

Identification and Matching of Sparse Delay-Doppler Spread Functions from High-Frequency Communications Signals

Ehsan Zamanizadeh

Institute for Systems and Robotics
Instituto Superior Técnico
Av. Rovisco Pais, Torre Norte 7.10
1049-001 Lisboa, Portugal
Email: ezamanizadeh@isr.ist.utl.pt

João Gomes

Institute for Systems and Robotics
Instituto Superior Técnico
Av. Rovisco Pais, Torre Norte 7.22
1049-001 Lisboa, Portugal
Email: jpg@isr.ist.utl.pt

José M. Bioucas-Dias

Instituto de Telecomunicações
Instituto Superior Técnico
Av. Rovisco Pais, Torre Norte 10
1049-001 Lisboa, Portugal
Email: jose.bioucas@lx.it.pt

Abstract—This work proposes methods for estimating the responses of time-varying underwater acoustic channels using $\ell_2 - \ell_1$ basis pursuit and quantifying their similarity. This is motivated by problems where environmental information is to be extracted from communications signals through high-frequency ocean acoustic tomographic techniques. We view underwater channels as time-varying linear systems and characterize them by their Delay-Doppler Spread Functions (DDSF), which have many adjustable parameters but are often very sparse. Basis pursuit estimation of DDSF coefficients is achieved with either the SpaRSA or TwIST algorithms, which can efficiently handle sparse unconstrained $\ell_2 - \ell_1$ minimization even for very large problem sizes, and directly operate on the complex signals of baseband communication models. We show that these algorithms offer better performance than matching pursuit methods, in terms of DDSF estimation accuracy and computational complexity, using both simulated and real data. We examine criteria based on Euclidean distance, correlation, and Kullback-Leibler divergence for quantifying the similarity of DDSFs, as needed in the envisaged tomographic systems for matching observations and model predictions. Their performance is found to be similar, with the correlation metric showing somewhat less consistent performance than the other two in simulation. Our matching approach also enables a simplification of DDSF estimation algorithms by tolerating the inherent biasing in $\ell_2 - \ell_1$ solutions, and thereby avoiding the need for further postprocessing.

I. INTRODUCTION

Many underwater acoustic channels exhibit extensive time dispersion, on the order of tens of milliseconds or more, due to multipath propagation, requiring powerful algorithms at the receiver and transmitter to properly convey digital messages. The relatively slow propagation speed of acoustic waves in the ocean induces challenging time variations, as even relatively slow motions of the transmitter/receiver or surface waves can cause waveform compression/expansion and broadband Doppler spreading [1]. Characterizing the statistics of fluctuating multipath arrivals in a way that is useful for communication systems designers remains a topic of great relevance today.

Underwater channel responses are often sparse, a property that has been successfully exploited to improve the perfor-

mance of adaptive receivers by zeroing small and jitter-prone estimated coefficients, e.g., in single-carrier [2] or OFDM systems [3]. In the presence of time variations a linear channel may be characterized by its Delay-Doppler Spread Functions (DDSF), which is conceptually obtained by computing the instantaneous impulse response as a function of delay and time, and then Fourier transforming along the time variable [4]. These 2D functions, defined in the delay-Doppler plane, provide physical insight, as the effect of a slowly moving scatterer in the underwater environment (over a sufficiently small time interval) that creates a delayed and Doppler-shifted replica of the transmitted signal at the receiver can be mathematically described by a single coefficient in the delay-Doppler plane, or a set of coefficients in a confined region. This makes DDSFs of typical underwater channels particularly sparse, with most of the energy concentrated on a few regions where there are contributions from multipath arrivals. These regions become broader as the medium becomes more dynamic, e.g., due to increased wave activity.

Switching from static impulse responses to DDSFs introduces an additional Doppler dimension and typically leads to a harder channel identification problem as, in the absence of prior information, there are many more coefficients to estimate. Exploiting the sparsity of DDSFs is key to computing them with moderate complexity, as shown in early work by Li and Preisig [5] using matching pursuit (MP) and orthogonal matching pursuit (OMP) methods. In this work we focus on alternative basis pursuit (BP) methods developed for compressive sensing [6], [7], particularly on recent algorithms for $\ell_2 - \ell_1$ pursuit, where the (unconstrained) cost function to be minimized is a weighted sum of a least-squares term that measures goodness of fit and an ℓ_1 term that acts as a regularizer to induce solutions where many coefficients become zero. The algorithms for handling these dictionaries can be highly efficient, and for large problem sizes their computational complexity compares favorably with search-based methods like MP/OMP [8]. We specifically examine SpaRSA [9] and TwIST [10], two algorithms that seamlessly

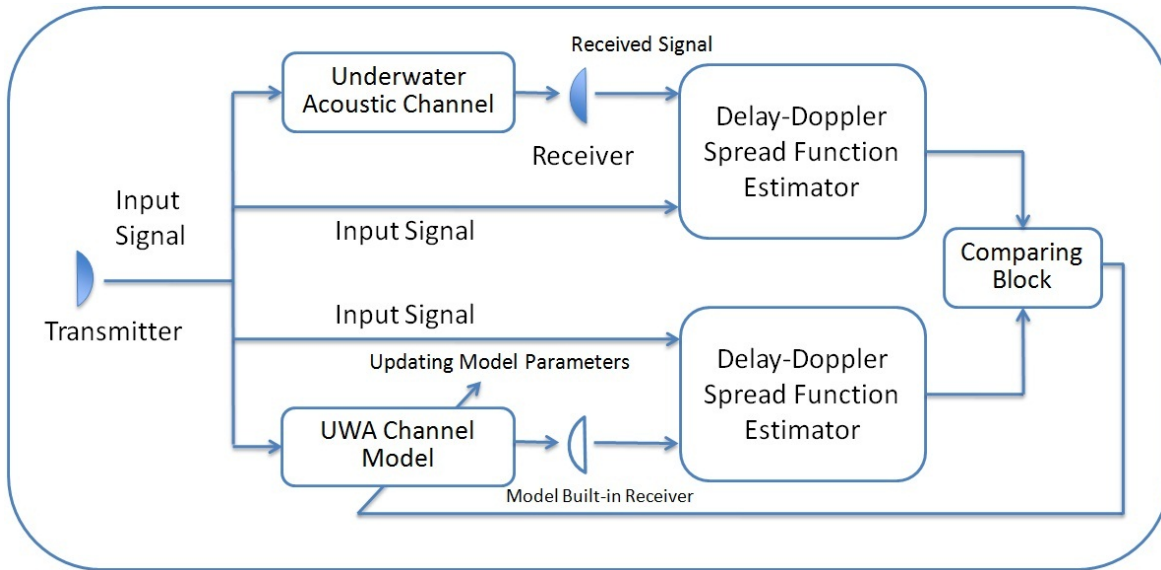


Fig. 1. Block diagram of the envisioned high-frequency tomographic approach. DDSFs are obtained from the observed response to a known transmission (top) or an acoustic propagation model (bottom). Model parameters (such as water depth or transmitter/receiver positions and relative speeds) are updated based on the matching analysis performed by the “Comparing Block”.

handle complex data and are therefore well suited to baseband communication models. We have found these algorithms to be robust in the sense that their adjustable parameters were set to values that provided suitable performance across our experimental and simulation data sets.

