oversampling factor to $\gamma = 20$. Since DOMP can be interpreted as a special case of NOMP, the convergence analysis presented here is valid for DOMP as well.

NOMP without Cyclic Refinements (NOMP-): If we skip the CYCLIC REFINEMENT step of NOMP, then we get an algorithm that emulates OMP over the continuum of atoms. Note that NOMP- does not have a feedback mechanism, hence it lies in the class of forward greedy methods. Our convergence analysis also holds here.

A. Proof of Convergence:

A trivial upper bound on the number of iterations of NOMP is the number of the observations N . This is a direct result of solving the least squares at Step 10 of the algorithm. After N iterations, X is a square full-rank matrix (no frequency

is detected twice), hence the residue is zero for $(N + 1)^{th}$ iteration and the algorithm terminates.

The following theorem states another upper bound on the number of iterations, obtained by characterizing the amount by which the residual energy decreases when adding a new frequency to the set of estimated sinusoids.

Theorem 1: The reduction of residual energy due to one iteration (adding a new sinusoid) is at least τ . Consequently, $\min\{N, \|\mathbf{y}\|^2/\tau\}$ is an upper bound on the number of iterations of the algorithm.

Proof: The residue at m^{th} iteration of the algorithm is given by $\mathbf{y}_r(\mathcal{P}_m) = \mathbf{y} - \sum_{l=1}^{l=m} g_l \mathbf{x}(\omega_l)$. The energy of the residue in each iteration of the algorithm satisfies the following,

$$\begin{aligned} \|\mathbf{y}_r(\mathcal{P}_{m-1})\|^2 &\stackrel{(a)}{=} \|\mathbf{y}_r(\mathcal{P}'_m)\|^2 + G_{\mathbf{y}_r(\mathcal{P}_{m-1})}(\hat{\omega}) \\ &\stackrel{(b)}{\geq} \|\mathbf{y}_r(\mathcal{P}'''_m)\|^2 + G_{\mathbf{y}_r(\mathcal{P}_{m-1})}(\hat{\omega}) \\ &\stackrel{(c)}{\geq} \|\mathbf{y}_r(\mathcal{P}_m)\|^2 + G_{\mathbf{y}_r(\mathcal{P}_{m-1})}(\hat{\omega}) \\ &\stackrel{(d)}{\geq} \|\mathbf{y}_r(\mathcal{P}_m)\|^2 + \tau. \end{aligned} \quad (16)$$

where (a) follows from Step 5 of the algorithm where we project $\mathbf{y}_r(\mathcal{P}_{m-1})$ orthogonal to the subspace spanned by $\mathbf{x}(\hat{\omega})$ to get $\mathbf{y}_r(\mathcal{P}'_m)$. Inequalities in (b) follow from (RAC) checks performed whenever the single frequency refinement algorithm is invoked and (c) from the fact that least squares UPDATE can only lead to a decrease in energy of the residual signal. (d) is a direct consequence of the stopping criteria of the algorithm.

Inequality (17) shows that the reduction of the residual energy due to detecting a new sinusoids is always greater than τ . This result bounds the number of iterations of the algorithm from above by $\|\mathbf{y}\|^2/\tau$, proving convergence. Combining this observation with the trivial upper-bound N completes the proof.

B. Rate of Convergence:

We first bound the maximum of the GLRT cost function $G_{\mathbf{y}}(\omega)$ over the continuum of frequencies, in terms of that over the oversampled grid. To this end, we borrow ideas used in [8] (Appendix C) for proving a similar property for the dual atomic norm. We briefly introduce the notion of atomic norm (also known as dictionary norm) [37] and specialize to the line spectral estimation problem.

The atomic set of unit norm sinusoids is given by $\mathcal{A} = \{e^{j\phi} \mathbf{x}(\omega) : \phi, \omega \in [0, 2\pi)\}$. The atomic norm for \mathbf{y} is defined by

$$\|\mathbf{y}\|_{\mathcal{A}} \triangleq \inf \{t > 0 : \mathbf{y} \in t \text{ conv}(\mathcal{A})\}. \quad (18)$$

where $\text{conv}(\mathcal{A})$ denotes the convex hull of the points in \mathcal{A} . Since the centroid of the $\text{conv}(\mathcal{A})$ is at the origin, the atomic norm can be rewritten as [48], [49]

$$\|\mathbf{y}\|_{\mathcal{A}} \triangleq \inf \left\{ \sum_l |g_l| : \mathbf{y} = \sum_l g_l \mathbf{x}(\omega_l), \mathbf{x}(\omega_l) \in \mathcal{A} \right\}. \quad (19)$$

Note that this is not the ℓ_1 norm of \mathbf{y} , but the ℓ_1 norm on the coefficients of the representation of \mathbf{y} by elements of \mathcal{A} . The

atomic norm is typically small when its argument has a good sparse approximation [2]. The dual norm of $\|\cdot\|_{\mathcal{A}}$ is defined by

$$\|\mathbf{y}\|_{\mathcal{A}}^* = \sup_{\mathbf{a} \in \mathcal{A}} \Re\{\langle \mathbf{a}, \mathbf{y} \rangle\}. \quad (20)$$

It is easy to see that

$$\begin{aligned} \|\mathbf{y}\|_{\mathcal{A}}^* &= \sup_{\omega \in [0, 2\pi)} \sup_{\phi \in [0, 2\pi)} \Re\{e^{i\phi} \langle \mathbf{x}(\omega), \mathbf{y} \rangle\} \\ &= \sup_{\omega \in [0, 2\pi)} |\langle \mathbf{x}(\omega), \mathbf{y} \rangle| \\ &= \sup_{\omega \in [0, 2\pi)} \sqrt{G_{\mathbf{y}}(\omega)}. \end{aligned} \quad (21)$$

Directly borrowing from [8] gives us the following Theorem.

Theorem 2: [8] Maximizing the GLRT cost function (for the dictionary of unit norm sinusoids) over $[0, 2\pi)$ is consistent with that over the oversampled grid Ω with oversampling factor γ . That is, we have

$$\max_{\omega \in \Omega} \sqrt{G_{\mathbf{y}}(\omega)} \leq \sup_{\omega \in [0, 2\pi)} \sqrt{G_{\mathbf{y}}(\omega)} \quad (22)$$

$$\leq \left(1 - \frac{2\pi}{\gamma}\right)^{-1} \max_{\omega \in \Omega} \sqrt{G_{\mathbf{y}}(\omega)}. \quad (23)$$

See ([8], Appendix C) for a proof.

We need the following lemma to prove Theorem 3, which provides a pessimistic characterization of the convergence rate.

Lemma 1: Assume $\{a_n\}_{n \geq 0}$ is a decreasing sequence of nonnegative numbers such that $a_0 \leq U$ and

$$a_n \leq a_{n-1} \left(1 - \frac{a_{n-1}}{U}\right), \quad \forall n > 0,$$

then we have $a_n \leq \frac{U}{n+1}$ for all $n \geq 0$.

