

Fast Transforms for Acoustic Imaging — Part I: Theory

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

Abstract—The classical approach for acoustic imaging consists of beamforming, and produces the source distribution of interest convolved with the array point spread function. This convolution smears the image of interest, significantly reducing its effective resolution. Deconvolution methods have been proposed to enhance acoustic images, and have produced significant improvements. Other proposals involve covariance fitting techniques, which avoid deconvolution altogether. However, in their traditional presentation, these enhanced reconstruction methods have very high computational costs, mostly because they have no means of efficiently transforming back and forth between a hypothetical image and the measured data. In this paper we propose the KAT (Kronecker Array Transform), a fast separable transform for array imaging applications. Under the assumption of a separable array, it allows the acceleration of imaging techniques by several orders of magnitude with respect to the fastest previously available methods, and enables the use of state of the art regularized least-squares solvers. Using the KAT, one can reconstruct images with higher resolutions than previously possible, and also use more accurate reconstruction techniques, opening new and exciting possibilities for acoustic imaging.

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

I. INTRODUCTION

Acoustic imaging refers to the problem of mapping the locations and intensities of sound sources over a region of interest using microphone arrays. For example, a microphone array can be positioned in a wind tunnel to determine the noise distribution over a model due to high velocity airflow [1], [2]. Such measurements are routinely used to design cars, trains and aircraft which are quieter to outside observers and to passengers. Microphone arrays have been employed to measure the noise generated by turbofan engines [3] and wind turbines [4], for similar noise reduction applications. Acoustic imaging has also been used to visualize the reverberant structure of concert halls [5]. We note that techniques for imaging the shapes of objects and structures via acoustic waves are sometimes called acoustic imaging, but we will not address such problems.

Array imaging is possible because sensor arrays can be electronically steered toward arbitrary directions. One can define a grid over a region of interest, electronically steer the array over all elements of the grid, and thus create a map of estimated sound pressure levels. Each point in the grid can be represented as a pixel. The value of the pixel can be chosen to represent the estimated sound pressure level, thus creating an acoustic image.

Array imaging differs from source localization techniques such as [6]–[10] because these usually produce a pseudospec-

trum of the wavefield, with maxima that indicate the location of dominant sources but with values that do not map to source powers. Thus, in this paper we will assume that accurate power estimates are desirable.

The simplest and most common method for imaging uses delay and sum beamforming. This technique consists of delaying and summing the signals arriving at each sensor such that the sources located at a direction of interest are reinforced, and sources located in other directions are attenuated. Beamforming is simple, but unfortunately produces the lowest quality images. Indeed, under the assumption that the sources are in the far-field of the array, beamforming produces the source distribution of interest convolved with the array point spread function (PSF). Since a typical acoustic array has a relatively small aperture with respect to its operating wavelengths, its PSF can be quite large, such that delay and sum beamforming produces very smeared images. Alternative beamforming techniques have been developed to improve resolution by using data-dependent methods and numerical optimization of the beampattern [11]–[13], but they do not overcome the fundamental limitation that beamforming produces convolved images.

To overcome this limitation, several deconvolution techniques have been proposed [14]–[17]. They use as inputs the image obtained with delay and sum beamforming and the array PSF, and generally produce a much better approximation of the original source distribution. Nevertheless, deconvolution is an ill conditioned inverse problem, and typically requires some knowledge of the solution to discriminate between different solutions which would be equally good fits for the measured data. The acoustic imaging methods proposed so far tend to use very simple types of regularization, such as low pass filtering between iterations, or no regularization at all.

Regularized signal reconstruction has been a topic of interest for many decades, and gained significant momentum with the popularity of compressive sensing [18]–[20]. Indeed, many image reconstruction problems can be recast as convex optimization problems, which can be solved with computationally efficient iterative methods. While many of these techniques were designed for imaging applications, they have remained limited to fields such as medical image reconstruction. Therefore, most of these developments have not been applied to acoustic imaging.

A major reason for this separation between fields has been the absence of computationally efficient transforms for aeroacoustic imaging. For example, consider the generic non-linear signal reconstruction problem given by

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

where \mathbf{x} is the measured signal, $\hat{\mathbf{y}}$ is the reconstructed signal, Ψ is a sparsifying transform, and Φ 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{x} would be a vectorized version of the array's sample covariance matrix.

Since in practice (1) is solved iteratively, one must be able to quickly evaluate $\Psi\mathbf{u}$ and $\Phi\mathbf{u}$ (and also $\Psi^H\mathbf{v}$ and $\Phi^H\mathbf{v}$, as we will see in Part II [21] of this paper) for arbitrary \mathbf{u} , \mathbf{v} . This is a very strong requirement, because the application of these transforms is the bottleneck of efficient convex optimization algorithms and completely determines their computational costs (regardless of the transform being fast or slow). While one can choose a convenient fast sparsifying transform Ψ , the transform Φ is determined by the physical measurement process. For example, for MRI applications we naturally have $\Phi = \mathbf{P}\mathcal{F}$, where \mathcal{F} is a fast Fourier transform and \mathbf{P} is a subsampling operator. Finite differences have been successfully used as the sparsifying transform Ψ with $\|\cdot\| = \|\cdot\|_1$ [22], [23].

While sparsity-enforcing approaches have been proposed for direction of arrival estimation [24] and acoustic imaging [25], to our knowledge no method of acoustic imaging uses a fast implementation of Φ . To motivate the need for a fast transform, consider a naive matrix representation of Φ . Given an array of N sensors and an image with M pixels, Φ has N^2 rows and M columns. For $N^2 = M = 256^2$, Φ has 4 billion elements and the products $\Phi\mathbf{u}$ and $\Phi^H\mathbf{v}$ are computationally very expensive, making convex optimization methods intractable with current desktop computers. Thus, the naive implementation of Φ is only practical for very small images and arrays, thus motivating the development of a fast transform.

In this paper we develop the KAT (Kronecker Array Transform), a fast transform which implements Φ , Φ^H and $\Phi^H\Phi$ for separable arrays. The KAT can be applied to many existing array imaging algorithms, with significant performance gains. It also allows the use of state of the art solvers for acoustic imaging problems, obviating ad-hoc solutions which typically produce worse results. Indeed, with a fast transform one can use most of the general purpose, state-of-the-art solvers developed for other imaging and compressive sensing applications. By combining the KAT with such methods, we can accelerate reconstruction times by several orders of magnitude with respect to the fastest previously available implementations. In practical terms, an image which would take minutes to reconstruct can be obtained in a few seconds. Finally, while this transform was motivated by applications in aeroacoustics, it also applies for generic wave fields and separable sensor arrays.

To our knowledge, previous proposals for accelerated acoustic imaging are all based on beamforming. Zimmerman and Studer [26] propose offloading delay-and-sum beamforming to an FPGA, which performs all the computation and draws acoustic images over a framebuffer. While this approach makes beamforming faster, it does not reduce its underlying computational cost. Huang [27] uses a state observer model to recursively obtain an approximation of the acoustic image

while acquiring data (in contrast to computing an image from a sample covariance matrix). While this method has the advantage of returning incremental results, it has the same computational cost as beamforming, and a comparable beampattern. In contrast, the KAT dramatically reduces the underlying computational costs of acoustic imaging, allowing more accurate reconstruction methods to be used instead of beamforming.

Part I is organized as follows: Section II gives several definitions and further motivates the need for fast transforms. Section III introduces the KAT, its adjoint and its direct-adjoint composition, under the assumption of far-field sources. Section IV presents connections with the fast Fourier transform, fast non-equispaced Fourier transform (NFFT) and fast non-equispaced in time and frequency Fourier transform (NNFFT) [28]. These connections are also a contribution, because to our knowledge, the NFFT and NNFFT have never been used for acoustic imaging. We show how the NFFT and NNFFT can also be used to accelerate acoustic imaging under a far-field approximation, despite being an order of magnitude slower than the KAT. Section V presents benchmarks comparing the KAT with the NFFT, NNFFT and explicit matrix representations. Section VI shows how to extend the KAT for near-field imaging, modeling spherical wavefronts instead of planar wavefronts. This generalization is unique to our proposal, and produces a transform which is orders of magnitude faster than direct matrix multiplication (which becomes the only alternative, since the FFT, NFFT and NNFFT require a far-field approximation). Section VII has our conclusions.

Part II [21] presents applications. Section I briefly reviews the results from Part I. In Section II we use the KAT to significantly accelerate existing techniques and to enable the use of general purpose solvers, obtaining more accurate reconstructions than possible with current state-of-the-art methods. Section III features examples, and compares several reconstruction methods with respect to computational cost and accuracy. Section IV compares separable arrays with multi-arm logarithmic spiral arrays, and shows that by requiring separable arrays we are not trading reconstruction quality for speed. Finally, Section V has our conclusions and final comments.

II. PRELIMINARIES

Consider a sensor array composed of N microphones at Cartesian coordinates $\mathbf{p}_0, \dots, \mathbf{p}_{N-1} \in \mathbb{R}^3$, and an arbitrary wavefield which we wish to estimate. Suppose that this wave field can be modeled as generated by the superposition of M point sources located at coordinates $\mathbf{q}_0, \dots, \mathbf{q}_{M-1} \in \mathbb{R}^3$, where M may be a large number in order to obtain an accurate model. Let

$$\mathbf{Q} = [\mathbf{q}_0 \ \mathbf{q}_1 \ \dots \ \mathbf{q}_{M-1}]. \quad (2)$$

The time-domain samples of each microphone are segmented into frames of K samples, and each frame is converted to the frequency domain using a fast Fourier transform (FFT). In the presence of additive noise, the $N \times 1$ array output vector for a single frequency on a single frame can be modeled as

$$\mathbf{x}(\omega_k) = \mathbf{V}(\mathbf{Q}, \omega_k) \mathbf{f}(\omega_k) + \boldsymbol{\eta}(\omega_k), \quad (3)$$

Table I
LIST OF SYMBOLS

Symbol	Meaning
T	matrix or vector transpose
H	matrix or vector Hermitian transpose
$*$	complex conjugate
$\text{mod}(a, b)$	remainder of a/b , for $a, b \in \mathbb{Z}_+$
$\lfloor x \rfloor$	round-off of $x \in \mathbb{R}$ towards $-\infty$
$\text{vec}\{\cdot\}$	vectorization operator
\otimes	Kronecker product
\odot	Hadamard product
$M = M_x M_y$	number of image pixels
$N = N_x N_y$	number of array microphones
ω	operating frequency in the Fourier domain
\mathbf{q}_m	source coordinates, for $0 \leq m < M$
\mathbf{u}_m	source coordinates in U-space, for $0 \leq m < M$
\mathbf{p}_n	array element coordinates, for $0 \leq n < N$
$\mathbf{v}_{\mathbf{u}_m}(\omega)$	$N \times 1$ array manifold vector for look direction \mathbf{u}_m
$\mathbf{v}_{x_n}(\omega)$	$N_x \times 1$ separable component of $\mathbf{v}_{\mathbf{u}_m}(\omega)$
$\mathbf{v}_{y_n}(\omega)$	$N_y \times 1$ separable component of $\mathbf{v}_{\mathbf{u}_m}(\omega)$
$v_{x_n}^i(\omega)$	i^{th} element of $\mathbf{v}_{x_n}(\omega)$
$\mathbf{Y}(\omega)$	$M_x \times M_y$ image at a frequency ω
$\mathbf{S}(\omega)$	$N \times N$ spectral matrix at a frequency ω
$\mathbf{B}_{i,j}$	(i, j) entry of matrix \mathbf{B}
$\mathbb{1}$	matrix of ones

where $0 \leq k < K/2$, $\mathbf{V}(\mathbf{Q}, \omega_k) = [\mathbf{v}(\mathbf{q}_0, \omega_k) \mathbf{v}(\mathbf{q}_1, \omega_k) \cdots \mathbf{v}(\mathbf{q}_{M-1}, \omega_k)]$ is the array manifold matrix, $\mathbf{f}(\omega_k) = [f_0(\omega_k) f_1(\omega_k) \cdots f_{M-1}(\omega_k)]^T$ is the frequency domain signal waveform and $\boldsymbol{\eta}(\omega_k)$ is the frequency domain noise waveform.