BP methods for wireless communications over sparse channels have been studied, e.g., in [11]–[13]. In [11] the emphasis is on deriving performance bounds for the channel estimates, whereas [12] focuses on designing the dictionary signals to maximize the sparsity of the estimated channel coefficients. In [13] the authors estimate time-varying channels in underwater OFDM systems using subspace, BP, and other methods, and then design frequency-domain equalizers to compensate for the intercarrier interference. BP methods in [13] were found to offer the most consistent performance over time and in several data sets.

Our work adopts a different viewpoint from those referenced above by emphasizing applications where environmental information is estimated from communications signals through tomographic techniques. Ocean acoustic tomography (OAT) is based on the principle that environmental properties such as sound speed profile, water depth, or bottom composition imprint a specific signature on acoustic signals, such that some of these properties can be recovered by solving an inverse problem. Conventional OAT uses acoustic signals whose frequencies are much lower than those used in underwater communications. However, the ability to infer environmental properties through digital communications signals, which is almost totally lacking in current underwater modems, is very appealing for a number of reasons:

- 1) To provide a clearer picture of the difficulties involved in a given spatial configuration of the acoustic link. This

could serve as a basis for tuning the receiver architecture (e.g., the filters used in equalization) to a particular ensemble of expected channel impulse responses.

- 2) To supply useful oceanographic information about the environment at no extra cost.
- 3) To provide useful localization information to be incorporated into the navigation system of a vehicle in which the receiver is installed.

Here, DDSFs are viewed as the link to assess how a given model of the environment conforms to experimental data (Fig. 1).

Matching DDSFs is envisaged for estimating geometric information such as source-receiver positions, velocities, and water depth, as modeling uncertainties at high frequencies seem too large to recover more subtle properties such as bottom composition. The framework is reminiscent of [14], [15], where the delays between selected multipath arrivals were used to invert for the source position through a model-based approach. In our work sparse DDSF estimation for single-carrier communications is used to obtain a “skeleton” of the acoustic field, including dynamic information on relative velocities from the Doppler shifts observed on different paths. This is then matched to predictions based on a ray propagation model, which provides sufficient accuracy at high acoustic frequencies in the tens of kHz. This paper specifically discusses criteria for measuring the discrepancy between DDSFs in the “comparing block” of Fig. 1 using Euclidean differences, correlation, and Kullback-Leibler divergence [16]. The actual statistical process of inferring model properties from matched DDSFs, denoted by the arrow from the “comparing block” to “UWA channel model”, is the subject of ongoing work and will be reported elsewhere.

Regarding DDSF estimation by BP methods, we assess the computational complexity of SpaRSA and TwIST as a function of problem size and compare it to MP and OMP, which have been used previously for similar purposes. It is found that BP methods have lower complexity than MP methods for suitable choices of the parameters that control the level of sparsity in the solution. Importantly, they also provide more physically plausible estimates with improved compact support in the delay-Doppler plane. Both SpaRSA and TwIST use a two-step process where the support region of the DDSF is first determined, as well as initial estimates of the DDSF coefficients. The second step, known as the *debiasing phase*, finds least-squares coefficient estimates restricted to the support region determined in the first step. We assess the relative computational complexity of these two steps and examine the feasibility of computing DDSF discrepancy measures without debiasing. We find that, while debiasing would likely be needed for receiver design purposes, it does not seem to offer significant advantages for OAT. Eliminating debiasing results in strong computational savings. The performance of the various methods is evaluated in simulation and using data from an at-sea experiment conducted in Norway, in September 2007.

The paper is organized as follows. Section II-A introduces DDSFs and formulates the DDSF estimation problem. Section II-B discusses criteria for matching DDSFs. Section III describes the BP framework for DDSF estimation and discusses some specific aspects of SpaRSA and TwIST. Section IV provides numerical results on DDSF estimation using both simulated and real data, as well as on DDSF matching measures. Finally, Section V outlines the main conclusions and provides directions for future research.

Notation: Superscripts $(\cdot)^T$, $(\cdot)^H$ stand for transpose and conjugate transpose (hermitian), respectively. ℓ_p norms are denoted by $\|\cdot\|_p$, and ℓ_2 is assumed when the argument p is omitted.

II. IDENTIFICATION AND MATCHING OF SPARSE DELAY-DOPPLER SPREAD FUNCTIONS

In this section we formulate the Delay-Doppler Spread Function estimation problem and introduce discrepancy measures for matching DDSFs. All communication signals and channel responses are represented by their complex baseband envelopes.

A. Problem Formulation

In a noiseless time-varying linear channel in continuous time the input signal, $x(t)$, and the output, $y(t)$, are related through the input delay spread function g [4]

$$y(t) = \int_{-\infty}^{\infty} x(t - \tau)g(t, \tau) d\tau. \quad (1)$$

The channel is represented as a continuum of scatterers, where $g(t, \tau)$ is the contribution at time t from a scatterer providing delays in the range $[\tau, \tau + d\tau]$. Variations in g along the time dimension are structured when caused by a small number of

Doppler shifts for a given delay, and Fourier transforming along the t variable yields a more compact representation of the system response through the DDSF U

$$y(t) = \iint_{\mathbb{R}^2} U(\tau, \nu) x(t - \tau) e^{j2\pi\nu(t - \tau)} d\tau d\nu, \quad (2)$$

$$U(\tau, \nu) = F_t \{g(t, \tau)\} e^{j2\pi\nu\tau}. \quad (3)$$

The argument in the complex exponential of (2) differs from the one appearing in [4], resulting in a DDSF definition that is different from the usual one, namely, $F_t \{g(t, \tau)\}$. This is more convenient for efficiently handling inner products between observed signals and multiple related dictionary waveforms through convolution in BP methods. The phase difference between (3) and the standard definition is of no consequence for the purposes of our work, which concentrates on the magnitude of U for tomographic inversion.

