

Fast Transforms for Acoustic Imaging — Part II: Applications

Flávio P. Ribeiro*, *Student Member, IEEE*, and Vítor H. Nascimento, *Member, IEEE*

Abstract—In Part I [“Fast Transforms for Acoustic Imaging—Part I: Theory,” *IEEE TRANSACTIONS ON IMAGE PROCESSING*, vol. XX, no. YY, Month 201x], we introduced the KAT (Kronecker Array Transform), a fast transform for imaging with separable arrays. Given a source distribution, the KAT produces the spectral matrix which would be measured by a separable sensor array. In this Part II, we establish connections between the KAT, beamforming and 2D convolutions, and show how these results can be used to accelerate classical and state of the art array imaging algorithms. We also propose using the KAT to accelerate general purpose regularized least-squares solvers. Using this approach, we avoid ill-conditioned deconvolution steps and obtain more accurate reconstructions than previously possible, while maintaining low computational costs. We also show how the KAT performs when imaging near-field source distributions, and illustrate the trade-off between accuracy and computational complexity. Finally, we show that separable designs can deliver accuracy competitive with multi-arm logarithmic spiral geometries, while having the computational advantages of the KAT.

Index Terms—array processing, fast transform, acoustic imaging, array imaging, regularized least-squares, sparse reconstruction.

I. INTRODUCTION

As described in Part I [1], array imaging requires solving the inverse problem of finding the best estimate for a source distribution, given wavefield statistics sampled by a sensor array. This is not a trivial problem, since in general one must rely on arrays with less than one hundred elements to reconstruct source distributions modeled with tens of thousands of point sources. To obtain accurate reconstructions, regularization is required to narrow the space of possible wavefields which result in essentially the same data at the sensors.

Let $\mathbf{S} \in \mathbb{C}^{N \times N}$ be a narrowband sample covariance matrix acquired using a planar sensor array with a separable geometry. Let $\mathbf{Y} \in \mathbb{R}^{M_y \times M_x}$ be a discretization of the source distribution at the same frequency. Assume for the sake of this argument that the true source distribution is represented exactly by \mathbf{Y} , and that the sources are uncorrelated. If $\mathbf{s} = \text{vec}\{\mathbf{S}\}$, $\mathbf{y} = \text{vec}\{\mathbf{Y}\}$ and \mathbf{A} is the KAT presented in Part I, in the absence of noise we have that $\mathbf{A}\mathbf{y} = \mathbf{s}$. The generic image reconstruction problem then becomes

$$\hat{\mathbf{y}} = \arg \min_{\mathbf{y}} \|\Psi\mathbf{y}\| \quad \text{such that } \mathbf{A}\mathbf{y} = \mathbf{s}, \quad (1)$$

where Ψ is a sparsifying transform for \mathbf{y} . For example, if \mathbf{Y} is known to be sparse in its canonical representation, then one could consider minimizing $\|\Psi\mathbf{y}\| = \|\mathbf{y}\|_1$, which turns (1) into an instance of basis pursuit [2].

In the presence of noise, the constraint $\mathbf{A}\mathbf{y} = \mathbf{s}$ no longer applies, motivating the formulation

$$\hat{\mathbf{y}} = \arg \min_{\mathbf{x}} \|\Psi\mathbf{y}\| + \mu \|\mathbf{A}\mathbf{y} - \mathbf{s}\|_2^2, \quad (2)$$

which is a regularized least-squares problem.

The problem of seeking sparse approximations to underdetermined systems has received significant attention in the recent years with the advent of compressive sensing [3]–[5]. Recently, many exact and approximate methods have been proposed for solving variations of (1) for specific instances of Ψ and $\|\cdot\|$, such as [6]–[11].

The computational bottleneck for solving (1) or (2) with efficient convex optimization methods lies exclusively in the implementations of \mathbf{A} , \mathbf{A}^H , Ψ or Ψ^H . For imaging applications, Ψ can be a fast wavelet transform, fast Fourier transform or a finite difference operator, which can all be evaluated quickly. Therefore, the potential bottleneck lies in the implementations of \mathbf{A} and \mathbf{A}^H . However, the KAT makes \mathbf{A} and \mathbf{A}^H orders of magnitude faster than competing transforms (and in particular, much faster than explicit matrix representations), allowing the use of regularized least-squares methods for acoustic imaging.

This Part II describes applications of the KAT for image reconstruction. In Section II we present several methods for acoustic imaging using a common language based on the transform. Using the KAT, we accelerate these techniques without compromising quality. We also propose applying the fast transforms to state-of-the-art, general purpose regularized least-squares solvers, and obtain more accurate reconstructions than what was possible with previous methods. Section III features examples, comparing the performance of the different approaches. Section IV compares the reconstruction accuracy using a separable array and a logarithmic spiral array. We show that by using regularized reconstruction methods, separable arrays can match logarithmic spiral arrays in terms of reconstruction accuracy, while allowing the computational benefits provided by the KAT. Finally, Section V has our conclusions and final comments.

II. IMAGE RECONSTRUCTION APPLICATIONS

A. Delay and sum imaging

Given a spectral matrix \mathbf{S} , its corresponding image is traditionally approximated using delay-and-sum beamforming with

$$|Y(u_{x_m}, u_{y_n})|^2 \approx \frac{\mathbf{v}^H(u_{x_m}, u_{y_n}) \mathbf{S} \mathbf{v}(u_{x_m}, u_{y_n})}{[\mathbf{v}^H(u_{x_m}, u_{y_n}) \mathbf{v}(u_{x_m}, u_{y_n})]^2}, \quad (3)$$

where the approximation is due to convolution effects.