The near-field array manifold vector for source m is given by [29]

$$\mathbf{v}(\mathbf{q}_m, \omega_k) = \left[\frac{e^{-j\frac{\omega_k}{c}\|\mathbf{p}_0 - \mathbf{q}_m\|}}{\|\mathbf{p}_0 - \mathbf{q}_m\|} \cdots \frac{e^{-j\frac{\omega_k}{c}\|\mathbf{p}_{N-1} - \mathbf{q}_m\|}}{\|\mathbf{p}_{N-1} - \mathbf{q}_m\|} \right]^T, \quad (4)$$

where c is the speed of sound.

Define $\mathbf{u}_m = \mathbf{q}_m / \|\mathbf{q}_m\|$, the look direction for source m . Under a far-field approximation (modeling a plane wave), the time differences of arrival are given by $-\mathbf{u}_m^T \mathbf{p}_n / c$, for $0 \leq n < N$. Since the wavefront is not expanding, the $1/\|\mathbf{p}_n - \mathbf{q}_m\|$ attenuation disappears, and the far-field array manifold vector for source m is given by

$$\mathbf{v}(\mathbf{u}_m, \omega_k) = \left[e^{j\frac{\omega_k}{c}\mathbf{u}_m^T \mathbf{p}_0} \cdots e^{j\frac{\omega_k}{c}\mathbf{u}_m^T \mathbf{p}_{N-1}} \right]^T. \quad (5)$$

Using spherical coordinates,

$$\mathbf{u}_m = \begin{bmatrix} \sin \phi \cos \theta \\ \sin \phi \sin \theta \\ \cos \phi \end{bmatrix}, \quad (6)$$

where θ and ϕ are the azimuth and elevation angles, respectively. One can reparameterize the unit half-sphere by defining

$$u_x(\theta, \phi) = \sin \phi \cos \theta \quad (7)$$

$$u_y(\theta, \phi) = \sin \phi \sin \theta, \quad (8)$$

such that

$$\mathbf{u}_m = \begin{bmatrix} u_x \\ u_y \\ \sqrt{1 - u_x^2 - u_y^2} \end{bmatrix} \quad (9)$$

for $u_x^2 + u_y^2 \leq 1$. Uniform sampling in U-space (where $\mathbf{U} = [-1, 1]^2$) is convenient in many applications, because under a far-field approximation it makes point-spread functions shift-invariant. In this paper it will enable us to decouple the x and y axes, producing the fast transform.

Since the optimizations presented in the following sections require Cartesian (not necessarily uniform) parameterizations in U-space and far-field approximations, we will assume that manifold vectors have the form (5). On Section VI we extend our results for near-field sources, which will allow us to approximate (4) with arbitrary accuracy.

Using the assumption of far-field sources, we rewrite (3) as

$$\mathbf{x}(\omega_k) = \mathbf{V}(\omega_k) \mathbf{f}(\omega_k) + \boldsymbol{\eta}(\omega_k), \quad (10)$$

where $\mathbf{V}(\omega_k) = [\mathbf{v}(\mathbf{u}_0, \omega_k) \mathbf{v}(\mathbf{u}_1, \omega_k) \cdots \mathbf{v}(\mathbf{u}_{M-1}, \omega_k)]$.

Let

$$\mathbf{S}_{\mathbf{x}}(\omega_k) = \mathbb{E} \{ \mathbf{x}(\omega_k) \mathbf{x}^H(\omega_k) \} \quad (11)$$

be the array's narrow-band cross spectral matrix for $0 \leq k < K/2$. If $\mathbf{x}_0(\omega_k), \dots, \mathbf{x}_{L-1}(\omega_k)$ correspond to L frequency domain frames (also known as snapshots), the spectral matrix can be estimated with

$$\mathbf{S}_{\mathbf{x}}(\omega_k) = \frac{1}{L} \sum_{l=0}^{L-1} \mathbf{x}_l(\omega_k) \mathbf{x}_l^H(\omega_k). \quad (12)$$

We assume that the statistics of the signal and noise are stationary over the measured period, such that (12) is an unbiased estimator.

Processing $\mathbf{S}_{\mathbf{x}}(\omega_k)$ instead of each $\mathbf{x}_l(\omega_k)$ is typically more convenient, because $\mathbf{S}_{\mathbf{x}}(\omega_k)$ carries only the relative phase shifts between microphones and is the result of averaging, such that it has less noise content. Indeed, for each $0 \leq l < L$, $\mathbf{x}_l(\omega_k)$ has a phase shift which is equal for every element but unknown, which disappears when computing $\mathbf{S}_{\mathbf{x}}(\omega_k)$. To save space, in the following we will assume narrow-band processing and omit the argument ω_k . Also, the \mathbf{x} subscript will be dropped, and $\mathbf{S}_{\mathbf{x}}(\omega_k)$ will be written as \mathbf{S} .

Substituting (10) into (11) and assuming that the noise is spatially white and uncorrelated with the sources of interest, we have

$$\mathbf{S}_{\mathbf{x}}(\omega_k) = \mathbf{V}(\omega_k) \mathbb{E} \{ \mathbf{f}(\omega_k) \mathbf{f}^H(\omega_k) \} \mathbf{V}^H(\omega_k) + \sigma^2 \mathbf{I}, \quad (13)$$

where $\sigma^2 = \mathbb{E} \{ \eta_i(\omega_k) \eta_i^*(\omega_k) \}$, $0 \leq i < N$.

Assume that the wavefield impinging on the array can be modeled as emitted by the superposition of uncorrelated point sources located in the array's far field. One can represent these sources by a collection of points at coordinates $\{\mathbf{u}_i\}_{i=0}^{M-1}$ located in a sufficiently fine grid in U-space. This representation is effectively a two-dimensional digital image, where the pixel coordinates correspond to locations in U-space, and the pixel values correspond to source intensities. Note that in (13), assuming that the sources are uncorrelated implies that $\mathbb{E} \{ \mathbf{f}(\omega_k) \mathbf{f}^H(\omega_k) \}$ is diagonal. Furthermore, the diagonal of $\mathbb{E} \{ \mathbf{f}(\omega_k) \mathbf{f}^H(\omega_k) \}$ is a vectorized version of the acoustic image.

Given an acoustic image where each pixel corresponds to a point source, one can easily obtain the array spectral matrix

as long as all point sources (pixels) are assumed to be pairwise uncorrelated. If there were cross-correlations, one would drop the assumption that $E \{ \mathbf{f}(\omega_k) \mathbf{f}^H(\omega_k) \}$ is diagonal, and require the correlation coefficient for each pair of sources, whose determination would be clearly impractical even for small images (for example, a 64×64 pixel image would have $64^2 = 4096$ pixels (sources) and $\frac{1}{2}4096 \cdot 4097 \approx 8$ million unique cross-correlations). Therefore, unless stated otherwise, we shall assume that sources are pairwise uncorrelated.

However, we note that a fast transform can also be obtained for correlated source distributions as a natural generalization of the KAT. It has special importance because one cannot apply Fourier methods to accelerate the reconstruction of correlated source maps. Nevertheless, since the number of cross-correlations scales quadratically, estimating every cross term is only viable for very simple source distributions. Thus, a proposal for imaging correlated sources should combine a fast transform and domain-specific regularization, the latter being an open problem. Due to space limitations, we will not address these topics in this text.

Recall that to solve (1) efficiently, one requires a fast method of obtaining \mathbf{S} from a hypothetical image. Consider an $M_x \times M_y$ pixel acoustic image, define $M = M_x M_y$ and let $\mathbf{u}_0, \dots, \mathbf{u}_{M-1}$ be an enumeration of all pixel coordinates in U-space. Let $\mathbf{v}(\mathbf{u}_m)$ be the array manifold vector when steered towards look direction \mathbf{u}_m . For a single source at \mathbf{u}_m radiating with power $|Y(\mathbf{u}_m)|^2$, the measured spectral matrix is $\mathbf{S} = |Y(\mathbf{u}_m)|^2 \mathbf{v}(\mathbf{u}_m) \mathbf{v}^H(\mathbf{u}_m)$. Given the source powers for $\mathbf{u}_0, \dots, \mathbf{u}_{M-1}$, one can reconstruct \mathbf{S} by superposition, so that

$$\mathbf{S} = \sum_{m=0}^{M-1} |Y(\mathbf{u}_m)|^2 \mathbf{v}(\mathbf{u}_m) \mathbf{v}^H(\mathbf{u}_m). \quad (14)$$

Unless the image is very sparse, this expression becomes computationally intractable. For instance, consider a 256 element array and a 256×256 acoustic image. Each $|Y(\mathbf{u}_m)|^2 \mathbf{v}(\mathbf{u}_m) \mathbf{v}^H(\mathbf{u}_m)$ outer product generates a 256-by-256 matrix. Neglecting the cost to scale $\mathbf{v}(\mathbf{u}_m)$ by $|Y(\mathbf{u}_m)|$, the outer product requires 2^{16} complex multiply-accumulate (MAC) instructions¹. This process must be repeated $M = 2^{16}$ times, resulting in 2^{32} complex MACs. Since each outer product has Hermitian symmetry, it suffices to determine its upper or lower triangular part (including the main diagonal), which reduces the total complex MAC count to approximately 2^{31} . Nevertheless, this computational cost is still excessive for a transform intended to be used in an iterative method. In the following, we describe how to implement an efficient transform to obtain \mathbf{S} from $\left\{ |Y(\mathbf{u}_m)|^2 \right\}_{0 \leq m < M}$.