In an ideal discrete path model, which approximates the characteristics of many real underwater acoustic channels, the DDSF is represented as a set of impulses in the delay-Doppler plane

$$U(\tau, \nu) = \sum_{p=1}^{N_p} \alpha_p \delta(\tau - \tau_p) \delta(\nu - \nu_p). \quad (4)$$

Practical systems have an essentially finite number of degrees of freedom that enables a sampled representation of (2), where the coarseness of the delay and Doppler grids, $\Delta\tau$ and $\Delta\nu$, is dictated by the reciprocal of the input signal bandwidth, and the reciprocal of its duration, respectively [4], [11], [17]. We thus adopt the discrete-time input-output model

$$y(n) = \sum_{k,l} u_{k,l} x_l(n - k), \quad x_l(n) = x(n) e^{j2\pi\nu_l n}, \quad (5)$$

where the sampling frequency, f_s , is a multiple of the input signal bandwidth and $\nu_l = \frac{l}{T f_s}$ for an input block of duration T . Our practical Doppler grids are usually coarser than this, and suitable ranges for the delay and Doppler indices in (5) are chosen from an empirical analysis of each data set. With correct delay and Doppler sampling the discrete-time DDSF u will remain sparse.

The channel model (5) is linear in the DDSF coefficients, and may be written in matrix form as $\mathbf{y} = \mathbf{X}\mathbf{u}$, where \mathbf{y} denotes a vector of M observed samples, \mathbf{u} holds the stacked DDSF coefficients to be determined, and \mathbf{X} is the known *dictionary matrix*. The p -th column of \mathbf{X} is a vector, \mathbf{x}_p , of M consecutive input samples, delayed relative to \mathbf{y} and Doppler shifted as required for the p -th element of \mathbf{u} according to (5). As in [5] we assume that the observed block is sufficiently short, on the order of 1 second, so that the channel coefficients in \mathbf{u} can be considered constant. The problem of DDSF estimation is to obtain the coefficient vector given the dictionary matrix and a noisy version of the observation vector \mathbf{y} .

This simple model is best suited for single-carrier transmissions, where a block of consecutive input symbols can be assumed known. In multicarrier communications the channel response would be estimated from a set of pilot symbols




Fig. 2. Sparse DDSF estimation for real channel. (a) MP (b) OMP (c) SpaRSA (d) TwIST. Although all methods detect some nonzero taps, SpaRSA and TwIST provide a better coverage in the delay-Doppler plane and a clearer picture of the DDSF.

spread across the signal bandwidth, in which case the problem formulation would be different. A linear system of equations for the DDSF coefficients would still be obtained, though [11], [13].

For realistic choices of the delay and Doppler grids under the assumed signal bandwidth the dictionary matrix will be very large, making it inconvenient and slow to operate on it explicitly. However, the sparse identification methods under consideration only require that matrix products of the form $\mathbf{X}\mathbf{u}$ or $\mathbf{X}^H\mathbf{y}$ be calculated and for blocks of contiguous observed samples this can be done very efficiently due to the special structure of \mathbf{X} , which is the (column-wise) concatenation of convolution matrices for the signals $x_l(n)$ in (5). Specifically, in **forward mapping** we generate $\mathbf{X}\mathbf{u}$ through (FFT-based) time-invariant filtering of $x_l(n)$ with the subset of elements of \mathbf{u} pertaining to Doppler frequency ν_l , then adding over l . In **adjoint mapping** we generate the vector $\mathbf{X}^H\mathbf{y}$ by crosscorrelating the sequences $y(n)$ and $x_l(n)$, for all l , restricted to the samples contained in \mathbf{y} and \mathbf{X} . Each element of the desired vector, $\mathbf{x}_p^H\mathbf{y}$, is given by the crosscorrelation for a specific Doppler index and lag. Usually this is also most efficiently done using the FFT.

The modified DDSF definition adopted in (2)–(3) directly supports the above procedures for forward and adjoint mapping with no need for further postprocessing. Fast algorithms for FFT-based forward and adjoint mapping can also be used for DDSF estimation in OFDM systems [13].

B. DDSF Discrepancy Measures

Propagation modeling at frequencies of tens of kHz, used in underwater modems, is inherently unreliable because uncertainties in bottom bathymetry, surface altimetry, or source/receiver positioning, among others, are of the same order of magnitude, or even much larger, than the acoustic wavelengths of communication signals. Under such conditions model mismatch induces significant errors in acoustic field computations, hence also in DDSF prediction, and it is generally believed that phase errors in the coefficients, in particular, are too large to justify meaningful comparisons between complex values in modeled and observed DDSFs. This contrasts with OAT at low frequencies, where prediction of the magnitude and phase of acoustic fields has been

successful with accurate propagation models and reasonably detailed environmental data.

In spite of large uncertainties in coefficient values, other properties of DDSFs or impulse responses of underwater channels, such as the arrival times of multipath components, are considered reasonably reliable. In fact, ray tracing is routinely used to get a basic understanding of the observed structure of arrivals in experimental data sets containing communications signals. This suggests that matching the *locations* with significant contributions in the delay-Doppler plane may be a reasonable criterion for quantifying the discrepancy between DDSFs, even at high frequencies. The approach proposed in [14], [15] for simple and robust inversion by matching the time difference between selected arrivals seems to be directly inspired by this. In our work we adopt an intermediate approach, discarding the phase information in DDSFs when evaluating discrepancies, but keeping the *relative* magnitudes of arrivals. The rationale for doing this is that absolute magnitudes are too dependent on uncertain environmental and technical factors (such as filter gains at the transmitter and receiver), but the relative strengths of arrivals still carry information that is usually considered useful in practice. For example, observed bottom-interacting arrivals are typically weaker than the direct arrival or the first surface-reflected arrival. Given a vector \mathbf{u} with DDSF coefficients, we first preprocess it by retaining only the magnitudes of coefficients and normalizing for unit ℓ_1 norm, yielding

$$u_i \leftarrow \frac{|u_i|}{\sum_j |u_j|}. \quad (6)$$

The preprocessed coefficients are therefore real, non-negative, and sum to one:

$$u_i \geq 0, \quad \sum_i u_i = \|\mathbf{u}\|_1 = 1. \quad (7)$$

We then evaluate different real-valued metrics $\varphi(\mathbf{u}, \mathbf{v}) \geq 0$ to quantify the discrepancy between two (normalized) DDSFs \mathbf{u} and \mathbf{v} :

a) *Euclidean distance*: We simply take the Euclidean distance between the vectors \mathbf{u} and \mathbf{v} :

$$\varphi_E(\mathbf{u}, \mathbf{v}) = \|\mathbf{u} - \mathbf{v}\|_2. \quad (8)$$




Fig. 3. Running time for DDSF estimation methods. (a) Variable sparsity, with fixed problem size. (b) Variable problem size, with fixed regularization parameters/stopping criteria.

