Frequency Estimation for a Mixture of Sinusoids: A Near-Optimal Sequential Approach

B. Mamandipoor, D. Ramasamy, U. Madhow

ECE Department, University of California, Santa Barbara, CA, USA
Emails: {bmamandi, dineshr, madhow}@ece.ucsb.edu

Abstract—**We propose a fast sequential algorithm for the fundamental problem of estimating continuous-valued frequencies and amplitudes using samples of a noisy mixture of sinusoids. Each step consists of two phases: detection of a new sinusoid, and refining the parameters of already detected sinusoids. The detection phase is performed on an oversampled DFT grid, while the refinement phase enables continuous-valued estimation, thus avoiding basis mismatch. By benchmarking against the Cramér Rao Bound, we show that the proposed algorithm achieves near-optimal performance under a variety of settings. We also compare our algorithm with the classical MUSIC, and more recent Lasso algorithms in terms of estimation accuracy and computational complexity.**

I. Introduction

Frequency estimation for 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 this paper, we propose an algorithm to estimate frequencies from N equi-spaced noisy samples in time, denoted by $\mathbf{y} \in \mathbb{C}^N$. We denote the unit norm sinusoid $[1 \ e^{j\omega} \ \cdots \ e^{j(N-1)\omega}]^T/\sqrt{N}$ of frequency ω by $\mathbf{x}(\omega)$. The signal is a mixture of K sinusoids:

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

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 $\{(g_l, \omega_l) : l = 1, 2, \ldots, K\}$ 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 $h(t) = \sum_{l=1}^{K} g_l \delta(t - \tau_l)$ (e.g., for communication, or for determining target ranges in radar), the channel transfer function $H(f) = \sum_{l=1}^{K} g_l e^{-j2\pi f \tau_l}$. It is easy to see that 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.

Contributions: Our key contributions are as follows:
(1) We propose a low-complexity sequential algorithm which employs a version of matching pursuit [1] for coarse detection on a grid, followed by Newton refinements, with the second phase being crucial for avoiding basis mismatch [2] and obtaining accuracies far better than would be possible by optimizing over a discrete grid. We do not require explicit estimates of model order, and provide a stopping criterion based on CFAR (constant false alarm rate). A freely downloadable software package implementing the proposed algorithm can be found in [20].
(2) We show that the algorithm is near-optimal by numerical comparisons against the Cramér Rao Bound (CRB) [19] in a variety of settings. Our numerical evaluations show that the proposed algorithm significantly outperforms classical MUSIC [3]. They also show that it outperforms recent sparse convex optimization techniques, with the gains (in terms of both mean squared error and computational complexity) becoming larger when the estimation problem is "more difficult."

Related work: Estimation of continuous-valued frequencies is a fundamental problem in statistical signal processing. Classical spectral methods such as MUSIC and ESPRIT algorithms [3], [4] exploit low-rank structure of the autocorrelation matrix to estimate the underlying signal subspace. While these methods are widely used, they perform poorly in noisy settings, and their performance is sensitive to model order estimation. Another family of classical methods are DFT-based methods [5], [6], which typically have lower computational complexity and similar level of estimation accuracy compared to subspace methods [6], [7]. The proposed algorithm handily outperforms such classical techniques in terms of estimation accuracy.

More recent techniques using convex optimization cast the frequency estimation problem as finding a sparse approximation for the received signal using an infinite-dimensional dictionary of sinusoids. It is shown in [8] 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 $4 \times \Delta_{\text{dft}}$ were $\Delta_{\text{dft}} = 2\pi/N$. The required minimum separation has been improved to $2.52 \times \Delta_{\text{dft}}$ in [9]. An extension to noisy scenarios is provided in [10]. Another approach is atomic norm denoising [11], [12], which provides theoretical guarantees of noise robustness in terms of mean squared error (MSE). Both total-variation norm and atomic norm are generalization of ℓ_1 norm to infinite-dimensional settings. While they outperform classical techniques in terms of accuracy, their computational complexity is prohibitive for most applications, especially when the number of observations gets large. A pragmatic approach is to use Lasso optimization on a highly oversampled grid as an approximation for atomic norm denoising [11], and this is what we compare our proposed algorithm against in our numerical experiments.