III. THE KRONECKER ARRAY TRANSFORM

Define $\mathbf{y} = [|Y(\mathbf{u}_0)|^2 \cdots |Y(\mathbf{u}_{M-1})|^2]^T$. Let us write (14) as a linear transform \mathbf{A} such that $\mathbf{s} = \mathbf{A}\mathbf{y}$, with $\mathbf{s} = \text{vec}\{\mathbf{S}\}$. To save space, we will write $\mathbf{v}(\mathbf{u}_m)$ as $\mathbf{v}_{\mathbf{u}_m}$,

¹Modern DSP architectures are able to implement a multiplication followed by an accumulation in the same clock cycle. This single cycle instruction is known as a multiply-accumulate, or MAC. Since the computational cost of performing a sum, product or MAC is the same, for the purposes of estimating computational complexity it suffices to estimate the total number of MACs.

and will denote its i^{th} element by $v_{\mathbf{u}_m}^i$ (elements of array manifold vectors will be indexed using superscripts). Let N be the number of microphones in the array. Note that

$$\text{vec}\{\mathbf{v}_{\mathbf{u}_m} \mathbf{v}_{\mathbf{u}_m}^H\} = \begin{bmatrix} \mathbf{v}_{\mathbf{u}_m} v_{\mathbf{u}_m}^{0*} \\ \mathbf{v}_{\mathbf{u}_m} v_{\mathbf{u}_m}^{1*} \\ \vdots \\ \mathbf{v}_{\mathbf{u}_m} v_{\mathbf{u}_m}^{(N-1)*} \end{bmatrix}. \quad (15)$$

Therefore,

$$\begin{aligned} \mathbf{s} &= \mathbf{A}\mathbf{y} \\ &= \begin{bmatrix} \mathbf{v}_{\mathbf{u}_0} v_{\mathbf{u}_0}^{0*} & \mathbf{v}_{\mathbf{u}_1} v_{\mathbf{u}_1}^{0*} & \cdots & \mathbf{v}_{\mathbf{u}_{M-1}} v_{\mathbf{u}_{M-1}}^{0*} \\ \mathbf{v}_{\mathbf{u}_0} v_{\mathbf{u}_0}^{1*} & \mathbf{v}_{\mathbf{u}_1} v_{\mathbf{u}_1}^{1*} & \cdots & \mathbf{v}_{\mathbf{u}_{M-1}} v_{\mathbf{u}_{M-1}}^{1*} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{v}_{\mathbf{u}_0} v_{\mathbf{u}_0}^{(N-1)*} & \mathbf{v}_{\mathbf{u}_1} v_{\mathbf{u}_1}^{(N-1)*} & \cdots & \mathbf{v}_{\mathbf{u}_{M-1}} v_{\mathbf{u}_{M-1}}^{(N-1)*} \end{bmatrix} \mathbf{y} \\ &= [\mathbf{v}_{\mathbf{u}_0}^* \otimes \mathbf{v}_{\mathbf{u}_0} \quad \mathbf{v}_{\mathbf{u}_1}^* \otimes \mathbf{v}_{\mathbf{u}_1} \quad \cdots \quad \mathbf{v}_{\mathbf{u}_{M-1}}^* \otimes \mathbf{v}_{\mathbf{u}_{M-1}}] \mathbf{y}, \end{aligned} \quad (16)$$

where \otimes is the Kronecker product.

Given a two-dimensional array, its array manifold vector $\mathbf{v}(\mathbf{u}) = \mathbf{v}(u_x, u_y)$ is said to be separable if there exist $\mathbf{a}(u_x)$ and $\mathbf{b}(u_y)$ such that $\mathbf{v}(u_x, u_y) = \mathbf{a}(u_x) \otimes \mathbf{b}(u_y)$ for all valid u_x, u_y . Note that $\mathbf{a}(u_x)$ and $\mathbf{b}(u_y)$ need not be submanifold vectors. We say that an array is separable if and only if it has a separable manifold vector. We will show below how the array geometry relates to its separability under a far-field assumption.

To simplify the notation that follows, let us specify the enumeration $\mathbf{u}_0, \dots, \mathbf{u}_{M-1}$ of look directions we are using. Suppose that \mathbf{Y} is a $M_x \times M_y$ digital image representing the acoustic image. The rows of \mathbf{Y} correspond to horizontal scan lines of arbitrarily sampled pixels, and the columns of \mathbf{Y} correspond to vertical scan lines of arbitrarily sampled pixels. Let $\{u_{x_m}\}_{0 \leq m < M_x}$ and $\{u_{y_n}\}_{0 \leq n < M_y}$ be points which sample U-space along the x and y axes, ordered from left to right and from top to bottom. We define $\mathbf{u}_0, \dots, \mathbf{u}_{M-1}$ such that

$$\mathbf{y} = \text{vec}\{\mathbf{Y}\} = \begin{bmatrix} |Y(\mathbf{u}_0)|^2 \\ |Y(\mathbf{u}_1)|^2 \\ \vdots \\ |Y(\mathbf{u}_{M-1})|^2 \end{bmatrix}. \quad (18)$$

Breaking \mathbf{u} into components, this implies that

$$\mathbf{u}_m = \begin{bmatrix} u_{x_{\lfloor m/M_y \rfloor}} \\ u_{y_{\text{mod}(m, M_y)}} \\ \sqrt{1 - (u_{x_{\lfloor m/M_y \rfloor}})^2 - (u_{y_{\text{mod}(m, M_y)}})^2} \end{bmatrix}. \quad (19)$$

Fig. 1 shows an example of how pixels are ordered and parameterized in U-space.

We now show that under the far-field parameterization given by (5), an array is separable if and only if it has elements positioned over a (potentially non-uniform) Cartesian grid.

To see this, consider an array with sensor coordinates $\mathbf{p}_i \in \mathbb{R}^3$, for $0 \leq i < N$, with x and y coordinates drawn from $\{p_{x_i}\}_{i=0}^{N_x-1}$ and $\{p_{y_i}\}_{i=0}^{N_y-1}$, respectively, such that

$$\mathbf{p}_i = [p_{x_{\lfloor i/N_y \rfloor}} \quad p_{y_{\text{mod}(i, N_y)}} \quad 0]^T.$$

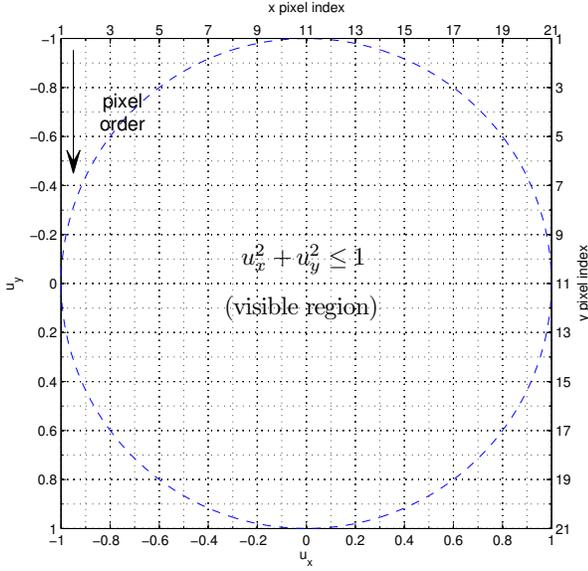


Figure 1. Example of pixel order and U-space parameterization for an acoustic image, for $M_x = M_y = 21$ and uniform sampling in U-space.

Let $N = N_x N_y$ be the number of array elements. Define a horizontal array with sensor coordinates $\mathbf{p}_{x_i} \in \mathbb{R}^3$, for $0 \leq i < N_x$ and a vertical array with sensor coordinates $\mathbf{p}_{y_j} \in \mathbb{R}^3$, for $0 \leq j < N_y$, such that

$$\mathbf{p}_{x_i} = [p_{x_i} \ 0 \ 0]^T \quad \mathbf{p}_{y_j} = [0 \ p_{y_j} \ 0]^T.$$

Let \mathbf{v}_x and \mathbf{v}_y be the $N_x \times 1$ and $N_y \times 1$ manifold vectors for these 1-dimensional arrays. Then for $0 \leq i < N_x$ and $0 \leq j < N_y$,

$$\begin{aligned} [\mathbf{v}(u_x, u_y)]_{i \cdot N_y + j} &= \\ &= e^{j \frac{\omega_k}{c} [u_x \ u_y \ \sqrt{1 - u_x^2 - u_y^2}]} \mathbf{p}_{i \cdot N_y + j} \\ &= e^{j \frac{\omega_k}{c} [u_x \ u_y \ \sqrt{1 - u_x^2 - u_y^2}]} (\mathbf{p}_{x_i} + \mathbf{p}_{y_j}) \\ &= e^{j \frac{\omega_k}{c} [u_x \ 0 \ \sqrt{1 - u_x^2}]} \mathbf{p}_{x_i} e^{j \frac{\omega_k}{c} [0 \ u_y \ \sqrt{1 - u_y^2}]} \mathbf{p}_{y_j} \\ &= [\mathbf{v}_x(u_x)]_i [\mathbf{v}_y(u_y)]_j, \end{aligned}$$

which by definition is equivalent to

$$\mathbf{v}(u_x, u_y) = \mathbf{v}_x(u_x) \otimes \mathbf{v}_y(u_y). \quad (20)$$

Thus, arrays with Cartesian geometries are separable under a U-space parameterization. To prove the converse, note that

$$\begin{aligned} [\mathbf{v}(u_x, u_y)]_i &= e^{j \frac{\omega_k}{c} [u_x \ u_y \ \sqrt{1 - u_x^2 - u_y^2}]} \mathbf{p}_i \\ &= e^{j \frac{\omega_k}{c} (u_x \bar{p}_{x_i} + u_y \bar{p}_{y_i})} \\ &= e^{j \frac{\omega_k}{c} u_x \bar{p}_{x_i}} e^{j \frac{\omega_k}{c} u_y \bar{p}_{y_i}} \end{aligned} \quad (21)$$

where $\mathbf{p}_i = [\bar{p}_{x_i} \ \bar{p}_{y_i} \ 0]^T$. By hypothesis, there exist $\mathbf{a}(u_x)$ and $\mathbf{b}(u_y)$ such that $\mathbf{v}(u_x, u_y) = \mathbf{a}(u_x) \otimes \mathbf{b}(u_y)$. The term $u_x \bar{p}_{x_i}$ from (21) must belong to $\mathbf{a}(u_x)$, since it is a function of u_x and \bar{p}_{x_i} is constant. It follows that $\mathbf{a}(u_x) = \mathbf{v}_x(u_x)$ and $\mathbf{b}(u_y) = \mathbf{v}_y(u_y)$, with $\mathbf{v}_x(u_x)$ and $\mathbf{v}_y(u_y)$ defined above, implying a Cartesian geometry.

