Identification of Sparse Time-Varying Underwater Channels Through Basis Pursuit Methods

Ehsan Zamanizadeh[1], João Gomes[1], José M. Bioucas-Dias[2]

[1]Institute for Systems and Robotics - Instituto Superior Tecnico, Av. Rovisco Pais, 1049-001 Lisboa, Portugal, {**ezamanizadeh, jpg**}**@isr.ist.utl.pt**
[2]Instituto de Telecomunicações - Instituto Superior Tecnico, Av. Rovisco Pais, 1049-001 Lisboa, Portugal, **jose.bioucas@lx.it.pt**

Abstract: Underwater acoustic channels often exhibit extensive time dispersion due to multipath, necessitating the development of powerful algorithms at the receiver and transmitter for reliable performance in digital communication systems. Such Impulse Responses (IR) are often sparse, a property that has been exploited to improve the performance of adaptive receivers by zeroing small and jitter-prone estimated coefficients. In time-varying channels, responses may be described by (2D) Delay-Doppler Spread Functions (DDSF), which have more parameters than IRs but are even sparser. Motivated by (i) demonstrated significant sparsification gains by simple coefficient truncation at receivers, and (ii) recent developments in compressive sensing algorithms, this work examines the performance of algorithms for $\ell_2 - \ell_1$ Basis Pursuit (SpaRSA, TwIST) as tools for estimating sparse DDSFs. Their ability to solve a large-scale regularized least-squares problem without explicitly building a dictionary matrix is key for efficiently handling DDSFs. Their performance is compared to matching pursuit approaches (MP, OMP), which have been used previously for similar purposes. The above basis pursuit algorithms are shown to provide better accuracy than MP/OMP with lower computational complexity in both simulated and real data, and therefore it is argued that they merit consideration for inclusion in the signal processing chains of digital receivers.

Index Terms: Delay-Doppler Spread Function (DDSF), SpaRSA, TwIST, Basis Pursuit, Matching Pursuit, Sparse Estimation, Underwater Acoustic Channels, Underwater Communications

1 Introduction

As in other wireless channels, digital transmission of information through the ocean by acoustic means must overcome challenges such as multipath propagation and channel time variations. However, the relatively slow propagation speed of acoustic waves in the ocean and significant interaction with the boundaries of the medium through surface and bottom reflections induce distortions in waveforms that can be much more severe than those observed in wireless radio channels. For example, delay spreads can easily reach several tens of milliseconds, and even relatively slow motions of the transmitter/receiver or surface waves can cause waveform compression/expansion and broadband Doppler spreading [1]. Moreover, the delays of surface-reflected arrivals can also fluctuate rapidly over intervals of a few seconds due to wave motion [2]. Finally, the approximately quadratic dependence of attenuation with frequency means that underwater acoustic channels are severely bandlimited over typical operating ranges [1, 3].

Time-varying channels may be described by Delay-Doppler Spread Functions (DDSF), a generalization of the concept of time-invariant channel impulse response to the time-frequency plane that adds a Doppler dimension [4]. The channel output is viewed as a sum of replicas of the input signal, each associated with a given delay and Doppler shift that are assumed constant over an averaging interval [2]. This type of black-box approach has the advantage of capturing the channel structure and its dynamics without physically modeling the channel, but it does lead to a large increase in the number of parameters to be estimated if no prior knowledge exists on where energetic regions are located in the delay-Doppler plane. Receivers which rely on DDSFs for improved performance in the presence of channel variations therefore require strategies for estimating relevant coefficients in the delay-Doppler plane, while simultaneously rejecting contributions from small coefficients whose statistical fluctuations would lead to unacceptably high output jitter.

For sufficiently large bandwidths many short-range and medium-range underwater channels exhibit sparse IRs or DDSFs, as the received signal is mostly formed by contributions from a small number of resolvable multipath arrivals that match the assumptions in the black-box models mentioned above [5]. In the time invariant case this property has been used, e.g., in [6] (single-carrier modulation) and [7] (OFDM) to obtain significant performance gains at the receiver simply by thresholding least-squares estimates of channel coefficients. Most DDSFs of real underwater channels are particularly sparse, featuring extended delay spreads, yet with most of the coefficient energy localized in several small regions [2]. Exploiting sparsity is a key insight for attaining a generic delay-Doppler representation of underwater channels with manageable complexity and relatively short observations, as shown in early work by Li and Preisig [2] using matching pursuit (MP) and orthogonal matching pursuit (OMP) methods.