The proof is by induction [37]. Suppose $a_{n-1} \leq \frac{U}{n}$. Either $a_{n-1} \leq \frac{U}{n+1}$, in which case $a_n \leq \frac{U}{n+1}$, or $a_{n-1} \geq \frac{U}{n+1}$, in which case

$$\begin{aligned} a_n &\leq a_{n-1} \left(1 - \frac{a_{n-1}}{U}\right) \\ &\leq \frac{U}{n} \left(1 - \frac{\frac{U}{n+1}}{U}\right) = \frac{U}{n+1}. \end{aligned}$$

Hence, $a_n \leq \frac{U}{n+1}$ for all $n \geq 0$.

Theorem 3: For all \mathbf{y} such that $\|\mathbf{y}\|_{\mathcal{A}} < \infty$, the residual energy of NOMP at the m^{th} iteration satisfies

$$\|\mathbf{y}_r(\mathcal{P}_m)\| \leq (m+1)^{-1/2} \left(1 - \frac{2\pi}{\gamma}\right)^{-1} \|\mathbf{y}\|_{\mathcal{A}}. \quad (24)$$

Proof: From (16), we have

$$\|\mathbf{y}_r(\mathcal{P}_m)\|^2 \leq \|\mathbf{y}_r(\mathcal{P}_{m-1})\|^2 - G_{\mathbf{y}_r(\mathcal{P}_{m-1})}(\hat{\omega}). \quad (25)$$

$\mathbf{y}_r(\mathcal{P}_{m-1})$ is the result of projecting \mathbf{y} orthogonal to the subspace spanned by \mathcal{P}_{m-1} , therefore

$$\begin{aligned} \|\mathbf{y}_r(\mathcal{P}_{m-1})\|^2 &= \langle \mathbf{y}_r(\mathcal{P}_{m-1}), \mathbf{y} \rangle \\ &\stackrel{(a)}{\leq} \|\mathbf{y}\|_{\mathcal{A}} \|\mathbf{y}_r(\mathcal{P}_{m-1})\|_{\mathcal{A}}^* \\ &= \|\mathbf{y}\|_{\mathcal{A}} \sup_{\omega \in [0, 2\pi)} \sqrt{G_{\mathbf{y}_r(\mathcal{P}_{m-1})}(\omega)} \\ &\stackrel{(b)}{\leq} \|\mathbf{y}\|_{\mathcal{A}} \left(1 - \frac{2\pi}{\gamma}\right)^{-1} \max_{\omega \in \Omega} \sqrt{G_{\mathbf{y}_r(\mathcal{P}_{m-1})}(\omega)}, \end{aligned}$$

where (a) follows by Hölder's inequality [50], and (b) is by Theorem 2. Let $\eta \triangleq \|\mathbf{y}\|_{\mathcal{A}} \left(1 - \frac{2\pi}{\gamma}\right)^{-1}$. From Step 5 of the algorithm we have $\hat{\omega} = \arg \max_{\omega \in \Omega} \sqrt{G_{\mathbf{y}_r(\mathcal{P}_{m-1})}(\omega)}$, hence,

$$\|\mathbf{y}_r(\mathcal{P}_{m-1})\|^2 \leq \eta \sqrt{G_{\mathbf{y}_r(\mathcal{P}_{m-1})}(\hat{\omega})}. \quad (26)$$

Combining (25) and (26), gives

$$\|\mathbf{y}_r(\mathcal{P}_m)\|^2 \leq \|\mathbf{y}_r(\mathcal{P}_{m-1})\|^2 \left(1 - \eta^{-2} \|\mathbf{y}_r(\mathcal{P}_{m-1})\|^2\right). \quad (27)$$

Using Lemma 1 and the fact that

$$\|\mathbf{y}_r(\mathcal{P}_0)\|^2 = \|\mathbf{y}\|^2 \leq \|\mathbf{y}\|_{\mathcal{A}}^2 \leq \eta^2, \quad (28)$$

we have

$$\|\mathbf{y}_r(\mathcal{P}_m)\|^2 \leq \frac{\eta^2}{m+1}. \quad (29)$$

In other words,

$$\|\mathbf{y}_r(\mathcal{P}_m)\| \leq (m+1)^{-1/2} \left(1 - \frac{2\pi}{\gamma}\right)^{-1} \|\mathbf{y}\|_{\mathcal{A}}. \quad (30)$$

This proves the Theorem.

For the simulations we set $\gamma = 4$, but it is worth mentioning that, since we employ FFTs for detection over the oversampled grid, increasing γ has marginal effect on the runtime of the algorithm. In fact setting $\gamma = 20$ leads to only about 5% increase in runtime. If we compare (30) to the rate of convergence of OMP over the continuum given by [37],

$$\|\mathbf{y}_r(\mathcal{P}_m)\|^2 \leq (m+1)^{-1} \|\mathbf{y}\|_{\mathcal{A}}^2, \quad (31)$$

we see that by choosing γ large enough, the bound on the convergence rate specified in (30) approaches that of OMP over the continuum. Note that, in the derivation of (30) we have not considered the effect of the refinement steps, but imposing RAC ensures that refinements can only help speed up convergence.

C. Empirical rate of convergence:

We use numerical simulations to highlight the convergence benefits of the refinement steps in NOMP compared to the bound specified in (30). We plot the mean residual energy (averaged over 1000 runs) as a function of the number of iterations in a noiseless setting. We set $K = 16$, and $\Delta\omega_{\min} = 2.5 \times \Delta\omega_{\text{fft}}$. Figure 5 shows that light oversampling ($\gamma = 4$) followed by SINGLE REFINEMENT step (NOMP-) leads to a better convergence rate than having a large oversampling factor ($\gamma = 20$) and no refinements (DOMP). On the other hand, NOMP enjoys an extremely fast convergence rate due to the CYCLIC REFINEMENT step. In fact we see that for the setting where $\Delta\omega_{\min}$ is fairly large, the residual energy essentially drops to machine precision after 16 iterations, which equals the number of sinusoids in the mixture.

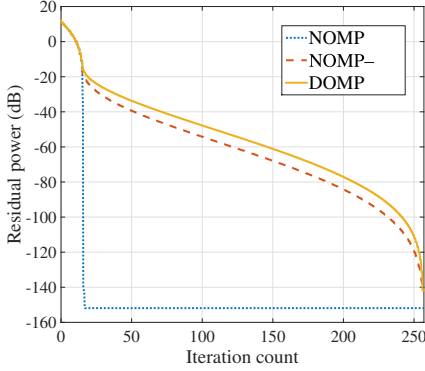


Fig. 5: Convergence rates at noiseless case

VI. SIMULATION RESULTS

Our performance measure is the mean squared error (MSE) of frequency estimation, and we compare the performance of NOMP against a number of benchmarks in various settings.

Benchmarks: The MUSIC algorithm is implemented using a modified version of the MATLAB routine `rootmusic`. The modifications are two-fold: (1) we use Minimum Description Length (MDL) criterion [51] for estimating the number of sinusoids in the mixture; (2) MUSIC operates by constructing an estimate of the autocorrelation matrix of the observed data vector. To this end, we use a sliding window of size W , to generate multiple snapshots from the observation vector $\mathbf{y} \in \mathbb{C}^N$. The choice of W has a significant impact on the performance of the algorithm, both in terms of estimation accuracy and run time: W too large leads to an inaccurate estimate of the autocorrelation matrix and the signal subspace, whereas too small a value effectively reduces the size of the observation window in time, degrading frequency estimation accuracy. We have found (empirically) that setting $W = 96$ results in the best estimation accuracy for the nominal settings in our simulations.