A. Fast direct transform

To save space, we will use the shorthand notation

$$\begin{aligned} \mathbf{v}_x(u_{x_m}) &= \mathbf{v}_{x_m} = \left[v_{x_m}^0 \ v_{x_m}^1 \ \dots \ v_{x_m}^{N_x-1} \right]^T \\ \mathbf{v}_y(u_{y_n}) &= \mathbf{v}_{y_n} = \left[v_{y_n}^0 \ v_{y_n}^1 \ \dots \ v_{y_n}^{N_y-1} \right]^T. \end{aligned} \quad (22)$$

Using the separability of the array in (17), we obtain

$$\begin{aligned} \mathbf{A} &= \left[(\mathbf{v}_{x_0}^* \otimes \mathbf{v}_{y_0}^*) \otimes (\mathbf{v}_{x_0} \otimes \mathbf{v}_{y_0}) \ \dots \right. \\ &\quad \left. \dots \ (\mathbf{v}_{x_{M-1}}^* \otimes \mathbf{v}_{y_{M-1}}^*) \otimes (\mathbf{v}_{x_{M-1}} \otimes \mathbf{v}_{y_{M-1}}) \right]. \end{aligned} \quad (23)$$

For $0 \leq m, n < N_x N_y$, the separability of the array also allows row $m \cdot N_x N_y + n$ of \mathbf{A} to be written as

$$\left[v_{x_0}^{i*} v_{x_0}^j \ \dots \ v_{x_{M_x-1}}^{i*} v_{x_{M_x-1}}^j \right] \otimes \left[v_{y_0}^{k*} v_{y_0}^l \ \dots \ v_{y_{M_y-1}}^{k*} v_{y_{M_y-1}}^l \right], \quad (24)$$

where $i = \lfloor \frac{m}{N_y} \rfloor$, $j = \lfloor \frac{n}{N_y} \rfloor$, $k = \text{mod}(m, N_y)$, $l = \text{mod}(n, N_y)$.

For $0 \leq i, j < N_x$ and $0 \leq k, l < N_y$, define

$$c_m(i, j) = v_{x_m}^{i*} v_{x_m}^j \quad d_n(k, l) = v_{y_n}^{k*} v_{y_n}^l. \quad (25)$$

For $0 \leq m, n < N_x N_y$, an arbitrary element $\mathbf{S}_{n,m}$ of \mathbf{S} can be written as the inner product of line $m \cdot N_x N_y + n$ of \mathbf{A} and $\text{vec}\{\mathbf{Y}\}$. Define

$$\begin{aligned} \mathbf{c}(i, j) &= \left[c_0(i, j) \ \dots \ c_{M_x-1}(i, j) \right]^T \\ &= \left[v_{x_0}^{i*} v_{x_0}^j \ \dots \ v_{x_{M_x-1}}^{i*} v_{x_{M_x-1}}^j \right]^T \\ \mathbf{d}(k, l) &= \left[d_0(k, l) \ \dots \ d_{M_y-1}(k, l) \right]^T \\ &= \left[v_{y_0}^{k*} v_{y_0}^l \ \dots \ v_{y_{M_y-1}}^{k*} v_{y_{M_y-1}}^l \right]^T. \end{aligned}$$

Using (24), we have

$$\mathbf{S}_{n,m} = [\mathbf{c}^T(i, j) \otimes \mathbf{d}^T(k, l)] \text{vec}\{\mathbf{Y}\} \quad (26)$$

$$= \mathbf{d}^T(k, l) \mathbf{Y} \mathbf{c}(i, j), \quad (27)$$

where $i = \lfloor \frac{m}{N_y} \rfloor$, $j = \lfloor \frac{n}{N_y} \rfloor$, $k = \text{mod}(m, N_y)$, $l = \text{mod}(n, N_y)$. Also, (26) and (27) are equivalent because $(\mathbf{A}^T \otimes \mathbf{B}) \text{vec}\{\mathbf{C}\} = \text{vec}\{\mathbf{BCA}\}$ whenever \mathbf{BCA} is defined [30].

Note that $\mathbf{c}(i, j)$ and $\mathbf{d}(k, l)$ completely model the response of the (n, m) pair of sensors, for all directions of arrival. This is already a more compact representation than before, since this model uses the separability of the array. All is left is how to efficiently compute all the responses for all pairs.

For $0 \leq i, j < N_x$ and $0 \leq k, l < N_y$, define

$$(i, j) \diamond (k, l) = \mathbf{d}^T(k, l) \mathbf{Y} \mathbf{c}(i, j) \quad (28)$$

and

$$\mathbf{T}_{j,i} = \begin{bmatrix} (i, j) \diamond (0, 0) & \dots & (i, j) \diamond (N_y - 1, 0) \\ (i, j) \diamond (0, 1) & \dots & (i, j) \diamond (N_y - 1, 1) \\ \vdots & & \vdots \\ (i, j) \diamond (0, N_y - 1) & \dots & (i, j) \diamond (N_y - 1, N_y - 1) \end{bmatrix} \quad (29)$$

From the results above, it is easy to show that

$$\mathbf{S} = \begin{bmatrix} \mathbf{T}_{0,0} & \mathbf{T}_{0,1} & \cdots & \mathbf{T}_{0,N_x-1} \\ \mathbf{T}_{1,0} & \mathbf{T}_{1,1} & \cdots & \mathbf{T}_{1,N_x-1} \\ \vdots & \vdots & & \vdots \\ \mathbf{T}_{N_x-1,0} & \mathbf{T}_{N_x-1,1} & \cdots & \mathbf{T}_{N_x-1,N_x-1} \end{bmatrix}. \quad (30)$$

Even though one could determine $(i, j) \diamond (k, l)$ for $0 \leq i, j < N_x$ and $0 \leq k, l < N_y$ by directly evaluating (28), one should organize the computations to eliminate redundancy. Also, since in modern computer architectures the arithmetic units can process data faster than the main memory can provide via random accesses, one should maximize locality of reference to ensure that the arithmetic operands are typically in the cache. In particular, the algorithm should promote sequential memory accesses so that the arithmetic units do not stall while waiting for a memory read. We will present such an implementation below.

Let

$$\mathbf{t}_{i,j} = \text{vec}\{\mathbf{T}_{i,j}\} \quad (31)$$

$$\mathbf{Z} = [\mathbf{t}_{0,0} \ \mathbf{t}_{1,0} \ \cdots \ \mathbf{t}_{N_x-1,N_x-1}]. \quad (32)$$

Given \mathbf{Z} , it is very easy to obtain \mathbf{S} , since every block $\mathbf{T}_{i,j}$ of \mathbf{S} can be obtained by unstacking $\mathbf{t}_{i,j}$.

Define

$$\mathbf{V}_x = \begin{bmatrix} c_0(0,0) & \cdots & c_{M_x-1}(0,0) \\ c_0(0,1) & \cdots & c_{M_x-1}(0,1) \\ \vdots & & \vdots \\ c_0(N_x-1, N_x-1) & \cdots & c_{M_x-1}(N_x-1, N_x-1) \end{bmatrix} \quad (33)$$

$$\mathbf{V}_y = \begin{bmatrix} d_0(0,0) & \cdots & d_{M_y-1}(0,0) \\ d_0(0,1) & \cdots & d_{M_y-1}(0,1) \\ \vdots & & \vdots \\ d_0(N_y-1, N_y-1) & \cdots & d_{M_y-1}(N_y-1, N_y-1) \end{bmatrix} \quad (34)$$

By comparison with (28), one can verify that

$$\mathbf{Z} = \mathbf{V}_y \mathbf{Y} \mathbf{V}_x^T. \quad (35)$$

Define Ξ such that $\text{vec}\{\mathbf{S}\} = \Xi \text{vec}\{\mathbf{Z}\}$ (note that Ξ is a permutation). Thus, $\text{vec}\{\mathbf{S}\} = \Xi (\mathbf{V}_x \otimes \mathbf{V}_y) \text{vec}\{\mathbf{Y}\}$ and

$$\mathbf{A} = \Xi (\mathbf{V}_x \otimes \mathbf{V}_y). \quad (36)$$

Since Ξ is a computationally efficient permutation and $(\mathbf{V}_x \otimes \mathbf{V}_y) \text{vec}\{\mathbf{Y}\} = \text{vec}\{\mathbf{V}_y \mathbf{Y} \mathbf{V}_x^T\}$, (36) can be implemented as a fast transform (see below).

From (30) it can be seen that each $\mathbf{T}_{i,j}$ contains the cross-covariance between two $N_y \times 1$ columns of sensors. Thus, \mathbf{Z} is a reorganization of \mathbf{S} which stacks these cross-covariances with a regularity that matches the row order of $\mathbf{V}_x \otimes \mathbf{V}_y$ (since $\text{vec}\{\mathbf{Z}\} = (\mathbf{V}_x \otimes \mathbf{V}_y) \text{vec}\{\mathbf{Y}\}$).

We now make some remarks regarding computational cost. The direct product $\mathbf{A}\mathbf{y}$ in (16) requires approximately $\frac{1}{2}M_x M_y N_x^2 N_y^2$ complex MACs when considering the Hermitian symmetry \mathbf{S} . Evaluating $(\mathbf{V}_y \mathbf{Y}) \mathbf{V}_x^T$ and $\mathbf{V}_y (\mathbf{Y} \mathbf{V}_x^T)$ requires $N_y^2 M_x M_y + N_x^2 N_y^2 M_x$ and $N_x^2 M_x M_y + N_x^2 N_y^2 M_y$ complex MACs, respectively. Since \mathbf{Y} is real-valued, the first

product can be optimized and the costs drop to $\frac{1}{2}N_y^2 M_x M_y + N_x^2 N_y^2 M_x$ and $\frac{1}{2}N_x^2 M_x M_y + N_x^2 N_y^2 M_y$ complex MACs, respectively. Using the first expression and neglecting the time to obtain \mathbf{S} from \mathbf{Z} , the relative speedup in terms of MACs is given by

$$\frac{\frac{1}{2}M_x M_y N_x^2 N_y^2}{\frac{1}{2}N_y^2 M_x M_y + N_x^2 N_y^2 M_x} = \frac{M_y N_x^2}{M_y + 2N_x^2}.$$

If the array geometry is symmetric with respect to the y axis, then \mathbf{V}_x has conjugate symmetry with respect to its middle row. An analogous statement applies to \mathbf{V}_y . If applicable, these symmetries can be used to further reduce the computational cost.

