

A Fast Interior-Point Method for Atomic Norm Soft Thresholding

Hansen, Thomas Lundgaard; Jensen, Tobias

Published in:
Signal Processing

DOI (link to publication from Publisher):
[10.1016/j.sigpro.2019.06.023](https://doi.org/10.1016/j.sigpro.2019.06.023)

Creative Commons License
CC BY-NC-ND 4.0

Publication date:
2019

Document Version
Accepted author manuscript, peer reviewed version

[Link to publication from Aalborg University](#)

Citation for published version (APA):
Hansen, T. L., & Jensen, T. (2019). A Fast Interior-Point Method for Atomic Norm Soft Thresholding. *Signal Processing*, 165, 7-19. <https://doi.org/10.1016/j.sigpro.2019.06.023>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal -

Take down policy

If you believe that this document breaches copyright please contact us at vbn@aub.aau.dk providing details, and we will remove access to the work immediately and investigate your claim.

Accepted Manuscript

A Fast Interior-Point Method for Atomic Norm Soft Thresholding

Thomas Lundgaard Hansen, Tobias Lindstrøm Jensen

PII: S0165-1684(19)30231-2
DOI: <https://doi.org/10.1016/j.sigpro.2019.06.023>
Reference: SIGPRO 7191

To appear in: *Signal Processing*

Received date: 19 November 2018
Revised date: 24 March 2019
Accepted date: 16 June 2019



Please cite this article as: Thomas Lundgaard Hansen, Tobias Lindstrøm Jensen, A Fast Interior-Point Method for Atomic Norm Soft Thresholding, *Signal Processing* (2019), doi: <https://doi.org/10.1016/j.sigpro.2019.06.023>

This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

Highlights

- Provide a fast methods to solve the atomic norm soft thresholding problem.
- The method is based on combining interior-point methods with fast methods for operations involving hermitian Toeplitz matrices.
- We measure the time-to-solve and compare with the theoretical obtained computational complexity.

ACCEPTED MANUSCRIPT

A Fast Interior-Point Method for Atomic Norm Soft Thresholding

Thomas Lundgaard Hansen¹, Tobias Lindstrøm Jensen*

Aalborg University, Fr. Bajersvej 7, DK-9220 Aalborg, Denmark

Abstract

The atomic norm provides a generalization of the ℓ_1 -norm to continuous parameter spaces. When applied as a sparse regularizer for line spectral estimation the solution can be obtained by solving a convex optimization problem. This problem is known as atomic norm soft thresholding (AST). It can be cast as a semidefinite program and solved by standard methods. In the semidefinite formulation there are $\mathcal{O}(N^2)$ dual variables which complicates the implementation of a standard primal-dual interior-point method based on symmetric cones. That has lead researchers to consider the alternating direction method of multipliers (ADMM) for the solution of AST, but this method is still somewhat slow for large problem sizes. To obtain a faster algorithm we reformulate AST as a non-symmetric conic program. That has two properties of key importance to its numerical solution: the conic formulation has only $\mathcal{O}(N)$ dual variables and the Toeplitz structure inherent to AST is preserved. Based on it we derive FastAST which is a primal-dual interior-point method for solving AST. Two variants are considered with the fastest one requiring only $\mathcal{O}(N^2)$ flops per iteration. Extensive numerical experiments demonstrate that both variants of FastAST solve AST significantly faster than a state-of-the-art solver based on ADMM.

Keywords: Atomic Norm Minimization, Atomic Norm Soft Thresholding, Line Spectral Estimation, Convex Optimization, Interior-Point Methods, Non-symmetric Conic Optimization

1. Introduction

It is well known that sparse estimation problems can be formulated as convex optimization problems using the ℓ_1 -norm. The ℓ_1 -norm can be generalized to continuous parameter spaces through the so-called atomic norm [1]. Convex modelling of sparsity constraints has two highly attractive traits: convex optimization problems can easily be solved both in theory [2] and in practice [3, 4], and, a number of recovery guarantees can be obtained within this framework. Such recovery guarantees are studied in signal processing under the name compressed sensing [5, 6, 7] and they generalize nicely to the atomic norm minimization approach [8, 9, 10, 11, 12].

*Corresponding author

Email address: tlj@its.aau.dk (Tobias Lindstrøm Jensen)

¹The work of T. L. Hansen is supported by the Danish Council for Independent Research under grant id DFF-4005-00549. The work of T. L. Jensen was partly supported by by Danich Council for Independent Research under grant id DFF-4005-00122

The most prominent example of estimation with the atomic norm is the application to line spectral estimation [9, 11, 12], in which case it is known as atomic norm soft thresholding (AST). The popularity of AST is, partly, due to the fact that it can be cast as a semidefinite programming (SDP) problem (we refer to Sec. 2 for a review of AST,)

$$\begin{aligned} & \text{minimize}_{v,x,u} \quad \|x - y\|_2^2 + \tau(v + w^T u) \\ & \text{subject to} \quad \begin{pmatrix} T(u) & x \\ x^H & v \end{pmatrix} \succeq 0, \end{aligned} \quad (1)$$

where $v \in \mathbb{R}$, $x \in \mathbb{C}^N$, $u \in \mathbb{R}^{2N-1}$ are the variables of the problem and $y \in \mathbb{C}^N$, $\tau \in \mathbb{R}$, $w \in \mathbb{R}^{2N-1}$ are fixed (known) parameters. The function $T(u) : \mathbb{R}^{2N-1} \rightarrow \mathbb{C}^{N \times N}$ outputs a complex Hermitian Toeplitz matrix constructed from u , such that the first row is $(2u_0, \dots, u_{N-1}) + j(0, u_N, \dots, u_{2N-2})$. To be precise, AST is obtained by selecting $w = 2e_0$ in (1), where e_0 is a vector with 1 in the first entry and zeros elsewhere.

The state-of-the-art method for solving (1) is via the alternating direction method of multipliers (ADMM) and used in [9, 13, 14, 15]. While this method is reasonably fast, it has some drawbacks. It requires the calculation of an eigenvalue decomposition in each iteration at cost $\mathcal{O}(N^3)$ floating-point operations (flops). This means that for large N it is exceedingly slow. As is often seen with proximal methods it also has slow convergence if a solution of high accuracy is requested.

Da Costa *et al.* [16] apply a low-dimensional projection of the observation vector to reduce the problem size and therefore the computational complexity of AST. In the noise-free case and under certain regularity conditions, it is shown that the estimation accuracy is not affected by doing so. However, it is clear that this approach discards observed data and the estimation accuracy will be degraded in the noisy case. Another attempt at a fast solver for AST is [17], but due to its non-SDP implementation utilizing a frequency grid and real, positive coefficients, this approach only allows for a fixed known phase for all components and hence cannot solve the line spectral problem as described in this paper. There has been other attempts to solve atomic norm problems efficiently in e.g. [18, 19], but without covering the AST problem.

The formulation of AST in (1) is casted as an SDP problem. SDP problems have been subject to intensive research since the 1990's and their solution using primal-dual interior-point methods (IPMs) is now understood well [2, 3, 20, 21, 22]. The Lagrangian dual of (1) has $\mathcal{O}(N^2)$ dual variables due to the semidefinite matrix constraint. The direct application of a standard primal-dual IPM thus requires $\mathcal{O}(N^6)$ flops per iteration at best (using direct methods for solving linear systems of equations), but can be reduced by eliminating the dual variables from the linear system (in general, the exact number flops per iteration will depend on the implementation). Compared to this approach, proximal methods (such as ADMM) which require $\mathcal{O}(N^3)$ flops per iteration are preferable, even if they converge much slower than primal-dual IPMs. That

explains why primal-dual IPMs have not gained traction for the solution of (1). In this work we reformulate the constraint in (1) as a non-symmetric conic constraint on the vector $(v, x^T, u^T)^T$. This formulation immediately reduces the number of dual variables to $\mathcal{O}(N)$ and sets the scene for a reintroduction of primal-dual IPMs as a very competitive class of algorithms for solution of AST.

Primal-dual IPMs for conic programming typically rely on a symmetry between the primal and dual problems. The formulation of such symmetric primal-dual IPMs requires the existence of a self-scaled barrier function for the cone involved in the constraint [23, 24]. Güler [25] showed that such barrier functions exist only for the class of homogeneous and self-dual cones. The cone in our formulation is not self-dual and so a symmetric primal-dual IPM cannot be formulated. Non-symmetric conic optimization has received some attention [26, 27, 28, 29]. These methods generally rely on the availability of a barrier function for the dual cone and possibly evaluation of its gradient and Hessian. An easy-to-calculate dual barrier is not available for the cone associated to the constraint of our formulation; only an oracle which can determine membership in the dual cone is available (part of our contribution is to show how such an oracle can be constructed.)

To derive a non-symmetric primal-dual IPM which does not rely on evaluating the dual barrier or its derivatives, we formulate the augmented Karush-Kuhn-Tucker conditions and devise a dedicated approach to solving these. This approach is shown to converge to a primal-dual feasible point. A lower bound on the objective function is calculated in those iterations where a dual feasible point (as determined by the oracle) is available. From the lower bound a duality gap can be evaluated, thus providing a method for dynamically updating the barrier parameter. We show that the proposed method enjoys global convergence.

Our focus is on obtaining an algorithm which has fast runtime *in practice*, i.e., it has both low per-iteration computational complexity and it exhibits reasonably fast convergence. Theoretical statements regarding for example convergence speed are left for future work. At the core of obtaining a practically fast algorithm lies the already mentioned conic formulation (which brings the number of dual variables down to $\mathcal{O}(N)$), along with techniques for fast evaluation of linear algebra in each step of the algorithm. These evaluations are based on fast algorithms [30, 31, 32, 33] for inversion of Toeplitz matrices. Related techniques are employed in [34, 35, 36, 37].

We dub the algorithm FastAST. Both Newton's method and a quasi-Newton method are considered for evaluation of the search direction in FastAST. When using Newton's method the algorithm requires $\mathcal{O}(N^3)$ flops per iteration, while the quasi-Newton variant only requires $\mathcal{O}(N^2)$ flops per iteration. The numerical experiments in Sec. 6 show that the quasi-Newton variant is faster in practice. Due to numerical inaccuracies in the calculation of the search direction the quasi-Newton variant is not able to obtain a solution of very high accuracy. Solving (1) to high accuracy makes a difference with very large signal-to-noise ratios and in these cases the variant of FastAST using Newton's method should be used. Both the Newton's and quasi-Newton variants

of FastAST are significantly faster than the ADMM-based solvers for (1).

Along with the primal-dual IPM presented here, we have also experimented with a primal-only version which is simpler to derive. The primal-only approach does not provide a good way to select the barrier parameter (which we denote t , see Sec. 3.2). This in turn forces the primal-only approach to use either overly conservative short-step [38] updates of the barrier parameter or it requires the barrier problem to be solved to high accuracy for each fixed t . Both scenarios lead to a significant increase in the number of iterations required by the primal-only algorithm, resulting in significantly increased runtime. On the contrary, the primal-dual version presented in this paper allows for evaluation of a duality gap in each iteration. The duality gap gives a natural way to select the barrier parameter and also provides a very precise stopping criterion.

The paper is outlined as follows. In Sec. 2 we begin with a brief review of atomic norm minimization and its application to line spectral estimation. Sec. 3 details the formulation of (1) as a non-symmetric conic optimization program along with the theory of its solution. In Sec. 4 we present our numerical algorithm along with implementation details. The exploitation of Toeplitz structure for fast evaluation of each step in the algorithm is discussed in Sec. 5. Numerical experiments which validate the practical efficacy of the proposed algorithm are presented in Sec. 6.