For sparse convex optimization, we consider Atomic norm Soft Thresholding (AST) [8], [52]. We use the alternating direction method of multipliers (ADMM) [53] to implement AST, as suggested in [8]. The updates in ADMM are iterative (typically 100 ~ 200 iterations for each simulation run), and each iteration includes an eigenvalue thresholding step for an $(N + 1) \times (N + 1)$ matrix. This $\mathcal{O}(N^3)$ step dominates the computational cost of the ADMM method, and becomes very expensive for large N .

The authors in [8] suggest solving Lasso as an alternative to the semi-definite program induced by AST. Lasso is solved on an oversampled frequency grid, using the highly optimized $\ell_2 - \ell_1$ software package SpaRSA [54]. We set the tolerance parameter to be 10^{-3} (other than this, we use the default parameters): smaller values of the tolerance parameter (e.g. 10^{-4}) increase runtime significantly, while providing marginal performance improvement. The regularization parameter in AST and Lasso formulation, suggested in [8], is set to

$$\text{reg} = \sigma \left(1 + \frac{1}{\log N} \right) \sqrt{\log N + \log(4\pi \log N)}.$$

The oversampling factor for the Lasso solver is set to 10 in our simulations. While increasing the oversampling factor improves Lasso performance, as mentioned in [8], for an oversampled grid, the frequencies estimated by Lasso cluster around the true frequencies. In order to avoid drastically overestimating model order, we implement a simple *clustering* scheme. In our simulations, we group frequencies into the largest number of clusters possible so that the frequency separation between any two sinusoids in different clusters is no smaller than $0.25 \times \Delta_{\text{dft}}$. After clustering, we update the gains of the cluster centers by solving the least squares problem $\min_{\{g_l\}} \|\mathbf{y} - \sum_l g_l \mathbf{x}(\omega_l)\|^2$.

Newtonized Lasso (NLasso): We also compare the results of NOMP algorithm with an extension of the Lasso formulation. In this scheme, we first apply the Lasso solver to identify the frequencies over the highly oversampled grid. Then we run the CYCLIC REFINEMENT step of the NOMP algorithm in order to refine the estimated frequencies, in order to prevent error floors caused by the off-grid effect. The parameters of Lasso are unchanged and the number of refinement steps is set to 5. Our simulations show that, for well-separated frequencies, the refinement step significantly improves estimation accuracy while incurring a small increase in runtime compared to Lasso. However, the benefit of refinement for Lasso diminishes as we increase the difficulty of the estimation problem (small $\Delta\omega_{\min}$).

Simulation set-up: We consider a mixture of $K = 16$ sinusoids of length $N = 256$. We perform 300 simulation runs for each of the four scenarios characterized by $\Delta\omega_{\min}$ and SNR values. The settings for different scenarios are summarized in Table I.


Scenarios	SNR (dB)	$\Delta\omega_{\min}/\Delta_{\text{dft}}$
1	SNR _{nom}	2.5
2	SNR _{nom}	0.5
3	Uniform[15, 35]	2.5
4	Uniform[15, 35]	0.5

TABLE I: Settings of different scenarios.

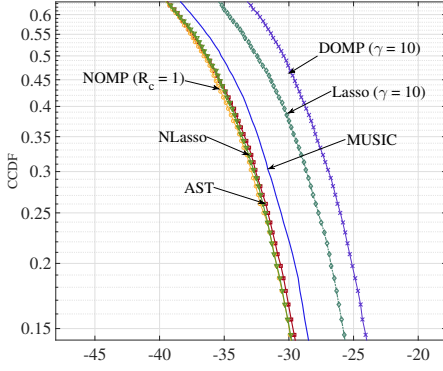
In Scenarios 1 and 2, the nominal SNR for each sinusoid is set as SNR_{nom} = 25 dB, whereas for Scenarios 3 and 4, the SNR values are chosen uniformly from [15, 35] dB, with mean equal to the nominal SNR of 25 dB. In each simulation run, the gain magnitudes are set to $|g_l| = \sigma\sqrt{\text{SNR}_l}$, while the phases $\{\angle g_l\}$ are chosen uniformly from $[0, 2\pi)$. The frequencies are chosen uniformly at random from $[0, 2\pi)^K$ while respecting the minimum separation constraints specified by $\Delta\omega_{\min}$ (if the minimum separation criterion is not met, we sample again from $[0, 2\pi)^K$). We plot the Complementary Cumulative Distribution Function (CCDF) of the squared frequency estimation error for all algorithms, along with the CRB (also a random variable, since it differs across realizations), and also compare against the DFT spacing, which is the resolution provided by coarse peak picking. See Appendix II for a quick overview of Cramér Rao Bound. The parameters of NOMP algorithm are set according to Table II for different scenarios.

NOMP	Scenario 1	Scenario 2	Scenario 3	Scenario 4
R_c	1	3	1	3
R_s	1	1	1	1
γ	4	4	4	4

TABLE II: NOMP parameters at different scenarios.



(a)



(b) Zoomed in

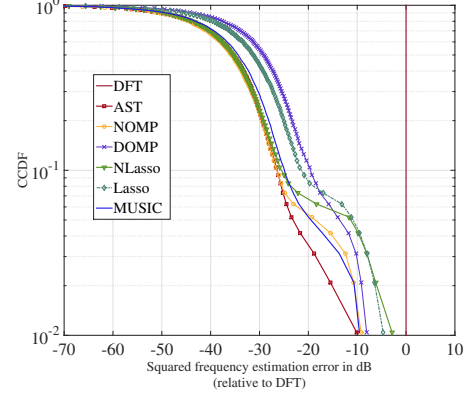
Fig. 6: CCDF of the frequency MSE for Scenario 1.

A. Frequency estimation accuracy

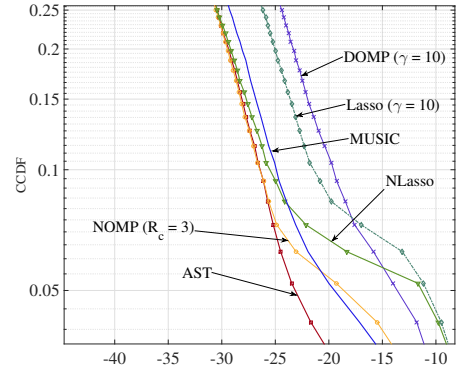
1) *Distribution of error:* Let us first examine the CCDF of squared frequency estimation error in each scenario. Figure 6 shows that NOMP, AST and NLasso lead to very similar error distributions in Scenario 1, while outperforming the other methods. Unlike Lasso and DOMP, which suffer from the off-grid effect, MUSIC picks frequencies over the continuum, and achieves better estimation accuracy. Another observation is that if the frequencies are well separated (as in Scenario 1), then adding a refinement stage at the output of Lasso leads to a significant improvement in estimation accuracy. As we move to more difficult Scenarios, however, the performance of NLasso degrades compared to NOMP and AST as shown in Figure 7 for Scenario 2. The reason is that the refinement stage of NLasso is able to improve the estimation accuracy only when the initial estimates provided by Lasso are close to the true frequencies. When Lasso fails in providing good initial estimates, there is no benefit in locally refining the frequencies.