b) *Correlation:*

$$\varphi_C(\mathbf{u}, \mathbf{v}) = \mathbf{u}^T \mathbf{v}. \quad (9)$$

Note that this metric is approximately zero when \mathbf{u} , \mathbf{v} are very dissimilar, and it takes on higher values when $\mathbf{u} \approx \mathbf{v}$. The converse happens for φ_E and φ_{KL} below.

c) *Kullback-Leibler divergence:* In light of (7), normalized DDSF vectors can be viewed as probability mass functions (PMF), and their similarity quantified through the Kullback-Leibler (KL) divergence. For PMFs \mathbf{p} and \mathbf{q} , defined over a discrete set \mathcal{I} , the KL divergence is defined as [16]

$$D(\mathbf{p} \parallel \mathbf{q}) = \sum_{i \in \mathcal{I}} p_i \log \frac{p_i}{q_i}, \quad (10)$$

with the convention that $0 \log \frac{0}{0} = 0$, $0 \log \frac{0}{q} = 0$, $p \log \frac{p}{0} = \infty$. The KL divergence is always nonnegative and is zero only if $\mathbf{p} = \mathbf{q}$, but unlike a true distance it is not symmetric in its arguments. As a discrepancy metric we adopt the symmetrized KL divergence

$$\varphi_{KL}(\mathbf{u}, \mathbf{v}) = \frac{D(\mathbf{u} \parallel \mathbf{v}) + D(\mathbf{v} \parallel \mathbf{u})}{2}. \quad (11)$$

Due to the nature of our sparse estimation methods, DDSF coefficient vectors typically have many zero entries. Conceivably, the vectors \mathbf{u} , \mathbf{v} for two relatively close propagation scenarios might end up having little or no intersection in the subset of indices of their nonzero entries, in which case discrepancy measures as proposed above would provide misleading results of high DDSF disparity. To avoid this potentially erratic behavior we propose to lowpass filter the DDSF estimates prior to calculating the metrics, thus artificially broadening the nonzero support region of coefficient vectors to improve the continuity of φ . In this work we use a truncated Gaussian FIR filter to smooth DDSFs along the delay dimension. This was found to be unnecessary along the Doppler dimension, as BP methods yield solutions that are already sufficiently smooth with the chosen grids for Doppler shifts.

III. BASIS PURSUIT METHODS

Recent achievements in compressive sensing (CS) have sparked enormous interest in such techniques for solving various linear, but ill-posed, inverse problems, where sparsity acts as a regularization criterion to stabilize the solution [7]. We are mainly concerned with solving unconstrained ℓ_2 - ℓ_1 optimization problems of the form

$$\min_{\mathbf{u}} \frac{1}{2} \|\mathbf{y} - \mathbf{X}\mathbf{u}\|_2^2 + \tau \|\mathbf{u}\|_1, \quad (12)$$

where the first term measures how well the candidate solution fits the observed data, in the least-squares (LS) sense, while the second one is a regularizer which acts as a surrogate for the intractable ℓ_0 norm, and tends to penalize more heavily vectors \mathbf{u} with many nonzero components. The so-called regularization parameter τ controls the relative weight of the two terms [10]. Many variants of (12) exist in the literature, e.g., keeping only one of the terms in the cost function and including the other one as a constraint under a prescribed bound [7].

In most reported applications of (12) the data matrix \mathbf{X} is fat, i.e., there are fewer observations than unknowns, and the regularizer is essential for obtaining a well-posed optimization problem. This is not necessarily the case in DDSF estimation, where data blocks may be large enough to enable even conventional LS estimation. The ℓ_1 regularizer then acts simply as a device for automatically setting to zero small coefficients, whose contribution to improve the fit to observations is marginal, but which would nonetheless be retained by a pure LS estimator.

1) *Sparse Reconstruction by Separable Approximation (SpaRSA):* SpaRSA is a general framework for numerically solving an unconstrained optimization problem of the form [9]

$$\min_{\mathbf{u}} f(\mathbf{u}) + \tau c(\mathbf{u}), \quad (13)$$

where f is a smooth function and c is the sparsity-inducing regularizer which, in state of the art CS methods, is non-




Fig. 4. Running times for DDSF estimation with (dashed) and without (solid) a debiasing postprocessing phase as a function of the number of nonzero taps. (a) TwIST (b) SpaRSA.

quadratic and nonsmooth (typically the ℓ_1 norm appearing in (12)). It is an iterative method that at each step solves an optimization subproblem with an approximation for f that is separable in the unknowns, interpolating the gradient information and using a diagonal approximation to the Hessian. Simple and efficient algorithms result when the regularizer is also separable, i.e., it is a sum of functions of the individual components of its argument, $c(\mathbf{u}) = \sum_i c_i(u_i)$, as is the case for the ℓ_1 norm. For the ℓ_1 regularizer SpaRSA repeatedly evaluates simple so-called *soft threshold* functions that transparently clip small entries in \mathbf{u} to exactly zero. In the complex case required for working with baseband communication signals the soft threshold function is given by

$$\text{soft}(u, a) = \frac{\max\{|u| - a, 0\}}{\max\{|u| - a, 0\} + a} u, \quad (14)$$

where the parameter a depends on the current step and regularization parameters. Experiments with CS problems show that SpaRSA is competitive with the fastest known methods for the standard $\ell_2 - \ell_1$ problem, as well as being efficient on problems with other separable regularization terms.

The regularization parameter τ , which is usually set by trial and error, provides an adjustable control to specify the desired level of sparsity in the solution. As a rule of thumb τ should approximately equal the maximum squared ℓ_2 norm of the dictionary matrix columns.

Both SpaRSA and TwIST, described below, only operate on the dictionary matrix through the forward and adjoint mappings introduced at the end of Sec. II-A. These mappings can be computed in a very efficient way for several types of dictionary, including the ones considered here for DDSF estimation (FFT-based). These BP methods therefore have the ability to handle very large estimation problems with reasonable computational complexity, which is extremely relevant for the proposed application in underwater communications. More specific issues related to real-time implementation will not be considered here.