The proposed sequential algorithm builds on the idea of coarse detection followed by Newton refinement first used in prior works on frequency estimation [13] and estimation of a single delay [14]. We have used similar sequential algorithms for spatial channel estimation with compressive measurements for millimeter wave communication [15], [16]. In the present paper, we provide the version of such algorithms which we have found to acheive near-optimal performance, together with a principled stopping criterion based on CFAR. Given the fundamental nature and widespread utility of the frequency estimation problem, our goal here is to present what we believe is the state of the art algorithm within an application-independent abstraction.

Outline: We explain the details of our algorithm in Section II. Simulation results are presented in Section III, and conclusions in Section IV.

Notation: abs$\{v\}$ denotes the absolute value of the elements of vector v. Complex conjugate transpose of v is denoted by v^H. $\Re\{a\}$ is the real part of complex number a. The DFT matrix with unit norm columns and the corresponding grid spacing are denoted by $\mathcal{F}$ and Δ_{dft}, respectively.

II. Algorithm

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

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 maximizing the function

$$S(g,\omega) = \Re\{\mathbf{y}^H g\mathbf{x}(\omega)\} - 0.5|g|^2||\mathbf{x}(\omega)||^2. \tag{2}$$

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 the frequencies 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 $S(g,\omega)$ give us the generalized likelihood ratio test (GLRT) estimate of ω (treating g as a nuisance parameter), as the solution to the following optimization problem

$$\hat{\omega} = \arg\max_{\omega}\ |\mathbf{x}(\omega)^H\ \mathbf{y}|^2 \Big/ \|\mathbf{x}(\omega)\|^2. \tag{3}$$

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 the frequencies to a finite discrete set denoted by $\Omega \triangleq \{k(2\pi/\gamma N) : k = 0, 1, \ldots, (\gamma N - 1)\}$, 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 $\omega_{\text{c}} \in \Omega$ that maximizes the GLRT cost function (3), and the corresponding gain $(\mathbf{x}(\omega_{\text{c}})^H\mathbf{y})/||\mathbf{x}(\omega_{\text{c}})||^2$.

Refinement: Since true frequencies can take any value in interval $[0, 2\pi)$, we add a Newton-based refinement stage for estimation on the continuum. Let $(\hat{g}_n, \hat{\omega}_n)$ denote the estimates after n rounds of refinement. The Newton step is given by

$$\hat{\omega}_{n+1} = \hat{\omega}_n - \dot{S}(\hat{g}_n,\hat{\omega}_n)/\ddot{S}(\hat{g}_n,\hat{\omega}_n) \tag{4}$$

where

$$\dot{S}(g,\omega) = \Re\{(\mathbf{y} - g\mathbf{x}(\omega))^H g(d\mathbf{x}(\omega)/d\omega)\} \tag{5}$$

$$\ddot{S}(g,\omega) = \Re\{(\mathbf{y} - g\mathbf{x}(\omega))^H g(d^2\mathbf{x}(\omega)/d\omega^2)\} - |g|^2||(d\mathbf{x}(\omega)/d\omega)||^2. \tag{6}$$

The gain update for the new frequency value is given by $\hat{g}_{n+1} = (\mathbf{x}(\omega_{n+1})^H\mathbf{y})/||\mathbf{x}(\omega_{n+1})||^2$. We perform multiple such refinement steps.

Refinement acceptance condition: We *accept* a refinement only if it leads to a decrease in the overall residual energy. This is a sufficient condition for the convergence of the proposed sequential algorithm.