2. A Review of Atomic Norm Soft Thresholding

2.1. Line Spectral Estimation

Consider an observation vector $y \in \mathbb{C}^N$,

$$y = x + \zeta, \quad x = \sum_{k=0}^{K-1} c_k a(\omega_k), \quad (2)$$

where $\zeta \in \mathbb{C}^N$ is a noise vector and $x \in \mathbb{C}^N$ is a signal of interest composed of K sinusoids, each with angular frequency $\omega_k \in [0, 2\pi)$ and complex coefficient $c_k \in \mathbb{C}$. The steering vector $a(\omega)$ has entries $(a(\omega))_n = \exp(jn\omega)$ for $n = 0, \dots, N-1$ and $j = \sqrt{-1}$. In line spectral estimation the task is to estimate the values $(K, c_0, \dots, c_{K-1}, \omega_0, \dots, \omega_{K-1})$. The crux of line spectral estimation lies in obtaining estimates of the model order K and the frequencies $\{\omega_k\}$. When these are available the coefficients $\{c_k\}$ can easily be estimated using a least-squares approach. The problem is ubiquitous in signal processing; examples include direction of arrival estimation using sensor arrays [39, 40], bearing and range estimation in synthetic aperture radar [41], channel estimation in wireless communications [42] and simulation of atomic systems in molecular dynamics [43].

2.2. Modelling Sparsity With the Atomic Norm

The atomic norm [1, 8, 9, 10] provides a tool for describing notions of sparsity in a general setting. It is defined in terms of the *atomic set* \mathcal{A} . Each member of \mathcal{A} is referred to as an *atom*.

The atoms are the basic building block of the signal and define the dictionary in which the signal has a sparse representation. The atomic norm induced by \mathcal{A} is defined as

$$\|x\|_{\mathcal{A}} = \inf\{\alpha > 0 : x \in \alpha \mathbf{conv} \mathcal{A}\}, \quad (3)$$

where $\mathbf{conv} \mathcal{A}$ is the convex hull of \mathcal{A} .

For line spectral estimation the atomic set is selected as the set of complex rotated Fourier atoms [8, 9, 10]

$$\mathcal{A} = \{a(\omega) \exp(j\phi) : \omega \in [0, 2\pi), \phi \in [0, 2\pi)\} \quad (4)$$

and the corresponding atomic norm can be described as

$$\|x\|_{\mathcal{A}} = \inf_{K, \{c_k, \omega_k\}} \left\{ \sum_{k=0}^{K-1} |c_k| : x = \sum_{k=0}^{K-1} c_k a(\omega_k) \right\}. \quad (5)$$

110 It is clear that the atomic norm provides a generalization of the ℓ_1 -norm to the continuous parameter space $\omega_k \in [0, 2\pi)$. Through the use of a dual polynomial characterization the atomic norm can be expressed as the solution of an SDP,

$$\begin{aligned} \|x\|_{\mathcal{A}} = \underset{v, u}{\text{minimize}} \quad & \frac{1}{2} \left(v + \frac{1}{N} \text{tr} T(u) \right) \\ \text{subject to} \quad & \begin{pmatrix} T(u) & x \\ x^H & v \end{pmatrix} \succeq 0. \end{aligned} \quad (6)$$

2.3. Atomic Norm Soft Thresholding

AST [9] is inspired by the least absolute shrinkage and selection operator (LASSO) [44] and
115 solves

$$\underset{x}{\text{minimize}} \quad \|x - y\|_2^2 + 2\tau \|x\|_{\mathcal{A}}, \quad (7)$$

where $\tau > 0$ is a regularization parameter to be chosen. It is clear that AST is recovered in (1) by selecting $w = 2e_0$.

Once a solution (v^*, x^*, u^*) of (7) has been found, estimates of the model order K and the frequencies $\{\omega_k\}$ can be obtained by examining a certain dual polynomial constructed from x^* .
120 This process determines the solution in (5) for the recovered signal x^* . Under a, somewhat restrictive, assumption concerning separation of the frequencies $\{\omega_k\}$ a number of theoretical statements can be given regarding signal and frequency recovery using AST. We refer to [8, 9, 10, 11, 12] for details.

We now consider the selection of the regularization parameter τ . Clearly the choice of τ crucially influences the estimation accuracy of AST. It is this parameter which determines the trade-off between fidelity and sparsity which is inherent in any estimator involving the model order K . With all else being equal, selecting larger τ gives estimates with smaller values of K . Let $\|\cdot\|_{\mathcal{A}}^*$ denote the dual norm of the atomic norm $\|\cdot\|_{\mathcal{A}}$. Then the theoretical analysis in

[9] requires $\tau \geq \mathbb{E}[\|\zeta\|_{\mathcal{A}}^*]$. For a white, zero-mean circularly symmetric complex Gaussian noise vector ζ with entry-wise variance σ^2 such an upper bound is given by [9],

$$\tau = \sigma \frac{\log(N) + 1}{\log(N)} \sqrt{N \log(N) + N \log(4\pi \log(N))}, \quad (8)$$

where $\log(\cdot)$ is the natural logarithm. This choice has been shown to perform well in practice and we also use it in our simulation study.

3. Non-symmetric Conic Optimization

We now return to our main focus: That of numerically solving the conic program (1). It can be written in the form

$$\begin{aligned} & \text{minimize} && f(\mu) \\ & \text{subject to} && \mu \in \mathcal{K}, \end{aligned} \quad (9)$$

where $f(\mu) = \|x - y\|_2^2 + \tau(v + w^T u)$ and \mathcal{K} is the cone defined by

$$\mathcal{K} \triangleq \left\{ \mu = \begin{pmatrix} v \\ x \\ u \end{pmatrix} : \begin{pmatrix} T(u) & x \\ x^H & v \end{pmatrix} \succeq 0 \right\}. \quad (10)$$

As a precursor to deriving a primal-dual IPM, we explore the properties of \mathcal{K} and its dual. It is easy to show that \mathcal{K} is a proper cone (convex, closed, solid and pointed; see [20]). The dual cone \mathcal{K}^* of \mathcal{K} is defined as

$$\mathcal{K}^* = \{ \lambda : \langle \lambda, \mu \rangle \geq 0 \ \forall \mu \in \mathcal{K} \}. \quad (11)$$

Since \mathcal{K} is proper, so is \mathcal{K}^* [20].

In this paper (primal) variables in the cone \mathcal{K} are denoted by $\mu = (v, x^T, u^T)^T$. The (dual) variables in the cone \mathcal{K}^* are denoted by $\lambda = (\rho, s^T, z^T)^T$, with $\rho \in \mathbb{R}$, $s \in \mathbb{C}^N$ and $z \in \mathbb{R}^{2N-1}$. The inner product between them is defined as $\langle \lambda, \mu \rangle = \rho v + \text{Re}(s^H x) + z^T u$.

3.1. Checking for Dual Cone Membership

In our proposed method, we need to check for $\lambda \in \mathcal{K}^*$. In order to characterize the dual cone \mathcal{K}^* , the cone of positive semidefinite Hermitian Toeplitz matrices is needed:

$$\mathcal{C} \triangleq \{ u \in \mathbb{R}^{2N-1} : T(u) \succeq 0 \}. \quad (12)$$

This cone is also proper. The corresponding dual cone \mathcal{C}^* is defined analogously to (11).

Let the function T^* be the adjoint² of the linear map T , i.e., $T^* : \mathbb{C}^{N \times N} \rightarrow \mathbb{R}^{2N-1}$ is such

² T^* is easy to calculate: Let B be Hermitian and let β_n denote the sum over the n th upper diagonal of B ,

$$\beta_n = \sum_{m=0}^{N-1-n} B_{m, m+n},$$

for $n = 0, \dots, N-1$. Then

$$T^*(B) = (2\beta_0, 2\text{Re}(\beta_1), \dots, 2\text{Re}(\beta_{N-1}), 2\text{Im}(\beta_1), \dots, 2\text{Im}(\beta_{N-1}))^T.$$

that for every Hermitian $B \in \mathbb{C}^{N \times N}$ we have $\text{tr}(T(u)^H B) = T^*(B)^T u$. Then we then have the following lemma.

Lemma 1. *The dual cone of \mathcal{K} can be characterized as*

$$\mathcal{K}^* = \left\{ \lambda = \begin{pmatrix} \rho \\ s \\ z \end{pmatrix} : \rho > 0, \left(z - \frac{1}{4\rho} T^*(ss^H) \right) \in \mathcal{C}^* \right\} \cup \{ \lambda : \rho = 0, s = 0, z \in \mathcal{C}^* \}. \quad (13)$$

PROOF. See the appendix.

It is clear that \mathcal{K} is not self-dual ($\mathcal{K} \neq \mathcal{K}^*$) and so (9) is a non-symmetric conic program.

The cone \mathcal{C} and its dual are defined in terms of real-valued vectors because this description simplifies the derivation of the method in Sec. 4. These sets are, however, more naturally understood from their corresponding complex-valued forms. We therefore define the vector $u_{\mathbb{C}} = (u_0, u_1 + ju_N, u_2 + ju_{N+1}, \dots, u_{N-1} + ju_{2N-2})^T$ and use a similar definition of $z_{\mathbb{C}}$.

The dual cone \mathcal{C}^* turns out to be the set of finite autocorrelation sequences. An excellent introduction to this set and a number of characterizations of it are given in [37] for the case of real-valued sequences. Here we extend the definition to the complex-valued case.

Definition 1. *A vector z is a finite autocorrelation sequence if there exists a vector $q \in \mathbb{C}^N$ such that³*

$$(z_{\mathbb{C}})_k = \sum_{n=0}^{N-1-k} \bar{q}_n q_{n+k}, \quad k = 0, \dots, N-1. \quad (14)$$

In other words, z is a finite autocorrelation sequence if

$$\dots, 0, 0, (\bar{z}_{\mathbb{C}})_{N-1}, (\bar{z}_{\mathbb{C}})_{N-2}, \dots, (\bar{z}_{\mathbb{C}})_1, (z_{\mathbb{C}})_0, (z_{\mathbb{C}})_1, \dots, (z_{\mathbb{C}})_{N-1}, 0, 0, \dots \quad (15)$$

is the autocorrelation sequence of some moving average process of order $N-1$ with filter coefficients q_1, \dots, q_{N-1} and input variance $|q_0|^2$. It is well known from the theory of linear time-invariant systems, that if (15) is a valid autocorrelation, then it can be represented by a moving average process (i.e., there exists a coefficient vector q such that (14) holds).

A sequence is a valid autocorrelation sequence if and only if its Fourier transform is non-negative [45]. The Fourier transform of (15) is

$$Z(\omega) = (z_{\mathbb{C}})_0 + 2 \sum_{k=1}^{N-1} \text{Re}((z_{\mathbb{C}})_k \exp(-j\omega k)), \quad (16)$$

for $\omega \in [0, 2\pi)$. Then $z \in \mathcal{C}^*$ if and only if $Z(\omega) \geq 0$ for all $\omega \in [0, 2\pi)$. The fast Fourier transform allows $Z(\omega)$ to be evaluated at a large number of points on $[0, 2\pi)$ in an efficient way. Using Lemma 1 we therefore have a low-complexity method of determining if $\lambda \in \mathcal{K}^*$.

³The complex conjugate of q is denoted \bar{q} .

This approach is approximate in the sense that $Z(\omega)$ is sampled at a finite number of points on $[0, 2\pi)$. The approximation can be made arbitrarily accurate by increasing the number of evaluated points. In our simulation study in Sec. 6 it is demonstrated that the approximation is of sufficient accuracy for our algorithm to be utilized in practice.

We still haven't shown that the dual of the cone \mathcal{C} is indeed the set of finite autocorrelation sequences. To that end, let $\tilde{\mathcal{C}}$ be the set of finite autocorrelation sequences and identify u with $u_{\mathbb{C}}$. Extending the approach of [37] to the complex-valued case, a vector u is in the dual of $\tilde{\mathcal{C}}$ if and only if $z^T u \geq 0$ for every $z \in \tilde{\mathcal{C}}$, or, in other words, if and only if

$$z^T u = \text{Re}(z_{\mathbb{C}}^H u_{\mathbb{C}}) = \text{Re}\left(\sum_{k=0}^{N-1} \sum_{n=0}^{N-1-k} (u_{\mathbb{C}})_k q_n \bar{q}_{n+k}\right) = \frac{1}{2} q^T T(u) \bar{q} \geq 0 \quad (17)$$

for every $q \in \mathbb{C}^N$. We can therefore identify \mathcal{C} with $\tilde{\mathcal{C}}^*$. Since $\tilde{\mathcal{C}}$ is a proper cone, we have $\mathcal{C}^* = \tilde{\mathcal{C}}^{**} = \tilde{\mathcal{C}}$.

