COMPUTATIONALLY EFFICIENT REGULARIZED ACOUSTIC IMAGING

Flávio P. Ribeiro, Vítor H. Nascimento

Electronic Systems Engineering Dept., Universidade de São Paulo
{fr,vitor}@lps.usp.br

ABSTRACT

Sparse recovery techniques have been shown to produce very accurate acoustic images, significantly outperforming traditional deconvolution approaches. However, so far these proposals have been computationally intractable for all but very small images, because they had no means of efficiently transforming back and forth between a hypothetical image under reconstruction and the measured data. In this paper we obtain a fast transform for planar array geometries using the fast non-equispaced Fourier transform (NFFT). We then apply it to accelerate general-purpose solvers by several orders of magnitude, enabling computationally-efficient regularized acoustic imaging. The proposed approach is not only tractable, but faster than competing deconvolution techniques, while delivering reconstructions with unprecedented accuracy.

Index Terms— array processing, microphone arrays, acoustic imaging, sparse reconstruction, regularized least-squares, fast transform, NFFT, deconvolution, covariance fitting.

1. INTRODUCTION

Acoustic imaging is a standard technique for mapping the location and intensity of sound sources with microphone arrays. It can provide insight into noise generating mechanisms, so it is routinely used to design quieter vehicles, machinery and aircraft [1]. It has also been used to analyze the acoustic properties of large rooms, by allowing one to visualize their late reverberation [2].

Due to its relatively low computational cost, beamforming remains the most popular method for acoustic imaging. Unfortunately, it produces the source distribution of interest convolved with the array point spread function (PSF). Since microphone arrays have at most a few hundred microphones, array PSFs are inevitably large, such that beamforming produces smeared images.

Deconvolution algorithms [3, 4] have been proposed to enhance images obtained with beamforming. More recently, regularized covariance fitting techniques [5] have been shown to deliver even better results. However, these methods have high computational costs, mainly because they have no way to efficiently transform back and forth between the image under reconstruction and the measured data. In the absence of fast transforms, covariance fitting is only tractable for low resolution images.

In this paper we show how to efficiently apply state-of-the-art iterative methods for convex optimization and sparse recovery to obtain high resolution acoustic images. In particular, we are interested in solving regularized least-squares problems of the form

$$\hat{\mathbf{y}} = \arg\min_{\mathbf{y}} \|\mathbf{\Psi}\mathbf{y}\| + \mu \|\mathbf{A}\mathbf{y} - \mathbf{s}\|_2^2 , \tag{1}$$

where $\mathbf{s}$ is the measured signal, $\hat{\mathbf{y}}$ is the reconstructed signal, $\mathbf{\Psi}$ is a sparsifying transform, and $\mathbf{A}$ is a transform which models the measurement process. For an acoustic image, $\mathbf{y}$ would be a vectorized version of the image describing the true source distribution, and $\mathbf{s}$ would be a vectorized version of the sample covariance matrix.

Since (1) is solved iteratively, one must be able to quickly and repeatedly evaluate $\mathbf{\Psi}$, $\mathbf{\Psi}^H$, $\mathbf{A}$ and $\mathbf{A}^H$ for arbitrary vectors. Indeed, the application of these transforms dominates the computational cost of solving (1). While one can usually choose a convenient fast sparsifying transform $\mathbf{\Psi}$, the measurement process determines $\mathbf{A}$. In practice, algorithms for solving (1) are strongly bottlenecked by the cost of evaluating $\mathbf{A}$, such that accelerating $\mathbf{A}$ by a given factor accelerates imaging by nearly the same factor.

While sparsity-enforcing approaches have been proposed for direction of arrival estimation [6] and acoustic imaging [5], to our knowledge these proposals have never been combined with a fast implementation of $\mathbf{A}$. Thus, previous formulations were only computationally tractable if $\mathbf{y}$ had a small number of dimensions, limiting their practical applications.