Multiple Frequencies: Let $\mathcal{P}_m = \{(g_l, \omega_l),\ l = 1, \ldots, m\}$ denote the set of estimates of the parameters of the sinusoids in the mixture. Let

$$\mathbf{y}_{\text{r}}(\mathcal{P}_m) = \mathbf{y} - \sum_{l=1}^{l=m} g_l\mathbf{x}(\omega_l) \tag{7}$$

denote the residual measurement corresponding to this estimate. The following procedure is a direct generalization of the single sinusoid estimation algorithm to multiple frequencies.

Procedure EXTRACTSPECTRUM($\mathbf{y}$, τ):
$m \leftarrow 0$, $\mathcal{P}_0 = \{\}$
while $\max\{\text{abs}\{\mathcal{F}\mathbf{y}_{\text{r}}(\mathcal{P}_m)\}\} > \sqrt{\tau}$ **do**
  $m \leftarrow m + 1$
  Find $\hat{\omega} = \arg\max_{\omega\in\Omega} \left|\left(\mathbf{x}(\omega)^H\mathbf{y}_{\text{r}}\left(\mathcal{P}_{m-1}\right)\right)\right|^2 \Big/ \|\mathbf{x}(\omega)\|^2$
  Compute corresponding gain
  $\hat{g} \leftarrow \left(\mathbf{x}(\hat{\omega})^H\mathbf{y}_{\text{r}}\left(\mathcal{P}_{m-1}\right)\right) \Big/ \|\mathbf{x}(\hat{\omega})\|^2$
  $\mathcal{P}_m \leftarrow \mathcal{P}_{m-1} \cup \{(\hat{g}, \hat{\omega})\}$
  Refine parameters in $\mathcal{P}_m$ *one at a time:*
  For each $(\hat{g}_l, \hat{\omega}_l) \in \mathcal{P}_m$ we treat $\mathbf{y}_{\text{r}}(\mathcal{P}_m \setminus \{(\hat{g}_l, \hat{\omega}_l)\})$ as the measurement $\mathbf{y}$, and apply single frequency Newton update algorithm.
return $\mathcal{P}_m$

Stopping Criterion: The stopping point of the algorithm is determined by the maximum of the absolute value of the residual signal in the DFT domain. This is a natural consequence of the proposed iterative algorithm: at the detection stage we are correlating the residual signal with normalized

$\mathbf{x}(\omega)$, which is equal to the output of the DFT operator for $\omega = k(2\pi/N)$, for $k \in \mathbb{Z}$. If all existing sinusoids have already been detected and subtracted out from $\mathbf{y}$, then the residual is $\mathbf{y}_r \approx \mathbf{z}$. Therefore, $\tilde{\mathbf{z}} = \mathcal{F}\mathbf{y}_r \sim \mathcal{CN}(0, \sigma^2 \mathbb{I}_N)$. Next we use the simple observation that for $w \sim \mathcal{CN}(0, \sigma^2)$, we have $|w|^2 \sim \exp(1/\sigma^2)$ and $\mathbb{E}[|w|^2] = \sigma^2$. For noise-only observations, the maximum magnitude squared of the DFT (not considering oversampling) is given by $M_N \triangleq \max\{w_i : i = 1, 2, \ldots, N\}$, where w_i's are iid exponential random variables with mean σ^2. We can now set a stopping criterion such that $\Pr\{M_N > \tau\} \leq p_0$, where p_0 is a nominal false alarm rate. This can be explicitly computed, since $\Pr\{M_N > \tau\} = 1 - \left(1 - \exp\left(-\tau/\sigma^2\right)\right)^N$, which gives $\tau = \log \frac{1}{1-(1-p)^{1/N}}$. However, a more transparent expression can be obtained by considering asymptotics for large N (which provide excellent approximations for even the moderate values of N considered in our numerical results). Specifically, 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

$$\Pr\{E \leq x\} = \exp(-\exp(-x/\sigma^2)). \tag{8}$$

We can now set $\tau = \sigma^2 \log(N) + e_0$, where $\Pr\{E > e_0\} = p_0$. This gives

$$e_0 = -\sigma^2 \log\log\left(1/(1-p_0)\right). \tag{9}$$

and

$$\tau = \sigma^2 \log(N) - \sigma^2 \log\log\left(1/(1-p_0)\right). \tag{10}$$

Thus, we terminate the algorithm whenever the magnitude of the residual in the DFT domain is smaller than $\sqrt{\tau}$. In our simulation results, we have found that setting the nominal false alarm rate to $p_0 = 10^{-2}$ works very well in both low and high SNR regimes.