2) *Two-step Iterative Shrinkage/ Thresholding Algorithm (TwIST)*: Iterative shrinkage/thresholding (IST) algorithms attempt to minimize the cost function (13) for $f(\mathbf{u}) = \frac{1}{2} \|\mathbf{y} - \mathbf{X}\mathbf{u}\|^2$ through the recursion

$$\mathbf{u}_{t+1} = (1 - \beta)\mathbf{u}_t + \beta\Psi_\tau(\mathbf{u}_t + \mathbf{X}^H(\mathbf{y} - \mathbf{X}\mathbf{u}_t)), \quad (15)$$

where Ψ_τ is the *shrinkage function*, a componentwise nonlinearity that reduces the range of the elements of \mathbf{u} and thus induces sparsity by setting small coefficients to zero [18]. When $c(\mathbf{u}) = \|\mathbf{u}\|_1$ in (13), Ψ_τ coincides with the soft threshold function defined in Section III-1. In (15) $0 < \beta \leq 1$ is a parameter which changes the convergence rate of the IST method. Setting $\beta = 1$ defines the original IST algorithm.

The convergence rate of IST algorithms depends on the linear observation operator $\mathbf{X}\mathbf{u}$, becoming very slow when it is ill-conditioned or ill-posed. Two-step iterative shrinkage/thresholding algorithms overcome this by implementing a modified version of IST where \mathbf{u}_{t+1} depends explicitly on \mathbf{u}_t and \mathbf{u}_{t-1} . The resulting algorithms exhibit a much faster convergence rate than IST for ill conditioned and ill-posed problems [10]. The TwIST recursion is given by

$$\mathbf{u}_{t+1} = (1 - \alpha)\mathbf{u}_{t-1} + (\alpha - \beta)\mathbf{u}_t + \beta\Psi_\tau(\mathbf{u}_t + \mathbf{X}^H(\mathbf{y} - \mathbf{X}\mathbf{u}_t)), \quad (16)$$

where α and β are algorithm parameters that determine the convergence rate and should be adjusted for each specific optimization problem. Like SpaRSA, this algorithm readily operates on complex data.

3) *Debiasing*: Solving (12) does not lead to a LS solution for \mathbf{u} over the set of identified nonzero taps. The practical implementations of SpaRSA and TwIST therefore provide an optional postprocessing step known as *debiasing*, where a true LS solution is calculated, restricted to the set of dictionary entries where nonzero coefficients were identified during the main processing phase. Debiasing solves a linear system of equations using the conjugate gradient method, which also