This paper offers two contributions: (i) we show how to use the non-equispaced fast Fourier transform (NFFT) [7] to obtain a fast transform for aeroacoustic imaging which outperforms explicit matrix multiplication by more than 100x, and (ii) propose a new approach for acoustic imaging by combining this transform with regularized least-squares solvers, making covariance fitting practical for large scale problems. Since all the domain-specific characteristics of array imaging are abstracted away in $\mathbf{A}$, this approach has the benefit of replacing ad-hoc techniques with state-of-the-art general purpose solvers. Furthermore, better regularization allow us to reconstruct more accurate representations than with previous methods.

The proposed transform has the advantage of modeling arbitrary planar array geometries. Thus, it allows one to perform imaging with random geometries, which provide better spatial diversity, resolution, operating bandwidth and reconstruction guarantees than geometries containing regularities.

This paper is organized as follows: Section 2 provides definitions and further motivates the need for fast transforms. Section 3 shows how to obtain a fast transform using the NFFT. Section 4 has applications, and Section 5 has our conclusions.

2. PRELIMINARIES

We use the following notation. Transposes and Hermitian transposes of vectors and matrices are denoted by $\cdot^T$ and $\cdot^H$, respectively. For $a, b \in \mathbb{Z}_+$, $\operatorname{mod}(a, b)$ denotes the remainder of a/b. $\lfloor x \rfloor$ represents rounding of $x \in \mathbb{R}$ towards $-\infty$.

Consider a planar array of N microphones lying on the horizontal plane at coordinates $\mathbf{p}_0, ..., \mathbf{p}_{N-1}$, such that $\mathbf{p}_i = \begin{bmatrix} p_{x_i} & p_{y_i} & 0 \end{bmatrix}^T$. Assume we are interested in the sound power coming from a set of M look directions $\{\mathbf{u}_i\}_{i=0}^{M-1}$, where in general M is a large number. The $N \times 1$ array output for a frequency ω is modeled as

$$\mathbf{x}(\omega) = \mathbf{V}(\omega)\,\mathbf{f}(\omega) + \boldsymbol{\eta}(\omega) , \tag{2}$$

where $\mathbf{V}(\omega) = \left[\, \mathbf{v}(\mathbf{u}_0, \omega) \;\; \mathbf{v}(\mathbf{u}_1, \omega) \;\; \cdots \;\; \mathbf{v}(\mathbf{u}_{M-1}, \omega) \,\right]$ is the array manifold matrix, $\mathbf{f}(\omega) = \left[\, f_0(\omega) \;\; f_1(\omega) \;\; \cdots \;\; f_{M-1}(\omega) \,\right]^T$ is the frequency domain signal waveform and $\boldsymbol{\eta}(\omega)$ is uncorrelated noise. The far-field array manifold vector for source m is given by

$$\mathbf{v}(\mathbf{u}_m, \omega) = \left[\, \mathrm{e}^{j\frac{\omega}{c}\mathbf{u}_m^T\mathbf{p}_0} \;\; \cdots \;\; \mathrm{e}^{j\frac{\omega}{c}\mathbf{u}_m^T\mathbf{p}_{N-1}} \,\right]^T. \tag{3}$$

If θ and ϕ are the azimuth and elevation angles, one can parameterize the look directions as

$$u_x(\theta, \phi) = \sin\phi\cos\theta, \qquad u_y(\theta, \phi) = \sin\phi\sin\theta,$$

such that

$$\mathbf{u} = \left[\, u_x \;\; u_y \;\; \sqrt{1 - u_x^2 - u_y^2} \,\right]^T \tag{4}$$

for $u_x^2 + u_y^2 \leq 1$. It is well-known that under a far-field approximation, uniform sampling in U-space (where $\mathrm{U} = [-1, 1]^2$) makes point-spread functions shift-invariant. Since in (3) the z coordinate of $\mathbf{u}_m$ is multiplied by 0, we can simplify notation by redefining $\mathbf{u}_m, \mathbf{p}_n$ as belonging to $\mathbb{R}^2$, using only the first two coordinates of each vector.