Figure 8 shows the distribution of error in Scenario 3.

We see that the overall gap in the performance of different algorithms is reduced compared to the first two scenarios, with AST, NOMP and NLasso still achieving the highest estimation accuracy. In Scenario 4, NOMP achieves superior performance compared to all of the other methods, as shown in Figure 9.



(a)



(b) Zoomed in

Fig. 7: CCDF of the frequency MSE for Scenario 2.

2) *Mean squared error:* Here we examine the performance of different algorithms in terms of frequency estimation accuracy by looking at the normalized Mean Squared Error (MSE), defined by $\mathbb{E}[(\omega_{true} - \omega_{est})^2] / \Delta_{dft}^2$, in different scenarios. In Scenarios 1 and 3, where frequencies are well-separated, one hopes to get to the estimation accuracy of a single sinusoid (as if the other sinusoids did not exist). We therefore use the CRB and ZZB corresponding to a *single sinusoid*, computed by (10) and (11), respectively, as measures of optimality. In Scenarios 2 and 4, where frequencies can get close to one another, we compute the CRB empirically for each realization of the problem, and employ the mean CRB as a measure of optimality.

Figure 10 shows the MSE of frequency estimation in Scenarios 1 and 2, for SNR_{nom} taking values from 13 dB to 35 dB. In Scenario 1, if the nominal SNR is high enough, AST, NOMP, and NLasso all achieve the CRB. As we decrease the nominal SNR, all of the algorithms exhibit threshold behavior, well predicted by the ZZB threshold. The threshold SNR of AST is lower than that of other methods, showing its noise

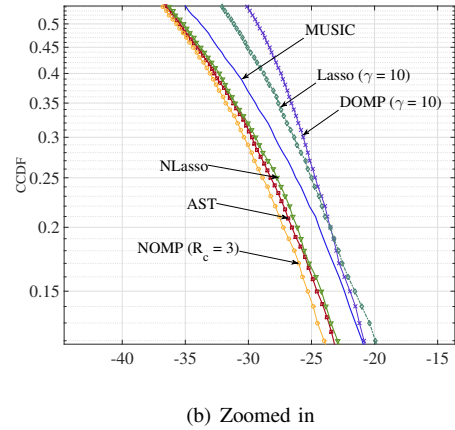
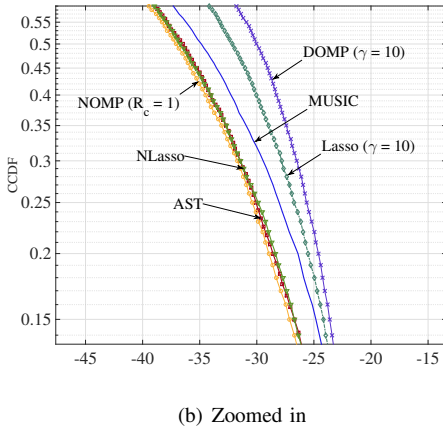




Fig. 8: CCDF of the frequency MSE for Scenario 3.

Fig. 9: CCDF of the frequency MSE for Scenario 4.

resilience. MUSIC does not achieve the CRB, but closely follows the bound for all SNR values in this scenario. DOMP and Lasso on the other hand, reach performance floors. We examine this floor more closely in Section VI-C to see whether it is a fundamental algorithmic limitation, or happens due to the off-grid effect. Figure 10-b corresponds to Scenario 2, where the separation between frequencies can be very small. We see that MUSIC is the only algorithm that benefits from increasing the nominal SNR. This goes back to the asymptotic optimality of MUSIC: for $SNR \rightarrow \infty$ and $K \ll N$, MUSIC is able to precisely determine the frequencies in the mixture, regardless of the separation between them [11].


In Scenarios 3 and 4, the SNRs are drawn independently and uniformly at random from the interval $[15, 35]$ dB. In order to evaluate the frequency MSE in these scenarios, we fix the SNR of one of the sinusoids in the mixture at a given value, while letting the other $(K - 1)$ SNRs to be realized randomly. Figure 11-a shows the frequency MSE curves corresponding to Scenario 3. As we expected from error distribution plots in Figure 8, the gap in the performance of different algorithms has decreased compared to the first scenario. Figure 11-b corresponds to Scenario 4, in which NOMP outperforms all of the other algorithms, and tightly follows the CRB. This indicates that NOMP is highly successful in exploiting the disparity in SNRs across sinusoids in the mixture, in order

to estimate closely-spaced frequencies. AST achieves the best performance at very low SNRs; however, its MSE curve stays bounded away from the CRB, with an expanding gap as we increase SNR.


3) *Number of cycles of Newton refinement (R_c):* We have seen that NOMP is able to achieve the CRB in Scenarios 1 and 3, with one cycle of refining the sinusoids in each iteration, i.e., $R_c = 1$. In this subsection, we want to highlight the effect of increasing R_c in improving the frequency estimation accuracy of NOMP in Scenarios 2 and 4 where $\Delta\omega_{\min} = 0.5 \times \Delta\omega_{\text{dft}}$. Figure 12 shows the frequency MSE of NOMP for $R_c \in \{1, 3, 5\}$. We see that NOMP enjoys the benefits of having more rounds of refinement, but exhibits diminishing return as R_c increases. In particular, increasing R_c beyond 3 cycles gives marginal improvement in estimation accuracy.

B. Computation time


Table III summarizes the time needed for running 300 simulations for each of the algorithms in different scenarios. We see that DOMP is extremely fast at the expense of estimation accuracy. NOMP is faster than all of the other methods in all four Scenarios, while achieving remarkable frequency estimation accuracy. As we discussed in Section II-C, the CYCLIC REFINEMENT step has the complexity $\mathcal{O}(R_c R_s K^2 N)$, and




(a) Scenario 1



(a) Scenario 3



(b) Scenario 2



(b) Scenario 4

Fig. 10: Normalized frequency MSE for Scenarios 1 and 2.

Fig. 11: Normalized frequency MSE for Scenarios 3 and 4.

dominates the computational cost of NOMP. Table IV shows the time needed for 300 simulation runs of NOMP for different values of R_c . Note that as we increase the difficulty of the estimation scenario, AST, Lasso and NLasso tend to take more time, while MUSIC and NOMP are unaffected.