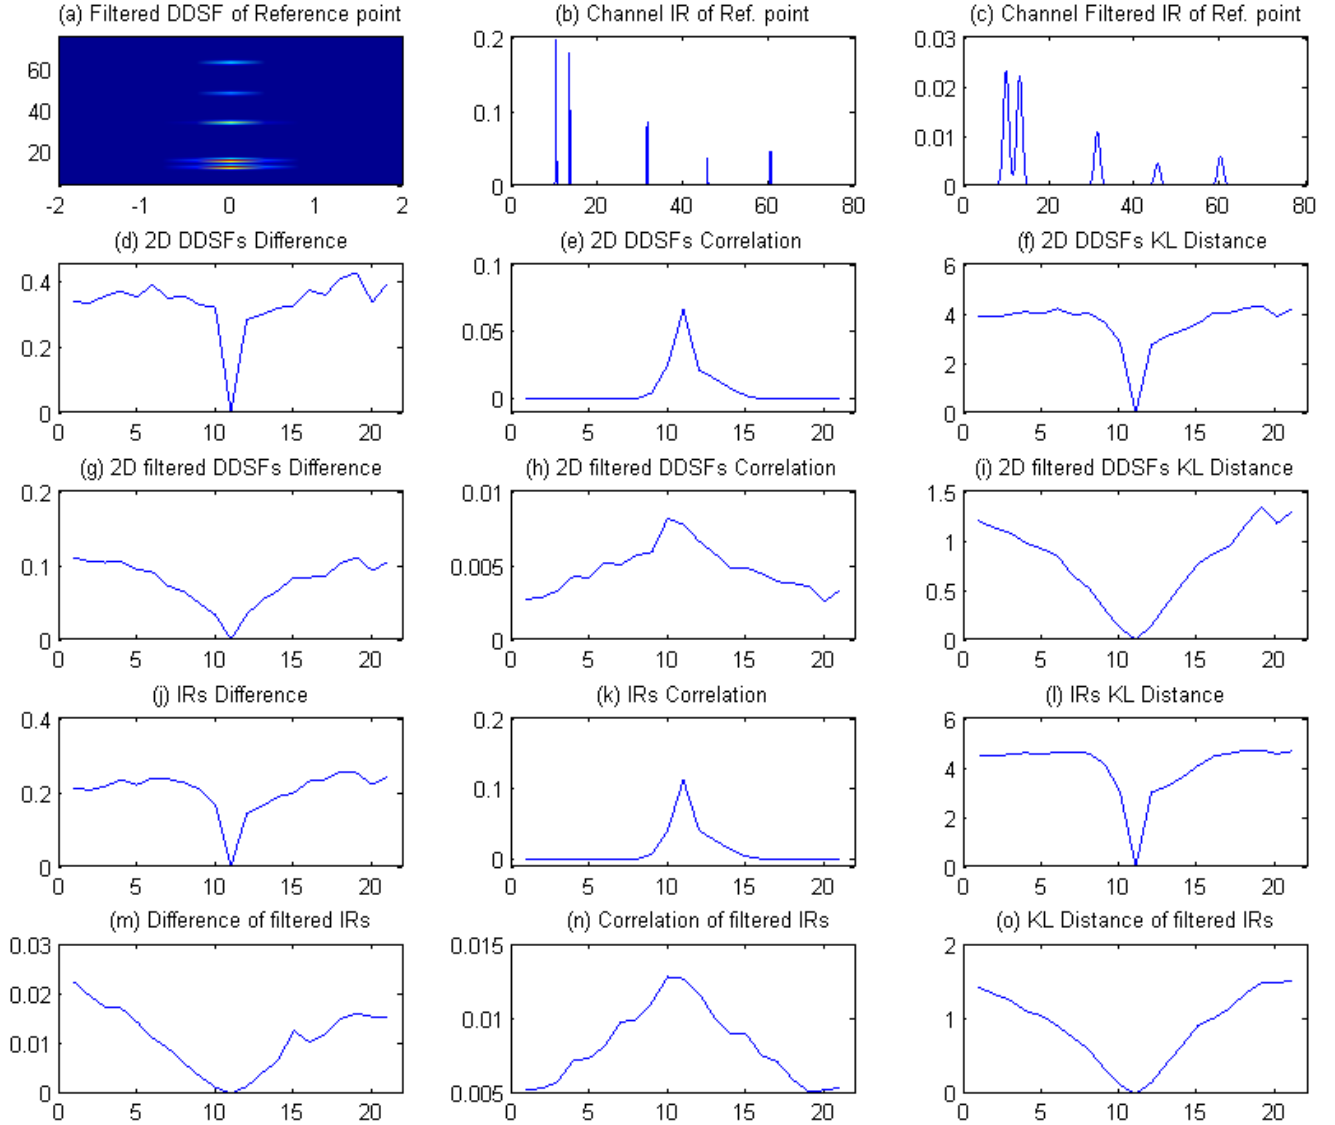


Fig. 5. Estimation and matching of sparse DDSFs. In this experiment the communication range is 400 m and the depth of the transmitter changes from 5 m to 45 m in 2 m steps. The depth of the transmitter for the reference configuration is 25 m. (a) Filtered DDSF, using a Gaussian filter along the delay axis, for the reference point (b) Channel impulse response (c) Channel impulse response after filtering by Gaussian filter (d) Difference between unfiltered DDSFs (e) Correlation between DDSFs (f) Kullback-Leibler distance between DDSFs (g)–(i) Discrepancy metrics for filtered DDSFs (j)–(o) Simplified discrepancy metrics for projection of unfiltered (top) and filtered (bottom) DDSFs onto the delay axis.

operates on the dictionary matrix only through the forward and adjoint mappings, thus retaining the ability to handle large problem sizes.

In spite of its efficiency, debiasing can account for a large fraction of the processing time in the current implementations of SpaRSA and TwIST. In this work we therefore investigate the possibility of comparing estimated DDSFs when the debiasing step is omitted. The rationale for doing so is the empirical observation that, in terms of magnitude, debiased solutions usually resemble slightly scaled-up versions of the solutions of (12), in which case the ℓ_1 normalization that

precedes the computation of disparity measures (Sec. II-B) will produce almost identical normalized DDSF coefficient vectors, and hence disparities will also be very similar in both cases. Note that debiasing would likely be needed if DDSFs were being used to design equalization filters at the receiver, where accurate knowledge of the coefficients is important.

IV. PERFORMANCE ASSESSMENT

In this section we compare the performances of sparse DDSF estimation and matching methods in QPSK transmissions over simulated and real underwater channels.

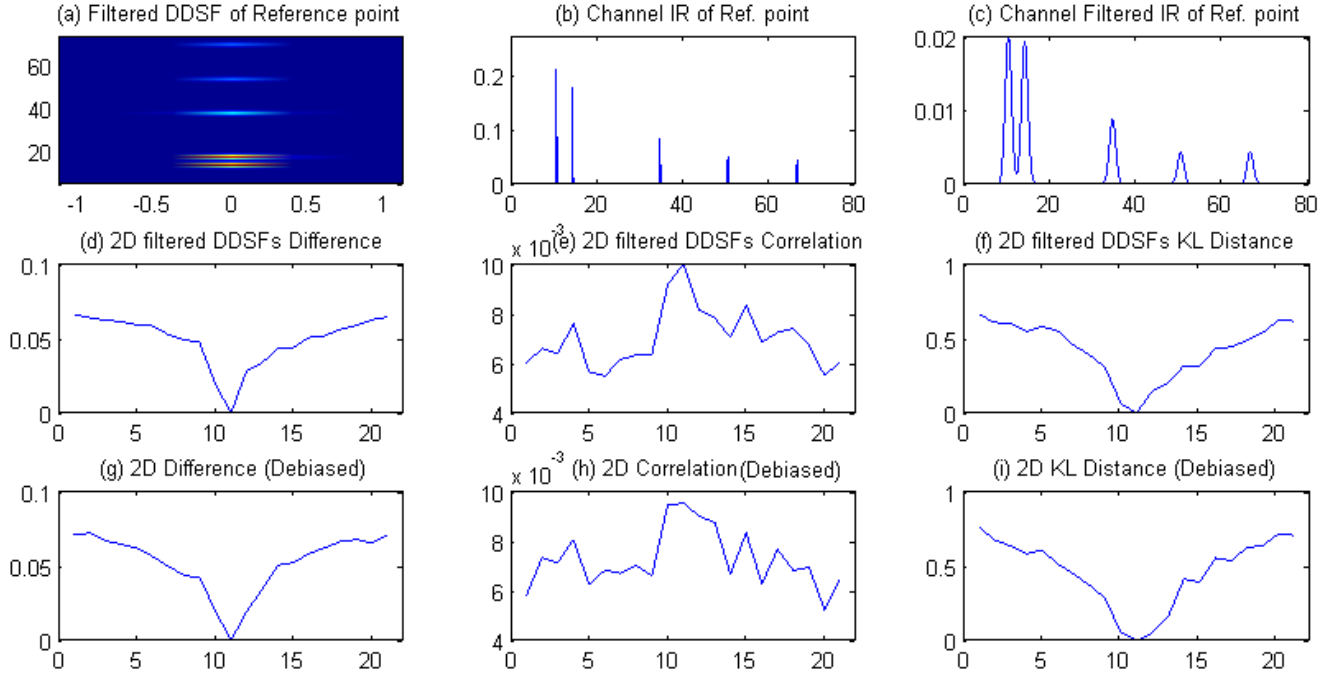


Fig. 6. Estimation and matching of sparse DDSFs. In this experiment the communication range changes from 300 m to 400 m in 5 m steps. The range is 350 m for the reference configuration. (a) Filtered DDSF, using a Gaussian filter, for the reference point (b) Channel impulse response (c) Channel impulse response after filtering by Gaussian filter (d) Difference between filtered DDSFs (e) Correlation between filtered DDSFs (f) Kullback-Leibler distance between filtered DDSFs (g)–(i) Discrepancy metrics using debiased and filtered DDSF estimates.

A. Experimental Results

The UAB’07 sea trial was carried out in Norway, in September 2007. In the particular experiment analyzed below the transmitter was suspended 10 m from shore, at a depth of about 5 m. The receiver (hydrophone #8 of a 16-element vertical array) was suspended from a drifting buoy at 35 m depth. The communication range was approximately 800 m, the bottom sloping from 10 m depth at the transmitter to about 100 m at the receiver. We focus on QPSK packets (type Q1) at 1 kbaud, with 1.5 kHz bandwidth, RRC pulse shapes (50% rolloff), 5.5 kHz carrier frequency, and total duration 3 s. Each packet is flanked by a pair of start/stop LFM markers for packet synchronization and coarse Doppler compensation through resampling. The Doppler axis in DDSF plots already accounts for this pre-processing. DDSFs are estimated based on a time window of 1 s centered on each received signal. The baseband sampling frequency is $f_s = 4$ kHz.