Let $\mathbf{S}_\mathbf{x}(\omega) = \mathrm{E}\left\{\mathbf{x}(\omega)\,\mathbf{x}^H(\omega)\right\}$ be the array cross spectral matrix. If $\mathbf{x}_0(\omega)$, ..., $\mathbf{x}_{L-1}(\omega)$ correspond to L frequency domain snapshots, the spectral matrix can be estimated with

$$\mathbf{S}_\mathbf{x}(\omega_k) = \frac{1}{L}\sum_{l=0}^{L-1} \mathbf{x}_l(\omega_k)\,\mathbf{x}_l^H(\omega_k), \tag{5}$$

which retains only the relative phase shifts between microphones. Each frequency is processed separately, so to simplify notation we omit the argument ω and write $\mathbf{S}_\mathbf{x}(\omega)$ simply as $\mathbf{S}$.

Assuming that the noise is spatially white and uncorrelated with the sources of interest, we have

$$\mathbf{S} = \mathbf{V}\,\mathrm{E}\left\{\mathbf{f}\mathbf{f}^H\right\}\mathbf{V}^H + \sigma^2\mathbf{I}, \tag{6}$$

where $\sigma^2 = \mathrm{E}\{\eta_i\eta_i^*\}, 0 \leq i < N$.

The sound power emanating from look directions $\{\mathbf{u}_i\}_{i=0}^{M-1}$ distributed along a grid in U-space defines a two-dimensional image, where pixel coordinates correspond to source locations, and pixel values correspond to source powers. A usual assumption in acoustic imaging is that the pixels are uncorrelated, i.e., that $\mathrm{E}\left\{\mathbf{f}\mathbf{f}^H\right\}$ is diagonal (note that the diagonal of $\mathrm{E}\left\{\mathbf{f}\mathbf{f}^H\right\}$ has the power radiated from each direction, i.e., a vectorized version of the image).

Given a $M_x \times M_y$ acoustic image, define $M = M_xM_y$ and let $\{\mathbf{u}_m\}_{m=0}^{M-1}$ be an enumeration of all pixel coordinates in U-space. Let $\mathbf{v}(\mathbf{u}_m)$ be the array manifold vector when steered towards $\mathbf{u}_m$. If the sound powers radiating from directions $\{\mathbf{u}_m\}_{m=0}^{M-1}$ are given by $\left\{|Y(\mathbf{u}_m)|^2\right\}_{m=0}^{M-1}$, in the absence of noise we have

$$\mathbf{S} = \sum_{m=0}^{M-1} |Y(\mathbf{u}_m)|^2\,\mathbf{v}(\mathbf{u}_m)\,\mathbf{v}^H(\mathbf{u}_m). \tag{7}$$

Covariance fitting algorithms iteratively compute a reconstructed image $\left\{\left|\hat{Y}(\mathbf{u}_m)\right|^2\right\}_{m=0}^{M-1}$ and compare the corresponding $\hat{\mathbf{S}}$ obtained through (7) with the measured values obtained from (5). Unless the image is very sparse, (7) becomes computationally intractable. For instance, for a 256 element array and a 256×256 acoustic image, (7) requires 2^{32} complex multiply-accumulate instructions to compute. This cost is excessive for a transform intended to be used in an iterative method. In the following section we describe how to obtain an efficient transform to compute (7) using the NFFT.

3. FAST TRANSFORMS WITH THE NFFT

Define $\mathbf{y} = \left[\, |Y(\mathbf{u}_0)|^2 \;\; \cdots \;\; |Y(\mathbf{u}_{M-1})|^2 \,\right]^T$. Note that we can rewrite (7) as $\mathbf{s} = \mathbf{A}\mathbf{y}$, where $\mathbf{A}$ is a linear transform and $\mathbf{s} = \mathrm{vec}\{\mathbf{S}\}$. Indeed, $\mathrm{vec}\left\{\mathbf{v}_{\mathbf{u}_m}\mathbf{v}_{\mathbf{u}_m}^H\right\} = \mathbf{v}_{\mathbf{u}_m}^* \otimes \mathbf{v}_{\mathbf{u}_m}$, where $\otimes$ is the Kronecker product, and

$$\mathbf{s} = \underbrace{\left[\, \mathbf{v}_{\mathbf{u}_0}^* \otimes \mathbf{v}_{\mathbf{u}_0} \;\; \cdots \;\; \mathbf{v}_{\mathbf{u}_{M-1}}^* \otimes \mathbf{v}_{\mathbf{u}_{M-1}} \,\right]}_{\mathbf{A}}\mathbf{y}. \tag{8}$$