Time [sec]	NOMP	AST	NLasso	Lasso	DOMP	MUSIC
Scenario 1	6.92	1.09e3	29.68	26.63	2.66	19.83
Scenario 2	14.26	1.02e3	29.63	27.62	2.81	20.15
Scenario 3	6.88	1.12e3	34.32	32.06	2.79	20.07
Scenario 4	14.19	1.18e3	36.37	33.71	2.80	19.69

TABLE III: Time [sec] for 300 runs of each algorithm. Parameters of NOMP are set by Table II

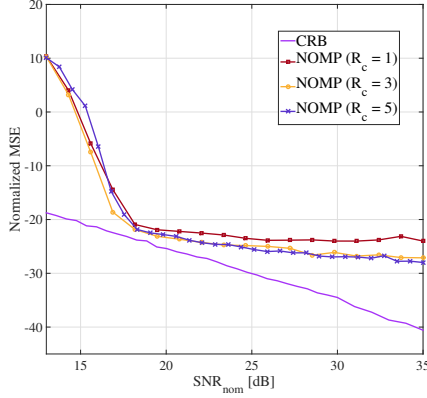
Time [sec]	$R_c = 1$	$R_c = 3$	$R_c = 5$
Scenario 1	6.92	14.24	21.67
Scenario 2	7.01	14.26	21.71
Scenario 3	6.88	14.21	21.61
Scenario 4	6.96	14.19	21.50

TABLE IV: Time [sec] for 300 runs of NOMP algorithm, for different numbers of cyclic refinements R_c .


C. Asymptotic regime

It is interesting to see the effect of increasing the oversampling factor for Lasso and DOMP. Figure 13 corresponds to Scenario 1, with the change being that the oversampling factor for these two algorithms is increased. We observe an improvement in the performance of both algorithms in terms of estimation accuracy, with that for Lasso being especially significant. Comparing these results to those in Figure 10-a, we see that the MSE performance of DOMP is marginally improved by increasing γ from 10 to 100. This performance plateau shows a fundamental *algorithmic limitation* of DOMP, and highlights the critical role of cyclic Newton refinements in NOMP. In other words, the performance limitation of DOMP is not just due to the off-grid error, but also a consequence of making “hard-decisions” at each iteration. The computational complexity of DOMP is insensitive to oversampling factor. For example, computation time increases from 3.25 to about 7.96 seconds as we go from $\gamma = 20$ to $\gamma = 100$.

In theory, Lasso is only limited by the oversampling factor γ , and as $\gamma \rightarrow \infty$, the result of Lasso converges to that of AST, as a consequence of the convergence of the corresponding atomic norms [8]. In practice, however, the performance of Lasso solved on a moderate size grid might be far from that of AST. Figure 13 shows that, at large enough oversampling



(a) Scenario 2



(b) Scenario 4

Fig. 12: Performance improvement of NOMP with increasing the number of cyclic refinements in Scenarios 2 and 4.

factors, Lasso approaches the CRB, but the computational cost becomes prohibitive. For example, the computation time of Lasso for 300 runs, increases from 72.16 seconds to 191.22 seconds as we go from $\gamma = 20$ to $\gamma = 50$.




Fig. 13: Frequency MSE for Scenario 1 and highly over-sampled grid for Lasso and DOMP.

D. Model order estimation

Estimating the model order K (the true number of non-zero atoms in the mixture) has significant importance in sparse approximation. Here we examine the Cumulative Distribution Function (CDF) of the estimated model order by different algorithms. As shown in Figure 14, both AST and MUSIC perform well, with MUSIC performance slightly degrading in scenarios with small $\Delta\omega_{\min}$. Note that the mean of the distributions are very close to the truth (small bias), and they have a small spread around the mean (small variance). Figure 15 shows the model order estimates for NOMP using both CFAR and BIC-based stopping criteria. We see that both criteria are very accurate in all four scenarios.

On the other hand, Lasso and DOMP perform poorly in estimating the model order, especially in scenarios with small $\Delta\omega_{\min}$. As shown in Figure 16, DOMP tends to overestimate the model order when some of the frequencies are placed too closely. This is again the result of a fundamental algorithmic limitation. DOMP does not allow for correcting errors that have happened in the previous iterations; instead, it tries to explain the residual energy by overestimating the number of non-zero atoms.

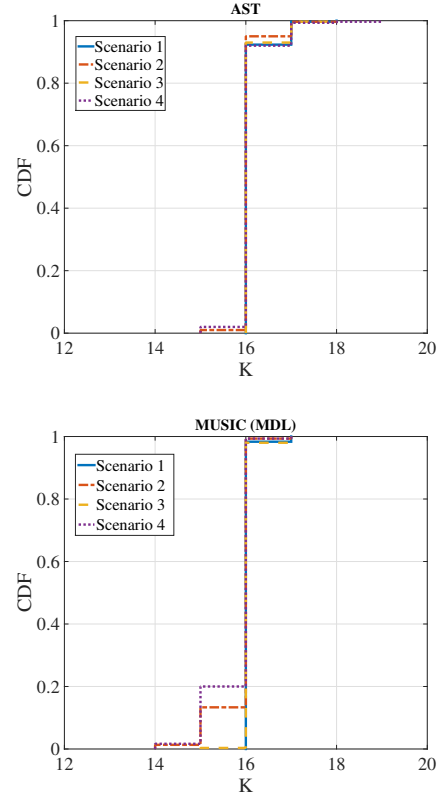


Fig. 14: CDF of the estimates of the model order for AST and MUSIC.

In Scenario 2, Lasso tends to underestimate the model order. The main reason is that for two closely spaced frequencies, Lasso generates two overlapping clusters of estimated frequencies, which are later replaced by a single frequency by our *clustering* algorithm. On the other hand, if we do not employ clustering, Lasso significantly overestimates the model order.

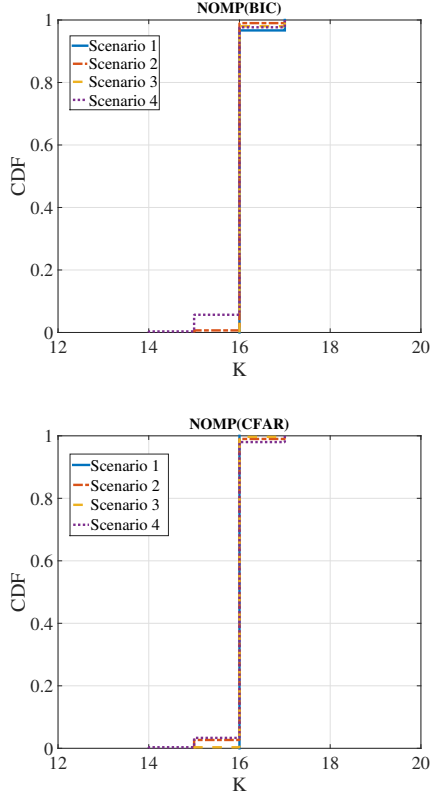


Fig. 15: CDF of the estimates of the model order for NOMP using CFAR and BIC criteria.

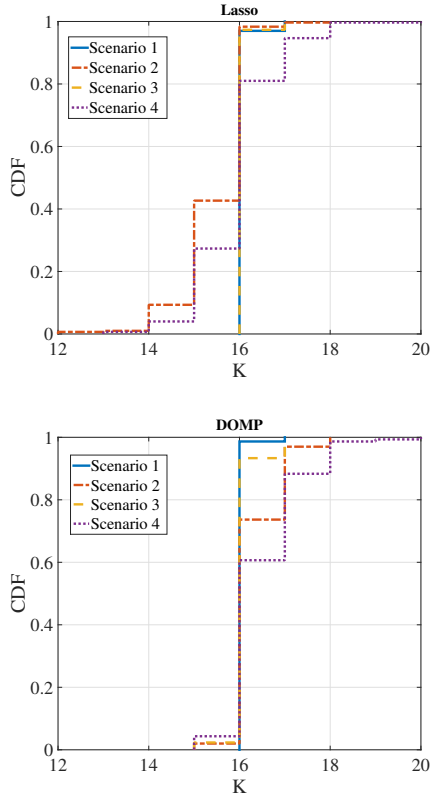


Fig. 16: CDF of the estimates of the model order for Lasso and DOMP with $\gamma = 20$.