We can rewrite

$$\begin{aligned} \mathbf{v}^H(u_{x_m}, u_{y_n}) \mathbf{S} \mathbf{v}(u_{x_m}, u_{y_n}) &= \\ &= [\mathbf{v}^T(u_{x_m}, u_{y_n}) \otimes \mathbf{v}^H(u_{x_m}, u_{y_n})] \mathbf{S} \end{aligned} \quad (4)$$

$$= [\mathbf{v}^*(u_{x_m}, u_{y_n}) \otimes \mathbf{v}(u_{x_m}, u_{y_n})]^H \text{vec}\{\mathbf{S}\}, \quad (5)$$

$$= [\mathbf{A}^H \text{vec}\{\mathbf{S}\}]_{m \cdot M_y + n} \quad (6)$$

where (4) is true because $(\mathbf{A}^T \otimes \mathbf{B}) \text{vec}\{\mathbf{C}\} = \text{vec}\{\mathbf{BCA}\}$ whenever \mathbf{BCA} is defined, and (6) follows by comparing (5) with (17) of Part I. Thus, delay-and-sum imaging can be implemented with the KAT adjoint.

It follows that the direct-adjoint composition $\mathbf{A}^H \mathbf{A}$ is a transform that obtains the delay-and-sum image from a clean (ideal) image. If we assume that the sources are in the far-field and that U-space is sampled uniformly, this delay-and-sum image is simply the clean image convolved with the beamformer's point spread function (PSF). Since it represents a convolution, under these assumptions $\mathbf{A}^H \mathbf{A}$ can also be accelerated with a 2D FFT. But as presented in Section V of Part I, the KAT can always be used to implement $\mathbf{A}^H \mathbf{A}$ more efficiently than an equivalent FFT-accelerated convolution.

B. MVDR imaging

Imaging using MVDR (minimum variance distortionless response) beamforming [12] is often preferable to delay-and-sum imaging, given that the MVDR beamformer can get very fine resolution for point sources (as long as the noise is not excessive, and the regularization parameter is chosen correctly). Recall that the MVDR processor steered towards $\mathbf{v}_T = \mathbf{v}(u_{x_T}, u_{y_T})$ is given by

$$\mathbf{w}_{\text{MVDR}}^H(u_{x_T}, u_{y_T}) = \frac{\mathbf{v}_T^H \mathbf{S}_n^{-1}}{\mathbf{v}_T^H \mathbf{S}_n^{-1} \mathbf{v}_T}, \quad (7)$$

where \mathbf{S}_n is the noise spectral matrix.

To obtain \mathbf{S}_n for acoustic imaging, one should perform a separate measurement (for example, with the model removed from the wind tunnel) [13]. If this is not possible, one can obtain the MPDR (minimum power distortionless response) processor [14] by using $\mathbf{S} + \lambda \mathbf{I}$ instead of \mathbf{S}_n , where λ is a suitably chosen regularization parameter, such that

$$\mathbf{w}_{\text{MPDR}}^H(u_{x_T}, u_{y_T}) = \frac{\mathbf{v}_T^H [\mathbf{S} + \lambda \mathbf{I}]^{-1}}{\mathbf{v}_T^H [\mathbf{S} + \lambda \mathbf{I}]^{-1} \mathbf{v}_T}, \quad (8)$$

where \mathbf{S} is the spectral matrix of the whole signal, including the sources of interest and noise.

Let \mathbf{x} be the frequency domain signal at the array output, such that $\mathbf{S} = \text{E}\{\mathbf{x}\mathbf{x}^H\}$. The acoustic image can be approximated by the power at the output of the MVDR beamformer, such that for $\mathbf{w} = \mathbf{w}_{\text{MVDR}}$,

$$\begin{aligned} |Y(u_{x_T}, u_{y_T})|^2 &\approx \text{E}\left\{|\mathbf{w}^H(u_{x_T}, u_{y_T}) \mathbf{x}|^2\right\} \\ &= \mathbf{w}^H(u_{x_T}, u_{y_T}) \text{E}\{\mathbf{x}\mathbf{x}^H\} \mathbf{w}(u_{x_T}, u_{y_T}) \\ &= \frac{\mathbf{v}_T^H \mathbf{S}_n^{-1} \mathbf{S} \mathbf{S}_n^{-1} \mathbf{v}_T}{[\mathbf{v}_T^H \mathbf{S}_n^{-1} \mathbf{v}_T]^2}. \end{aligned}$$

From the results of the previous section, one can obtain $\mathbf{v}_T^H \mathbf{S}_n^{-1} \mathbf{S} \mathbf{S}_n^{-1} \mathbf{v}_T$ simultaneously for all look directions by

evaluating $\mathbf{A}^H \text{vec}\{\mathbf{S}_n^{-1} \mathbf{S} \mathbf{S}_n^{-1}\}$. Likewise, one can compute $[\mathbf{v}_T^H \mathbf{S}_n^{-1} \mathbf{v}_T]^2$ for all directions with the pointwise square of $\mathbf{A}^H \text{vec}\{\mathbf{S}_n^{-1}\}$. By dividing one by the other, one can efficiently perform imaging with an MVDR beamformer. MPDR imaging follows similarly.