Recall that we introduced \mathbf{Y} as having scan lines which realize an arbitrary Cartesian sampling of U-space. If $\{u_{x_i}\}$ and $\{u_{y_i}\}$ uniformly sample U-space, then \mathbf{V}_x and \mathbf{V}_y can be interpreted as DFT matrices for non-uniform frequency sampling (this fact can be verified by explicitly writing \mathbf{V}_x and \mathbf{V}_y in terms of complex exponentials). Therefore, for sufficiently large values of N_x and N_y , a further optimization consists of using a fast non-equispaced Fourier transform (NFFT) [28] instead of each matrix product in (35). A rule of thumb obtained from numerical experiments is to use the NFFT for $N_x > 8$ or $N_y > 8$ and $M_x > 2^8$ or $M_y > 2^8$. Details regarding the performance with and without the NFFT are presented in Section V.

B. Fast adjoint transform

As we present in Section II of Part II, given a measured spectral matrix \mathbf{S} , many computationally efficient image reconstruction methods require only fast implementations of \mathbf{A} and \mathbf{A}^H to estimate a source distribution $\hat{\mathbf{Y}}$. A computationally efficient reconstruction algorithm must have fast implementations of both, otherwise the slow transform becomes the bottleneck for the solver.

Let $\bar{\mathbf{S}} \in \mathbb{C}^{N \times N}$ and $\bar{\mathbf{Y}} \in \mathbb{R}^{M_y \times M_x}$ such that $\text{vec}\{\bar{\mathbf{Y}}\} = \mathbf{A}^H \text{vec}\{\bar{\mathbf{S}}\}$. It follows from (36) that

$$\mathbf{A}^H = (\mathbf{V}_x^H \otimes \mathbf{V}_y^H) \Xi^H. \quad (37)$$

Since Ξ is a permutation, $\Xi^{-1} = \Xi^T = \Xi^H$. If $\text{vec}\{\bar{\mathbf{Z}}\} = \Xi^T \text{vec}\{\bar{\mathbf{S}}\}$, then

$$\bar{\mathbf{Y}} = \mathbf{V}_y^H \bar{\mathbf{Z}} \mathbf{V}_x^*, \quad (38)$$

which is the fast implementation of \mathbf{A}^H (note that it has the same computational cost as the direct transform).

If the U-space sampling is symmetric with respect to the y axis, then \mathbf{V}_x has conjugate symmetry with respect to its center column. An analogous statement applies to \mathbf{V}_y . If applicable, this symmetry can be used to further reduce the computational cost.

For separable arrays which are uniformly sampled in U-space, multiplication by \mathbf{V}_x and \mathbf{V}_y can again be optimized using NFFTs, under the same considerations presented for the direct transform.

Table II
APPLICABILITY OF THE FFT, NFFT, NNFFT AND KAT

Transform	Geometry	U-space sampling	Exact?
FFT	Uniform	Uniform	Yes
NFFT	Arbitrary	Uniform	No
NNFFT	Arbitrary	Arbitrary	No
KAT	Separable	Separable	Yes

C. Fast direct-adjoint transform

Given the direct transform \mathbf{A} and its adjoint \mathbf{A}^H , consider the transform given by $\mathbf{A}^H \mathbf{A}$. This composition will be used in Section II of Part II for image reconstruction, and in this section we present a method of accelerating it further. Using that $\Xi^H = \Xi^{-1}$, it follows from the previous results that $\text{vec}\{\bar{\mathbf{Y}}\} = \mathbf{A}^H \mathbf{A} \text{vec}\{\mathbf{Y}\}$ can be implemented as

$$\bar{\mathbf{Y}} = \mathbf{V}_y^H \mathbf{V}_y \mathbf{Y} \mathbf{V}_x^T \mathbf{V}_x^* \quad (39)$$

This implementation is especially interesting when N_x and N_y are sufficiently large in comparison to M_x and M_y , because it can be evaluated as

$$\bar{\mathbf{Y}} = (\mathbf{V}_y^H \mathbf{V}_y) \mathbf{Y} (\mathbf{V}_x^T \mathbf{V}_x^*), \quad (40)$$

with precomputed versions of $\mathbf{V}_y^H \mathbf{V}_y$ and $\mathbf{V}_x^T \mathbf{V}_x^*$, which are real-valued.

Implementing the direct-adjoint transform with (39) can be much faster than using a composition of the direct and adjoint KAT, because for large problems one can precompute $\mathbf{V}_y^H \mathbf{V}_y$ and $\mathbf{V}_x^T \mathbf{V}_x^*$. Furthermore, (39) can be parallelized more effectively, since it avoids applying Ξ .

The implementations which use the NFFT for further acceleration are at a disadvantage for the direct-adjoint transform, since one cannot precompute the equivalent of $\mathbf{V}_y^H \mathbf{V}_y$ and $\mathbf{V}_x^T \mathbf{V}_x^*$. Therefore, one is forced to use a composition of the previously presented transforms.

IV. CONNECTIONS

In this section we briefly describe how the KAT relates to the two-dimensional FFT, NFFT and NNFFT. To our knowledge, the NFFT and NNFFT have never been applied to acoustic imaging. With the exception of the NNFFT, each transform is only suitable for specific array geometries or U-space sampling patterns. Transforms which make more restrictive assumptions about the array geometry and U-space pattern can generally be more computationally efficient, so the choice of which transform to use depends on a series of tradeoffs, summarized in Table II.

Assume that the sampled wavefield is a zero-mean random process which is stationary in time and homogenous in space². Consider a sensor array consisting of N microphones at coordinates $\mathbf{p}_0, \dots, \mathbf{p}_{N-1} \in \mathbb{R}^3$. For a fixed frequency ω , the cross spectral matrix $\mathbf{S}(\omega) = \mathbb{E}\{\mathbf{x}(\omega) \mathbf{x}^H(\omega)\}$ is by definition a covariance matrix. For $0 \leq m, n < N$,

²The random process model follows naturally from the fact that we do not know a priori what the source waveforms are. We model this process using its second-order statistics given by the cross-spectral matrix. Stationarity in time and homogeneity in space let us estimate the power spectral density of the process as a function of ω and \mathbf{k} , which is the acoustic image of interest.

$[\mathbf{S}(\omega)]_{m,n}$ holds the frequency-domain cross-covariance of the wavefield between any two points whose coordinates differ by $\mathbf{p}_m - \mathbf{p}_n$. Let $P(\omega, \mathbf{k}) = |Y(\omega, \mathbf{k})|^2$ be the power spectral density when parameterized as a function of the wavenumber $\mathbf{k} = -\frac{\omega}{c} \mathbf{u} \in \mathbb{R}^3$, and $S(\omega, \Delta \mathbf{p})$ be the spectral covariance between two points whose coordinates differ by $\Delta \mathbf{p}$. It can be shown [29] that

$$P(\omega, \mathbf{k}) = \int_{\mathbb{R}^3} S(\omega, \Delta \mathbf{p}) e^{+j\mathbf{k}^T \Delta \mathbf{p}} d\Delta \mathbf{p} \quad (41)$$

$$S(\omega, \Delta \mathbf{p}) = \frac{1}{(2\pi)^3} \int_{\mathbb{R}^3} P(\omega, \mathbf{k}) e^{-j\mathbf{k}^T \Delta \mathbf{p}} d\mathbf{k}, \quad (42)$$

which is essentially a generalization of the relationship between cross-covariance and cross-spectral density for wide-sense stationary spatial-temporal processes, and is expressed as a Fourier transform.

Therefore, the knowledge of $S(\omega, \Delta \mathbf{p})$, limited to a finite set of baselines $\Delta \mathbf{p}$, allows us to approximate a discrete space version of $P(\omega, \mathbf{k})$, which is the image of interest. The following connections arise naturally from different ways of sampling these relations, in order to evaluate them numerically for discrete space and discrete U-space.

A. NFFT imaging

A d-dimensional NDFT [28] (non-equispaced discrete Fourier transform) 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 $\mathcal{X} = \left\{x_i \in \left[-\frac{1}{2}, \frac{1}{2}\right]^d : 0 \leq i < N\right\}$ such that $|\mathcal{X}| = N$, where $|\cdot|$ indicates set cardinality. The index set

$$I_{\mathbf{N}} = \mathbb{Z}^d \cap \prod_{t=0}^{d-1} \left[-\frac{M_t}{2}, \frac{M_t}{2} \right), \quad (43)$$

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

$$\hat{h}_i = \sum_{\mathbf{k} \in I_{\mathbf{N}}} h_{\mathbf{k}} e^{-j2\pi \mathbf{k}^T \mathbf{x}_i}, \quad (44)$$

for $0 \leq i < N$.

The NFFT is a fast implementation of the NDFT obtained by interpolating an oversampled FFT. It is an approximate method which provides a very good compromise between accuracy and computational complexity.

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

$$\tilde{\mathbf{U}} = \left\{ \frac{2i}{M_x} \right\}_{i=-M_x/2}^{M_x/2-1} \times \left\{ \frac{2j}{M_y} \right\}_{j=-M_y/2}^{M_y/2-1}. \quad (45)$$

We now show that \mathbf{S} can be obtained from \mathbf{Y} using a 2D NFFT. Due to the linearity of the NFFT, it suffices to show that this transform is exact for an image containing one arbitrary unit impulse at coordinates \mathbf{u}_0 , which must be in the U-space sampling grid.