Estimation of sparse time-varying channels through basis pursuit (BP) techniques developed for compressive sensing has recently drawn interest in wireless communications scenarios [8, 9]. According to the original BP principle, a signal is decomposed into a superposition of possibly highly redundant dictionary signals, and an optimal set of weights is found such that the resulting coefficient vector has minimum ℓ_1 norm [10]. Among several variations of BP that have been proposed [11], we examine recently developed compressive sensing algorithms for $\ell_2 - \ell_1$ pursuit, where the 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. Specifically, we focus on Sparse Reconstruction by Separable Approximation (SpaRSA) [12] and Two-step Iterative Shrinkage/Thresholding (TwIST) [13], two elegant methods for solving unconstrained $\ell_2 - \ell_1$ optimization problems with complex variables and data. SpaRSA is a framework for the general problem of iteratively minimizing the sum of a smooth convex function and a nonsmooth regularizer. TwIST is an improved version of the Iterative Shrinkage/Thresholding (IST) algorithm, exhibiting a much faster convergence rate for ill conditioned and ill-posed problems. Both methods avoid explicitly building the dictionary matrix, but rather use operators (typically FFT-based) for projecting onto dictionary elements or generating an output vector from a set of weights. This allows them to efficiently handle very large-scale problems, a property that is extremely valuable for DDSF estimation.

BP methods, among others, have been considered in [14] for sparse channel estimation in underwater OFDM systems. The resulting channel estimates were then used to design frequency-domain equalizers to counter the effect of intercarrier interference caused by time variations. Equalizers based on sparse methods consistently outperformed those relying on conventional least-squares (LS) estimates, and BP methods were found to be the most robust in experimental data collected over diverse environmental conditions. Our work focuses specifically on DDSF estimation for single-carrier modulation, not on receiver design issues. Compared to previously used methods for DDSF estimation such as MP and OMP, it is shown that both SpaRSA and TwIST can provide a clear image of the DDSF faster. We examine the accuracy/complexity tradeoff as the desired level of sparseness is modified through the weight of the ℓ_1 regularizer. BP methods are found to be faster and produce more compact DDSF representations in the delay-Doppler plane than MP/OMP. 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 2 introduces the sparse DDSF estimation problem which we want to solve. Section 3 describes the MP and OMP methods for sparse estimation. Section 4 describes the BP framework and discusses some specific aspects of SpaRSA and TwIST. Section 5 provides numerical results on DDSF estimation performance and complexity using both simulated and real data. Finally, Section 6 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.

2 Problem Formulation

In this section we formulate the DDSF estimation problem. All signals and channel responses are represented by their complex baseband envelopes.

In a noiseless time-varying continuous-time channel the input signal, $x(t)$, and the output, $y(t)$, may be 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 nonmoving scintillating scatterers, where $g(t,\tau)$ is the contribution at time t from a scatterer providing delays in the range $[\tau, \tau + d\tau]$. Time variations in g are quite structured when caused by a small number of Doppler shifts, such that it becomes more convenient to Fourier transform g along the t variable and thus express (1) through the DDSF[1] U

$$y(t) = \iint_{\mathbb{R}^2} U(\tau,\nu)\, x(t-\tau) e^{j2\pi\nu(t-\tau)}\, d\tau d\nu \,, \quad (2)$$
$$U(\tau,\nu) = F_t\left\{g(t,\tau)\right\} e^{j2\pi\nu\tau} \,. \quad (3)$$

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. The reader is referred to [4, 15, 8] for details. 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 $\nu_l = \frac{l}{Tf_s}$ for an input block of duration T. Below, 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 retain the sparsity properties of its continuous-time counterpart.

[1] Our DDSF definition is slightly different from Bello's [4] to simplify the specification of the forward and adjoint operators in MP/BP methods.

The channel model (5) is linear in the DDSF coefficients, and may be written in matrix form as

$$\mathbf{y} = \mathbf{X}\mathbf{u}\,, \tag{6}$$

where $\mathbf{y}$ denotes a vector of M observed samples, $\mathbf{u}$ holds the 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 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 [2] we assume that the observed block is sufficiently short (lasting for about 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 $\mathbf{y}$ in (6).

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:

- [**Forward mapping**] To generate $\mathbf{X}\mathbf{u}$ perform time-invariant filtering of $x_l(n)$ with the subset of elements of $\mathbf{u}$ pertaining to Doppler frequency ν_l, then add over l. Naturally, the FFT can be used for filtering.