VII. EXTENSIONS OF THE ALGORITHM

In this section we point out an immediate extension of NOMP algorithm. Specifically, we can replace the manifold of sinusoids $\{\mathbf{x}(\omega) : \omega \in [0, 2\pi)\}$ by $\{A\mathbf{x}(\omega)/\|A\mathbf{x}(\omega)\| : \omega \in [0, 2\pi)\}$, where $A \in \mathbb{C}^{M \times N}$ is a known measurement matrix. This is motivated by the measurement setup $\mathbf{y} = \sum_{l=1}^K g_l A\mathbf{x}(\omega_l)$.

Compressive measurements

We consider a compressive measurement model in which the number of measurements $M \ll N$. As in the bulk of literature on compressive sensing, we assume that the elements of A are chosen i.i.d from appropriate zero-mean distributions (with variance conveniently scaled to $1/N$ such as $\text{Uniform}\{\pm 1/\sqrt{N}\}$, $\text{Uniform}\{\pm 1/\sqrt{N}, \pm j/\sqrt{N}\}$, $\mathcal{N}(0, 1/N)$, etc.), so that certain concentration results hold. It has been shown in [41] that when A satisfies certain isometry conditions (related to the estimation problem at hand), the CRB and ZZB are approximately preserved for compressive estimation, except for an SNR degradation of M/N due to randomly projecting down the signal to a smaller space. The number of compressive measurements needed to estimate continuous frequencies scales as $M = \mathcal{O}(K \log N)$ [41]. In order to get concrete numerical intuition, we set $M = N/4 = 64$ with the elements of A chosen uniformly and independently at random from $\{\pm 1/\sqrt{N}, \pm j/\sqrt{N}\}$ and run NOMP with atoms $\{\mathbf{s}(\omega) = A\mathbf{x}(\omega)/\|A\mathbf{x}(\omega)\| : \omega \in [0, 2\pi)\}$. We consider Scenario 1 with $R_c = 3$, with number of sinusoids set to $K = 13$ and $K = 16$. Our algorithm approaches the CRB in the setting where $K = 13$, whereas we incur large estimation errors when $K = 16$; see Figure 17. The large estimation errors for $K = 16$ occur because the compressive measurement matrix A fails to preserve the structure of the estimation problem: $M = 64$ compressive measurements is too few for $K = 16$ sinusoids.




Fig. 17: CCDF of the frequency MSE for Scenario 1 with Compressive measurements.

VIII. CONCLUSIONS

We have shown that NOMP is fast and near-optimal for frequency estimation in AWGN. It performs better than classical

methods such as MUSIC, and more recent convex optimization based methods such as AST, in terms of both estimation accuracy and run time. The algorithm uses a fundamental element of OMP, ensuring that the residue is orthogonal to the signal space spanned by the current set of frequencies. However, NOMP avoids the error floors of naively discretized OMP by refinement over the continuum. Specifically, it searches for a signal subspace in the “neighborhood” of our current signal space which can better explain the observed measurements. The algorithm has a natural “decision feedback” interpretation in which it gives the already detected sinusoids a chance to adjust their frequencies in light of new evidence which is presented in the form of an updated residue after we add the next sinusoid.

We believe that the ideas underlying NOMP are broadly applicable to problems of sparse approximation over continuous dictionaries, but further work is required to justify this assertion. An open problem is to go beyond the pessimistic convergence rate estimates provided here by analytically quantifying the benefits of refinement.

APPENDIX I: BAYESIAN INFORMATION CRITERIA

The BIC balances an increase in the likelihood with the number of parameters used to achieve that increase. Namely,

$$\text{BIC} = 2 \ln \frac{L_2(\theta_1, \dots, \theta_{m_2})}{L_1(\phi_1, \dots, \phi_{m_1})} - (m_2 - m_1) \ln(N), \quad (32)$$

where $L_1(\phi_1, \dots, \phi_{m_1})$ and $L_2(\theta_1, \dots, \theta_{m_2})$ are likelihood functions and $\{\phi_i\}_{i=1}^{m_1}$ and $\{\theta_j\}_{j=1}^{m_2}$ are their corresponding parameters. For our measurement model, $\mathbf{y} = \sum_{l=1}^K g_l \mathbf{x}(\omega_l) + \mathbf{z}$, where $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbb{I}_N)$, we have

$$L(\{g_l, \omega_l\}_{l=1}^K) = \frac{1}{(\pi\sigma^2)^N} \exp(-\|\mathbf{y}_r\|^2/\sigma^2),$$

where $\mathbf{y}_r = \mathbf{y} - \sum_{l=1}^K g_l \mathbf{x}(\omega_l)$ is the residual. Therefore, $\text{BIC} = \frac{2}{\sigma^2} \Delta \|\mathbf{y}_r\|^2 - 2 \ln(N)$, where $\Delta \|\mathbf{y}_r\|^2 = \|\mathbf{y}_r(\text{old})\|^2 - \|\mathbf{y}_r(\text{new})\|^2$ is the reduction in residual energy after detecting a new sinusoid. If $\text{BIC} > 10$, it is a strong evidence that the most recent reduction in the residual energy corresponds to a newly detected sinusoid. Therefore, when $\text{BIC} < 10$, we stop the algorithm.

APPENDIX II: CRAMÉR RAO BOUND

In this Appendix we first review the Cramér Rao Bound [7] for an unbiased estimator in a general setting, then specialize to the frequency estimation problem. Let $\mathbf{a} \in \mathbb{R}^\lambda$. If $\mathbf{a}^T \hat{\theta}(\mathbf{y})$ is an unbiased estimator of $\mathbf{a}^T \theta$, then the variance of the estimator given by $\mathbb{E}_{\mathbf{y}|\theta} \left[\left(\mathbf{a}^T \hat{\theta}(\mathbf{y}) - \mathbf{a}^T \theta \right)^2 \right]$ is lower bounded by $\mathbf{a}^T F^{-1}(\theta) \mathbf{a}$, where $F(\theta)$ is the Fisher Information Matrix. The $(m, n)^{\text{th}}$ element of $F(\theta)$ is given by

$$F_{m,n}(\theta) = \mathbb{E}_{\mathbf{y}|\theta} \left\{ \frac{\partial \ln p(\mathbf{y}|\theta)}{\partial \theta_m} \frac{\partial \ln p(\mathbf{y}|\theta)}{\partial \theta_n} \right\}. \quad (33)$$

For parameter estimation in additive white Gaussian noise i.e. $\mathbf{y} = \mathbf{s}(\theta) + \mathbf{z}$, $\mathbf{z} \sim \mathcal{CN}(0, \sigma^2 \mathbb{I})$, equation (33) simplifies to,

$$F_{m,n}(\theta) = \frac{2}{\sigma^2} \Re \left\{ \left(\frac{\partial \mathbf{s}(\theta)}{\partial \theta_m} \right)^H \frac{\partial \mathbf{s}(\theta)}{\partial \theta_n} \right\}. \quad (34)$$

For our frequency estimation problem in (2), θ is the vector of all parameters, namely $\{|g_l|, \angle g_l, \omega_l : l = 1, \dots, K\}$, and $\mathbf{s}(\theta) = \sum_{l=1}^K |g_l| e^{j\angle g_l} \mathbf{x}(\omega_l)$. We form $F(\theta)$ for this measurement model then choose those diagonal elements of $F^{-1}(\theta)$ that correspond to the frequencies $\{\omega_l : l = 1, \dots, K\}$.