Using (9), for arbitrary $-\frac{M_x}{2} \leq m_0 < \frac{M_x}{2}$ and $-\frac{M_y}{2} \leq n_0 < \frac{M_y}{2}$, define

$$\mathbf{u}_0 = \begin{bmatrix} \frac{2m_0}{M_x} \\ \frac{2n_0}{M_y} \\ \sqrt{1 - \left(\frac{2m_0}{M_x}\right)^2 - \left(\frac{2n_0}{M_y}\right)^2} \end{bmatrix} \quad (46)$$

and

$$\mathbf{Y}_{n,m} = \begin{cases} 1 & \text{if } (m, n) = (m_0 + \frac{M_x}{2}, n_0 + \frac{M_y}{2}) \\ 0 & \text{otherwise.} \end{cases} \quad (47)$$

By definition,

$$\mathbf{S}_{r,s} = [\mathbf{v}(\mathbf{u}_0) \mathbf{v}^H(\mathbf{u}_0)]_{r,s} \quad (48)$$

$$= e^{j\frac{\omega}{c} \mathbf{u}_0^T (\mathbf{p}_r - \mathbf{p}_s)} \quad (49)$$

$$= e^{j2\pi \mathbf{u}_0^T \frac{(\mathbf{p}_r - \mathbf{p}_s)}{\lambda}} \quad (50)$$

$$= e^{j2\pi \frac{\mathbf{u}_0^T}{\lambda} \cdot \frac{2(\mathbf{p}_r - \mathbf{p}_s)}{\lambda}}. \quad (51)$$

To obtain (51) using the NFFT, rectangular U-space sampling and an arbitrary geometry of N microphones we use

$$\mathbf{M} = [M_x \ M_y]^T \quad (52)$$

$$I_{\mathbf{N}} = \mathbb{Z}^2 \cap \left[-\frac{M_x}{2}, \frac{M_x}{2} \right) \times \left[-\frac{M_y}{2}, \frac{M_y}{2} \right) \quad (53)$$

$$\mathcal{X} = \left\{ \mathbf{x}_i = \frac{2}{\lambda} (\mathbf{p}_{\lfloor i/N \rfloor} - \mathbf{p}_{\text{mod}(i,N)}) \odot \begin{bmatrix} M_x^{-1} \\ M_y^{-1} \end{bmatrix} : 0 \leq i < N^2 \right\} \quad (54)$$

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 the direct transform.

Once again,

$$\begin{aligned} \mathbf{S}_{r,s} &= e^{j2\pi \frac{\mathbf{u}_0^T}{\lambda} \cdot \frac{2(\mathbf{p}_r - \mathbf{p}_s)}{\lambda}} \\ &= e^{j2\pi (\mathbf{u}_0^T \odot [\frac{M_x}{2} \ \frac{M_y}{2}]) \left(\frac{2}{\lambda} (\mathbf{p}_r - \mathbf{p}_s) \odot [M_x^{-1} \ M_y^{-1}]^T \right)}, \end{aligned} \quad (55)$$

where \mathbf{u}_0 is also represented only by its x and y components.

Comparing (55) with (44), the first term in parentheses clearly belongs to $I_{\mathbf{N}}$. Since for $0 \leq i < N^2$, $(\mathbf{p}_{\lfloor i/N \rfloor} - \mathbf{p}_{\text{mod}(i,N)})$ spans all possible baselines, the second term in parentheses belongs to \mathcal{X} . The enumeration given by $(\lfloor i/N \rfloor, \text{mod}(i, N))$ indexes the elements of \mathbf{S} row by row. Given the Hermitian symmetry of \mathbf{S} , this is equivalent to conjugating (55) and indexing the elements of \mathbf{S} column by column (in the order of $\text{vec}\{\cdot\}$), making (55) equivalent to (44).

It is possible to 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), this NFFT reduces to 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, which have their upper operating frequency constrained by the Nyquist-Shannon sampling theorem.

The NFFT has the advantage of allowing arbitrary array geometries, but as we will see in Section V, it is one order of magnitude slower than the KAT. Furthermore, as will be shown in Section VI, the KAT can be generalized to approximate the spherical wave fronts due to near-field sources. In contrast, the FFT, NFFT and NFFT require a far-field assumption.

Finally, the KAT has the advantage of allowing separable (as opposed to uniform) U-space sampling grids. Acoustic images are often formed by clusters of distributed sources (for example, located over a model in a windtunnel) and large regions with no significant sources. Thus, the KAT allows one to oversample the regions which are expected to have sources and undersample quiet regions, while maintaining low computational requirements.

B. NNFFT imaging

By dropping the uniform sampling constraint (45) one obtains the far-field array transform in its fullest generality. This transform can be accelerated with the non-equispaced in time and frequency fast Fourier transform (NNFFT). While the NNFFT is significantly slower than the KAT and the NFFT, it requires much less memory than the matrix representation of \mathbf{A} , which makes it useful for smaller problems that can be solved offline.

Given as input a set of samples $h_l \in \mathbb{C}$ for $0 \leq l < L$, the NNFFT is defined as

$$\hat{h}_i = \sum_{l=0}^{L-1} h_l e^{-j2\pi (\mathbf{v}_l \odot \mathbf{M})^T \mathbf{x}_i}, \quad (56)$$

for $0 \leq i < N$, and arbitrary $\mathbf{v}_l, \mathbf{x}_i \in [-\frac{1}{2}, \frac{1}{2}]^d$. The NNFFT is a fast approximation of the NNFFT.

To obtain (51) using the NNFFT, arbitrary U-space sampling and an arbitrary geometry of N microphones we use $x_j \in \mathcal{X}$ and $v_l \in \mathcal{V}$, with

$$\mathbf{M} = [M_x \ M_y]^T \quad (57)$$

$$\mathcal{V} = \left\{ v_l = \left(\frac{u_{x_l}}{2}, \frac{u_{y_l}}{2} \right) \in \left[-\frac{1}{2}, \frac{1}{2} \right)^2 : 0 \leq l < M \right\} \quad (58)$$

$$\mathcal{X} = \left\{ \mathbf{x}_i = \frac{2}{\lambda} (\mathbf{p}_{\lfloor i/N \rfloor} - \mathbf{p}_{\text{mod}(i,N)}) \odot \begin{bmatrix} M_x^{-1} \\ M_y^{-1} \end{bmatrix} : 0 \leq i < N^2 \right\} \quad (59)$$

which has the same form as (52)-(54), but allows arbitrary U-space sampling.

V. COMPUTATIONAL COST

A. Asymptotic complexity

To simplify the following formulas, we will assume that $N_x = N_y$ and $M_x = M_y$. We will present the asymptotic complexity for the direct and adjoint transforms.

Table III
ASYMPTOTIC COMPLEXITY OF THE KAT, FFT, NFFT, NNFFT AND EXPLICIT MATRIX REPRESENTATION.

Transform	Complexity
KAT with matrix multiplication	$O(MN + M^{1/2}N^2)$
KAT with 1-D NFFTs	$O(M \log M + M^{1/2}N)$
2-D FFT	$O(M \log M)$
2-D NFFT	$O(M \log M + N^2)$
2-D NNFFT	$O(M \log M + N^2)$
Explicit matrix representation	$O(MN^2)$

The product $\mathbf{A}\mathbf{y}$ in (16) requires approximately $\frac{1}{2}MN^2$ complex MACs to compute, and thus has complexity $O(MN^2)$. For the KAT, the cost of computing Ξ can be neglected. Evaluating $(\mathbf{V}_y\mathbf{Y})\mathbf{V}_x^T$ with matrix multiplication requires $\frac{1}{2}MN + M^{1/2}N^2$ complex MACs, and thus has complexity $O(MN + M^{1/2}N^2)$.

One can also evaluate $(\mathbf{V}_y\mathbf{Y})\mathbf{V}_x^T$ by interpreting each matrix product as a 1-D NFFT. Evaluating $\mathbf{V}_y\mathbf{Y}$ requires M_x 1-D NFFTs, each with cost $O(M_y \log M_y + N_y^2)$ [28]. The second product requires N_y^2 1-D NFFTs, each with cost $O(M_x \log M_x + N_x^2)$. Assuming that $M^{1/2} > N$, the total complexity becomes $O(M \log M + M^{1/2}N)$. The direct NFFT and NNFFT implementations both have complexity $O(M \log M + N^2)$ [28].

Table III summarizes these results. Note that since the asymptotic complexity is similar for most of the fast transforms, it simply guarantees that these methods will scale about as well as an FFT. Nevertheless, the constants hidden in the $O(\cdot)$ notation are significant. As we show next, the direct NNFFT implementation is much slower than an KAT with 1-D NFFTs, despite having a similar asymptotic complexity. Also, since the \mathbf{A} matrix is very large, memory bandwidth becomes the limiting factor for the explicit matrix representation. Thus, the constant hiding in the $O(\cdot)$ notation for the explicit matrix representation is greater than the MAC count suggests. Furthermore, for practical problem sizes one does not have enough memory to store a full matrix representation and is forced to recompute the rows of \mathbf{A} every time a matrix-vector product is required. This can dramatically increase the computational cost of the explicit matrix representation.

B. Numerical benchmarks

This section presents experiments to assess the execution times for the KAT, the NFFT and the NNFFT. Even though the relative performance of algorithms based on matrix multiplication can be easily estimated in terms of MACs, actual runtimes can deviate significantly from these estimates for certain problem sizes. Indeed, for modern architectures, performance is strongly dependent on the interaction of parallel arithmetic units, memory bandwidth, cache size and branch prediction, such that the number of floating point operations only serves as an approximate measure of computational complexity.

The runtimes presented in Figs. 2 and 3 are averages collected over 10 seconds for each algorithm and problem size. All simulations were run on an Intel Core 2 Duo T9400 processor in 64-bit mode, using only one core. The permutation Ξ which obtains \mathbf{S} from \mathbf{Z} was written in ANSI C, the NFFT

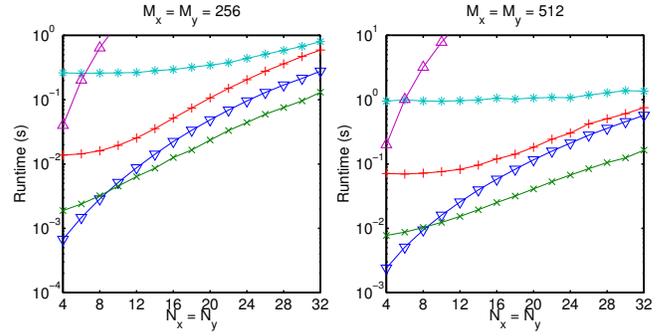


Figure 2. Runtimes for the direct transform. ∇ : KAT implemented with matrix multiplication, \times : KAT implemented with 1-D NFFTs replacing matrix multiplication, $+$: direct NFFT implementation with (52)-(54), $*$: direct NNFFT implementation with (57)-(59), \triangle : explicit matrix representation.

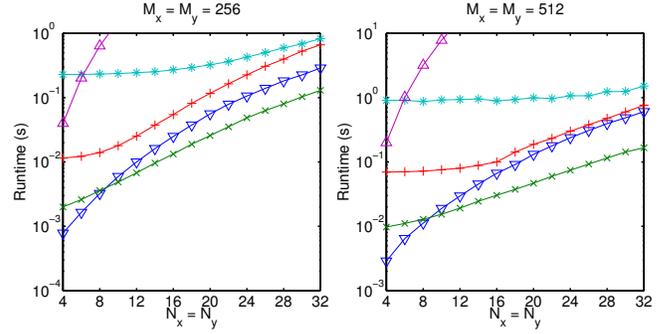


Figure 3. Runtimes for the adjoint transform. ∇ : KAT implemented with matrix multiplication, \times : KAT implemented with 1-D NFFTs replacing matrix multiplication, $+$: direct NFFT implementation with (52)-(54), $*$: direct NNFFT implementation with (57)-(59), \triangle : explicit matrix representation.

library was compiled with default optimizations as used by its authors, and all other functions were written in M-code for MATLAB R2008b. Since the code does not feature time consuming loops and MATLAB uses the Intel Math Kernel Library for matrix and vector arithmetic, the proposed transforms run very much like machine-specific tuned code. MATLAB and the NFFT use FFTW [31] for computing FFTs, such that they also run like machine-specific code. Thus, having the code written in MATLAB actually incurs negligible computational overhead when compared to an optimized implementation in C and machine-specific assembly code.