- [**Adjoint mapping**] To generate the vector $\mathbf{X}^H\mathbf{y}$ cross-correlate 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. The FFT should be used to efficiently compute these crosscorrelations, except when the range of delays considered in the DDSF is extremely short.

The modified DDSF definition adopted in (2)–(3) directly supports the above procedures for forward and adjoint mapping with no need for further postprocessing.

3 Basic and Orthogonal Matching Pursuit Methods

Matching Pursuit is an elegant method introduced by Mallat and Zhang [16], which iteratively decomposes a signal into a linear expansion of waveforms that are selected from a redundant dictionary. This section reviews the algorithms for DDSF estimation through basic MP and OMP, following a similar notation to [2]. Both MP and OMP sequentially select dominant taps of the DDSF that maximize the projection of the residual observation vector onto the corresponding symbol vector and then calculate tap coefficients according to some criteria. The difference is that MP calculates each tap coefficient directly from the projection regardless of the previous history, while OMP obtains a joint LS solution for the coefficients of all the selected taps [2].

At each iteration, t, MP selects one column of $\mathbf{X}$ that correlates best with the approximation residual from the previous iteration, $\mathbf{r}_{t-1}$, [16]

$$p = \arg\max_{s\notin I_{t-1}} \frac{\left|\mathbf{x}_s^H\mathbf{r}_{t-1}\right|^2}{\left\|\mathbf{x}_s\right\|^2}\,, \tag{7}$$

where I_{t-1} is the index set of all previously selected columns. The initial residual is the observation vector, $\mathbf{r}_0 = \mathbf{y}$. Then u_p, which is the element of $\mathbf{u}$ associated with the dictionary column $\mathbf{x}_p$, can be computed as

$$u_p = \frac{\mathbf{x}_p^H\mathbf{r}_{t-1}}{\left\|\mathbf{x}_p\right\|^2}\,. \tag{8}$$

Finally, the residual vector and index set are updated for the next iteration as

$$\mathbf{r}_t = \mathbf{r}_{t-1} - u_p\mathbf{x}_p\,, \qquad I_t = I_{t-1}\cup p\,. \tag{9}$$

The stopping criterion can be based on thresholding of $|u_p|$ or $\|\mathbf{r}_t\|$. When the same dictionary matrix is repeatedly applied to different observation vectors, the online computation of inner products can be eliminated by precomputing a table with all inner products between columns of the dictionary matrix.

OMP improves upon MP by recognizing that the correlation between selected columns of $\mathbf{X}$ as iterations progress should be taken into account to produce a LS estimate for the corresponding subset of coefficients in $\mathbf{u}$ [2]. A basic approach would be to explicitly minimize $\|\mathbf{y} - \mathbf{X}_{I_t}\mathbf{u}_{I_t}\|$ at each iteration, where $\mathbf{X}_{I_t}$ denotes the set of selected columns of $\mathbf{X}$, but this would be computationally burdensome. A better option is to recursively update a QR factor of $\mathbf{X}_{I_t}$ and only solve for the DDSF coefficients at the end, as described below. The same criterion of MP is used to choose a new column of $\mathbf{X}$ at each iteration.

1. Initialize $\mathbf{r}_0 = \mathbf{y}$ and $\mathbf{Q}_{I_0} = \mathbf{R}_{I_0} = \varnothing$
2. While the stopping criterion is not met:
   a) Select column p as in (7)
   b) Update the QR factor[2] of $\mathbf{X}_{I_t} = \mathbf{Q}_{I_t}\mathbf{R}_{I_t}$
   c) Project the residual onto $\mathbf{q}$, the new column of $\mathbf{Q}_{I_t}$, as $z_p = \mathbf{q}^H\mathbf{r}_{t-1}$ and store the coefficient
   d) Update I_t and $\mathbf{r}_t = \mathbf{r}_{t-1} - z_p\mathbf{q}$
3. Return $\mathbf{u}_{I_t} = \mathbf{R}_{I_t}^{-1}\mathbf{z}_{I_t}$, where $\mathbf{z}_{I_t}$ holds all sequentially stored residual projections

[2] We first build $\mathbf{Q}_{I_t} = \left[\mathbf{Q}_{I_{t-1}}\ \mathbf{x}_p\right]$ and $\mathbf{R}_{I_t} = \begin{bmatrix}\mathbf{R}_{I_{t-1}} & \mathbf{0}\\ \mathbf{0}^T & 1\end{bmatrix}$, then apply the Gram-Schmidt procedure to $\mathbf{x}_p$ to reorthogonalize the basis.