C. DAMAS2

DAMAS2 [15] is a state of the art deconvolution method for aeroacoustic imaging. By using a far-field approximation, it assumes that the convolved image produced by delay-and-sum beamforming is equal to the clean image convolved with the beamformer's PSF. These convolutions are the bottleneck of the algorithm, but if uniform U-space sampling is used, they can be significantly accelerated with 2D FFTs.

Let $\hat{\mathbf{Y}}$ be the image obtained with delay-and-sum beamforming, \mathbf{P} the array PSF for delay-and-sum imaging, \mathbf{Y} the clean image and $\hat{\mathbf{Y}}^{(k)}$ the reconstructed image at iteration k . By definition, $\hat{\mathbf{Y}} = \mathbf{P} * \mathbf{Y}$, where $*$ represents 2D convolution.

DAMAS2 solves for \mathbf{Y} by iterating

$$\hat{\mathbf{Y}}^{(k+1)} = \max\left\{\hat{\mathbf{Y}}^{(k)} + \frac{1}{a} \left[\check{\mathbf{Y}} - (\mathbf{P} * \hat{\mathbf{Y}}^{(k)})\right], \mathbf{0}\right\}, \quad (9)$$

where $\max\{\cdot, \cdot\}$ returns the pointwise maximum, $a = \sum_{i,j} |\mathbf{P}|_{i,j}$, $\check{\mathbf{Y}}^{(0)} = \mathbf{0}$ and the convolution is implemented with a 2D FFT and zero-padding.

Given the fast transform, it is possible to implement a faster version of the already FFT-accelerated DAMAS2. Indeed, from Section II-A we have that $\check{\mathbf{y}} = \text{vec}\{\check{\mathbf{Y}}\} = \text{vec}\{\mathbf{P} * \mathbf{Y}\} = \mathbf{A}^H \mathbf{A} \text{vec}\{\mathbf{Y}\} = \mathbf{A}^H \mathbf{A} \mathbf{y}$. Similarly, $\text{vec}\{\mathbf{P} * \hat{\mathbf{Y}}^{(k)}\} = \mathbf{A}^H \mathbf{A} \hat{\mathbf{y}}^{(k)}$, where $\mathbf{A}^H \mathbf{A}$ can be implemented with the fast direct-adjoint KAT, described in Section III-C of Part I.

Thus, (9) becomes

$$\hat{\mathbf{y}}^{(k+1)} = \max\left\{\hat{\mathbf{y}}^{(k)} + \frac{1}{a} \left[\check{\mathbf{y}} - \mathbf{A}^H \mathbf{A} \hat{\mathbf{y}}^{(k)}\right], \mathbf{0}\right\}, \quad (10)$$

where a has the same definition as before and $\hat{\mathbf{y}}^{(k)} = \text{vec}\{\hat{\mathbf{Y}}^{(k)}\}$.

Since convolutions are the bottleneck of DAMAS2, the performance improvement of (10) with the fast transform with respect to (9) as conventionally implemented is given by the runtime of $\mathbf{A}^H \mathbf{A}$ when compared to that of an FFT accelerated convolution. By referring to Figure 4 of Part I, one can see that significant improvements can be obtained for all problem sizes. In particular, for the examples shown in Section III, the KAT is 8 times faster than an FFT accelerated convolution.

Even though DAMAS2 is considered to be a state-of-the-art method for computationally efficient acoustic imaging, it does not use any regularization other than forcing pointwise non-negativity. Thus, it does not incorporate a prior model of the source distribution. Furthermore, DAMAS2 is a deconvolution approach that relies on restoring detail from very smeared delay and sum images. We have shown that delay-and-sum imaging is equivalent to the application of $\mathbf{A}^H \mathbf{A}$ or to convolution by the array PSF, which is a low pass filter. The low-pass characteristic implies that $\mathbf{A}^H \mathbf{A}$ has many small singular values. Applying $\mathbf{A}^H \mathbf{A}$ significantly attenuates input basis vector components corresponding to these small singular

values, such that solving $\check{y} = \mathbf{A}^H \mathbf{A} \mathbf{y}$ for \mathbf{y} (as proposed by DAMAS) is not trivial. On the other hand, the singular values of \mathbf{A} are the square roots of the singular values of $\mathbf{A}^H \mathbf{A}$. Thus, the application of \mathbf{A} only attenuates input basis vector components by the square root of the previous factors, making it preferable to solve $\mathbf{s} = \mathbf{A} \mathbf{y}$ for \mathbf{y} (as proposed by least-squares formulations). For these two reasons, we favor regularized least-squares methods.

D. ℓ_1 -regularized least-squares

To avoid deconvolution, [16] proposes a covariance fitting technique. Since in the absence of noise, $\text{vec}\{\mathbf{S}\} = \mathbf{A} \text{vec}\{\mathbf{Y}\}$, the authors propose solving

$$\min_{\hat{\mathbf{Y}}, \sigma^2} \left\| \text{vec}\{\mathbf{S}\} - \mathbf{A} \text{vec}\{\hat{\mathbf{Y}}\} - \sigma^2 \text{vec}\{\mathbf{I}\} \right\|_2^2, \quad (11)$$

subject to $\hat{\mathbf{Y}}_{i,j} \geq 0$, $\sigma^2 \geq 0$ and $\left\| \text{vec}\{\hat{\mathbf{Y}}\} \right\|_1 \leq \lambda$, where σ^2 is the white noise power and $\left\| \text{vec}\{\hat{\mathbf{Y}}\} \right\|_1 \leq \lambda$ is a sparsity constraint. This method assumes that the source distribution is sparse and that only a small number of U-space points have radiating sources. (11) is a convex optimization problem, and can be solved with reasonably efficient numerical methods.