The computational cost of efficient convex optimization methods (and in particular, of the regularized least-squares methods presented in Part II) is completely dependent on the cost of applying \mathbf{A} , \mathbf{A}^H , Ψ and Ψ^H . Since Ψ depends on the regularization method and can be chosen to be very fast, the bottleneck is on applying \mathbf{A} , \mathbf{A}^H and possibly $\mathbf{A}^H\mathbf{A}$. Thus, from the runtimes of \mathbf{A} , \mathbf{A}^H and $\mathbf{A}^H\mathbf{A}$ presented in this section, one can assume with good approximation that an K -fold decrease in computational time translates to an algorithm which reconstructs an image K times faster.

It is clear that the KAT with the NFFT optimization is the fastest transform for arrays with more than 64 elements ($N_x > 8$ and $N_y > 8$). This is the case because as \mathbf{V}_x and \mathbf{V}_y grow, the NFFT scales better than matrix multiplication. The direct implementation with the NFFT is useful if one must have an

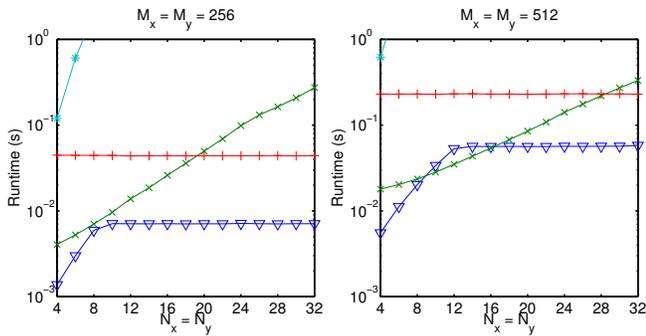


Figure 4. Runtimes for the direct-adjoint composition. ∇ : KAT implemented with (39), \times : KAT implemented with the composition \mathbf{A} and \mathbf{A}^H , with the 1-D NFFT optimization, $+$: 2D FFT-accelerated convolution, $*$: composition of the explicit matrix representation.

arbitrary array geometry, but it has the drawback of being around an order of magnitude slower and requiring a far-field approximation (as we show in Section VI, the KAT can be extended for near-field imaging). The direct implementation with the NNFFT is by far the slowest. For all implementations, the direct and adjoint transforms perform similarly.

Fig. 4 presents runtimes for the direct-adjoint composition. The implementation using (39) analyzes the problem size and automatically selects the optimal order for matrix multiplication. For large values of N_x and N_y , it also uses precomputed versions of $\mathbf{V}_y^H \mathbf{V}_y$ and $\mathbf{V}_x^T \mathbf{V}_x^*$, which makes the computational complexity depend only on M_x and M_y . The NFFT implementation uses a composition of the direct and adjoint transforms, without any additional optimizations.

As is shown in Part II, under a far-field assumption and uniform U-space sampling, the direct-adjoint composition reduces to a 2D convolution of the input image with the array point-spread function. Thus, it can be accelerated with a 2D FFT (with zero padding to prevent edge effects). Fig. 4 shows that the direct-adjoint composition implemented with the KAT also outperforms 2D FFT accelerated convolutions. As we show in the following section, the KAT can be generalized to near-field scenarios, allowing us to drop the far-field assumption. Note that for near-field cases, the direct-adjoint composition no longer reduces to a convolution, and KAT becomes the only fast transform suitable for imaging.

Finally, the KAT has the additional advantage of being easy to implement and parallelize, since it only requires relatively small matrix multiplications and computationally efficient permutations.

VI. NEAR-FIELD IMAGING

Up to this point, we have assumed that the sources were located in the far-field. Thus, we used a plane wave model. In this section, we show how to generalize the KAT and address near-field scenarios, where one has spherical wavefronts.

Note that the KAT does not impose any structure onto the array manifold vector other than its separability. The specific far-field representation was only chosen for convenience, since for any Cartesian geometry, the far-field array manifold vector is separable. Nothing prevents us from choosing a different

separable representation that is more suitable for the near-field case. In this section, we show that the problem of finding the best separable representation can be recast as a rank-1 approximation of a rearranged version of $\mathbf{\Xi}^{-1} \mathbf{A} = \mathbf{\Xi}^T \mathbf{A}$. By using a rank- K approximation (for $K > 1$), one is able to obtain an arbitrarily accurate model for near-field propagation, while maintaining low computational requirements.

To simplify the language in this section, we will use the following notation. Given suitably sized matrices \mathbf{A} , \mathbf{B} , \mathbf{C} , we use $\mathbf{C} = \underline{\mathbf{A}}(\mathbf{B})$ to denote $\text{vec}\{\mathbf{C}\} = \mathbf{A} \text{vec}\{\mathbf{B}\}$.

We approximate \mathbf{A} by $\check{\mathbf{A}} = \mathbf{\Xi} \left(\sum_{k=1}^K \mathbf{C}_k \otimes \mathbf{D}_k \right)$, for small values of K . Note that $\mathbf{S} = \check{\mathbf{A}}(\mathbf{Y})$ can be efficiently implemented as $\mathbf{S} = \mathbf{\Xi} \left(\sum_{k=1}^K \mathbf{D}_k \mathbf{Y} \mathbf{C}_k^T \right)$. Comparing with (36), we are approximating the near-field transform \mathbf{A} with a series of K separable transforms, to which we can apply the KAT.

Let us consider the problem of approximating a generic $\mathbf{B} \in \mathbb{C}^{m \times n}$ with $m = m_1 m_2$ and $n = n_1 n_2$ with a sum of Kronecker products, such that

$$\min_{\{\mathbf{C}_k\}, \{\mathbf{D}_k\}} \left\| \mathbf{B} - \sum_{k=1}^K \mathbf{C}_k \otimes \mathbf{D}_k \right\|_F,$$

where $\mathbf{C}_k \in \mathbb{C}^{m_1 \times n_1}$ and $\mathbf{D}_k \in \mathbb{C}^{m_2 \times n_2}$ for $1 \leq k \leq K$. This problem is addressed in [32], where it is shown to be equivalent to

$$\min_{\{\mathbf{C}_k\}, \{\mathbf{D}_k\}} \left\| \mathcal{R}(\mathbf{B}) - \sum_{k=1}^K \text{vec}\{\mathbf{C}_k\} \otimes \text{vec}\{\mathbf{D}_k\} \right\|_F, \quad (60)$$

where $\mathcal{R}(\cdot)$ is a matrix rearrangement operator such that $\mathcal{R}(\mathbf{B}) \in \mathbb{C}^{m_1 n_1 \times m_2 n_2}$. This is a low-rank approximation problem which can be solved with the SVD of $\mathcal{R}(\mathbf{B})$.

For our purposes, we approximate $\mathbf{B} = \mathbf{\Xi}^{-1} \mathbf{A}$. We note that $\mathbf{\Xi}$ is the key to a successful low-rank decomposition. As we show later in this section, using $\mathbf{B} = \mathbf{A}$ is not useful, since $\mathcal{R}(\mathbf{A})$ has too many significant singular values.

Computing the dominant singular values and vectors of $\mathcal{R}(\mathbf{B})$ is not trivial, since in practice $\mathcal{R}(\mathbf{B})$ is too large to be stored explicitly in memory. Nevertheless, one can use Lanczos methods [32], [33] which only require the implementation of the matrix-vector products $\mathcal{R}(\mathbf{B}) \alpha$ and $\mathcal{R}(\mathbf{B})^H \beta$ for arbitrary α, β . One can also use approximate SVD methods which are designed to require a small number of passes over $\mathcal{R}(\mathbf{B})$ (e.g. [34], [35]).

Using the definition of $\mathcal{R}(\cdot)$ from [32], it can be shown that

$$\mathcal{R}(\mathbf{B})^T = \begin{bmatrix} \mathbf{Z}_{0,0} & \cdots & \mathbf{Z}_{M_x-1,0} \\ \vdots & & \vdots \\ \mathbf{Z}_{0,M_y-1} & \cdots & \mathbf{Z}_{M_x-1,M_y-1} \end{bmatrix} \quad (61)$$

$$\mathbf{Z}_{m,n} = \mathbf{\Xi} \left(\mathbf{v}(u_{x_m}, u_{y_n}) \mathbf{v}^H(u_{x_m}, u_{y_n}) \right). \quad (62)$$

Since $\mathbf{v}(u_{x_m}, u_{y_n})$ can be precomputed for $0 < m < M_x$ and $0 < n < M_y$ and $\mathbf{\Xi}$ is a very fast permutation, $\mathcal{R}(\mathbf{B}) \alpha$ and $\mathcal{R}(\mathbf{B})^H \beta$ can be evaluated with relative efficiency. Indeed, using the Lanczos method from [33], $N = 64$ and $M_x = M_y = 256$, we can solve (60) for $K = 8$ in 8 minutes on

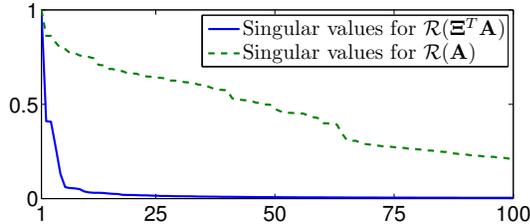


Figure 5. First 100 singular values for $\mathcal{R}(\mathbf{A})$ and $\mathcal{R}(\Xi^T \mathbf{A})$, out of a total of 16384, normalized to 1.

an Intel Core 2 Duo 2.4 GHz processor, using only one core. Note that this procedure only has to be performed once.

The decomposition obtained with (60) is especially useful for $K > 1$. Indeed, even in the presence of strong near field effects, $\mathcal{R}(\mathbf{B})$ can be well approximated by a low-rank decomposition. Even though the transform cost grows linearly with K , due to the Kronecker representation, the cost of applying each $\mathbf{C}_k \otimes \mathbf{D}_k$ is very small, so a transform with $K = 8$ is still very fast.