In this section we describe how to use the NFFT to implement $\mathbf{A}$.

Let $\mathbf{Y} \in \mathbb{C}^{M_y \times M_x}$ be an image obtained by a uniform rectangular sampling of U-space, for even M_x, M_y and with sampling coordinates drawn from

$$\tilde{\mathrm{U}} = \left\{\frac{2i}{M_x}\right\}_{i=-M_x/2}^{M_x/2-1} \times \left\{\frac{2j}{M_y}\right\}_{j=-M_y/2}^{M_y/2-1}. \tag{9}$$

We enumerate $\mathbf{u}_0, ..., \mathbf{u}_{M-1} \in \tilde{\mathrm{U}}$ such that $\mathbf{y} = \mathrm{vec}\{\mathbf{Y}\}$.

One can show that for uniform rectangular arrays with horizontal and vertical inter-element spacings $d_x = d_y = \lambda/2$ (where λ is the wavelength of the signal of interest), $\mathbf{A}$ can be implemented with a 2D FFT. This implementation is not convenient for aeroacoustic imaging, since: (i) the constraint $d_x = d_y = \lambda/2$ can only be satisfied for one frequency, and we are interested in wideband operation; (ii) the 2D FFT is inefficient, since it ignores that image pixels significantly outnumber array sensors, and determines covariances for sensors that do not exist; (iii) the 2D FFT requires uniform rectangular geometries, severely limiting the spatial resolution of the reconstructed images.

A d-dimensional NDFT is defined by a set of arbitrary spatial nodes $\mathcal{X}$ and a frequency bandwidth vector $\mathbf{M} \in \mathbb{N}^d$. Each node x_j belongs to the sampling set $\mathcal{X} = \left\{x_i \in \left[-\frac{1}{2}, \frac{1}{2}\right)^d : 0 \leq i < N\right\}$. The index set $I_\mathbf{N} = \mathbb{Z}^d \cap \prod_{t=0}^{d-1}\left[-\frac{M_t}{2}, \frac{M_t}{2}\right)$ defines a rectangular grid over which a function of interest is sampled.

Given as input a set of samples $h_\mathbf{k} \in \mathbb{C}$ for $\mathbf{k} \in I_\mathbf{N}$, the NDFT is defined as

$$\hat{h}_i = \sum_{\mathbf{k}\in I_\mathbf{N}} h_\mathbf{k} e^{-j2\pi\mathbf{k}^T\mathbf{x}_i}, \tag{10}$$

for $0 \leq i < N$. The NFFT is a fast approximate implementation of the NDFT which interpolates an oversampled FFT, and obtains a very good compromise between accuracy and computational complexity.

To obtain $\mathbf{A}$ using the NFFT, rectangular U-space sampling and an arbitrary geometry of N microphones we use

$$\mathbf{M} = \left[\, M_x \;\; M_y \,\right]^T \tag{11}$$

$$I_\mathbf{N} = \mathbb{Z}^2 \cap \left[-\frac{M_x}{2}, \frac{M_x}{2}\right) \times \left[-\frac{M_y}{2}, \frac{M_y}{2}\right) \tag{12}$$

$$\mathcal{X} = \left\{\mathbf{x}_i = \frac{2}{\lambda}\left(\mathbf{p}_{\lfloor i/N \rfloor} - \mathbf{p}_{\mathrm{mod}(i,N)}\right) \odot \begin{bmatrix} M_x^{-1} \\ M_y^{-1} \end{bmatrix} : i \in \left[0, N^2\right)\right\} \tag{13}$$

where $\odot$ represents the pointwise (Hadamard) product and the baselines $\mathbf{p}_r - \mathbf{p}_s$ are represented only by their x and y coordinates.