The ℓ_1 constraint serves to regularize the problem, and to permit the inversion of an otherwise ill-conditioned system. Thanks to the ℓ_1 regularization, the authors of [16] show using numerical examples that by solving (11) one can indeed reconstruct sparse images with very high accuracy. Their proposal outperforms DAMAS regarding reconstruction accuracy due to the use of regularization and because no deconvolution was involved.

However, as we have detailed, \mathbf{A} can be a very large matrix, such that solving (11) with a matrix representation of \mathbf{A} (as implemented previously) is very computationally intensive. Of course, the KAT replaces the multiplications by \mathbf{A} and \mathbf{A}^H , which is all that most convex optimization algorithms require.

In order to obtain a fast formulation that is amenable to existing solvers, we propose recasting (11) as a basis pursuit with denoising problem (BPDN), which has the form

$$\min_{\hat{\mathbf{Y}}} \left\| \hat{\mathbf{Y}} \right\|_1 \quad \text{subject to} \quad \left\| \text{vec}\{\mathbf{S}\} - \mathbf{A} \text{vec}\{\hat{\mathbf{Y}}\} \right\|_2 \leq \sigma, \quad (12)$$

and has been studied in detail in the compressive sensing literature. In the examples, we solve (12) with SPGL1 [10], which is a state-of-the-art solver designed for large scale problems. The use of the fast transform not only makes this problem tractable, but makes it competitive with our already very efficient variation of DAMAS2, despite using a more robust method for reconstruction. Note that the FFT acceleration is not applicable to (12).

E. Total variation regularized least-squares

To address scenarios where the acoustic images are not sparse in their canonical representations, we propose reconstructing acoustic images with total variation (TV) regularization.




Figure 1. Simulated array geometry.

Given $\mathbf{Y} \in \mathbb{C}^{M_y \times M_x}$, define its isotropic total variation as

$$\|\mathbf{Y}\|_{BV} = \sum_{i,j} \sqrt{[\nabla_x \mathbf{Y}]_{i,j}^2 + [\nabla_y \mathbf{Y}]_{i,j}^2}, \quad (13)$$

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

$$\min_{\hat{\mathbf{Y}}} \left\| \hat{\mathbf{Y}} \right\|_{BV} + \mu \left\| \text{vec}\{\mathbf{S}\} - \mathbf{A} \text{vec}\{\hat{\mathbf{Y}}\} \right\|_2^2, \quad (14)$$

subject to $\hat{\mathbf{Y}}_{i,j} \geq 0$. The first term measures how much an image oscillates. Therefore, it is smallest for images with plateaus and monotonic transitions, and tends to privilege simple solutions with small amounts of noise. The second term ensures a good fit between the reconstructed image and the measured data. This formulation was first proposed for image denoising by Rudin, Osher and Fatemi [17], for $\mathbf{A} = \mathbf{I}$. It was later generalized and applied successfully to many image reconstruction problems.

To solve (14) we have chosen TVAL3 [11], which uses the augmented Lagrangian method and variable splitting to decouple the TV-minimization and covariance fitting problems. TVAL3 compares very favorably to other solvers in terms of processing time and reconstruction quality, and with the fast transform it becomes practically as efficient as our accelerated version of DAMAS2, while providing more accurate and stable reconstructions with guaranteed convergence.

III. RECONSTRUCTION EXAMPLES

In the following we show image reconstruction examples illustrating the use of delay and sum beamforming, DAMAS2, ℓ_1 regularization and TV regularization, all implemented with the KAT. We simulate a 64-element separable array, with $N_x = N_y = 8$, and with horizontal and vertical apertures of 30 cm. Each $N_x \times 1$ and $N_y \times 1$ linear subarray is chosen to be a nonredundant array with minimum missing lags [18], with interelement spacing .1.3.5.6.7.10.2. (where the dots represent elements, and the numbers represent interelement distances). This geometry is plotted in Fig. 1.

In this section we present results comparing delay and sum beamforming, DAMAS2, ℓ_1 -regularized reconstruction with SPGL1 [10] solving (12), and TV-regularized reconstruction with TVAL3 [11] solving (14). All methods were accelerated with exact versions of the KAT (not using the NFFT), and

the 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. Thanks to the KAT, the reconstruction times for delay-and-sum, DAMAS2, ℓ_1 regularized reconstruction and TV regularized reconstruction were approximately 5 ms, 1.5 s, 8 s and 4 s per image, respectively. DAMAS2 requires no parameters. SPGL1 used $\sigma = 0.01 \|\mathbf{S}\|_F$ in (12). TVAL3 used $\mu = 10^3$ in (14). The signal model is given by $\mathbf{S} = \mathbf{V} \mathbf{E} \{\mathbf{f}\mathbf{f}^H\} \mathbf{V}^H + \sigma^2 \mathbf{I}$, with σ^2 set to obtain 20 dB SNR. Since the intent of these simulations is not to analyze the noise sensitivity of each method, only one SNR is used.