REFERENCES

- [1] “NOMP Software,” <https://bitbucket.org/wcs/spectralestimation/continuous-frequency-estimation/src/NOMP>, 2016.
- [2] J. A. Tropp and S. J. Wright, “Computational methods for sparse solution of linear inverse problems,” *Proceedings of the IEEE*, vol. 98, no. 6, pp. 948–958, 2010.
- [3] J. A. Tropp and A. C. Gilbert, “Signal recovery from random measurements via orthogonal matching pursuit,” *Information Theory, IEEE Transactions on*, vol. 53, no. 12, pp. 4655–4666, 2007.
- [4] Y. Chi, L. L. Scharf, A. Pezeshki, and A. R. Calderbank, “Sensitivity to basis mismatch in compressed sensing,” *Signal Processing, IEEE Transactions on*, vol. 59, no. 5, pp. 2182–2195, 2011.
- [5] L. Jacques and C. De Vleeschouwer, “A geometrical study of matching pursuit parametrization,” *Signal Processing, IEEE Transactions on*, vol. 56, no. 7, pp. 2835–2848, 2008.
- [6] S. G. Mallat and Z. Zhang, “Matching pursuits with time-frequency dictionaries,” *Signal Processing, IEEE Transactions on*, vol. 41, no. 12, pp. 3397–3415, 1993.
- [7] H. L. Van Trees, *Detection, estimation, and modulation theory*. John Wiley & Sons, 2004.
- [8] B. N. Bhaskar, G. Tang, and B. Recht, “Atomic norm denoising with applications to line spectral estimation,” *Signal Processing, IEEE Transactions on*, vol. 61, no. 23, pp. 5987–5999, 2013.
- [9] R. O. Schmidt, “Multiple emitter location and signal parameter estimation,” *Antennas and Propagation, IEEE Transactions on*, vol. 34, no. 3, pp. 276–280, 1986.
- [10] R. Roy and T. Kailath, “Esprit-estimation of signal parameters via rotational invariance techniques,” *Acoustics, Speech and Signal Processing, IEEE Transactions on*, vol. 37, no. 7, pp. 984–995, 1989.
- [11] H. Teutsch, *Modal Array Signal Processing: Principles and Applications of Acoustic Wavefield Decomposition*, ser. Lecture Notes in Control and Information Sciences. Springer Berlin Heidelberg, 2007.
- [12] K. Duda, L. B. Magalas, M. Majewski, and T. P. Zieliński, “Dft-based estimation of damped oscillation parameters in low-frequency mechanical spectroscopy,” *Instrumentation and Measurement, IEEE Transactions on*, vol. 60, no. 11, pp. 3608–3618, 2011.
- [13] T. Zieliński and K. Duda, “Frequency and damping estimation methods—an overview,” *Metrology and Measurement Systems*, vol. 18, no. 4, pp. 505–528, 2011.
- [14] D. Vu, L. Xu, M. Xue, and J. Li, “Nonparametric missing sample spectral analysis and its applications to interrupted sar,” *Selected Topics in Signal Processing, IEEE Journal of*, vol. 6, no. 1, pp. 1–14, 2012.
- [15] E. J. Candès and C. Fernandez-Granda, “Towards a mathematical theory of super-resolution,” *Communications on Pure and Applied Mathematics*, vol. 67, no. 6, pp. 906–956, 2014.
- [16] C. Fernandez-Granda, “Super-resolution of point sources via convex programming,” *arXiv preprint arXiv:1507.07034*, 2015.
- [17] E. J. Candès and C. Fernandez-Granda, “Super-resolution from noisy data,” *Journal of Fourier Analysis and Applications*, vol. 19, no. 6, pp. 1229–1254, 2013.
- [18] G. Tang, B. N. Bhaskar, and B. Recht, “Sparse recovery over continuous dictionaries—just discretize,” in *Signals, Systems and Computers, 2013 Asilomar Conference on*. IEEE, 2013, pp. 1043–1047.
- [19] J. Friedman, T. Hastie, H. Höfling, and R. Tibshirani, “Pathwise coordinate optimization,” *The Annals of Applied Statistics*, vol. 1, no. 2, pp. 302–332, 2007.
- [20] T. T. Wu and K. Lange, “Coordinate descent algorithms for lasso penalized regression,” *The Annals of Applied Statistics*, pp. 224–244, 2008.
- [21] A. Panahi, M. Viberg, and B. Hassibi, “A numerical implementation of gridless compressed sensing,” in *Acoustics, Speech and Signal Processing (ICASSP), 2015 IEEE International Conference on*. IEEE, 2015, pp. 3342–3346.
- [22] D. Needell and J. A. Tropp, “Cosamp: Iterative signal recovery from incomplete and inaccurate samples,” *Applied and Computational Harmonic Analysis*, vol. 26, no. 3, pp. 301–321, 2009.