We now show that this parameterization of the NFFT produces $\mathbf{A}$. Due to the linearity of the NFFT, it suffices to show that this parameterization produces $\mathbf{A}$ for an image containing a unit impulse at arbitrary coordinates $\bar{\mathbf{u}} \in \tilde{\mathrm{U}}$. For arbitrary $-\frac{M_x}{2} \leq m_0 < \frac{M_x}{2}$ and $-\frac{M_y}{2} \leq n_0 < \frac{M_y}{2}$, let $\bar{\mathbf{u}} = \begin{bmatrix} \frac{2m_0}{M_x} & \frac{2n_0}{M_y} \end{bmatrix}^T$ and

$$\mathbf{Y}_{n,m} = \begin{cases} 1 & \text{if } (m,n) = \left(m_0 + \frac{M_x}{2}, n_0 + \frac{M_y}{2}\right) \\ 0 & \text{otherwise.} \end{cases}$$

Let $\mathrm{vec}\{\mathbf{S}\} = \mathbf{A}\mathrm{vec}\{\mathbf{Y}\}$. It follows from (7) that

$$\begin{aligned} \mathbf{S}_{r,s} &= \left[\mathbf{v}(\bar{\mathbf{u}})\,\mathbf{v}^H(\bar{\mathbf{u}})\right]_{r,s} \\ &= \mathrm{e}^{j\frac{\omega}{c}\mathbf{u}_0^T(\mathbf{p}_r - \mathbf{p}_s)} = \mathrm{e}^{j2\pi\frac{\mathbf{u}_0^T}{2}\frac{2}{\lambda}(\mathbf{p}_r - \mathbf{p}_s)} \\ &= \mathrm{e}^{j2\pi\left(\mathbf{u}_0^T \odot \left[\begin{array}{cc} \frac{M_x}{2} & \frac{M_y}{2} \end{array}\right]\right)\left(\frac{2}{\lambda}(\mathbf{p}_r - \mathbf{p}_s) \odot \left[\begin{array}{cc} M_x^{-1} & M_y^{-1} \end{array}\right]^T\right)}. \end{aligned} \tag{14}$$

Comparing (14) with (10), the first term in parentheses clearly belongs to $I_{\mathbf{N}}$. Since for $0 \leq i < N^2$, $\left(\mathbf{p}_{\lfloor i/N \rfloor} - \mathbf{p}_{\mathrm{mod}(i,N)}\right)$ spans all possible baselines, the second term in parentheses belongs to $\mathcal{X}$. The enumeration given by $(\lfloor i/N \rfloor, \mathrm{mod}\,(i,N))$ indexes the elements of $\mathbf{S}$ row by row. Given the Hermitian symmetry of $\mathbf{S}$, this is equivalent to conjugating (14) and indexing the elements of $\mathbf{S}$ column by column (in the order of $\mathrm{vec}\{\cdot\}$), making (14) equivalent to (10).

Thus, with (11)-(13) one can implement $\mathbf{A}$ using the NFFT. The fast implementation of $\mathbf{A}^H$ is provided automatically by the NFFT implementation, allowing us to proceed with applications.

4. APPLICATIONS

Due to limited space, we only illustrate two approaches for regularized imaging, which do not exhaust the range of possible applications. For other applications, we refer the reader to [8].