A. Checkerboard patterns

Fig. 2 shows reconstructed checkerboard patterns. This reconstruction clearly shows the deficiencies of delay-and-sum imaging. The images are quite smeared, and all images present artifacts due to sidelobes. DAMAS2 produces very good results, reproducing the checkerboard patterns correctly, with small artifacts outside the checkerboards. ℓ_1 -regularized reconstruction shows better defined edges, but some artifacts, since the patterns are not sparse. Note that some checkerboard squares have dots where the sound pressure level has been underestimated. Artifacts of this kind are common when using ℓ_1 regularization to reconstruct signals with plateaus, because they present a compromise between sparsity and fitness to the measured data. TV regularization produces reconstructions similar to DAMAS2, but with some smearing around the edges of the visible region.

B. Sparse patterns

Fig. 3 presents reconstruction results for a test image with 17 unit impulses at U-space coordinates $(\pm n/6, \pm n/6)$, for $0 \leq n \leq 4$. This test is designed to evaluate the equivalent PSF for the reconstruction methods, as well as detect the presence of aliasing artifacts. As expected, delay and sum has the lowest spatial resolution and shows significant sidelobes. DAMAS2 shows some artifacts, which result from its lack of regularization. ℓ_1 -regularized reconstruction presents the best results, with very small sources and no artifacts. This is a reasonable result, since the image of interest is indeed very sparse. TV-regularization also presents good results, with no discernible artifacts but with larger sources than ℓ_1 regularization.

C. Non-sparse test pattern

Fig. 4 shows reconstruction results for a non-sparse test pattern designed for this experiment. Once again, delay-and-sum has low resolution and rectangular smearing due to the separable geometry's sidelobes. DAMAS2 produces much better results, but still shows some artifacts, especially for high frequency images. The artifacts are gone with ℓ_1 -regularized reconstruction, which also has better resolution than DAMAS2. Nevertheless, it does not represent smooth transitions well, since they are not sparse. Finally, TV-regularized reconstruction produces the most accurate representations, with correct shapes and low noise.

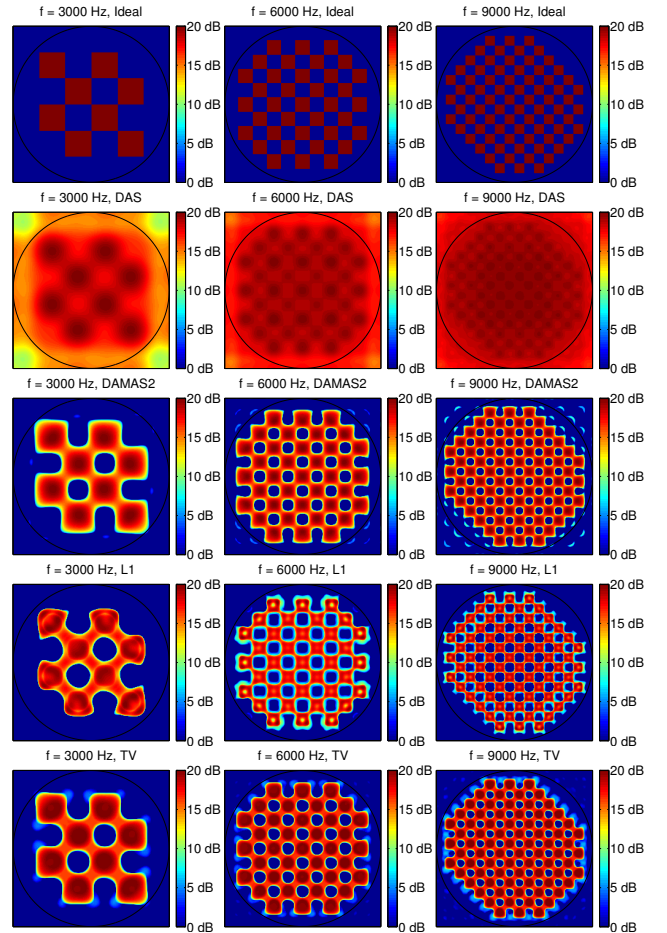


Figure 2. Reconstruction of the checkerboard patterns, for $M_x = M_y = 256$. First row: ideal distribution, second row: delay-and-sum; third row: DAMAS2; fourth row: ℓ_1 -regularized least-squares; fifth row: TV-regularized least-squares.

D. Near-field imaging

In this section, we show how the far-field assumption can break down, and how the near-field extension of the KAT can be used to model near-field effects. We simulate the checkerboard source distribution positioned over a hemispherical shell with a radius of 1.0 m. Both the array and the shell are centered at $(0, 0, 0)$.

The top row of Fig. 5 presents the checkerboard images reconstructed with the exact (slow) near-field transform. Reconstruction results are very similar to the far-field ones, indicating that the transform did not degenerate. The second row shows the reconstruction using a far-field approximation. The estimated distributions are very smeared and show significant artifacts. The other rows show reconstruction results for $K = 1, 4$ and 8 , as prescribed in Section VI of Part I. The artifacts are essentially gone, and the smearing has been significantly reduced. Note that the computational cost for implementing a rank- K KAT is K times larger than implementing a far-field KAT. Nevertheless, even for $K = 8$ this approach is about as fast as a direct NFFT implementation (which cannot be used in this case), while accurately modeling strong near-field effects. Indeed, the reconstruction times for Fig. 5 (using TVAL) were 4.0, 4.9, 6.4 and 10.0 s for $K = 1, 2, 4$ and 8 . In contrast, explicit matrix multiplication requires approximately 2000 s.