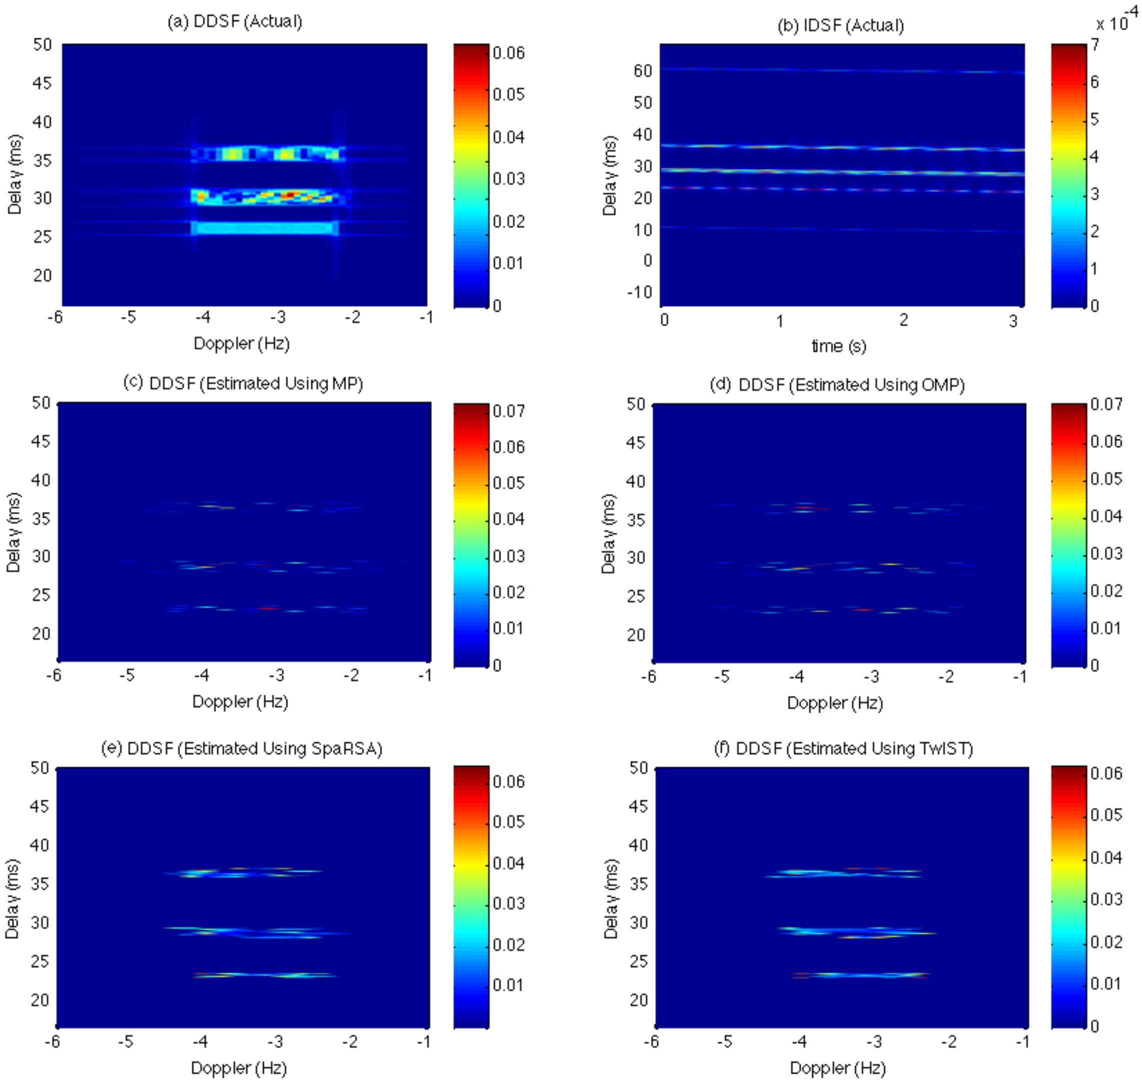

Figure 1: Sparse DDSF estimation for simulated channel using various pursuit algorithms. (a)–(b) True DDSF and time-varying channel impulse response. (c) MP (d) OMP (e) SpaRSA (f) TwIST. Although all methods detect some nonzero taps, SpaRSA and TwIST provide a clearer picture of the DDSF. Their running time is also significantly lower than for MP and OMP.