3.2. Barrier Functions

IPMs are built on the idea of a barrier function $F : \text{int } \mathcal{K} \rightarrow \mathbb{R}$ associated to the cone \mathcal{K} ($\text{int } \mathcal{K}$ denotes the interior of \mathcal{K}). The barrier function must be a smooth and strictly convex⁴ function with $F(\mu_k) \rightarrow \infty$ for every sequence of points $\mu_k \in \text{int } \mathcal{K}$ with limit point $\tilde{\mu} \in \text{bd } \mathcal{K}$, where $\text{bd } \mathcal{K}$ denotes the boundary of \mathcal{K} . The typical approach to IPMs also assumes that the barrier function is logarithmically homogeneous (LH). F is a LH barrier function for the cone \mathcal{K} if there exists a $\theta_F > 0$ such that $F(\alpha\mu) = F(\mu) - \theta_F \log(\alpha)$ for all $\alpha > 0$, $\mu \in \text{int } \mathcal{K}$. The value θ_F is called the degree of the barrier.

We will use the following well-known properties of a LH barrier function F for \mathcal{K} [20, 24, 23]: If $\mu \in \text{int } \mathcal{K}$, then

$$\langle -\nabla_{\mu} F(\mu), \mu \rangle = \theta_F, \quad (18)$$

$$-\nabla_{\mu} F(\mu) \in \text{int } \mathcal{K}^*, \quad (19)$$

where the gradient operator is defined as $\nabla_{\mu} F = (\nabla_v F, \nabla_x F^T, \nabla_u F^T)^T$. The gradient with respect to the complex vector $x = a + jb$ is to be understood as⁵ $\nabla_x F = \nabla_a f + j \nabla_b f$.

The usefulness of barrier functions is clear when considering their role in path-following methods. A primal-only path-following method finds a solution to (9) by iteratively solving

$$\begin{aligned} & \text{minimize} && f(\mu) + t^{-1} F(\mu) \\ & \text{subject to} && \mu \in \text{int } \mathcal{K} \end{aligned} \quad (20)$$

for an increasing sequence of values $t > 0$. In each step μ is initialized with the solution of the previous step. This approach is desirable because each step can be solved by an algorithm for unconstrained optimization such as Newton's method.

⁴Hessian positive definite everywhere.

⁵This is actually twice the Wirtinger derivative of F with respect to \bar{x} .

In this paper we use the standard log-determinant barrier function for \mathcal{K} :

$$F(\mu) = -\log \left| \begin{pmatrix} T(u) & x \\ x^H & v \end{pmatrix} \right| \quad (21)$$

$$= -\log |T(u)| - \log(v - x^H T^{-1}(u)x), \text{ for } \mu \in \text{int } \mathcal{K}.$$

175 It is easy to show that it is LH with degree $\theta_F = N + 1$.

3.3. Solvability

We now consider conditions for the problem (9) to be solvable. An optimization problem is solvable when a feasible point exists and its objective is bounded below on the feasible set.

Lemma 2. *The function $f(\mu)$ is bounded below on $\mu \in \mathcal{K}$ if and only if $\tau = 0$ or $\tau > 0$ and $w \in \mathcal{C}^*$.*
180

PROOF. The case $\tau = 0$ is trivial. Assume $\tau \neq 0$ in the following.

If $\tau < 0$ or $w \notin \mathcal{C}^*$ there exists $\mu \in \mathcal{K}$ with $x = 0$ such that $\tau v + \tau w^T u < 0$. Then $\alpha\mu \in \mathcal{K}$ for any $\alpha \geq 0$ and $\lim_{\alpha \rightarrow \infty} f(\alpha\mu) = -\infty$, so $f(\mu)$ is unbounded below on $\mu \in \mathcal{K}$.

Conversely, if $\tau > 0$ and $w \in \mathcal{C}^*$, we have $\tau v \geq 0$ and $\tau w^T u \geq 0$ for every $\mu \in \mathcal{K}$. So $f(\mu) \geq 0$
185 for $\mu \in \mathcal{K}$.

Since a primal feasible point always exists (take for example $v = 1, x = 0, u = e_0$), the problem (9) is solvable if and only if $\tau = 0$ or the conditions in Lemma 2 are fulfilled. These conditions can easily be checked prior to executing the algorithm and we assume that the problem is solvable in the following.

3.4. Optimality Conditions

With the conic modelling machinery in place we can begin to analyze the solution of (1) by considering the non-symmetric conic formulation (9). The Lagrangian of (9) is

$$L(\mu, \lambda) = \|x - y\|_2^2 + \tau(v + w^T u) - \langle \lambda, \mu \rangle \quad (22)$$

and the dual is

$$\begin{aligned} &\text{maximize} && -\frac{1}{4} \|s\|_2^2 - \text{Re}(y^H s) \\ &\text{subject to} && \lambda \in \mathcal{K}^*, \rho = \tau, z = \tau w. \end{aligned} \quad (23)$$

Notice that by taking the dual of (9) instead of (1), the number of dual variables is reduced from $\mathcal{O}(N^2)$ to $\mathcal{O}(N)$ (see [9] for an explicit formulation of the dual of (1)). This is the reason why, from a computational point of view, it is beneficial to work with the form (9) instead of
195 (1).

Since f is convex, the Karush-Kuhn-Tucker (KKT) are necessary and sufficient [20, Sec. 5.9] for variables (μ^*, λ^*) to be solutions of the primal and dual problems (9) and (23). The KKT conditions are

$$\begin{cases} \nabla_{\mu} L(\mu^*, \lambda^*) = 0 \\ \mu^* \in \mathcal{K} \\ \lambda^* \in \mathcal{K}^* \\ \langle \lambda^*, \mu^* \rangle = 0 \end{cases}. \quad (24)$$

Instead of directly solving the KKT conditions, our primal-dual IPM finds solutions $(\mu^{(t)}, \lambda^{(t)})$ of the augmented KKT conditions [20, 46]

$$\begin{cases} \nabla_{\mu} L(\mu^{(t)}, \lambda^{(t)}) = 0 \\ \mu^{(t)} \in \text{int } \mathcal{K} \\ \lambda^{(t)} \in \text{int } \mathcal{K}^* \\ \lambda^{(t)} = -t^{-1} \nabla_{\mu} F(\mu^{(t)}) \end{cases} \quad (25)$$

for an increasing sequence of values $t > 0$. It is easy to realize that $(\mu^{(t)}, \lambda^{(t)})$ solves (25) only if $\mu^{(t)}$ is a solution of the barrier problem (20). This observation provides the link between primal-only barrier methods and primal-dual IPMs. The set of values $\{(\mu^{(t)}, \lambda^{(t)}) : t > 0\}$ is known as the primal-dual central path. The primal-dual central path converges to the desired solution in the sense that $\lim_{t \rightarrow \infty} (\mu^{(t)}, \lambda^{(t)}) = (\mu^*, \lambda^*)$ [20, 46].

The last condition in (25) is known as the augmented complementary slackness condition. It follows from (19) that the second and fourth condition in (25) together imply $\lambda^{(t)} \in \text{int } \mathcal{K}^*$, so the third condition can be dropped.

From (18) it follows that the duality gap for the primal-dual problems (9) and (23) at a point on the primal-dual central path is $\langle \lambda^{(t)}, \mu^{(t)} \rangle = \theta_F/t$ [20, 23]. So solving the augmented KKT gives a primal feasible solution $\mu^{(t)}$ which is no more than $(N+1)/t$ suboptimal. Consequently, an arbitrarily accurate solution can be obtained by solving (25) for sufficiently large t .

3.5. Obtaining a Solution of the Augmented KKT Conditions

We now define $v^{(t)}, x^{(t)}, u^{(t)}, \rho^{(t)}, s^{(t)}$ and $z^{(t)}$ as the entries of $\mu^{(t)}$ and $\lambda^{(t)}$. By solving the first equation in (25) (the stationarity condition) we get

$$\rho^{(t)} = \tau, \quad z^{(t)} = \tau w, \quad s^{(t)} = 2(x^{(t)} - y). \quad (26)$$

We continue by writing out the last condition of (25). Solve for $v^{(t)}$ and $x^{(t)}$ and insert the relations above to get

$$v^{(t)} = (\tau t)^{-1} + \left(x^{(t)}\right)^H T^{-1} \left(u^{(t)}\right) x^{(t)} \quad (27)$$

$$x^{(t)} = T \left(u^{(t)}\right) T^{-1} \left(u^{(t)} + 2^{-1} \tau e_0\right) y. \quad (28)$$

Finally, solve $z^{(t)} = -t^{-1}\nabla_u F(\mu^{(t)})$ for $u^{(t)}$ and insert the above to obtain

$$\tau w - \tau T^*(\phi\phi^H) - t^{-1}T^*(T^{-1}(u^{(t)})) = 0, \quad (29)$$

where $\phi = T^{-1}(u^{(t)} + 2^{-1}\tau e_0)y$.

For a given $t > 0$ the corresponding point on the primal-dual central path can be obtained as follows: First a solution $u^{(t)}$ of (29) that fulfills $u^{(t)} \in \text{int } \mathcal{C}$ is found (existence of such a solution is shown below). Then the point $(\mu^{(t)}, \lambda^{(t)})$ is obtained by inserting into (26), (27) and (28). It is easy to show from $u^{(t)} \in \text{int } \mathcal{C}$ that $\mu^{(t)} \in \text{int } \mathcal{K}$ and so $(\mu^{(t)}, \lambda^{(t)})$ solves (25) and it is a primal-dual central point.

How can we obtain a solution $u^{(t)} \in \text{int } \mathcal{C}$ of (29)? The left-hand side of (29) is recognized as the gradient of $h_t(u) = g(u) + t^{-1}G(u)$, with

$$g(u) = \tau w^T u + \tau y^H T^{-1}(u + 2^{-1}\tau e_0)y \quad (30)$$

$$G(u) = -\log |T(u)|. \quad (31)$$

Now consider the barrier problem

$$\begin{aligned} & \text{minimize} && h_t(u) \\ & \text{subject to} && u \in \text{int } \mathcal{C}. \end{aligned} \quad (32)$$

The gradient of h_t vanishes at the solution of (32) because G is a LH barrier function for \mathcal{C} . So solving (29) with $u^{(t)} \in \text{int } \mathcal{C}$ is equivalent to solving (32). Since we have assumed that the problem (9) is solvable, then so is (32) (thus proving that there exists a $u^{(t)} \in \text{int } \mathcal{C}$ that solves (29)).

The idea of the primal-dual IPM presented in the following section is to use an iterative algorithm for unconstrained optimization (either Newton's method or a quasi-Newton method) to solve (32). However, we do not need to exactly solve (32) for a sequence of values $t > 0$. In each iteration of the solver the value of t can be updated in a dynamic manner based on the duality gap.

4. The Primal-Dual Interior-Point Method

We now outline FastAST, a primal-dual IPM for the solution of (9). Let $(\mu_i, u_i, \lambda_i, t_i)$ denote (μ, u, λ, t) in iteration i . The proposed method is given in Algorithm 1. In the remainder of this section, each step of the algorithm is discussed in detail.

Low-complexity evaluation of the steps in FastAST are presented in Sec. 5. With these approaches, the computational complexity is dominated by the evaluation of the search direction. For this step we propose to use either Newton's method or a quasi-Newton method. The quasi-Newton method has much lower computational complexity per iteration and is also faster in practice. It is, however, not able to obtain a solution of high accuracy. If a highly accurate

Algorithm 1: Primal-dual IPM for fast atomic norm soft thresholding (FastAST).