Figure 3. Reconstruction of the impulsive patterns, for $M_x = M_y = 256$. First row: ideal distribution, second row: delay-and-sum; third row: DAMAS2; fourth row: ℓ_1 -regularized least-squares; fifth row: TV-regularized least-squares.

IV. HOW GOOD ARE CARTESIAN ARRAYS?

Multi-arm logarithmic spiral arrays [19] have been shown to have low sidelobes over a wide range of frequencies. Since the low sidelobe characteristic is crucial when performing imaging with beamforming, these geometries have found widespread use. Nevertheless, sidelobes have little relevance if one can efficiently use deconvolution or regularized least-squares methods. In this case, ideal geometries become the ones with zero redundancy and minimum missing lags (which give highest bandwidth and some reconstruction artifacts) or minimum redundancy and zero missing lags (which theoretically allow ideal reconstruction up to a given frequency, under a far-field assumption and in the absence of noise). In general, these geometries do not produce low sidelobes, but the sidelobes are low enough to allow nonambiguous reconstruction.

In this section, we compare the Cartesian geometry presented in Fig. 1 and the 63-element logarithmic spiral geometry presented in Fig. 6. This spiral array has an aperture of 50×50 cm, which was chosen to produce images with resolution similar to those of our separable array (which has a 30×30 cm aperture). Furthermore, its parameters were carefully chosen to produce optimal reconstruction for the frequencies of interest. Fig. 7 shows reconstruction results for this logarithmic spiral geometry, under the same conditions as

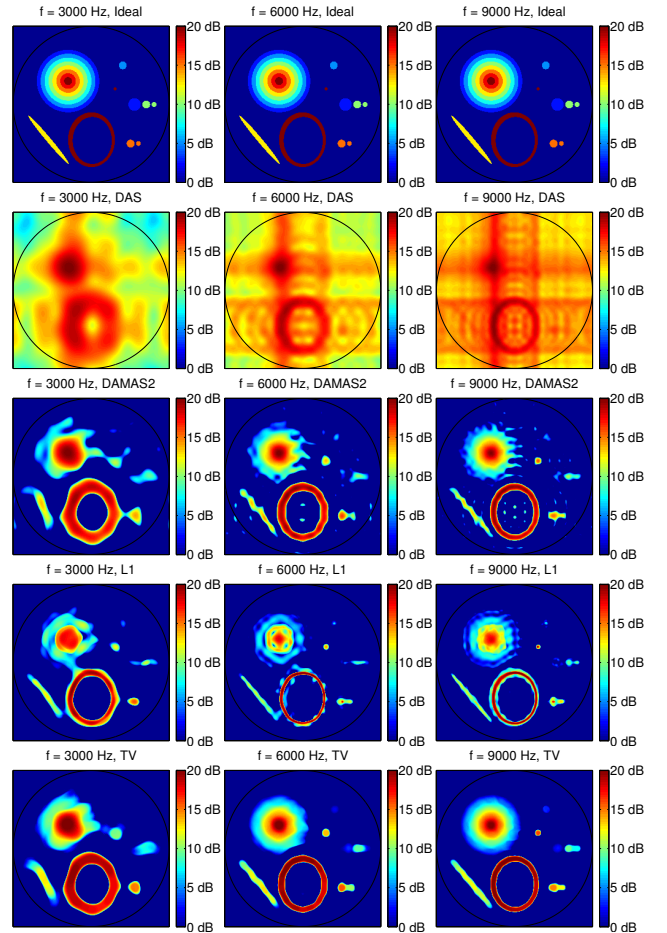


Figure 4. Reconstruction of the non-sparse test pattern, for $M_x = M_y = 256$. First row: ideal distribution, second row: delay-and-sum; third row: DAMAS2; fourth row: ℓ_1 -regularized least-squares; fifth row: TV-regularized least-squares.

Fig. 4.

While the logarithmic spiral geometry produces better results for delay-and-sum, the other techniques produce results of comparable quality. In particular, TV-regularized least-squares produces very similar results for both geometries. This is not surprising, since the Cartesian geometry was chosen to have optimal characteristics. While this example is by no means exhaustive, it is meant to convince the reader that given appropriate image reconstruction techniques, Cartesian geometries can be nearly as powerful as more traditional logarithmic spiral geometries. Of course, with Cartesian arrays one can apply the KAT and obtain extremely fast and accurate near-field reconstruction for arbitrary focal surfaces. With logarithmic spiral geometries, one has no such option.