4.1. ℓ_1-regularized least-squares imaging

If only a small number of U-space points have radiating sources, then one should enforce sparsity in the canonical basis, as proposed in [5] using an objective function similar to LASSO [9]. Here we use a basis pursuit with denoising (BPDN) formulation, given by

$$\min_{\hat{\mathbf{Y}}} \left\|\hat{\mathbf{Y}}\right\|_1 \text{ subject to } \left\|\mathrm{vec}\{\mathbf{S}\} - \mathbf{A}\mathrm{vec}\{\hat{\mathbf{Y}}\}\right\|_2 \leq \sigma, \tag{15}$$

which has been studied in detail in the compressive sensing literature. Minimizing $\left\|\hat{\mathbf{Y}}\right\|_1$ yields a sparse approximation to a solution of the underdetermined system given by $\mathrm{vec}\{\mathbf{S}\} = \mathbf{A}\mathrm{vec}\{\hat{\mathbf{Y}}\}$. Furthermore, (15) is a convex optimization problem which can be solved iteratively with good convergence properties.

In the examples, we solve (15) with SPGL1 [9], which is a state-of-the-art solver designed for large scale ℓ_1-regularized least-squares problems. The fast transform not only makes this solution tractable, but faster than deconvolution methods such as the FFT-accelerated DAMAS2 [4].

4.2. Total variation regularized least-squares imaging

To address cases where the source distribution cannot be assumed to be sparse in its canonical representation, we propose reconstructing acoustic images with total variation (TV) regularization.

Given $\mathbf{Y} \in \mathbb{C}^{M_y \times M_x}$, define its isotropic total variation as $\|\mathbf{Y}\|_{BV} = \sum_{i,j} \sqrt{[\nabla_x \mathbf{Y}]_{i,j}^2 + [\nabla_y \mathbf{Y}]_{i,j}^2}$, where ∇_x and ∇_y are the first difference operators along the x and y dimensions with periodic boundaries, for $0 \leq i < M_y$ and $0 \leq j < M_x$. $\|\cdot\|_{BV}$ is called the bounded variation (BV) semi-norm.

We propose solving

$$\min_{\hat{\mathbf{Y}}} \left\|\hat{\mathbf{Y}}\right\|_{BV} + \mu \left\|\mathrm{vec}\{\mathbf{S}\} - \mathbf{A}\mathrm{vec}\{\hat{\mathbf{Y}}\}\right\|_2^2, \tag{16}$$

subject to $\hat{\mathbf{Y}}_{i,j} \geq 0$. The first term measures how much an image oscillates, and privileges solutions with small amounts of noise. The second term ensures a good fit between the reconstructed image and the measured data.

To solve (16) we use TVAL3 [10], which is a state-of-the-art solver for TV-regularized least-squares problems. TVAL3 compares very favorably to other solvers in terms of convergence rate and reconstruction quality. With the fast transform it becomes significantly faster than DAMAS2, while providing more accurate and stable reconstructions with guaranteed convergence.

4.3. Examples

In this section we present results comparing delay-and-sum beamforming, DAMAS2 implemented with the FFT2, ℓ_1-regularized reconstruction with SPGL1 [9] solving (15), and TV-regularized reconstruction with TVAL3 [10] solving (16).

To sample a wave field with maximum diversity, we use a 64-element random array with a 50×50 cm aperture. This geometry was obtained by positioning array elements uniformly at random inside a unit square, which was then rescaled to put the outermost elements on the edge of a 50×50 cm square.

All algorithms were implemented in MATLAB R2008b, and the NFFT was compiled with default optimizations. All images were reconstructed with $M_x = M_y = 256$. DAMAS2, SPGL1 and TVAL3 used 1000, 200 and 100 iterations, respectively, which provide a good compromise between computational cost and image quality.

Reconstruction times for DAMAS2, SGPL1 and TVAL3 were approximately 40s, 15s and 8s per image, respectively, on an Intel Core 2 Duo T9400 processor using only one core. In contrast, SPGL1 and TVAL3 without a fast transform would be more than 100 times slower, requiring approximately 2350s and 1250s per image, respectively.

We used $\sigma = 0.01\,\|\mathbf{S}\|_F$ in (15) for SPGL1 and $\mu = 10^3$ in (16) for TVAL3. The signal model is given by (6), with σ^2 set to obtain a 20 dB SNR. Since the intent of these simulations is not to analyze the noise sensitivity of each method, only one SNR is used. Due to limited space, we only present results for 2, 5 and 8 kHz.