Parameters: $\gamma > 1$.

Input: Initial values $u_0 \in \mathcal{C}$ and $t_1 > 0$.

Set objective lower bound $f_{\text{LB}} = -\infty$.

for $i = 1, 2, \dots$ **do**

 Determine the search direction Δu .

 Perform a line search along Δu to obtain the step size α .

 Update estimate $u_i = u_{i-1} + \alpha \Delta u$.

 Form primal-dual variables (μ_i, λ_i) using (26), (27) and (28).

if $\lambda_i \in \mathcal{K}^*$ **then**

 Update lower bound on objective $f_{\text{LB}} = \max\left(f_{\text{LB}}, -\frac{1}{4} \|s_i\|_2^2 - \text{Re}(y^H s_i)\right)$.

end

 Evaluate duality gap $\eta_i = f(\mu_i) - f_{\text{LB}}$.

 Terminate if the stopping criterion is satisfied.

 Update barrier parameter $t_{i+1} = \max\left(t_i, \gamma \frac{N+1}{\eta_i}\right)$.

end

Output: Primal-dual solution (μ_i, λ_i) .

solution is required, Newton's method is preferred. We refer to the numerical evaluation in
235 Sec. 6 for a detailed discussion thereof.

4.1. Determining the Search Direction Using Newton's Method

Applying Newton's method to solve (32) we get the search direction

$$\Delta u = -(\nabla_u^2 h_{t_i}(u_{i-1}))^{-1} \nabla_u h_{t_i}(u_{i-1}), \quad (33)$$

where $\nabla_u^2 h_{t_i}(u_{i-1})$ denotes the Hessian of h_{t_i} evaluated at u_{i-1} . As discussed in Sec. 5 the Hessian can be evaluated in $\mathcal{O}(N^3)$ flops and the same cost is required for solution of the system (33).

4.2. Determining the Search Direction Using L-BFGS

In scenarios with large N the computation time for evaluation of the Newton search direction can become prohibitively large. In these cases we propose to use the limited-memory Broyden-Fletcher-Goldfarb-Shanno (L-BFGS) algorithm [47] for the solution of (32). L-BFGS possesses two key properties that are instrumental in obtaining an algorithm with low per-iteration
245 computational complexity: 1) it uses only gradient information and the gradient of h_{t_i} can be evaluated with low computational complexity;⁶ and 2) by appropriately modifying the L-BFGS

⁶To speed up convergence, our implementation also uses an approximation of the diagonal of the Hessian of h_{t_i} .

Algorithm 2: Modified L-BFGS two-loop recursion for calculation of the search direction.

Parameters: Number of saved difference vectors M .

Input: Current iteration number i and parameter t_i . Saved difference vectors r_k, q_k, Q_k for $k = i - 1, i - 2, \dots, \max(i - M, 1)$. Current gradient vector $\nabla_u h_{t_i}(u_{i-1})$ and initial Hessian approximation \hat{H}_i .

```

 $d \leftarrow -\nabla_u h_{t_i}(u_{i-1})$ 
for  $k = i - 1, i - 2, \dots, \max(i - M, 1)$  do
     $\psi_k \leftarrow q_k + t_i^{-1} Q_k$ 
     $\sigma_k \leftarrow \frac{r_k^T d}{r_k^T \psi_k}$ 
     $d \leftarrow d - \sigma_k \psi_k$ 
end
 $d \leftarrow \hat{H}_i^{-1} d$ 
for  $k = \max(i - M, 1), \max(i - M, 1) + 1, \dots, i - 1$  do
     $\beta_k \leftarrow \frac{\psi_k^T d}{\psi_k^T r_k}$ 
     $d \leftarrow d + r_k (\sigma_k - \beta_k)$ 
end
Output: Search direction  $\Delta u = d$ .

```

two-loop recursion, it can be used for the solution of (32) in a computational efficient manner even though t is increased in every iteration. It is this property, and not the limited memory requirements, that makes L-BFGS preferable over other quasi-Newton methods (such as vanilla BFGS) for our purposes.

In relation to the second property, note that since $t_i \neq t_{i-1} \neq \dots$, the normal formulation of L-BFGS does not apply. A simple modification of the L-BFGS two-loop recursion [47] overcomes this limitation. At the end of the i th iteration, the following difference vectors are calculated and saved for later use:

$$r_i = u_i - u_{i-1} \quad (34)$$

$$q_i = \nabla_u g(u_i) - \nabla_u g(u_{i-1}) \quad (35)$$

$$Q_i = \nabla_u G(u_i) - \nabla_u G(u_{i-1}). \quad (36)$$

This set of vectors is retained for M iterations. The modified two-loop recursion in Algorithm 2 can then be used to calculate the search direction Δu . This algorithm calculates the normal L-BFGS search direction for minimization of h_{t_i} , as if $t_i = t_{i-1} = \dots$. That can be achieved because L-BFGS only depends on t_i through $\nabla_u h_{t_i}(u_k) = \nabla_u g(u_k) + t_i^{-1} \nabla_u G(u_k)$, for $k = i - 1, \dots, i - M - 1$. The gradients $\nabla_u g(u_k)$ and $\nabla_u G(u_k)$ need only be calculated once to allow $\nabla_u h_{t_i}(u_k)$ to be calculated for any value of t_i .

In each iteration the initial Hessian \hat{H}_i should be chosen as an approximation of the Hessian of h_{t_i} evaluated at u_{i-1} . It is the matrix upon which L-BFGS successively applies rank-2 updates

to form the Hessian approximation that is used for calculating the search direction. An easy, and popular, choice for the initial Hessian is the identity matrix $\hat{H}_i = I$. Through numerical experiments we have seen that this choice leads to slow convergence. It turns out that the slow convergence is caused by the scaling of the Hessian, leading to non-acceptance of a full Newton step (i.e., α is selected much smaller than 1). Using a diagonal approximation of the true Hessian remedies this, but, unfortunately, it cannot be calculated with low computational complexity. (Our best attempt at devising a fast evaluation of the Hessian diagonal yielded cubic complexity $\mathcal{O}(N^3)$, the same as evaluation of the full Hessian.) Instead our algorithm uses the following heuristic approximation of the diagonal Hessian

$$\hat{H}_i = \text{diag}\left(1, \frac{N-1}{2N}, \dots, \frac{1}{2N}, \frac{N-1}{2N}, \dots, \frac{1}{2N}\right) (\nabla_u^2 h_{t_i}(u_{i-1}))_{0,0}, \quad (37)$$

where $(\nabla_u^2 h_{t_i}(u_{i-1}))_{0,0}$ is the $(0,0)$ th entry of the true Hessian evaluated at u_{i-1} . This approximation can be calculated with low computational complexity as demonstrated in Sec. 5. The approximation is motivated as follows: The diagonal entries are scaled according to the number of times the corresponding entry of u appears in $T(u)$. This scaling resembles that in the biased autocorrelation estimate (except for a factor of 2 caused by the scaling of the diagonal in the definition of $T(u)$). In our numerical experiments, we have observed the above approximation to be fairly accurate; each entry typically takes a value within $\pm 50\%$ of the true value. To this end we note that only a crude approximation is needed, since the role of \hat{H}_i is to account for the scaling of the problem. Our numerical investigation suggests that using the approximation (37) leads to only marginally slower convergence, compared to using a diagonal Hessian approximation using the diagonal of the true Hessian.

A final note on our adaptation of L-BFGS is that the usual observations regarding positive definiteness of the approximated Hessian remain valid. First note that the objective upon which L-BFGS is applied (h_{t_i}) is a strictly convex function for $u \in \text{int } \mathcal{C}$. It follows that the initial Hessian approximation \hat{H}_i is positive definite. Also, the curvature condition $r_k^T \psi_k > 0$ is valid for all k . Then the approximated Hessian is positive definite and the calculated search direction Δu is a descent direction [47, 4].

4.3. Line Search

The line search along the search direction Δu is a simple backtracking line search starting at $\alpha = 1$. A step size is accepted if the new point is strictly feasible, i.e., if $u_{i-1} + \alpha \Delta u \in \text{int } \mathcal{C}$. It is then easy to show that the primal solution μ_i calculated from inserting u_i into (27) and (28) is strictly primal feasible ($\mu_i \in \text{int } \mathcal{K}$).

To guarantee that the objective is sufficiently decreased, the Armijo rule is also required for acceptance of a step size α :

$$h_{t_i}(u_{i-1} + \alpha \Delta u) \leq h_{t_i}(u_{i-1}) + c \alpha \Delta u^T \nabla_u h_{t_i}(u_{i-1}), \quad (38)$$

where $0 < c < 1$ is a suitably chosen constant.

4.4. The Duality Gap and Update of t

The line search guarantees that the primal solution is strictly feasible in all iterations, i.e., that $\mu_i \in \text{int } \mathcal{K}$. Dual feasibility of a solution λ_i obtained from (26) is not guaranteed. The algorithm therefore checks for $\lambda_i \in \mathcal{K}^*$ using the approximate approach described in Sec. 3.1.

Let f^* denote the optimal value of the problem (9). If λ_i is dual feasible, the objective of the dual (23) provides a lower bound on the optimal value, i.e.,

$$f^* \geq -\frac{1}{4} \|s_i\|_2^2 - \text{Re}(y^H s_i). \quad (39)$$

The algorithm always retains the largest lower bound it has encountered in f_{LB} . From the lower bound, a duality gap η_i can be evaluated in each iteration:

$$\eta_i = f(\mu_i) - f_{\text{LB}}. \quad (40)$$

This value gives an upper bound on the sub optimality of the solution μ_i , i.e., $f(\mu_i) - f^* \leq \eta_i$.

Recall that the algorithm is “aiming” for a solution of the augmented KKT conditions (25). At this solution, the duality gap is θ_F/t_{i+1} . The next value of t can then be determined so that the algorithm is aiming for a suitable (not too large, not too small) decrease in the duality gap, i.e., we select t_{i+1} such that $\eta_i/\gamma = \theta_F/t_{i+1}$ for some preselected $\gamma > 1$. To guarantee convergence it is also imposed that t_i is a non-decreasing sequence.

4.5. Termination

The duality gap provides a natural stopping criterion. The proposed algorithm terminates based on either the duality gap ($\eta_i < \varepsilon_{\text{abs}}$) or the relative duality gap ($\eta_i/f(\mu_i) < \varepsilon_{\text{rel}}$). The relative duality gap is a sensible stopping criterion because $f(\mu) \geq 0$ as is seen in the proof of Lemma 2.

Algorithm 1 is guaranteed to terminate at a point that fulfills either of the two stopping criteria listed above. To see why that is the case, consider a scenario where t_i converges to some finite constant \tilde{t} as $i \rightarrow \infty$. Then, as $i \rightarrow \infty$, the algorithm implements L-BFGS with a backtracking line search to minimize $h_{\tilde{t}}$. Thus u_i converges to the minimizer $u^{(\tilde{t})}$ of $h_{\tilde{t}}$. Let $(\mu^{(\tilde{t})}, \lambda^{(\tilde{t})})$ denote the corresponding primal and dual variables calculated from (27), (28) and (26).

Now, $(\mu^{(\tilde{t})}, \lambda^{(\tilde{t})})$ constitute a solution to (25) with $t = \tilde{t}$. Then $\lambda^{(\tilde{t})} \in \text{int } \mathcal{K}^*$ follows from (19). Further, we have from (40), (26) and (18) that the duality gap η_i converges to $\langle \mu^{(\tilde{t})}, \lambda^{(\tilde{t})} \rangle = \theta_F/\tilde{t}$ as $i \rightarrow \infty$. However, that implies $t_{i+1} = \gamma\theta_F/\eta_i = \gamma\tilde{t} > \tilde{t}$ in the limit, a contradiction to the assumption that t_i converges to \tilde{t} . This means that t_i does not converge to a finite value and, as it is non-decreasing, it must diverge to $+\infty$. It is also evident that the duality gap $\eta_i \rightarrow 0$ as $t_i \rightarrow \infty$, and so either of the stopping criteria are eventually fulfilled.