4 Basis Pursuit Methods

Finding sparse approximate solutions to large underdetermined linear systems of equations is relevant in several problems in signal/image processing and statistics [11]. In signal processing recent achievements in compressive sensing (CS) have sparked enormous interest in such techniques for solving various linear inverse problems. We are mainly concerned with solving unconstrained ℓ_2-ℓ_1 optimization problems of the form

$$\min_{\mathbf{u}} \frac{1}{2} \left\|\mathbf{y} - \mathbf{X}\mathbf{u}\right\|_2^2 + \tau \left\|\mathbf{u}\right\|_1 , \qquad (10)$$

where the first term measures how well the candidate solution fits the observed data, in the 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 [13]. We note that many variants of (10) 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 [11].

In most reported applications of (10) 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

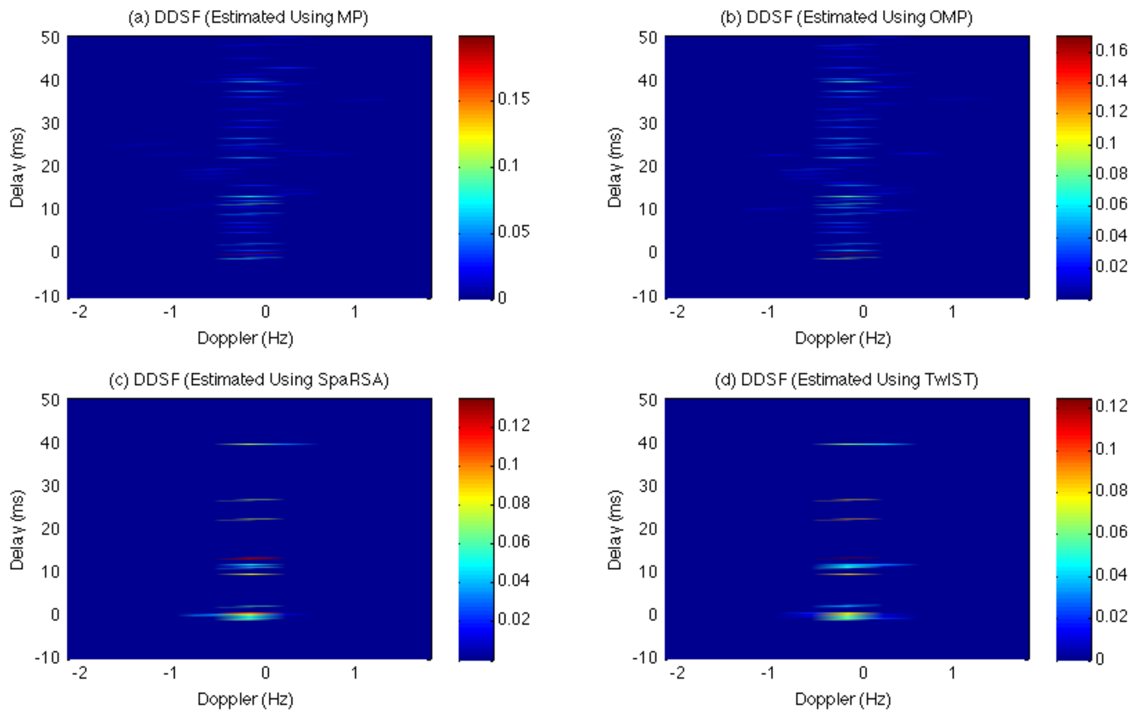

Figure 2: Sparse DDSF estimation for real channel. (a) MP (b) OMP (c) SpaRSA (d) TwIST. In this case all solutions have adequate support in the delay-Doppler plane.

observations is marginal, but which would nonetheless be retained by a pure LS estimator.

Unlike OMP, (10) does not lead to a LS solution for $\mathbf{u}$ over the set of identified nonzero taps, although the overall shapes are usually similar up to a scaling factor. The practical implementations for the BP methods described in Sections 4.1 and 4.2 therefore provide an option for a postprocessing *debiasing step*, where a true OMP-type LS solution is calculated. Debiasing is included in the running times reported in Section 5.

4.1 Sparse Reconstruction by Separable Approximation (SpaRSA)

SpaRSA is a general framework for numerically solving an unconstrained optimization problem of the form [12]

$$\min_{\mathbf{u}} f(\mathbf{u}) + \tau c(\mathbf{u}) \,, \tag{11}$$

where f is a smooth function and c is the sparsity-inducing regularizer which, in state of the art CS methods, is non-quadratic and nonsmooth (typically the ℓ_1 norm appearing in (10)). 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 in the real field SpaRSA repeatedly evaluates simple so-called *soft threshold* functions of the form[3] $\text{soft}(u, a) = \text{sign}(u) \max\{|u| - a, 0\}$ for each component of $\mathbf{u}$, where a depends on the current step and regularization parameters. It is through the $\max$ operation that small elements of $\mathbf{u}$ are set to zero.