III. Simulation Results

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

Benchmarks: The MUSIC algorithm is implemented using the MATLAB routine "`rootmusic`". This method needs an initial estimate of the number of the sinusoids in the observation. In order to give it the best chance of success, we provide the *true value of* K as the input to the MUSIC algorithm. For sparse convex optimization, we consider Lasso on an oversampled frequency grid [11], using the highly optimized $\ell_2 - \ell_1$ software package SpaRSA [18]. We set the tolerance parameter to be 10^{-4} (other than that, we use the default parameters). The regularization parameter in Lasso formulation, suggested in [11], 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 and for the proposed sequential algorithm are set to 10 and 4, respectively.

Minimum Frequency Separation: When two frequencies, say ω_1 and ω_2 come "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 $\Delta\omega_{\min} = \min_{k \neq l} \text{dist}(\omega_k, \omega_l)$, where

$$\text{dist}(\omega_k, \omega_l) \triangleq \min_{a \in \mathbb{Z}} |\omega_k - \omega_l + 2\pi a|.$$

We would like our algorithms to work well even for small values of $\Delta\omega_{\min}$.

Simulation set-up: We consider a mixture of $K = 16$ sinusoids of length $N = 256$. We perform 300 simulation runs for each of three 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	25	2.5
2	25	0.5
3	Uniform$[15, 35]$	0.5

TABLE I: Settings of different Scenario

For Scenarios 1 and 2, the SNR for each sinusoid is set as 25 dB, whereas for Scenarios 3 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 $[0, 2\pi)$ while respecting the minimum separation constraints specified by $\Delta\omega_{\min}$ (if the minimum separation criterion is not met, we sample again from $[0, 2\pi)^K$). The number of refinement steps for the sequential algorithm is set to 8. We plot the complementary CDF of the MSE 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.

In all three scenarios, the proposed sequential algorithm outperforms both Lasso and MUSIC in terms of frequency MSE. In fact, the algorithm is nearly optimal and closely follows the CRB in all settings. Table II summarizes the time needed for running 300 simulations for each of the algorithms in different scenarios. We see that MUSIC algorithm is extremely fast but does not provide enough estimation accuracy. On the other hand, as we increase the difficulty of the estimation scenario, Lasso tends to take more time, while the sequential algorithm is unaffected.

Time [sec]	Sequential	Lasso	MUSIC
Scenario 1	30.90	37.78	3.26
Scenario 2	31.09	37.95	3.27
Scenario 3	30.91	49.39	3.36

TABLE II: Time [sec] for 300 runs of each algorithm

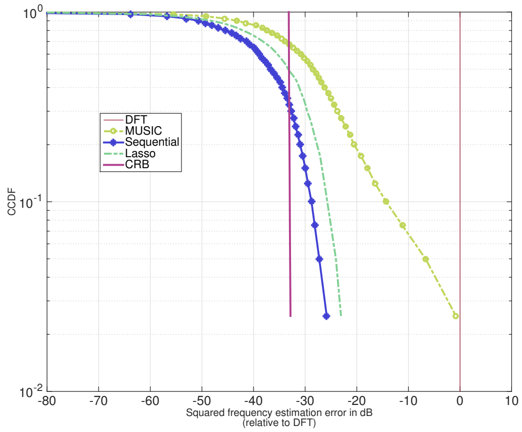

Fig. 1: CCDF of the frequency MSE for Scenario 1.