4.6. Initialization

FastAST must be initialized with a primal variable $u_0 \in \mathcal{C}$ and a barrier parameter $t_1 > 0$. To determine a suitable value of the initial barrier parameter t_1 we first identify a primal-dual
 310 feasible point from which the duality gap can be evaluated. A primal-dual feasible point can be obtained by assuming⁷ $w \in \mathbf{int} \mathcal{C}^*$ and iterating these steps:

1. Set $u = 10\|y\|_2^2/N, 0, \dots, 0)^T$.
2. Calculate (μ, λ) from u based on (26), (27) and (28).
3. If $\lambda \in \mathcal{K}^*$, terminate, otherwise double the first entry of u and go to step 2.

315 The value of u in Step 1 has been chosen heuristically. It is easy to see that u stays primal feasible throughout. It is guaranteed that a dual feasible point is reached because $u \rightarrow (\infty, 0, \dots, 0)^T$. Then, following (28), we have $x \rightarrow y$ and so $s \rightarrow 0$. Considering the result in Lemma 1 and the assumption $w \in \mathbf{int} \mathcal{C}^*$ we get that λ converges to a point $\tilde{\lambda} \in \mathbf{int} \mathcal{K}^*$.

When a primal-dual feasible point (μ, λ) has been found the corresponding duality gap is $\eta_0 =$
 320 $\langle \mu, \lambda \rangle$. The initial value of the barrier parameter is selected as $t_1 = \gamma\theta_F/\eta_0$. The corresponding value of u is used as the initial value of the primal variable u_0 .

5. Fast Computations

For brevity iteration indices are dropped in the following. The computationally demanding steps of FastAST (Alg. 1) all involve the determinant or the inverse of Toeplitz matrices $T(u)$
 325 and $T(u + 2^{-1}\tau e_0)$. In this section we demonstrate how the Toeplitz structure can be exploited to significantly reduce the computational complexity of these evaluations. The use of such structure for fast solution of optimization problems have previously been seen [34, 35], including for evaluation of the gradient and Hessian of the barrier function G [36, 37].

5.1. Fast Algorithms for Factorizing a Toeplitz Inverse

Our computational approach is based on the following factorizations of Toeplitz inverses. The Gohberg-Semencul formula [30, 48] gives a factorization of the inverse of a Toeplitz matrix $T(u)$,

$$T^{-1}(u) = \delta_{N-1}^{-1}(U^H U - V V^H), \quad (41)$$

where the entries of Toeplitz matrices U and V are

$$U_{n,m} = \rho_{N-1+n-m}, \quad (42)$$

$$V_{n,m} = \rho_{n-m-1}, \quad (43)$$

⁷The problem (1) is solvable if and only if $w \in \mathcal{C}^*$. The restriction to the interior has no practical effect.

for $n, m = 0, \dots, N-1$. Note that $\rho_n = 0$ for $n < 0$ and $n > N-1$; thus U is unit upper triangular ($\rho_{N-1} = 1$) and V is strictly lower triangular.

The values δ_n and ρ_n for $n = 0, \dots, N-1$ can be computed with a generalized Schur algorithm in $\mathcal{O}(N \log^2 N)$ flops [30]. Alternatively, the Levinson-Durbin algorithm can be used to obtain the decomposition in $\mathcal{O}(N^2)$ flops. The latter algorithm is significantly simpler to implement and is faster for small N . In [31] it is concluded that the Levinson-Durbin algorithm requires fewer total operations than the generalized Schur algorithm for $N \leq 256$.

We will also use a Cholesky factorization of $T^{-1}(u)$,

$$T^{-1}(u) = PDP^H \quad (44)$$

where P is unit upper triangular and D is diagonal. The matrix $D = \text{diag}(\delta_0^{-1}, \dots, \delta_{N-1}^{-1})$ is inherently computed when the generalized Schur algorithm is executed [30]. The generalized Schur algorithm does not compute the matrix P . The Levinson-Durbin algorithm inherently computes both P and D , a property which we exploit for evaluation of the Hessian of the barrier function G .

In the following we let $\rho_0, \dots, \rho_{N-1}$ and $\delta_0, \dots, \delta_{N-1}$ be the entries obtained by executing the generalized Schur or Levinson-Durbin algorithm with either $T^{-1}(u)$ or $T^{-1}(u + 2^{-1}\tau e_0)$; which one will be clear from the context.

5.2. Evaluating the Objective and the Primal Variables

We first discuss evaluation of the objective $h_t(u) = g(u) + t^{-1}G(u)$. Since P in (44) has unit diagonals, it is easy to obtain

$$G(u) = -\log |T(u)| = -\sum_{n=0}^{N-1} \log \delta_n. \quad (45)$$

To evaluate $g(u)$ insert (41) into (30) and realize that all matrix-vector products involve Toeplitz matrices. Vector multiplication onto a Toeplitz matrix can be performed using the fast Fourier transform (FFT) in $\mathcal{O}(N \log N)$ flops (such products are convolutions, see e.g. [37] for details). In conclusion, the dominant cost of evaluating $h_t(u)$ is the execution of the generalized Schur (or Levinson-Durbin) algorithm.

Evaluating the primal variables $v^{(t)}$ and $x^{(t)}$ in (27) – (28) similarly amounts to vector products onto Toeplitz matrices.

The line search in Algorithm 1 must check for $u \in \mathcal{C}$, i.e., if $T(u) \succ 0$. The generalized Schur (or Levinson-Durbin) algorithm can again be used here, as $T(u) \succ 0$ if and only if $\delta_n > 0$ for $n = 0, \dots, N-1$.

5.3. Evaluating the Gradients

The following gradients must be evaluated in each iteration of Algorithm 1:

$$\nabla_u g(u) = \tau w - \tau T^*(\phi\phi^H) \quad (46)$$

$$\nabla_u G(u) = -T^*(T^{-1}(u)). \quad (47)$$

We first consider the term $T^*(\phi\phi^H)$. The vector ϕ can be evaluated with low complexity (confer the evaluation of primal variables, above). Let $\beta_n \in \mathbb{C}$ denote the sum over the n th upper diagonal of $\phi\phi^H$ for $n = 0, \dots, N-1$, i.e.,

$$\beta_n = \sum_{m=0}^{N-1-n} (\phi\phi^H)_{m,m+n} = \sum_{m=0}^{N-1-n} \phi_m \bar{\phi}_{m+n}. \quad (48)$$

It is recognized that the values $\beta_0, \dots, \beta_{N-1}$ can be calculated as a correlation, which can be implemented using FFTs in $\mathcal{O}(N \log N)$ flops. Then $T^*(\phi\phi^H)$ can be obtained by concatenating and scaling the real and imaginary parts of β ,

$$T^*(\phi\phi^H) = (2\beta_0, 2\operatorname{Re}(\beta_1), \dots, 2\operatorname{Re}(\beta_{N-1}), 2\operatorname{Im}(\beta_1), \dots, 2\operatorname{Im}(\beta_{N-1}))^T. \quad (49)$$

Now consider evaluation of the term $T^*(T^{-1}(u))$. The sum over the n th upper diagonal of $T^{-1}(u)$ is denoted as $\tilde{\beta}_n$ and can be rewritten as

$$\tilde{\beta}_n = \sum_{m=0}^{N-1-n} (T^{-1}(u))_{m,m+n} = \delta_{N-1}^{-1} \sum_{k=0}^{N-1} (n - N + 2(k+1)) \rho_k \bar{\rho}_{k+n}, \quad (50)$$

see [34, 35] for details. The above is recognized as two correlations, thus allowing a low-complexity evaluation. The vector $T^*(T^{-1}(u))$ is found by concatenating and scaling the real and imaginary parts of $\tilde{\beta}$, analogously to (49).

5.4. Evaluating the Full Hessian

When Newton's method is used to determine the search direction, the Hessian of h_t must be evaluated. We now derive an approach to calculate the Hessians of g and G , from which the required Hessian is easily found.

The (n, m) th entry of the Hessian of g is

$$(\nabla_u^2 g(u))_{n,m} = 2\tau \phi^H (E_n + E_n^H) T^{-1}(u + 2^{-1} \tau e_0) (E_m + E_m^H) \phi, \quad (51)$$

where

$$E_n = \begin{cases} I & n = 0 \\ \tilde{E}^n & 1 \leq n \leq N-1 \\ -jE_{n-N+1} & N \leq n \leq 2N-1. \end{cases} \quad (52)$$

The matrix \tilde{E} is the lower shift matrix, i.e., it has ones on the lower subdiagonal and zeros elsewhere. Note that $T(e_n) = E_n + E_n^H$. The m th column of the Hessian is then

$$(\nabla_u^2 g(u))_m = \tau T^*(d_m \phi^H + \phi d_m^H), \quad (53)$$

where we let d_m denote a vector $d_m = T^{-1}(u + 2^{-1}\tau e_0)(E_m + E_m^H)\phi$. Assuming the decomposition (41) is available, a column of the Hessian can be calculated in $\mathcal{O}(N \log N)$ flops by explicitly forming d_m and performing sums over diagonals (as in (48)). The full Hessian of g is then obtained in $\mathcal{O}(N^2 \log N)$ flops.

To evaluate the Hessian of the barrier function G we generalize the approach of [37] to the complex-valued case. The (n, m) th entry of the Hessian is

$$(\nabla_u^2 G(u))_{n,m} = \text{tr}(T^{-1}(u)(E_n + E_n^H)T^{-1}(u)(E_m + E_m^H)). \quad (54)$$

Define the $N \times N$ matrices A and B with entries

$$A_{n,m} = 2 \text{tr}(T^{-1}(u)E_n T^{-1}(u)E_m) \quad (55)$$

$$B_{n,m} = 2 \text{tr}(T^{-1}(u)E_n T^{-1}(u)E_m^T). \quad (56)$$

Then the Hessian can be written in the form

$$\nabla_u^2 G(u) = \begin{pmatrix} \text{Re}(A + B) & \text{Re}(-jAJ^T) \\ \text{Re}(-jJA - jJB) & \text{Re}(-JAJ^T + JBJ^T) \end{pmatrix}, \quad (57)$$

where J is a matrix that removes the first row, i.e., $J = (0, I)$, where 0 is a column of zeros and I is the $(N-1) \times (N-1)$ identity matrix.

At this point, we need a fast way of evaluating matrices A and B . Define the discrete Fourier transform matrix $W \in \mathbb{C}^{N_{\text{FFT}} \times N}$ with entries

$$W_{n,m} = \exp(-j2\pi nm/N_{\text{FFT}}), \quad (58)$$

where N_{FFT} is chosen such that $N_{\text{FFT}} \geq 2N - 1$. Recall that the Levinson-Durbin algorithm gives the decomposition $T^{-1}(u) = PDP^H$, from which $T^{-1}(u) = RR^H$ is obtained by calculating $R = PD^{\frac{1}{2}}$. Let S_n denote the discrete Fourier transform of the n th column of R (denote this column R_n), i.e., $S_n = WR_n$. Then by straight-forward generalization of the derivation in [37] to the complex-valued case, we get that A and B can be written in the forms

$$A = \frac{2}{N_{\text{FFT}}^2} W^T \left(\left(\sum_{l=0}^{N-1} S_l S_l^H \right) \odot \left(\sum_{l=0}^{N-1} S_l S_l^H \right) \right) W \quad (59)$$

$$B = \frac{2}{N_{\text{FFT}}^2} W^T \left(\left(\sum_{l=0}^{N-1} S_l S_l^H \right) \odot \left(\sum_{l=0}^{N-1} S_l S_l^H \right) \right) \bar{W}, \quad (60)$$

with \odot denoting the Hadamard (entrywise) product. Using (59) – (60) the Hessian of G can be evaluated in $\mathcal{O}(N^3)$ flops.