Experimental results are presented in Figure 2. We use sub-symbol delay resolution $\Delta\tau = 1/f_s = 2.5 \times 10^{-4}$ s and Doppler step $\Delta\nu = 0.1$ Hz, for a total of 244×41 DDSF coefficients. Results demonstrate a case with strong Doppler spread, where MP and OMP identify many scattered points that make it difficult to grasp the shape of the DDSF, while BP methods provide a clearer picture.

The run times for the four methods as a function of the

number of nonzero taps¹ for estimating a fixed-size DDSF are shown in Figure 3a. The debiasing tolerance parameter, which defines the desired accuracy for the debiasing phase discussed in Section III-3, is set to 10^{-3} for SpaRSA and TwIST. Since MP methods search for the highest projection for each new selected tap, the run time for MP and OMP is directly dependent on the number of nonzero taps. BP methods estimate the sparse vector jointly by solving a single optimization problem, which makes them less dependent on the number of nonzero taps. It is clear that BP algorithms are significantly faster than MP methods with no sacrifice in accuracy. Figure 3b shows the dependency of elapsed time for variable-size DDSF estimation as the Doppler resolution increases. Although getting higher DDSF resolution obviously implies greater computational cost in all algorithms, the elapsed time for MP and OMP grows significantly faster than for SpaRSA and TwIST.

Figures 4a and 4b respectively show the running times of TwIST and SpaRSA with and without debiasing as a function of the number of nonzero taps. The debiasing tolerance is set to 10^{-5} for both algorithms. These results demonstrate that, depending on the debiasing tolerance and number of non-zero taps, debiasing may account for a large fraction of the running time in both algorithms.

¹These were obtained by varying the number of iterations in MP/OMP or the value of the regularization parameter τ in BP methods.

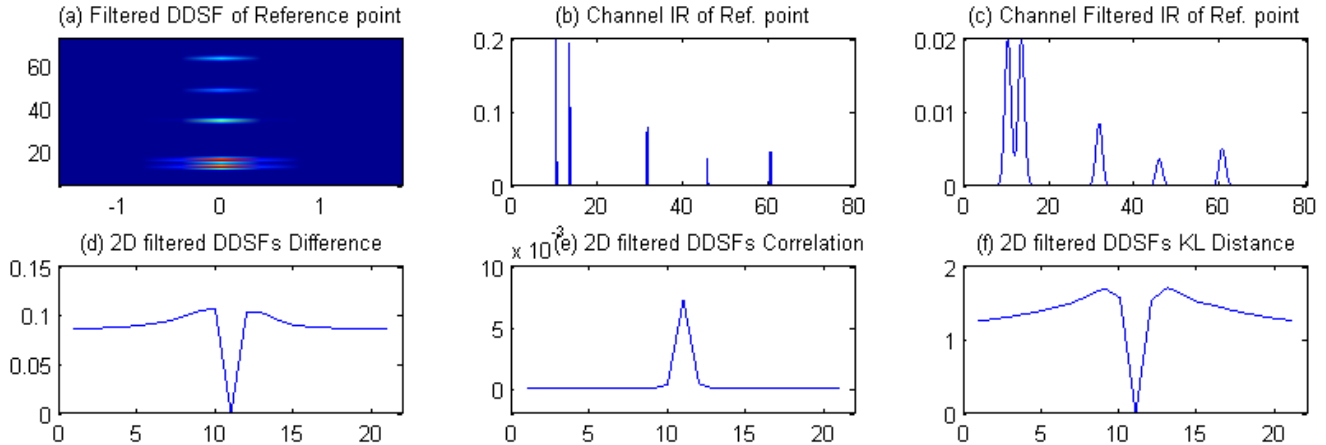


Fig. 7. Estimation and matching of sparse DDSFs. In this experiment the communication range is 400 m. The source is moving towards the receiver with changing speed from 2 m/s to -2 m/s in 0.2 m/s steps. The speed is 0 m/s for the reference configuration. (a) Filtered DDSF, using a Gaussian filter, for the reference point (b) Channel impulse response (c) Channel impulse response after filtering by Gaussian filter (d) Difference between filtered DDSFs (e) Correlation between filtered DDSFs (f) Kullback-Leibler distance between filtered DDSFs.

B. Simulation Results for DDSF Discrepancy Metrics

Simulation results were obtained using an online underwater acoustic simulator developed by the University of Algarve². The transmitted signal is a QPSK packet at 2.4 kbaud, with 5.5 kHz carrier frequency, 4.5 kHz bandwidth, root-raised-cosine (RRC) pulse shape (88% rolloff), and total duration 1 s. The baseband received signal is sampled at 4 times the symbol rate, $f_s = 9.6$ kHz. The bottom is sandy (1600 m/s, 2 g/cm³, 0.8 dB/λ). DDSFs are estimated using TwIST.

Figures 5, 6, and 7 show results for the discrepancy metrics as the transmitter depth, range to receiver, and speed are changed from a nominal configuration (depth 25 m, range 350 m, and speed 0 m/s). The receiver is stationary, at a depth of 50 m. Figure 5 shows metrics computed for raw DDSFs, and also for lowpass filtered versions with smoother behavior, as discussed in Section II-B (see the figure caption for details). In this experiment all metrics have similar performance. In addition to the discrepancy metrics calculated for the full 2D DDSF, the bottom two rows in the figure show the same metrics calculated on the projection of the DDSF along the delay axis. The price to be paid for this projection method is a moderate decrease in resolution, particularly when operating on filtered DDSFs.

Figures 6 and 7 only present metrics calculated on 2D filtered DDSFs. The shapes of the three metrics are similar in all cases, except for the correlation in Figure 6 that shows a more erratic evolution. Comparable results were observed in other simulations that are not reported here.

All the discrepancy metrics in Figures 5, 6, and 7 are obtained without debiasing of DDSF estimates, except for those in the bottom row of Figure 6. The similarity between the latter and their biased counterparts suggests that debiasing

does not play a major role in evaluating those metrics.

V. CONCLUSION

This paper examined issues related to the estimation of time-varying underwater channel responses using basis pursuit methods and quantification of their disparity. The work is intended as a contribution to the development of high-frequency tomographic methods that can be used to extract useful environmental information from communication signals, in addition to the digital messages themselves. The DDSF is envisaged here as the core (time-varying) channel characterization on which environmental inference is performed.