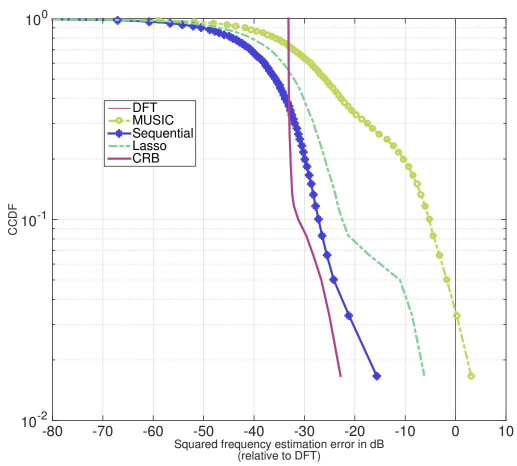

Fig. 2: CCDF of the frequency MSE for Scenario 2.

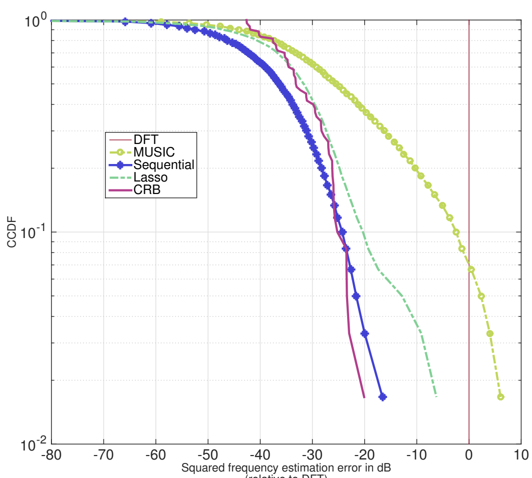

Fig. 3: CCDF of the frequency MSE for Scenario 3.

It is also interesting to see the effect of increasing the oversampling factor for Lasso on the estimation accuracy and computational complexity. Fig. 4 corresponds to Scenario 1, when the oversampling factor for Lasso is increased to 20 (that of our sequential algorithm remains at 4). We observe that sequential algorithm is still better in terms of MSE, while the computation time for Lasso increases significantly (to about 89.60 seconds for 300 runs).

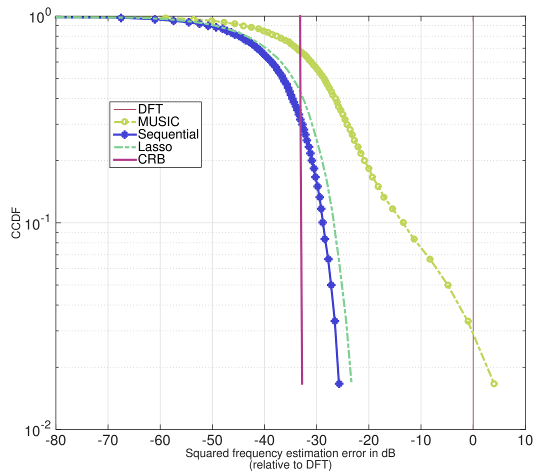

Fig. 4: CCDF of the frequency MSE for Scenario 1 and highly over-sampled grid for Lasso ($\gamma = 20$).

IV. Conclusions

We have presented a fast, near-optimal algorithm for the fundamental problem of line spectral estimation. The algorithm has a natural "decision feedback" interpretation, while providing MSE performance far superior to classical methods such as MUSIC, somewhat superior to Lasso techniques, and approaching fundamental estimation-theoretic bounds. The complexity of the proposed algorithm is smaller than that even highly optimized implementations of Lasso, especially as the setting for estimation becomes more challenging. We note that our sequential detection plus refinement approach is compatible with compressive measurements, as illustrated by earlier versions of the algorithm applied for compressive spatial channel estimation [15], [16], as well as according to a general theory of compressive estimation [17].

There are a number of topics for future work. While our stopping criterion is based on a nominal false alarm probability, the actual false alarm and miss probabilities are sensitive to the definition of detecting a frequency (i.e., the size of the bin around a true frequency that is considered adequate), and further study is needed to characterize ROCs.

V. Acknowledgment

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.