The SpaRsa framework also yields efficient solution techniques for other regularizers, such as the nonconvex ℓ_0 norm and group-separable regularizers. It readily generalizes to the case in which the data is complex rather than real, which is highly desirable for working with complex baseband representations of digital communications signals. Experiments with CS problems show that this approach 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. Section 5 assesses the computational complexity (running time) vs. accuracy as a function of this parameter.

[3] In the complex case this becomes $\text{soft}(u, a) = \frac{\max\{|u| - a, 0\}}{\max\{|u| - a, 0\} + a} u$.

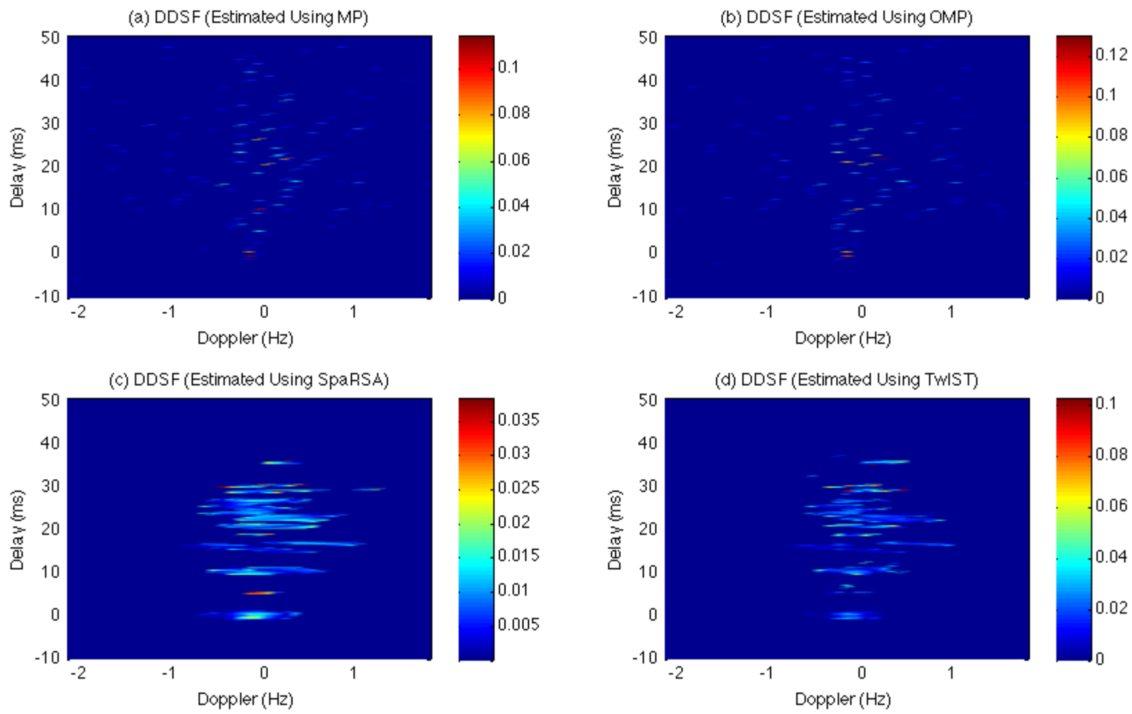

Figure 3: Sparse DDSF estimation for real channel. (a) MP (b) OMP (c) SpaRSA (d) TwIST. In this case the coverage of MP and OMP in the delay-Doppler plane seems too disperse.

4.2 Two-step Iterative Shrinkage/ Thresholding Algorithm (TwIST)

Iterative shrinkage/thresholding (IST) algorithms attempt to minimize the cost function (11) for $f(\mathbf{u}) = \frac{1}{2}\left\|\mathbf{y} - \mathbf{X}\mathbf{u}\right\|^2$ through the recursion

$$\mathbf{u}_{t+1} = (1-\beta)\mathbf{u}_t + \beta\Psi_\tau\Big(\mathbf{u}_t + \mathbf{X}^H(\mathbf{y} - \mathbf{X}\mathbf{u}_t)\Big), \quad (12)$$

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 [17]. When $c(\mathbf{u}) = \left\|\mathbf{u}\right\|_1$ in (11), Ψ_τ coincides with the soft threshold function defined in Section 4.1. In (12) $0 < \beta \leq 1$ 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 shortcoming 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 [13]. The TwIST recursion is given by

$$\mathbf{u}_{t+1} = (1-\alpha)\mathbf{u}_{t-1} + (\alpha-\beta)\mathbf{u}_t + \beta\Psi_\tau\Big(\mathbf{u}_t + \mathbf{X}^H(\mathbf{y} - \mathbf{X}\mathbf{u}_t)\Big), \quad (13)$$

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 can readily be applied to the complex case.