We present experimental results for DDSF estimation using the SpaRSA and TwIST algorithms for efficient $\ell_2 - \ell_1$ basis pursuit of sparse systems in the complex domain. These BP methods achieved more compact support in the delay-Doppler plane and an overall clearer picture of the DDSF than popular matching pursuit algorithms. Accuracy improvements over MP and OMP were attained with both significantly lower computational effort and more favorable scaling of complexity as a function of the problem size and degree of sparsity of the solution. Our focus is not on real-time implementation, but for run times of plain Matlab code usually under 10 seconds using 1-second data blocks this seems a reasonable goal. SpaRSA and TwIST attained essentially identical accuracy, the latter having slightly lower computational complexity.

Regarding DDSF discrepancies, we assessed metrics based on Euclidean distance, correlation, and KL distance. DDSF coefficient vectors were normalized to magnitude-only and unit ℓ_1 norm prior to applying those metrics to build tolerance to intrinsic environmental and modeling uncertainties in tomography, particularly at high frequencies. Simulation results for Euclidean and KL distances were very similar, with correlation performing less consistently than these two in some of the

²<http://www.ua-net.eu/projects/simulator/>

scenarios. Discrepancies calculated in the full delay-Doppler plane provided the best sensitivity to mismatches in range, depth, or speed. Not surprisingly, alternative simplified metrics operating on the 1D delay or Doppler projections of DDSFs exhibited lower resolution to changes in the environment. Our discrepancy metrics depend only weakly on the optional debiasing step after the main $\ell_2 - \ell_1$ execution of SpARSA and TwIST, and eliminating it when matching observations to a large number of candidate environments could result in substantial computational savings.

Future work will improve the strategy for DDSF matching, integrating it more tightly with ray tracing results to obtain more useful criteria for adjusting the environmental parameters when closing the tomographic loop. Another relevant topic is to improve the efficiency of the debiasing step for our specific basis pursuit problem, rather than simply eliminating it.

ACKNOWLEDGMENT

This work was supported by Fundação para a Ciência e a Tecnologia through project PTDC/EEA-TEL/71263/2006 and ISR/IST plurianual funding. The authors would like to thank CINTAL and ISR - Universidade do Algarve for collecting the experimental data used in this work, the Norwegian University of Science and Technology (NTNU) for the use of the Trondheim marine system research infrastructure (Hydralab III), and R/V Gun-nerus master and crew.

REFERENCES

- [1] T. H. Eggen, A. B. Baggeroer, and J. C. Preisig, "Communication over Doppler spread channels — Part I: Channel and receiver presentation," *IEEE Journal of Oceanic Engineering*, vol. 25, no. 1, pp. 62–71, January 2000.
- [2] M. Stojanovic, "Efficient processing of acoustic signals for high rate information transmission over sparse underwater channels," *Elsevier Journal on Physical Communication*, vol. 1, no. 2, pp. 146–161, June 2008.
- [3] —, "OFDM for underwater acoustic communications: Adaptive synchronization and sparse channel estimation," in *Proceedings of the International Conference on Acoustics, Speech and Signal Processing (ICASSP'08)*, Las Vegas, Nevada, USA, March 2008.
- [4] P. A. Bello, "Characterization of randomly time-variant linear channels," *IEEE Transactions on Communications Systems*, vol. CS-11, pp. 360–393, December 1963.
- [5] W. Li and J. C. Preisig, "Estimation of rapidly time-varying sparse channels," *IEEE Journal of Oceanic Engineering*, vol. 32, no. 4, pp. 927–939, October 2007.
- [6] S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic decomposition by basis pursuit," *SIAM Review*, vol. 43, no. 1, pp. 129–159, March 2001.
- [7] A. M. Bruckstein, D. L. Donoho, and M. Elad, "From sparse solutions of systems of equations to sparse modeling of signals and images," *SIAM Review*, vol. 51, no. 1, pp. 34–81, February 2009.
- [8] E. Zamanizadeh, J. Gomes, and J. Bioucas-Dias, "Identification of sparse time-varying underwater channels through basis pursuit methods," in *Proceedings of the 10th European Conference on Underwater Acoustics (ECUA'10)*, Istanbul, Turkey, July 2010.
- [9] S. J. Wright, R. D. Nowak, and M. A. T. Figueiredo, "Sparse reconstruction by separable approximation," *IEEE Transactions on Signal Processing*, vol. 57, no. 7, pp. 2479–2493, July 2009.
- [10] J. M. Bioucas-Dias and M. A. T. Figueiredo, "A new TwIST: Two-step iterative shrinkage/thresholding algorithms for image restoration," *IEEE Transactions on Image Processing*, vol. 16, no. 2, pp. 2992–3004, December 2007.
- [11] W. U. Bajwa, J. Haupt, A. M. Sayeed, and R. Nowak, "Compressed channel sensing: A new approach to estimating sparse multipath channels," *Proceedings of the IEEE*, vol. 98, no. 6, pp. 1058–1076, June 2010.
- [12] G. Tauböck, F. Hlawatsch, D. Eiwien, and H. Rauh, "Compressive estimation of doubly selective channels in multicarrier systems: Leakage effects and sparsity-enhancing processing," *IEEE Journal of Selected Topics in Signal Processing*, vol. 4, no. 2, pp. 255–271, April 2010.
- [13] C. R. Berger, S. Zhou, J. C. Preisig, and P. Willett, "Sparse channel estimation for multicarrier underwater acoustic communication: From subspace methods to compressed sensing," *IEEE Transactions on Signal Processing*, vol. 58, no. 3, pp. 1708–1721, March 2010.
- [14] X. Ma and Z.-H. Michalopoulou, "Matched arrival processing for efficient inversion in underwater acoustics," in *Proceedings of MTS/IEEE OCEANS'99*, vol. 3, Seattle, USA, September 1999, pp. 1577–1580.
- [15] Z.-H. Michalopoulou, X. Ma, M. Picarelli, and U. Ghosh-Dastidar, "Fast matching methods for inversion with underwater sound," in *Proceedings of MTS/IEEE OCEANS'00*, vol. 1, Providence, RI, USA, September 2000, pp. 647–651.
- [16] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. John Wiley & Sons, 1991.
- [17] W. U. Bajwa, A. M. Sayeed, and R. Nowak, "Learning sparse doubly-selective channels," in *Proceedings of the 46th Annual Allerton Conference on Communication, Control, and Computing*, Urbana-Champaign, IL, USA, September 2008, pp. 575–582.
- [18] M. A. T. Figueiredo and R. D. Nowak, "An EM algorithm for wavelet-based image restoration," *IEEE Transactions on Image Processing*, vol. 12, no. 8, pp. 906–916, August 2003.