Fig. 1 presents reconstruction results for images with unit impulses at U-space coordinates $(\pm n/6, \pm n/6)$ for $0 \leq n \leq 4$. As expected, delay-and-sum has the lowest spatial resolution and shows significant sidelobes. DAMAS2 shows good results and a few small artifacts, but due to the absence of regularization, it spreads the point sources. ℓ_1-regularized reconstruction presents significantly more accurate results. This is not a surprise, since the source distribution is indeed very sparse.

Fig. 2 show reconstruction results for a non-sparse test pattern designed for this experiment. Once again, delay-and-sum produces very obvious smearing, as evidenced by the grey background that spreads outside the visible region. DAMAS2 produces much better results, but the plateaus are noisy due to the lack of regularization.

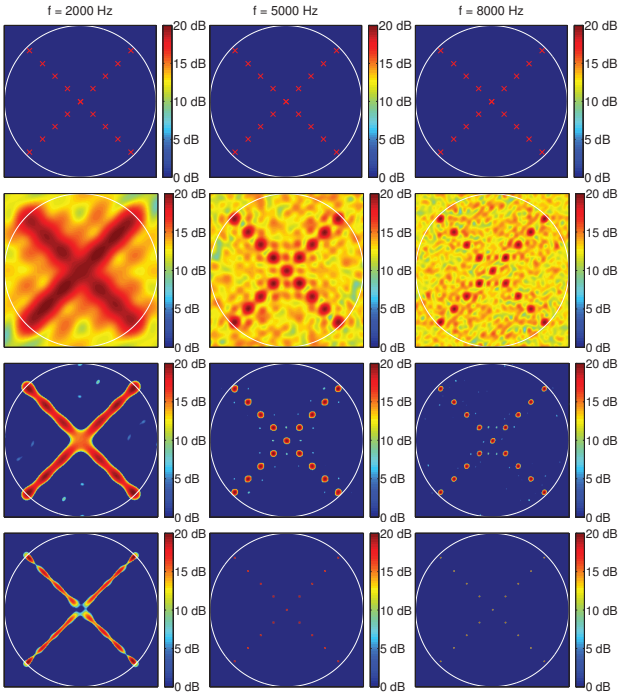

Fig. 1. Reconstructions of an impulsive pattern, for $M_x = M_y = 256$. Line 1: ideal; line 2: delay-and-sum; line 3: DAMAS2; line 4: ℓ_1-regularized least-squares. The outer circle represents the array's visible region.

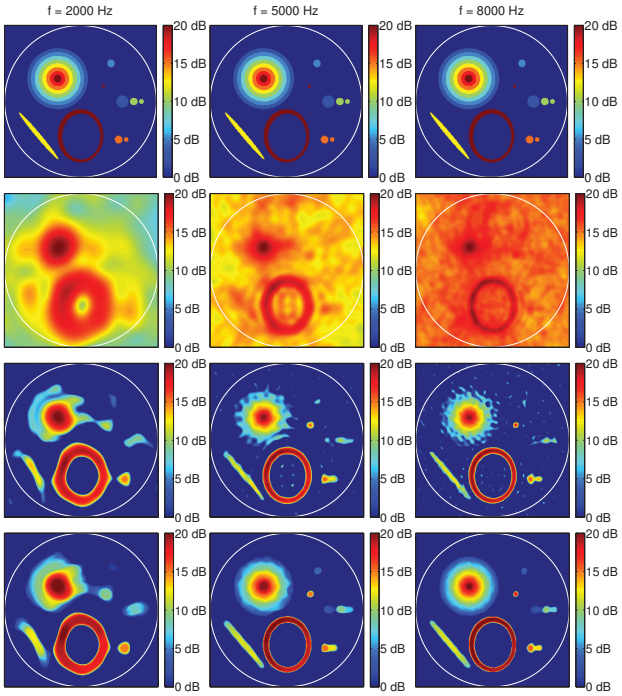

Fig. 2. Reconstructions of a non-sparse test pattern, for $M_x = M_y = 256$. Line 1: ideal; line 2: delay-and-sum; line 3: DAMAS2; line 4: TV-regularized least-squares. The outer circle represents the array's visible region.