REFERENCES

[1] S. Mallat and Z. Zhang, “Matching pursuits with time-frequency dictionaries,” IEEE Trans. Signal Process., vol. 41, no. 12, pp. 33973415, Dec. 1993.

[2] Y. Chi, L. L. Scharf, A. Pezeshki and R. Calderbank, “Sensitivity to basis mismatch in compressed sensing,” IEEE Trans. Signal Process., vol. 59, no. 5, pp. 2182-2195, 2011.

[3] R. Schmidt, ”Multiple emitter location and signal parameter estimation,” IEEE Trans. on Antennas and Propagation, vol. 34, no. 3, pp. 276-280, 1986.

[4] R. Roy and T. Kailath, “ESPRIT - estimation of signal parameters via rotational invariance techniques,” IEEE Trans. on Acoustics, Speech and Signal Processing, vol. 37, no. 7, pp. 984-995, 1989.

[5] K. Duda, T. P. Zieliński, L. B. Magalas, M. Majewski, “DFT-based Estimation of Damped Oscillation Parameters in Low-frequency Mechanical Spectroscopy,” IEEE Trans. Instrum. Meas., 60(11), 3608-3618, 2011.

[6] T. Zieliński and K. Duda, “Frequency and damping estimation methods – an overview,” Metrol. Meas. Syst., vol. 8, pp.505-528, 2011.

[7] D. Vu, L. Xu, M. Xue, J. Li, “Nonparametric missing sample spectral analysis and its applications to interrupted sar,” IEEE Journal of Selected Topics in Signal Processing, vol.6, no.1, pp.1-14, Feb. 2012.

[8] E. Candes and C. Fernandez-Granda, “Towards a mathematical theory of super-resolution,” arXiv Preprint 1203.5871, 2012.

[9] C. Fernandez-Granda, “Super-Resolution of Point Sources via Convex Programming,” arXiv Preprint 1507.07034, 2015.

[10] E. Candes and C. Fernandez-Granda, “Super-resolution from noisy data,” arXiv Preprint 1211.0290v3, 2013.

[11] B. N. Bhaskar, G. Tang, and B. Recht, “Atomic norm denoising with applications to line spectral estimation,” arXiv preprint arXiv:1204.0562, 2012.

[12] 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.

[13] T. J. Abatzoglou, “A Fast Maximum Likelihood Algorithm for Frequency Estimation of a Sinusoid Based on Newton’s Method,” IEEE Transactions on Acoustics, Speech and Signal Proc., vol. 33, no. 1, pp. 77-89, February 1985.

[14] P. Bidigare, U. Madhow, R. Mudumbai, D. Scherber, “Attaining fundamental bounds on timing synchronization,” **invited paper,** *Proc. 2012 IEEE International Conf. on Acoustics, Speech and Signal Processing (ICASSP 2012),* pp. 5229-5232, Kyoto, Japan, March 2012.

[15] D. Ramasamy, S. Venkateswaran, U. Madhow, “Compressive adaptation of large steerable arrays,” *Proc. 2012 Information Theory and Applications Workshop (ITA 2012),* pp. 234-239, University of California, San Diego, February 2012.

[16] ——, “Compressive tracking with 1000-element arrays: A framework for multi-Gbps mm wave cellular downlinks,” *Proc. 2012 50th Annual Allerton Conference on Communication, Control, and Computing (Allerton 2012),* pp. 690-697, Monticello, IL, October 2012.

[17] ——, “Compressive parameter estimation in AWGN,” *IEEE Transactions on Signal Processing,* vol. 62, no. 8, pp. 2012-2027, April 2014.

[18] S. J. Wright, R. D. Nowak, and M. A. T. Figueiredo, “Sparse reconstruction by separable approximation,” IEEE Trans. on Signal Processing, vol. 57, no. 7, pp. 2479-2493, 2009.

[19] H. L. Van Trees, K. L. Bell, and Z. Tian, Detection, Estimation, and Modulation Theory. John Wiley & Sons, 2013.

[20] https://bitbucket.org/wcslspectralestimation/continuous-frequency-estimation/src