5.5. Evaluating the Diagonal Hessian Approximation

The L-BFGS variant of FastAST uses the approximation of the Hessian diagonal (37) which requires calculation of the first entry of the Hessian

$$(\nabla_u^2 h_t(u))_{0,0} = (\nabla_u^2 g(u))_{0,0} + \frac{1}{t} (\nabla_u^2 G(u))_{0,0}. \quad (61)$$

An $\mathcal{O}(N \log N)$ evaluation of the first term is easily obtained from (51). The second term can be evaluated based on (57), but a more efficient way is as follows: From (54) we have

$$(\nabla_u^2 G(u))_{0,0} = 4 \operatorname{tr}(T^{-1}(u)T^{-1}(u)). \quad (62)$$

The matrix $T^{-1}(u)$ can be formed explicitly in $\mathcal{O}(N^2)$ flops using the Trench algorithm [49, 50]. However, since the decomposition (41) is already available in our setting it is much easier to form $T^{-1}(u)$ from it by writing for $n = 0, \dots, N-1$ and $m = 0, \dots, N-1-n$

$$T^{-1}(u)_{m,m+n} = \delta_{N-1}^{-1} \left(\sum_{k=0}^m \bar{\rho}_{N-1-k} \rho_{N-1-(k+n)} - \rho_{k-1} \bar{\rho}_{k+n-1} \right),$$

i.e., $T^{-1}(u)$ is “formed along the diagonals”. By implementing the above sum as a cumulative sum, the complete matrix $T^{-1}(u)$ is formed in $\mathcal{O}(N^2)$ flops. Note that since $T(u)$ is both Hermitian and persymmetric, then so is $T^{-1}(u)$. This means that only one “wedge” of the matrix, about $N/4$ entries, must be calculated explicitly [50].

The trace in (62) is evaluated by taking the magnitude square of all entries in $T^{-1}(u)$ and summing them.

5.6. Analysis of Computational Complexity

To summarize the computational complexity of an implementation of Alg. 1 based on the low-complexity evaluations above, consider each of the two variants for determining the search direction.

- FastAST Newton: The computation time is asymptotically dominated by evaluation and inversion of the Hessian, i.e., $\mathcal{O}(N^3)$ flops.
- FastAST L-BFGS: The computation time is asymptotically dominated by the $\mathcal{O}(MN)$ modified L-BFGS two-loop recursion in Alg. 2 or by the $\mathcal{O}(N^2)$ evaluation of the diagonal Hessian approximation.

When using the Newton search direction, the decomposition (44) is required and the Levinson-Durbin algorithm must therefore be used to evaluate the factorization of the Toeplitz inverse.

When using the L-BFGS search direction either the generalized Schur or the Levinson-Durbin algorithm can be used. The choice does not affect the asymptotic computational complexity, but one may be faster than the other in practice.

Variant	L-BFGS	Newton
Number of saved difference vectors M	$2N - 1$	-
Armijo parameter c	0.05	0.05
Barrier parameter multiplier γ	2	10
Absolute tolerance ε_{abs}	10^{-4}	10^{-7}
Relative tolerance ε_{rel}	10^{-4}	10^{-7}

Table 1: Algorithm parameters.

6. Numerical Experiments

6.1. Setup & Algorithms

In our experiments we use the signal model (2). The frequencies $\omega_0, \dots, \omega_{K-1}$ are drawn randomly on $[0, 2\pi)$, such that the minimum separation⁸ between any two frequencies is $4\pi/N$. The coefficients c_0, \dots, c_{K-1} are generated independently random according to a circularly symmetric standard complex Gaussian distribution. After generating the set of K frequencies and coefficients the variance of the noise vector ζ is selected such that the desired signal-to-noise ratio (SNR) is obtained. The regularization parameter τ is selected from (8) based on the true noise variance. We assess the algorithms based on their ability to solve AST, which is obtained by selecting $w = 2e_0$ in (1).

We show results for both the L-BFGS and Newton's variants of FastAST⁹. For $N \leq 512$ our implementation uses the Levinson-Durbin algorithm for Toeplitz inversion, while for $N > 512$ it uses the generalized Schur algorithm where applicable. The parameters of the algorithm are listed in Table 1. It is worth to say a few words about the number of saved difference vectors M in L-BFGS. On the one hand, selecting larger values of M can decrease the total number of iterations required (by improving the Hessian approximation), but on the other hand doing so increases the number of flops required per iteration. In our numerical experiments we have found that setting it equal to the size of u ($M = 2N - 1$) provides a good trade-off. Loosely speaking this choice allows L-BFGS to perform a full-rank update of the Hessian approximation, while it does not increase the asymptotic per-iteration computational complexity. With this choice the algorithm asymptotically requires $\mathcal{O}(N^2)$ flops per iteration.

Performance of the ADMM algorithm¹⁰ [9] is also shown along with that of CVX [51] applied with both the SeDuMi[21] and Mosek¹¹ solvers.

6.2. Solution Accuracy Per Iteration

For this investigation a ground-truth solution of (1) is obtained using CVX+SeDuMi with the precision setting set to "best". We denote this value as μ^* . Fig. 1 shows the normalize

⁸The wrap-around distance on $[0, 2\pi)$ is used for all frequency differences.

⁹Our code is publicly available at github.com/thomaslundgaard/fast-ast.

¹⁰We use the implementation from github.com/badrinarayan/astlinespec.

¹¹mosek.com

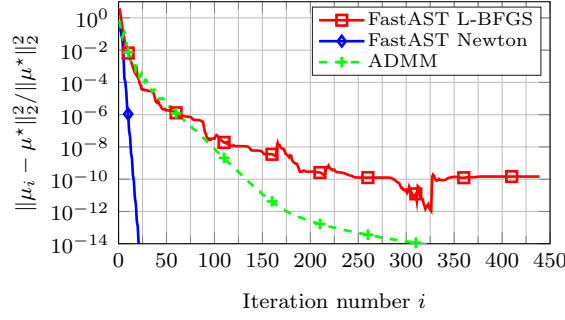


Figure 1: Solution accuracy versus iteration. The signal length is $N = 64$, the number of sinusoids is $K = 6$ and the SNR is 20 dB.

squared error between μ^* and the solution in each iteration of the algorithms. The algorithms ignore the stopping criteria and run until no further progress can be made towards the solution.

FastAST Newton converges very fast and a solution of very high accuracy is obtained within 25 iterations. This is due to the well-known quadratic convergence of Newton's method. While FastAST L-BFGS converges significantly slower it requires only $\mathcal{O}(N^2)$ flops per iteration versus the $\mathcal{O}(N^3)$ flops per iteration of FastAST Newton. We therefore cannot, at this point, conclude which version of FastAST is faster in practice. Note that ADMM on the other hand requires $\mathcal{O}(N^3)$ flops per iteration, the same as FastAST Newton, but requires significantly more iterations.

It is seen that FastAST L-BFGS seems to not make progress after approx. 300 iterations. This happens due to numerical challenges in evaluating the L-BFGS search direction. It is well-known that Woodbury's matrix identity, upon which L-BFGS is based, has limited numerical stability. For this reason FastAST L-BFGS is unable to obtain a solution of the same accuracy as the SeDuMi and Mosek solvers. Despite of this, as seen in the following sections, the solution accuracy of FastAST L-BFGS is sufficiently high in all cases but those with very high SNR. The tolerance values of FastAST L-BFGS are selected larger than for FastAST Newton (Table 1) because of the mentioned numerical issues with obtaining a high-accuracy solution.

FastAST Newton does not suffer from this problem and can obtain a solution of about the same accuracy as SeDuMi and Mosek. ADMM can also obtain a solution of high accuracy but, as can be seen in Fig. 1, it has slow convergence starting around iteration number 175. It therefore takes a large number of iterations to obtain a solution of the same accuracy as SeDuMi/Mosek or FastAST Newton.

6.3. Metrics

In the following we perform a Monte Carlo simulation study. Four metrics of algorithm performance and behaviour are considered: normalized mean-square error (NMSE) of the reconstructed signal x ; mean-square error (MSE) of the frequencies $\{\omega_k\}$ conditioned on successful recovery; number of iterations and algorithm runtime. The NMSE of the reconstructed signal is obtained by estimating the frequencies from the dual polynomial as described in [9] and using

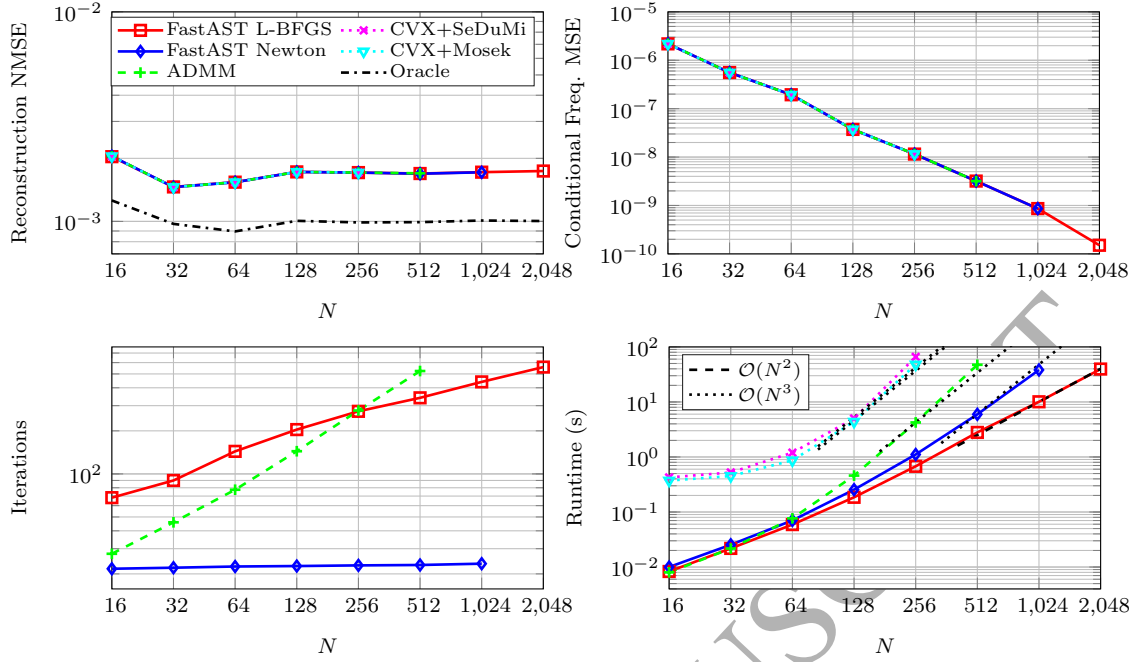


Figure 2: Simulation results for varying problem size N . The SNR is 20 dB and the number of sinusoids K is selected as $N/10$ rounded to the nearest integer. Results are averaged over 100 Monte Carlo trials. The legend applies to all plots; only the NMSE of Oracle is shown. In the figure with runtime the asymptotic per-iteration computational complexity is also plotted.

these to obtain the least-squares solution for the coefficients. An estimate of x is then obtained by inserting into (2). This estimate is also known as the *debiased* solution and it is known to have smaller NMSE than the estimate of x directly obtained as the solution of (1) [9]. In the evaluation of the signal reconstruction the performance of an Oracle estimator is also shown. The Oracle estimator knows the true frequencies and estimates the coefficients using least-squares.

To directly assess the accuracy with which the frequencies are estimated we present the MSE of the frequency estimates obtained from the dual polynomial. The MSE of the frequency estimates is only calculated based on these Monte Carlo trials in which the set of frequencies is successfully recovered. Successful recovery is understood as correct estimation of the model order K and that all frequency estimates are within a distance of π/N from their true value. The association of the estimated to the true frequencies is obtained by minimizing the frequency MSE using the Hungarian method [52].