Finally, TV-regularized reconstruction produces the most accurate representations, with accurate shapes and low noise.

5. CONCLUSIONS

This paper describes a method for computationally efficient and accurate acoustic imaging. We combine covariance fitting and sparse recovery by recasting acoustic imaging as ℓ_1- or TV-regularized least-squares problems. While these formulations would be ordinarily intractable for all but small problems (and thus of little practical interest), we implement them with a fast transform, which is the key to this approach. Indeed, this transform accelerates solvers by two orders of magnitude (in our examples, by more than 100x).

Given the good convergence properties and numerical stability of the regularized least-squares methods, our proposal is not only tractable, but significantly faster than competing deconvolution methods. Furthermore, by using regularization, we can incorporate a priori information about the source distribution. Thus, the proposed methods also deliver more accurate reconstructions than deconvolution algorithms.

We note that fast transforms can also naturally implement and accelerate other aeroacoustic imaging methods, including beamforming and DAMAS2. Furthermore, other fast transforms can be designed to better fit certain applications (for example, by obtaining optimizations for constrained geometries). Thus, the $\mathbf{A}$ transform offers a useful language for describing array processing methods required to deliver many source estimates at once. By recasting algorithms in terms of $\mathbf{A}$ and $\mathbf{A}^H$, one can naturally eliminate redundancies and obtain optimizations. For more details and other applications, we refer the reader to [8].

6. REFERENCES

[1] W.C. Home, K.D. James, T.K. Arledge, P.T. Sodermant, N. Burnside, and S.M. Jaeger, "Measurements of 26%-scale 777 Airframe Noise in the NASA Ames 40- by 80-Foot Wind Tunnel," in *Proc. of the 11th AIAA/CEAS Aeroacoustics Conference*, 2005.

[2] A. O'Donovan, R. Duraiswami, and D. Zotkin, "Imaging concert hall acoustics using visual and audio cameras," in *Proc. of ICASSP*, 2008, pp. 5284–5287.

[3] Y. Wang, J. Li, P. Stoica, M. Sheplak, and T. Nishida, "Wideband RELAX and wideband CLEAN for aeroacoustic imaging," *The Journal of the Acoustical Society of America*, vol. 115, pp. 757, 2004.

[4] R.P. Dougherty, "Extensions of DAMAS and Benefits and Limitations of Deconvolution in Beamforming," in *Proc. of the 11th AIAA/CEAS Aeroacoustics Conference*, 2005.

[5] T. Yardibi, J. Li, P. Stoica, and L.N. Cattafesta III, "Sparsity constrained deconvolution approaches for acoustic source mapping," *The Journal of the Acoustical Society of America*, vol. 123, pp. 2631, 2008.

[6] D. Malioutov, M. Cetin, and A.S. Willsky, "A sparse signal reconstruction perspective for source localization with sensor arrays," *IEEE Trans. Signal Process.*, vol. 53, no. 8, pp. 3010–3022, 2005.

[7] J. Keiner, S. Kunis, and D. Potts, "Using NFFT 3—A Software Library for Various Nonequispaced Fast Fourier Transforms," *ACM Transactions on Mathematical Software (TOMS)*, vol. 36, no. 4, pp. 19, 2009.

[8] F.P. Ribeiro and V.H. Nascimento, "A fast transform for acoustic imaging with separable arrays," in *Proc. of ICASSP*, 2011.

[9] E. van den Berg and M.P. Friedlander, "Probing the Pareto frontier for basis pursuit solutions," *SIAM Journal on Scientific Computing*, vol. 31, no. 2, pp. 890–912, 2008.

[10] C. Li, "An efficient algorithm for total variation regularization with applications to the single pixel camera and compressive sensing," M.S. thesis, Rice University, 2009.