5 Performance Assessment

In this section we compare the performances of the sparse DDSF estimation methods in single-carrier (QPSK) transmissions over simulated and real underwater channels.

5.1 Simulation Results

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 shapes (88% rolloff), and total duration 1 s. The source and receiver are located 1.1 km apart, at 30 m depth, and approaching each other with relative speed 1 m/s. The bottom is sandy (1600 m/s, 2 g/cm^3, 0.8 dB/λ), at 110 m depth. The received signal, sampled at 4 times the symbol rate ($f_s = 9.6$ kHz), is computed with an online ocean acoustic simulator developed by the University of Algarve[4].

Figure 1 shows the DDSF estimation results for this channel. As in [14] we use sub-symbol delay resolution $\Delta\tau = 1/f_s \approx 10^{-4}$ s, whereas the Doppler step is $\Delta\nu = 0.2$

[4] http://www.ua-net.eu/projects/simulator/

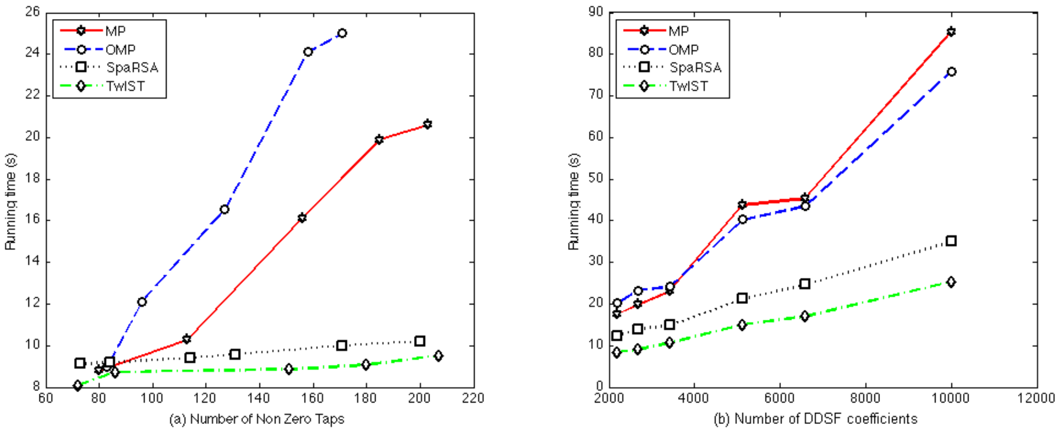

Figure 4: Running time for DDSF estimation methods. (a) Variable sparsity, with fixed problem size. (b) Variable problem size, with fixed regularization parameters/stopping criteria.

Hz. The DDSF has a total of 324×26 coefficients, so the (implicit) dictionary matrix $\mathbf{X}$ has 9600×8424 entries. From the figure it is clear that all methods capture the effective support region for the DDSF reasonably well, although SpaRSA and TwIST provide more compact representations that better match the true DDSF shown in Figure 1a. Besides, the running times for BP methods are significantly shorter than those form MP and OMP to estimate the same number of non-zero taps.

5.2 Experimental Results

The UAB'07 sea trial was carried out in Norway during the first two weeks of September 2007. The transmitter was suspended from a fixed platform 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 a depth of about 35 m. The communication range was approximately 800 m, the bottom depth gradually increasing from 10 m at the transmitter to about 100 m at the receiver location. Several modulation formats were transmitted with carrier frequency 5.5 kHz and variable bandwidth. In this work we focus on QPSK packets (type Q1) at 1 kbaud, with 1.5 kHz bandwidth, RRC pulse shapes (50% rolloff), and total duration 3 s. Each packet is flanked by a pair of start/stop LFM markers to be detected by cross-correlation for packet synchronization and coarse Doppler compensation through resampling. The Doppler axis in DDSF plots already accounts for this preprocessing. Each DDSF is estimated based on a time window of 1 s centered on the received signal.