V. CONCLUSION

Using the assumption of a separable array geometry, in Part I we presented the Kronecker array transform (KAT), which can efficiently transform back and forth between a hypothetical source distribution and its corresponding spectral matrix, under the assumption of a separable array. This transform is orders of magnitude faster than explicit matrix multiplication, and one order of magnitude faster than NFFT-based approaches. Furthermore, the KAT can be generalized

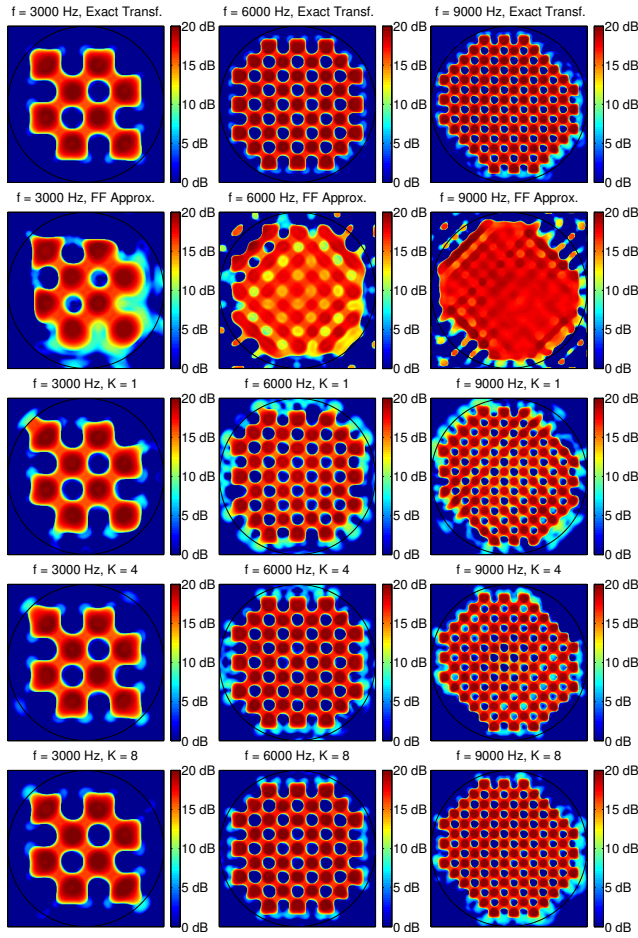


Figure 5. TV-regularized reconstruction of the checkerboard patterns. The real source distribution is located over a hemispherical shell with a radius of 1 m. From top to bottom: reconstruction using the exact (slow) transform; reconstruction using a far-field approximation; reconstructions using the best Kronecker approximation for varying values of K .




Figure 6. Logarithmic spiral array geometry with 63 elements, inner radius $r_0 = 1.5$ cm, outer radius $r_{max} = 25$ cm, 9 concentric circles, 7 arms and having each arm perform 2 full rotations.

for near-field imaging, while the NFFT must use a far-field approximation.

In this Part II, we have recast delay-and-sum beamforming, MVDR beamforming and the DAMAS2 deconvolution algorithm [15] as applications of our transform. Thus, we have shown how the computational benefits from Part I can be realized in practice. Furthermore, the KAT allows the efficient use of general purpose regularized least-squares solvers. To demonstrate this application, we used it to recast acoustic imaging as least-squares problems with ℓ_1 and total variation regularization. With the KAT, regularized reconstruction be-

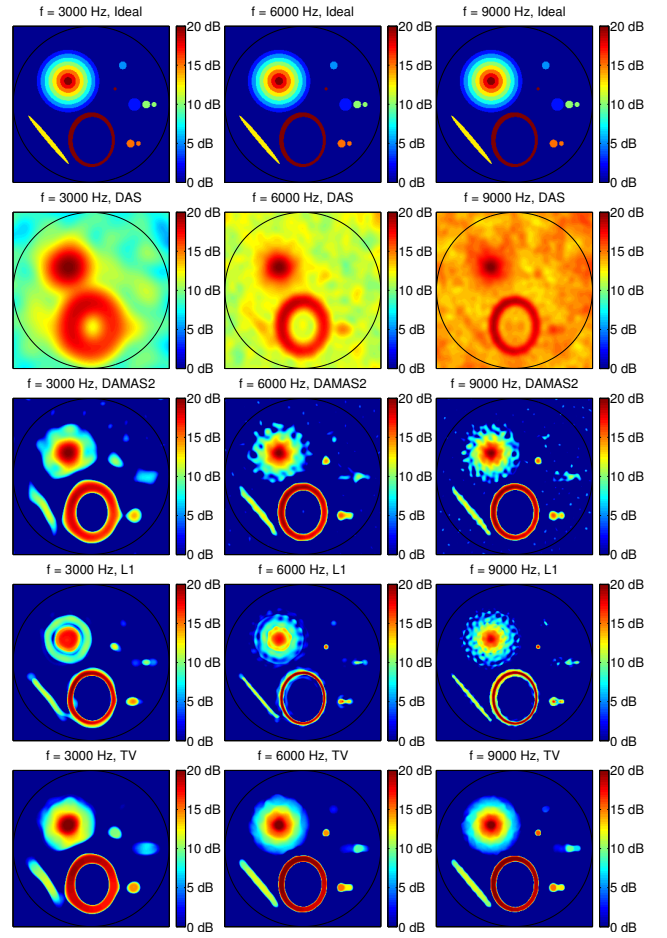


Figure 7. Reconstruction of the non-sparse test pattern with a logarithmic spiral array, for $M_x = M_y = 256$. First row: ideal distribution, second row: delay-and-sum; third row: DAMAS2; fourth row: ℓ_1 -regularized least-squares; fifth row: TV-regularized least-squares.

comes straightforward, elegant and computationally efficient. With this approach, image reconstruction can be delegated to third party solvers, saving the time and effort of developing ad hoc methods.

Finally, we have shown that by using carefully chosen separable arrays one does not have to compromise on reconstruction quality. Thus, the KAT does not require a trade-off between accuracy and reconstruction time.