- [23] T. Zhang, "Adaptive forward-backward greedy algorithm for learning sparse representations," *Information Theory, IEEE Transactions on*, vol. 57, no. 7, pp. 4689–4708, 2011.
- [24] T. Blumensath and M. E. Davies, "Iterative thresholding for sparse approximations," *Journal of Fourier Analysis and Applications*, vol. 14, no. 5-6, pp. 629–654, 2008.
- [25] P. Jain, A. Tewari, and I. S. Dhillon, "Orthogonal matching pursuit with replacement," in *Advances in Neural Information Processing Systems*, 2011, pp. 1215–1223.
- [26] J. Li and P. Stoica, "Efficient mixed-spectrum estimation with applications to target feature extraction," *Signal Processing, IEEE Transactions on*, vol. 44, no. 2, pp. 281–295, 1996.
- [27] T. Yardibi, J. Li, P. Stoica, M. Xue, and A. B. Baggeroer, "Source localization and sensing: A nonparametric iterative adaptive approach based on weighted least squares." Institute of Electrical and Electronics Engineers, 2010.
- [28] A. Fannjiang and W. Liao, "Coherence pattern-guided compressive sensing with unresolved grids," *SIAM Journal on Imaging Sciences*, vol. 5, no. 1, pp. 179–202, 2012.
- [29] J. A. Tropp, "Greed is good: Algorithmic results for sparse approximation," *Information Theory, IEEE Transactions on*, vol. 50, no. 10, pp. 2231–2242, 2004.
- [30] M. A. Davenport and M. B. Wakin, "Analysis of orthogonal matching pursuit using the restricted isometry property," *Information Theory, IEEE Transactions on*, vol. 56, no. 9, pp. 4395–4401, 2010.
- [31] T. J. Abatzoglou, "A fast maximum likelihood algorithm for frequency estimation of a sinusoid based on newton's method," *Acoustics, Speech and Signal Processing, IEEE Transactions on*, vol. 33, no. 1, pp. 77–89, 1985.
- [32] P. Bidigare, U. Madhow, R. Mudumbai, and D. Scherber, "Attaining fundamental bounds on timing synchronization," in *Acoustics, Speech and Signal Processing (ICASSP), 2012 IEEE International Conference on*. IEEE, 2012, pp. 5229–5232.
- [33] D. Ramasamy, S. Venkateswaran, and U. Madhow, "Compressive adaptation of large steerable arrays," in *Information Theory and Applications Workshop (ITA), 2012*. IEEE, 2012, pp. 234–239.
- [34] —, "Compressive tracking with 1000-element arrays: A framework for multi-gbps mm wave cellular downlinks," in *Communication, Control, and Computing (Allerton), 2012 50th Annual Allerton Conference on*. IEEE, 2012, pp. 690–697.
- [35] B. Mamandipoor, D. Ramasamy, and U. Madhow, "Frequency estimation for a mixture of sinusoids: A near-optimal sequential approach," in *3rd IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, 2015.
- [36] T. L. Hansen, M. A. Badiu, B. H. Fleury, and B. Rao, "A sparse bayesian learning algorithm with dictionary parameter estimation," in *Sensor Array and Multichannel Signal Processing Workshop (SAM), 2014 IEEE 8th*. IEEE, 2014, pp. 385–388.
- [37] A. R. Barron, A. Cohen, W. Dahmen, and R. A. DeVore, "Approximation and learning by greedy algorithms," *The annals of statistics*, pp. 64–94, 2008.
- [38] D. Slepian, "Prolate spheroidal wave functions, fourier analysis, and uncertainty: The discrete case," *Bell System Technical Journal*, vol. 57, no. 5, pp. 1371–1430, 1978.
- [39] F. B. Hildebrand, *Introduction to numerical analysis*. Courier Corporation, 1987.
- [40] H. L. Van Trees and K. L. Bell, "Bayesian bounds for parameter estimation and nonlinear filtering/tracking," *AMC*, vol. 10, p. 12, 2007.
- [41] D. Ramasamy, S. Venkateswaran, and U. Madhow, "Compressive parameter estimation in awgn," *Signal Processing, IEEE Transactions on*, vol. 62, no. 8, pp. 2012–2027, 2014.
- [42] P. P. Gandhi and S. A. Kassam, "Analysis of cfar processors in homogeneous background," *Aerospace and Electronic Systems, IEEE Transactions on*, vol. 24, no. 4, pp. 427–445, 1988.
- [43] H. Rohling, "Radar cfar thresholding in clutter and multiple target situations," *Aerospace and Electronic Systems, IEEE Transactions on*, no. 4, pp. 608–621, 1983.
- [44] B. Eisenberg, "On the expectation of the maximum of iid geometric random variables," *Statistics & Probability Letters*, vol. 78, no. 2, pp. 135–143, 2008.
- [45] G. Schwarz, "Estimating the dimension of a model," *The annals of statistics*, vol. 6, no. 2, pp. 461–464, 1978.
- [46] H. Akaike, "A new look at the statistical model identification," *Automatic Control, IEEE Transactions on*, vol. 19, no. 6, pp. 716–723, 1974.
- [47] V. S. Ryaben'kii and S. V. Tsynekov, *A theoretical introduction to numerical analysis*. CRC Press, 2006.
- [48] V. Chandrasekaran, B. Recht, P. A. Parrilo, and A. S. Willsky, "The convex geometry of linear inverse problems," *Foundations of Computational mathematics*, vol. 12, no. 6, pp. 805–849, 2012.
- [49] F. F. Bonsall, "A general atomic decomposition theorem and banach's closed range theorem," *The Quarterly Journal of Mathematics*, vol. 42, no. 1, pp. 9–14, 1991.
- [50] P. Shah, B. N. Bhaskar, G. Tang, and B. Recht, "Linear system identification via atomic norm regularization," in *Decision and Control (CDC), 2012 IEEE 51st Annual Conference on*, 2012, pp. 6265–6270.
- [51] M. Wax and T. Kailath, "Detection of signals by information theoretic criteria," *Acoustics, Speech and Signal Processing, IEEE Transactions on*, vol. 33, no. 2, pp. 387–392, 1985.
- [52] G. Tang, B. N. Bhaskar, and B. Recht, "Near minimax line spectral estimation," *Information Theory, IEEE Transactions on*, vol. 61, no. 1, pp. 499–512, 2015.
- [53] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed optimization and statistical learning via the alternating direction method of multipliers," *Foundations and Trends® in Machine Learning*, vol. 3, no. 1, pp. 1–122, 2011.
- [54] S. J. Wright, R. D. Nowak, and M. A. Figueiredo, "Sparse reconstruction by separable approximation," *Signal Processing, IEEE Transactions on*, vol. 57, no. 7, pp. 2479–2493, 2009.



Babak Mamandipoor received his B.Sc. degree in electrical engineering from Iran University of Science and Technology (IUST), Tehran, Iran in 2010, and the M.A.Sc. degree in electrical and computer engineering from University of Waterloo, Canada, in 2013. He started his Ph.D. in communications and signal processing at the University of California, Santa Barbara (UCSB), USA, in Fall 2013. His main research interests lie within the areas of millimeter wave communications and imaging systems, sparse approximation and information theory.



Dinesh Ramasamy received his Bachelor's degree in electrical engineering from the Indian Institute of Technology, Madras, in 2009 and his Ph.D. degree from the University of California, Santa Barbara (UCSB) in 2014. He is interested in machine learning, millimeter wave communication and inference from large scale networks.



Upamanyu Madhow is Professor of Electrical and Computer Engineering at the University of California, Santa Barbara. His research interests broadly span communications, signal processing and networking, with current emphasis on millimeter wave communication, and on distributed and bio-inspired approaches to networking and inference. He received his bachelor's degree in electrical engineering from the Indian Institute of Technology, Kanpur, in 1985, and his Ph. D. degree in electrical engineering from the University of Illinois, Urbana-Champaign in

1990. He has worked as a research scientist at Bell Communications Research, Morristown, NJ, and as a faculty at the University of Illinois, Urbana-Champaign. Dr. Madhow is a recipient of the 1996 NSF CAREER award, and co-recipient of the 2012 IEEE Marconi prize paper award in wireless communications. He has served as Associate Editor for the IEEE Transactions on Communications, the IEEE Transactions on Information Theory, and the IEEE Transactions on Information Forensics and Security. He is the author of two textbooks published by Cambridge University Press, *Fundamentals of Digital Communication* (2008) and *Introduction to Communication Systems* (2014).