The simulations are performed on a T470p Lenovo, with an Intel(R) Core(TM) i7-7820HQ CPU @ 2.90GHz, using MATLAB R2018b. MATLAB is restricted to only use a single CPU core, such that the runtime of the algorithms can be compared without differences in the parallelism achieved in the implementations. The computationally heavy steps of FastAST and ADMM are implemented in native code using the automatic code generation (“codegen”) feature of MATLAB.

6.4. Performance Versus Problem Size

The performance versus problem size N is depicted in Fig. 2. First note that all algorithms give the same estimation accuracy at all problem sizes, providing strong evidence that they correctly solve (1).

The number of iterations of FastAST L-BFGS increases with N . It is then expected that the total runtime asymptotically scales at a rate above the per-iteration cost of $\mathcal{O}(N^2)$ flops. Even still, the runtime for N up to 2,048 scales at a rate of about $\mathcal{O}(N^2)$. The number of iterations of FastAST Newton is practically independent of N . We then expect the total runtime to scale asymptotically as $\mathcal{O}(N^3)$. In practice it scales a little better for the values of N considered here. The number of iterations of ADMM increases significantly with N (doubling N roughly doubles the number of iterations). This in turn means that the runtime scales faster than the asymptotic per-iteration cost of $\mathcal{O}(N^3)$ flops.

In conclusion both variants of FastAST are faster than ADMM already at $N = 128$ and their runtime scales at a rate much slower than ADMM. This means that they are significantly faster than ADMM for large values of N . For large N it is also clear that the L-BFGS variant of FastAST is significantly faster than the Newton variant.

6.5. Performance Versus Signal-to-Noise Ratio

Fig. 3 shows performance versus the SNR level. Note that the conditional MSE of the frequency estimates is not shown for 0 dB SNR because there are no Monte Carlo trials with successful recovery of the frequencies at this SNR.

At SNR up to 30 dB all the algorithms perform the same in terms of NMSE of x and conditional MSE of the frequency estimates. This means that all algorithms have found a sufficiently accurate solution of (1) (relative to the SNR). In SNR larger than 30 dB FastAST L-BFGS shows a degraded solution accuracy compared to the remaining algorithms. This is due to the mentioned numerical issues and the consequently larger tolerances selected (cf. Table 1).

In terms of number of iterations and runtime note that both variants of FastAST show roughly unchanged behaviour with different SNR. ADMM on the other hand requires more iterations and has larger runtime for large SNR. In large SNR it is evident that FastAST Newton is preferred due to lower runtime than ADMM and higher estimation accuracy than FastAST L-BFGS.

7. Conclusions

The FastAST algorithm presented in this paper provides a fast approach to solving the atomic norm soft thresholding problem (1). The L-BFGS variant provides a reasonably accurate solution and is much faster than any other algorithm for large problem size N . If a solution of high accuracy is requested, which may be desirable in very high SNR, a variant of FastAST based on Newton's method is also provided. This variant can find a solution of high accuracy in

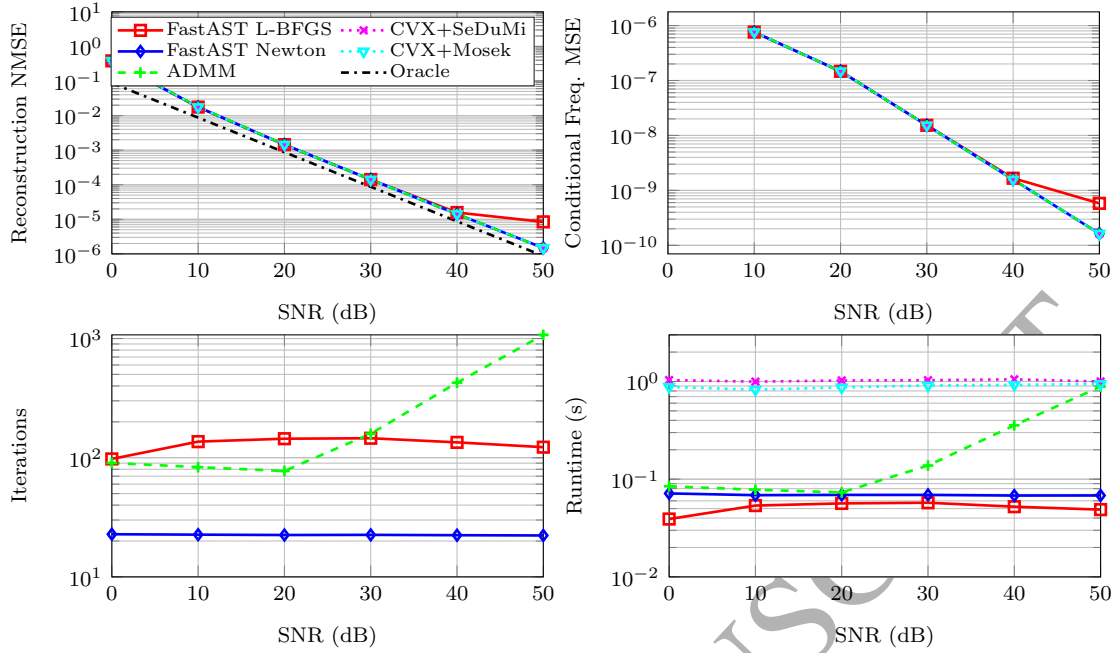


Figure 3: Simulation results for varying SNR. The signal length is $N = 64$ and the number of sinusoids is $K = 6$. Results are averaged over 100 Monte Carlo trials. The legend applies to all plots; only the NMSE of Oracle is shown.

a small number of iterations. While it is slower than FastAST L-BFGS, it is significantly faster than the state-of-the-art method based on ADMM.

The FastAST algorithm is obtained by reformulating the semidefinite program (1) as a non-symmetric conic program (9). This reformulation is of key importance in obtaining a fast algorithm. This work has provided an example of an optimization problem where it is beneficial to formulate it as a non-symmetric conic program instead of the standard, and much better understood, formulation as a symmetric conic program. We have also provided an implementation of a non-symmetric conic solver, thereby demonstrating the practical feasibility of this class of methods.

We have demonstrated how the L-BFGS two-loop recursion can be modified to allow a quasi-Newton solution of the barrier problem (32) even when the barrier parameter t is updated in every iteration. This approach can directly be applied in other algorithms based on the barrier method, including primal-only methods.

Finally note that there are many examples of optimization problems of practical interest which involve a constraint in either the cone of finite autocorrelation sequences \mathcal{C}^* or the cone \mathcal{K} . An example is the gridless SPICE method [15] for line spectral estimation; or frequency-domain system identification and filter design as summarized in [37]. We expect that equally fast primal-dual IPMs can be derived for all of these problems using the techniques of this paper. We also expect that it is fairly straight-forward to extend FastAST to atomic norm minimization with partial observations [8] or multiple measurement vectors [14]. An interesting, but less obvious, extension is to the multi-dimensional harmonic retrieval problem [53]; for that purpose

the work [54] may contain some useful insights.

Acknowledgements

We would like to thank Lieven Vandenberghe and Martin Skovgaard Andersen for providing valuable input to the work and pointing us to some important references.

Appendix A. Characterization of \mathcal{K}^*

To characterize the dual cone \mathcal{K}^* a number of lemmas are needed.

Lemma 3. *Let \mathcal{K} be a proper cone and assume $\lambda \neq 0$. If $\langle \lambda, \mu \rangle \geq 0$ for every $\mu \in \text{int } \mathcal{K}$, then $\langle \lambda, \mu \rangle \geq 0$ for every $\mu \in \mathcal{K}$.*

PROOF. Let $\tilde{\mu} \in \mathcal{K}$ and let $\{\mu_i\}$ be a sequence which converges to $\tilde{\mu}$ with $\mu_i \in \text{int } \mathcal{K}$. Then $\langle \lambda, \mu_i \rangle \geq 0$ and so $\langle \lambda, \tilde{\mu} \rangle = \lim_{i \rightarrow \infty} \langle \lambda, \mu_i \rangle \geq 0$, completing the proof.

Lemma 4. *Let \mathcal{K} be a proper cone. The interior of its dual is given by*

$$\text{int } \mathcal{K}^* = \{\lambda : \langle \lambda, \mu \rangle > 0 \ \forall \mu \in \mathcal{K}\}. \quad (\text{A.1})$$

PROOF. See [20], exercise 2.31.

To formulate the next lemma, the dual barrier of F is introduced:

$$F^*(\lambda) = \sup \{-\langle \lambda, \mu \rangle - F(\mu) : \mu \in \text{int } \mathcal{K}\}. \quad (\text{A.2})$$

This function is a slight modification ($-\langle \lambda, \mu \rangle$ replaces $\langle \lambda, \mu \rangle$) of the convex conjugate of F . It turns out that F^* is a LH barrier function for the dual cone \mathcal{K}^* [20, 23]. Its usefulness for our purposes lies in the following property.

Lemma 5. *Assume $\lambda \neq 0$ and let \mathcal{K} be a proper cone with corresponding LH barrier function F . Then $\lambda \in \text{int } \mathcal{K}^*$ if and only if $F^*(\lambda) < \infty$ (i.e., F^* is bounded above.)*

PROOF. We first prove the direct implication. Reasoning by contradiction, assume that $F^*(\lambda) < \infty$ and that there exists a $\mu \in \text{int } \mathcal{K}$ such that $\langle \lambda, \mu \rangle < 0$. Then $\alpha\mu \in \text{int } \mathcal{K}$ for all $\alpha > 0$. But $\lim_{\alpha \rightarrow \infty} -\langle \lambda, \alpha\mu \rangle - F(\alpha\mu) = \lim_{\alpha \rightarrow \infty} -\alpha \langle \lambda, \mu \rangle - F(\mu) + \theta_F \log(\alpha) = \infty$, a contradiction, so $\langle \lambda, \mu \rangle \geq 0$ for every $\mu \in \text{int } \mathcal{K}$. By Lemma 3 we have $\langle \lambda, \mu \rangle \geq 0$ for every $\mu \in \mathcal{K}$, thus $\lambda \in \mathcal{K}^*$. Since F^* is a LH barrier function for \mathcal{K}^* , it is easy to show that $F^*(\lambda) < \infty$ implies $\lambda \notin \text{bd } \mathcal{K}^*$, so $\lambda \in \text{int } \mathcal{K}^*$.

To prove the converse assume $\lambda \in \text{int } \mathcal{K}^*$. Then by Lemma 4, we have $\langle \lambda, \mu \rangle > 0$ for all $\mu \in \mathcal{K}$. It follows that there exists an $\varepsilon > 0$ such that $\langle \lambda, \tilde{\mu} \rangle \geq \varepsilon$ for every $\tilde{\mu} \in \mathcal{K}$ with $\|\tilde{\mu}\|_2 = 1$.

By continuity of F it can also be shown that there exists a δ such that $F(\tilde{\mu}) \geq \delta$ for every $\tilde{\mu} \in \mathcal{K}$ with $\|\tilde{\mu}\|_2 = 1$. With $\tilde{\mu} = \mu / \|\mu\|_2$, the objective in (A.2) obeys

$$\begin{aligned} -\langle \lambda, \mu \rangle - F(\mu) &= -\|\mu\|_2 \langle \lambda, \tilde{\mu} \rangle - F(\|\mu\|_2 \tilde{\mu}) \\ &= -\|\mu\|_2 \langle \lambda, \tilde{\mu} \rangle - F(\tilde{\mu}) + \theta_F \log(\|\mu\|_2) \\ &\leq -\|\mu\|_2 \varepsilon - \delta + \theta_F \log(\|\mu\|_2). \end{aligned}$$

The second equality follows from logarithmic homogeneity of F . This function is bounded above and so $F^*(\lambda) < \infty$.

We are now ready to give the desired proof.