REFERENCES

- [1] F. Ribeiro and V. Nascimento, “Fast Transforms for Acoustic Imaging – Part I: Theory,” 2010, submitted.
- [2] S. Chen, D. Donoho, and M. Saunders, “Atomic decomposition by basis pursuit,” *SIAM review*, vol. 43, no. 1, pp. 129–159, 2001.
- [3] E. Candès, J. Romberg, and T. Tao, “Stable signal recovery from incomplete and inaccurate measurements,” *Communications on Pure and Applied Mathematics*, vol. 59, no. 8, p. 1207, 2006.
- [4] E. Candès, J. Romberg, and T. Tao, “Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information,” *IEEE Transactions on Information Theory*, vol. 52, no. 2, pp. 489–509, 2006.
- [5] D. Donoho, “Compressed sensing,” *IEEE Transactions on Information Theory*, vol. 52, no. 4, pp. 1289–1306, 2006.
- [6] D. Donoho and J. Tanner, “Sparse nonnegative solution of underdetermined linear equations by linear programming,” *Proceedings of the National Academy of Sciences of the United States of America*, vol. 102, no. 27, p. 9446, 2005.

- [7] E. Hale, W. Yin, and Y. Zhang, "A Fixed-Point Continuation Method for ℓ_1 -Regularized Minimization with Applications to Compressed Sensing," Rice University, Tech. Rep. CAAM R07-07, 2007.
- [8] M. Figueiredo, R. Nowak, and S. Wright, "Gradient Projection for Sparse Reconstruction: Application to Compressed Sensing and Other Inverse Problems," *IEEE Journal of Selected Topics in Signal Processing*, vol. 1, no. 4, pp. 586–597, 2007.
- [9] S. Kim, K. Koh, M. Lustig, S. Boyd, and D. Gorinevsky, "An Interior-Point Method for Large-Scale ℓ_1 -Regularized Least Squares," *IEEE Journal of Selected Topics in Sig. Proc.*, vol. 1, no. 4, pp. 606–617, 2007.
- [10] E. van den Berg and M. Friedlander, "Probing the Pareto frontier for basis pursuit solutions," *SIAM Journal on Scientific Computing*, vol. 31, no. 2, pp. 890–912, 2008.
- [11] C. Li, "An Efficient Algorithm For Total Variation Regularization with Applications to the Single Pixel Camera and Compressive Sensing," Master's thesis, Rice University, 2009.
- [12] J. Capon, "High-resolution frequency-wavenumber spectrum analysis," *Proceedings of the IEEE*, vol. 57, no. 8, pp. 1408–1418, 1969.
- [13] D. Blacodon, "Spectral estimation noisy data using a reference noise," in *Proc. BeBeC*, 2010.
- [14] H. L. Van Trees, *Optimum Array Processing: Part IV of Detection, Estimation, and Modulation Theory*. New York, NY: John Wiley & Sons, 2002.
- [15] R. Dougherty, "Extensions of DAMAS and Benefits and Limitations of Deconvolution in Beamforming," in *Proc. of the 11th AIAA/CEAS Aeroacoustics Conference*, 2005.
- [16] T. Yardibi, J. Li, P. Stoica, and L. Cattafesta III, "Sparsity constrained deconvolution approaches for acoustic source mapping," *The Journal of the Acoustical Society of America*, vol. 123, p. 2631, 2008.
- [17] L. Rudin, S. Osher, and E. Fatemi, "Nonlinear total variation based noise removal algorithms," *Physica D: Nonlinear Phenomena*, vol. 60, no. 1-4, pp. 259–268, 1992.
- [18] E. Vertatschitsch and S. Haykin, "Nonredundant arrays," *Proceedings of the IEEE*, vol. 74, no. 1, p. 217, 1986.
- [19] J. Underbrink and R. Dougherty, "Array design for non-intrusive measurement of noise sources," in *Proc. of NOISE-CON*, 1996, pp. 757–762.



Flávio P. Ribeiro received the B.S. degree in electrical engineering from Escola Politécnica, University of São Paulo in 2005, and the B.S. degree in mathematics from the Institute of Mathematics and Statistics, University of São Paulo in 2008. He currently pursuing the Ph.D. degree in electrical engineering, also at Escola Politécnica, University of São Paulo.

From 2007 to 2009, he was a hardware engineer at Licht Labs, where he developed controllers for power transformers and substations. On the summers of 2009 and 2010, he was a research intern at

Microsoft Research Redmond. He won the best student paper award at ICME 2010. His research interests include array signal processing, multimedia signal processing and computational linear algebra.



Vítor H. Nascimento was born in São Paulo, Brazil. He obtained the B.S. and M.S. degrees in Electrical Engineering from the University of São Paulo, in 1989 and 1992, respectively, and the Ph.D. degree from the University of California, Los Angeles, in 1999. From 1990 to 1994 he was a Lecturer at the Univ. of São Paulo, and in 1999 he joined the faculty at the same school, where he is now an Associate Professor. In 2002, he received a IEEE SPS Best Paper Award. He served as an Associate Editor for *Signal Processing Letters* from 2003 to

2005, for the *Transactions on Signal Processing* from 2005 to 2008 and for the *EURASIP Journal on Advances in Signal Processing* from 2006 to 2009. He is currently a member of the IEEE-SPS Signal Processing Theory and Methods Technical Committee. His research interests include signal processing theory and applications, robust and nonlinear estimation, and applied linear algebra.