Since we discretize the focal surface with Cartesian sampling grids, some images may have invalid regions. For example, one can parameterize a spherical half-shell with radius r_0 using $u_x = r_0 \sin \phi \cos \theta$ and $u_y = r_0 \sin \phi \sin \theta$ for $\|\mathbf{u}\| \leq r_0$, such that source locations for $\|\mathbf{u}\| > r_0$ are invalid. Note that this implies that some $\mathbf{Z}_{m,n}$ in $\mathcal{R}(\mathbf{B})$ are not defined. Referring to (60), the optimal $\mathbf{C}_k, \mathbf{D}_k$ are now given by

$$\min_{\{\mathbf{C}_k\}, \{\mathbf{D}_k\}} \left\| \mathbf{W} \odot \left(\mathcal{R}(\mathbf{B}) - \sum_{k=1}^K \text{vec}\{\mathbf{C}_k\} \otimes \text{vec}\{\mathbf{D}_k\} \right) \right\|_F,$$

where \odot is the pointwise (Hadamard) product and \mathbf{W} is a binary mask set to 1 for valid elements and 0 for invalid elements of $\mathcal{R}(\mathbf{B})$. This masked-SVD problem was considered in [36], and can be solved by iterating

$$\mathbf{E}^{(i)} = \text{LRA}_K \left(\mathbf{W} \odot \mathcal{R}(\mathbf{B}) + (\mathbf{1} - \mathbf{W}) \odot \mathbf{E}^{(i-1)} \right),$$

where $\text{LRA}_K(\cdot)$ is a rank- K approximation as computed by the SVD, $\mathbf{E}^{(i)} = \sum_{k=1}^K \text{vec}\{\mathbf{C}_k^{(i)}\} \otimes \text{vec}\{\mathbf{D}_k^{(i)}\}$ and $\mathbf{E}^{(0)} = \mathbf{0}$. In our experiments, 1 or 2 iterations have shown to be sufficient for a good fit.

Fig. 5 compares the first 100 (out of a total of 16384) singular values for $\mathcal{R}(\mathbf{A})$ and $\mathcal{R}(\Xi^T \mathbf{A})$. In this experiment, \mathbf{A} models 256×256 sources radiating at 9 kHz, located over a spherical half-shell with a radius of 1.0 m, as measured by an 8×8 separable array with a $30 \text{ cm} \times 30 \text{ cm}$ aperture (this configuration is the same used in the examples from Part II). The sharp decay of the curve for $\mathcal{R}(\Xi^T \mathbf{A})$ highlights the importance of Ξ in enabling an accurate low-rank approximation. For lower frequencies, the singular values show an even sharper roll-off. Also, this transform models the complete hemisphere. By modeling a smaller field of view, near-field effects are not as severe and the singular values also decay faster.

Note that by using a rank- K approximation, we obtain a transform with a computational cost that is K times larger than the far-field KAT presented in the previous sections. Nevertheless, as we show in Part II, K will be small enough

that this penalty is not significant. In fact, we will show that it is possible to compensate for strong near-field effects with $K = 8$, which makes the KAT about as fast as the NFFT, while being able to model arbitrary near-field focal surfaces.

VII. CONCLUSION

This paper presents the Kronecker array transform (KAT), which was designed to enable computationally efficient and accurate acoustic imaging. To obtain a fast implementation, we assumed a separable microphone array, source parameterization in U-space and far-field sources. The KAT transforms back and forth between a hypothetical source distribution and the corresponding covariance matrix which would be measured by the array. Another contribution is the novel application of the NFFT and NNFFT to acoustic imaging, which we used as baselines against which our proposal was compared.

The KAT is orders of magnitude faster than equivalent implementations using explicit matrix representations. Despite the fact that the NFFT and NNFFT are already fast transforms, we have shown in Section V that our proposal is at least one order of magnitude faster than them. In contrast with FFT- and NFFT-based transforms, the KAT allows arbitrary Cartesian samplings of the source distributions, which let one oversample regions with sources and undersample silent areas without performance degradation. Also in contrast with the NFFT, the KAT makes no numerical approximations, and can be more easily implemented and parallelized, since it only requires relatively small matrix products and simple permutations.

Even though the KAT was motivated with the far-field assumption, it does not impose any structure onto the array manifold vector other than its separability. We have used this fact to extend it for near-field imaging, providing a computationally efficient approximation of the exact near-field transform.

Future work involves developing fast transforms for other array geometries. In contrast with the KAT, which is exact for far-field sources, transforms for other geometries will most likely require approximations to obtain good performance. Furthermore, the KAT can be generalized for correlated distributions, which we also intend to address in future work.

REFERENCES

- [1] W. Home, K. James, T. Arledge, P. Sodermant, N. Burnside, and S. 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] W. Humphreys and T. Brooks, "Noise Spectra and Directivity For Scale-Model Landing Gear," in *Proc. of the 13th AIAA/CEAS Aeroacoustics Conference*, 2007.
- [3] S. Lee, "Phased-Array Measurement of Modern Regional Aircraft Turbofan Engine Noise," in *Proc. of the 12th AIAA/CEAS Aeroacoustics Conference*, 2006.
- [4] S. Oerlemans, P. Sijtsma, and B. Mendez Lopez, "Location and quantification of noise sources on a wind turbine," *Journal of Sound and Vibration*, vol. 299, no. 4-5, pp. 869–883, 2007.
- [5] 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.
- [6] M. Wax and T. Kailath, "Optimum localization of multiple sources by passive arrays," *IEEE Transactions on Acoustics, Speech and Signal Processing*, vol. 31, no. 5, pp. 1210–1217, 1983.

- [7] R. Schmidt, "Multiple emitter location and signal parameter estimation," *IEEE Transactions on Antennas and Propagation*, vol. 34, no. 3, pp. 276–280, 1986.
- [8] R. Roy and T. Kailath, "ESPRIT-estimation of signal parameters via rotational invariance techniques," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 37, no. 7, pp. 984–995, 1989.
- [9] C. Zhang, D. Florencio, D. Ba, and Z. Zhang, "Maximum Likelihood Sound Source Localization and Beamforming for Directional Microphone Arrays in Distributed Meetings," *IEEE Transactions on Multimedia*, vol. 10, no. 3, pp. 538–548, 2008.
- [10] M. Brandstein and H. Silverman, "A robust method for speech signal time-delay estimation in reverberant rooms," in *Proc. of ICASSP*, vol. 1, 1997.
- [11] Z. Wang, J. Li, P. Stoica, T. Nishida, and M. Sheplak, "Constant-beamwidth and constant-powerwidth wideband robust Capon beamformers for acoustic imaging," *The Journal of the Acoustical Society of America*, vol. 116, p. 1621, 2004.
- [12] S. Yan, Y. Ma, and C. Hou, "Optimal array pattern synthesis for broadband arrays," *The Journal of the Acoustical Society of America*, vol. 122, p. 2686, 2007.
- [13] J. Li, Y. Xie, P. Stoica, X. Zheng, and J. Ward, "Beampattern synthesis via a matrix approach for signal power estimation," *IEEE Transactions on Signal Processing*, vol. 55, no. 12, pp. 5643–5657, 2007.
- [14] 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, p. 757, 2004.
- [15] R. Dougherty and R. Stoker, "Sidelobe suppression for phased array aeroacoustic measurements," in *Proc. of the 4th AIAA/CEAS Aeroacoustics Conference*, 1998, pp. 235–245.
- [16] R. Dougherty, "Extensions of DAMAS and Benefits and Limitations of Deconvolution in Beamforming," in *Proc. of the 11th AIAA/CEAS Aeroacoustics Conference*, 2005.
- [17] K. Ehrenfried and L. Koop, "Comparison of iterative deconvolution algorithms for the mapping of acoustic sources," *AIAA Journal*, vol. 45, no. 7, p. 1584, 2007.
- [18] 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.
- [19] 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.
- [20] D. Donoho, "Compressed sensing," *IEEE Transactions on Information Theory*, vol. 52, no. 4, pp. 1289–1306, 2006.
- [21] F. Ribeiro and V. Nascimento, "Fast Transforms for Acoustic Imaging – Part II: Applications and Extensions," 2010, submitted.
- [22] L. He, T. Chang, S. Osher, T. Fang, and P. Speier, "MR image reconstruction by using the iterative refinement method and nonlinear inverse scale space methods," *UCLA CAM Report*, vol. 6, p. 35, 2006.
- [23] M. Lustig, D. Donoho, and J. Pauly, "Sparse MRI: The application of compressed sensing for rapid MR imaging," *Magnetic Resonance in Medicine*, vol. 58, no. 6, pp. 1182–1195, 2007.
- [24] D. Malioutov, M. Cetin, and A. Willsky, "A sparse signal reconstruction perspective for source localization with sensor arrays," *IEEE Transactions on Signal Processing*, vol. 53, no. 8, pp. 3010–3022, 2005.
- [25] 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.
- [26] B. Zimmermann and C. Studer, "FPGA-based real-time acoustic camera prototype," in *Proc. of ISCAS*, 2010.
- [27] X. Huang, "Real-time algorithm for acoustic imaging with a microphone array," *The Journal of the Acoustical Society of America*, vol. 125, no. 5, 2009.
- [28] 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, p. 19, 2009.
- [29] H. L. Van Trees, *Optimum Array Processing: Part IV of Detection, Estimation, and Modulation Theory*. New York, NY: John Wiley & Sons, 2002.
- [30] R. Horn and C. Johnson, *Matrix analysis*. Cambridge University Press, 1990.
- [31] M. Frigo and S. Johnson, "The design and implementation of FFTW 3," *Proceedings of the IEEE*, vol. 93, no. 2, pp. 216–231, 2005.
- [32] C. Van Loan and N. Pitsianis, "Approximation with Kronecker products," in *Linear Algebra for Large Scale and Real Time Applications*, M. Moonen and G. Golub, Eds. Kluwer Publications, 1992, pp. 293–314.
- [33] J. Baglama and L. Reichel, "Augmented implicitly restarted Lanczos bidiagonalization methods," *SIAM Journal on Scientific Computing*, vol. 27, no. 1, pp. 19–42, 2006.
- [34] P. Drineas, R. Kannan, and M. Mahoney, "Fast Monte Carlo algorithms for matrices II: Computing a low-rank approximation to a matrix," *SIAM Journal on Computing*, vol. 36, no. 1, pp. 158–183, 2007.
- [35] N. Halko, P. Martinsson, and J. Tropp, "Finding structure with randomness: Stochastic algorithms for constructing approximate matrix decompositions," California Inst. Tech., Tech. Rep. ACM 2009-05, Sep. 2009.
- [36] N. Srebro and T. Jaakkola, "Weighted low-rank approximations," in *Proc. ICML*, 2003, pp. 720–727.



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.