PROOF (PROOF OF LEMMA 1). It is easy to show the following:

1. If $\rho < 0$, then $\lambda \notin \mathcal{K}^*$.
2. If $\rho = 0$ and $s \neq 0$, then $\lambda \notin \mathcal{K}^*$.
3. If $\rho = 0$ and $s = 0$, then $\lambda \in \mathcal{K}^*$ if and only if $z \in \mathcal{C}^*$.

The first and second property are shown by constructing a $\mu \in \mathcal{K}$ such that $\langle \lambda, \mu \rangle < 0$. The third property is shown by writing $\langle \lambda, \mu \rangle = z^T u \geq 0$ for all $\mu \in \mathcal{K}$ if and only if $z^T u \geq 0$ for all $u \in \mathcal{C}$.

The only case we have not considered so far is $\rho > 0$. By Lemma 5 and (21), we have $\lambda \in \text{int } \mathcal{K}^*$ if and only if $F^*(\lambda) < \infty$, i.e., when

$$h(\mu) = -\rho v - \text{Re}(s^H x) - z^T u + \log |T(u)| + \log(v - x^H T^{-1}(u)x)$$

is bounded above on the domain $\mu \in \text{int } \mathcal{K}$. The function h is concave and by setting the gradient equal to zero we get optimal points

$$\begin{aligned} v^* &= \rho^{-1} + (2\rho)^{-2} s^H T(u)s \\ x^* &= -(2\rho)^{-1} T(u)s. \end{aligned}$$

It is easy to show that if $u \in \text{int } \mathcal{C}$, then $(v^*, x^*, u)^T \in \text{int } \mathcal{K}$. Inserting into $h(\mu)$ we obtain

$$\begin{aligned} h(\mu) &\leq -z^T u + \frac{1}{4\rho} s^H T(u)s - 1 - \log(\rho) + \log |T(u)| \\ &= -c(\lambda)^T u - 1 - \log(\rho) + \log |T(u)|, \end{aligned}$$

with $c(\lambda) = z - \frac{1}{4\rho} T^*(ss^H)$. For each $u \in \text{int } \mathcal{C}$ there exists some corresponding $\mu \in \text{int } \mathcal{K}$ such that the above holds with equality.

If $c(\lambda) = 0$, the function $h(\mu)$ is unbounded above on the domain $\mu \in \text{int } \mathcal{K}$ and so $\lambda \notin \text{int } \mathcal{K}^*$. If $c(\lambda) \neq 0$ we can use Lemma 5 because $-\log |T(u)|$ is a LH barrier function for \mathcal{C} . So $h(\mu)$ is bounded above on the domain $\mu \in \text{int } \mathcal{K}$ if and only if $c(\lambda) \in \text{int } \mathcal{C}^*$. Tracing back our steps above we have (for $\rho > 0$) that $\lambda \in \text{int } \mathcal{K}^*$ if and only if $c(\lambda) \in \text{int } \mathcal{C}^*$. Since both of the dual cones are closed sets and $c(\cdot)$ is a continuous function, we have $\lambda \in \mathcal{K}^*$ if and only if $c \in \mathcal{C}^*$.

That completes the proof.

References

- [1] V. Chandrasekaran, B. Recht, P. A. Parrilo, A. S. Willsky, The convex geometry of linear inverse problems, *Found. Comp. Math.* 12 (6) (2012) 805–849.
- [2] Y. Nesterov, A. Nemirovskii, *Interior-Point Polynomial Methods in Convex Programming*, SIAM, 1994.
- [3] S. J. Wright, *Primal-Dual Interior-Point Methods*, SIAM, 1997.
- [4] J. Nocedal, S. Wright, *Numerical Optimization*, Springer, 1999.
- [5] E. Candès, J. Romberg, T. Tao, Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information 52 (2) (2006) 489–509.
- [6] D. Donoho, Compressed sensing 52 (4) (2006) 1289–1306.
- [7] E. J. Candès, M. B. Wakin, An introduction to compressive sampling, *IEEE Signal Process. Mag.* 25 (2) (2008) 21 – 30.
- [8] G. Tang, B. N. Bhaskar, P. Shah, B. Recht, Compressed sensing off the grid 59 (11) (2013) 7465–7490.
- [9] B. Bhaskar, G. Tang, B. Recht, Atomic norm denoising with application to line spectral estimation 61 (23) (2013) 5987–5999.
- [10] G. Tang, B. N. Bhaskar, B. Recht, Near minimax line spectral estimation 61 (1) (2015) 499–512.
- [11] E. J. Candès, C. Fernandez-Granda, Towards a mathematical theory of super-resolution, *Commun. Pure Appl. Math.* 67 (6) (2014) 906–956.
- [12] E. J. Candès, C. Fernandez-Granda, Super-resolution from noisy data, *J. Fourier Anal. Applicat.* 19 (6) (2013) 1229–1254.
- [13] M. Cho, K. V. Mishra, J. F. Cai, W. Xu, Block iterative reweighted algorithms for super-resolution of spectrally sparse signals, *IEEE Signal Process. Lett.* 22 (12) (2015) 2319–2313.
- [14] Y. Li, Y. Chi, Off-the-grid line spectrum denoising and estimation with multiple measurement vectors, *IEEE Trans. Signal Process.* 64 (5) (2016) 1257–1269.
- [15] Z. Yang, L. Xie, On gridless sparse methods for line spectral estimation from complete and incomplete data 63 (12) (2015) 3139–3153.
- [16] M. F. Da Costa, W. Dai, Low dimensional atomic norm representations in line spectral estimation, in: *IEEE Int. Symp. Inform. Theory*, 2017, pp. 226–230.

- [17] N. Rao, P. Shah, S. Wright, Forward-backward greedy algorithms for atomic norm regularization, *IEEE Trans. on Signal Process.* 63 (21) (2015) 5798–5811.
- [18] N. Boyd, G. Schiebinger, B. Recht, The alternating descent conditional gradient method for sparse inverse problems, *SIAM J. on Optim.* 27 (2) (2017) 616–639.
- 595 [19] M. Vinyes, G. Obozinski, Fast column generation for atomic norm regularization, in: A. Singh, J. Zhu (Eds.), *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, Vol. 54 of *Proceedings of Machine Learning Research*, PMLR, Fort Lauderdale, FL, USA, 2017, pp. 547–556.
- [20] S. Boyd, L. Vandenberghe, *Convex Optimization*, Cambridge University Press, 2004.
- 600 [21] J. F. Sturm, Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones, *Optim. Methods Softw.* 11-12 (1999) 625–653.
- [22] K. Toh, M. Todd, R. Tutuncu, SDPT3 — a Matlab software package for semidefinite programming, *Opt. Methods Softw.* 11 (1999) 545–581.
- [23] Y. E. Nesterov, M. J. Todd, Self-scaled barriers and interior-point methods for convex programming, *Math. Oper. Res.* 22 (1) (1997) 1–42.
- 605 [24] Y. E. Nesterov, M. J. Todd, Primal-dual interior-point methods for self-scaled cones, *SIAM J. Opt.* 8 (2) (1998) 324–364.
- [25] O. Güler, Barrier functions in interior point methods, *Math. Oper. Res.* 21 (4) (1996) 860–885.
- 610 [26] Y. Nesterov, Towards non-symmetric conic optimization, *Opt. Methods Softw.* 27 (4-5) (2012) 893–917.
- [27] A. Skajaa, J. B. Jørgensen, P. C. Hansen, On implementing a homogeneous interior-point algorithm for nonsymmetric conic optimization, *Tech. Rep. IMM-2011-02*, Technical University of Denmark (2011).
- 615 [28] A. Skajaa, Y. Ye, A homogeneous interior-point algorithm for nonsymmetric convex conic optimization, *Math. Prog.* 150 (2) (2015) 391–422.
- [29] L. Tunçel, Generalization of primal-dual interior-point methods to convex optimization problems in conic form, *Found. Comp. Math.* 1 (3) (2001) 229–254.
- [30] G. S. Ammar, W. B. Gragg, The generalized Schur algorithm for the superfast solution of Toeplitz systems, in: *Rational Approx. Applicat. Math. Phys.*, Springer, 1987, pp. 315–330.
- 620

- [31] G. S. Ammar, W. B. Gragg, Numerical experience with a superfast real Toeplitz solver, *Linear Algebra Applicat.* 121 (1989) 185–206.
- [32] N. Levinson, The Wiener (root mean square) error criterion in filter design and prediction, *Stud. Appl. Math.* 25 (1-4) (1946) 261–278.
- 625 [33] J. Durbin, The fitting of time-series models, *Revue de l’Institut Int. de Statistique* (1960) 233–244.
- [34] T. L. Hansen, B. H. Fleury, B. D. Rao, Superfast line spectral estimation, in press, *IEEE Trans. Signal Process.* (2018).
- [35] B. R. Musicus, Fast MLM power spectrum estimation from uniformly spaced correlations
630 33 (5) (1985) 1333–1335.
- [36] Y. Genin, Y. Hachez, Y. Nesterov, P. Van Dooren, Optimization problems over positive pseudopolynomial matrices, *SIAM J. Matrix Anal. Applicat.* 25 (1) (2003) 57–79.
- [37] B. Alkire, L. Vandenberghe, Convex optimization problems involving finite autocorrelation sequences, *Math. Prog.* 93 (3) (2002) 331–359.
- 635 [38] A. Nemirovski, Lecture notes: Interior point polynomial time methods in convex programming, Georgia Institute of Technology (2004).
- [39] D. Malioutov, M. Cetin, A. S. Willsky, A sparse signal reconstruction perspective for source localization with sensor arrays 53 (2005) 3010–3022.
- [40] B. Ottersten, M. Viberg, T. Kailath, Analysis of subspace fitting and ML techniques for
640 parameter estimation from sensor array data 40 (1992) 590–600.
- [41] R. Carriere, R. L. Moses, High resolution radar target modeling using a modified Prony estimator 40 (1992) 13–18.
- [42] W. Bajwa, A. Sayeed, R. Nowak, Compressed channel sensing: A new approach to estimating sparse multipath channels 98 (2010) 1058–1076.
- 645 [43] X. Andrade, J. N. Sanders, A. Aspuru-Guzik, Application of compressed sensing to the simulation of atomic systems, *Proc. Nat. Academy Sciences* 109 (35) (2012) 13928–13933.
- [44] R. Tibshirani, Regression shrinkage and selection via the lasso, *J. R. Stat. Soc., Ser. B* 58 (1994) 267–288.
- [45] M. G. Krein, A. A. Nudelman, The Markov moment problem and extremal problems, Vol. 50
650 of *Translations of Mathematical Monographs*, American Mat. Soc., 1977.

- [46] A. Nemirovski, M. J. Todd, Interior-point methods for optimization, *Acta Numerica* 17 (2008) 191–234.
- [47] J. Nocedal, Updating quasi-Newton matrices with limited storage, *Math. Computation* 35 (151) (1980) 773–782.
- 655 [48] I. Gohberg, I. A. Feldman, Convolution equations and projection methods for their solution, Vol. 41 of *Translations of Mathematical Monographs*, American Mat. Soc., 2005.
- [49] W. F. Trench, An algorithm for the inversion of finite Toeplitz matrices, *J. Soc. Ind. Appl. Math.* 12 (3) (1964) 515–522.
- [50] G. H. Golub, C. F. Van Loan, *Matrix computations*, Vol. 3, JHU Press, 2012.
- 660 [51] M. Grant, S. Boyd, CVX: Matlab software for disciplined convex programming, <http://cvxr.com/cvx/> (Jan. 2018).
- [52] H. W. Kuhn, The Hungarian method for the assignment problem, *Nav. Res. Logist. Q.* 2 (1-2) (1955) 83–97.
- [53] Y. Chi, Y. Chen, Compressive two-dimensional harmonic retrieval via atomic norm mini-
665 mization 63 (4) (2015) 1030–1042.
- [54] Z. Yang, L. Xie, P. Stoica, Vandermonde decomposition of multilevel Toeplitz matrices with application to multidimensional super-resolution 62 (6) (2016) 3685–3701.