Experimental results are presented in Figures 2 and 3. Despite the fact that the most energetic taps of the DDSF are detected by all methods in the packet shown in Figure 2, the results for BP methods are more concentrated. Similarly to the simulation results of Figure 1, these seem more physically plausible. Results for a second packet, in Figure 3, demonstrate a case with stronger Doppler spread, where MP and OMP identify many scattered points that make it even more difficult than previously to grasp the shape of the DDSF, while BP methods provide a much more clear picture.

The run times for the four methods as a function of the number of nonzero taps[5] for estimating a fixed-size DDSF are shown in figure 4a. 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. On the other hand, 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 in all situations BP algorithms are significantly faster than MP methods with no sacrifice in accuracy. Figure 4b 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.

6 Conclusion

In this paper we assessed the performance of two basis pursuit algorithms (SpaRSA and TwIST) for sparse DDSF estimation, comparing them with well established matching pursuit methods (MP and OMP). Numerical results using both simulated and real data showed that BP methods lead to more compact support in the delay-Doppler plane and a clearer picture of the DDSF. Moreover, this improvement over MP and OMP was attained with both significantly lower computational effort and more favorable scaling of complexity as a function of the problem size and number of nonzero entries in the DDSF.

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

Run times for BP methods on the order of 10 seconds seem amenable to improvements in algorithmic and software efficiencies that could ensure real-time or near real-time operation. No major differences were encountered between algorithms in the same class; MP and OMP performed similarly, possibly because the columns of our dictionary matrix are not strongly correlated. SpaRSA and TwIST also showed comparable performance, such that a choice between them for practical implementation could end up being dictated by considerations such as parallelizability and suitability to mapping onto low-level processor operations. Future work will address the problem of quantifying the disparity between DDSFs.

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 Gunnerus master and crew.

References

[1] T. Eggen, A. Baggeroer, and J. Preisig, "Communication over Doppler spread channels — Part I: Channel and receiver presentation," *IEEE J. of Oceanic Engineering*, vol. 25, no. 1, pp. 62–71, Jan. 2000.

[2] W. Li and J. Preisig, "Estimation of rapidly time-varying sparse channels," *IEEE J. of Oceanic Engineering*, vol. 32, no. 4, pp. 927–939, Oct. 2007.

[3] M. Stojanovic, "On the relationship between capacity and distance in an underwater acoustic communication channel," *ACM SIGMOBILE Mobile Computing and Communications Review (MC2R)*, vol. 11, no. 4, pp. 34–43, Oct. 2007.

[4] P. Bello, "Characterization of randomly time-variant linear channels," *IEEE Trans. on Communications Systems*, vol. CS-11, pp. 360–393, Dec. 1963.

[5] W. Li, "Estimation and tracking of rapidly time-varying broadband acoustic communication channels," Ph.D. dissertation, Massachusetts Institute of Technology and Woods Hole Oceanographic Institution, Feb. 2006.

[6] M. Stojanovic, "Efficient processing of acoustic signals for high rate information transmission over sparse underwater channels," *Elsevier J. on Physical Communication*, vol. 1, no. 2, pp. 146–161, June 2008.

[7] ——, "OFDM for underwater acoustic communications: Adaptive synchronization and sparse channel estimation," in *Proc. Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP'08)*, Las Vegas, Nevada, USA, March 2008.

[8] W. Bajwa, J. Haupt, A. Sayeed, and R. Nowak, "Compressed channel sensing: A new approach to estimating sparse multipath channels," *Proc. IEEE*, vol. 98, no. 6, pp. 1058–1076, June 2010.

[9] G. Tauböck, F. Hlawatsch, D. Eiwen, and H. Rauhu, "Compressive estimation of doubly selective channels in multicarrier systems: Leakage effects and sparsity-enhancing processing," *IEEE J. of Selected Topics in Signal Processing*, vol. 4, no. 2, pp. 255–271, April 2010.

[10] S. Chen, D. Donoho, and M. Saunders, "Atomic decomposition by basis pursuit," *SIAM Review*, vol. 43, no. 1, pp. 129–159, March 2001.

[11] A. Bruckstein, D. 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, Feb. 2009.

[12] S. Wright, R. Nowak, and M. Figueiredo, "Sparse reconstruction by separable approximation," *IEEE Trans. on Signal Processing*, vol. 57, no. 7, pp. 2479–2493, July 2009.

[13] J. Bioucas-Dias and M. Figueiredo, "A new TwIST: Two-step iterative shrinkage/thresholding algorithms for image restoration," *IEEE Trans. on Image Processing*, vol. 16, no. 2, pp. 2992–3004, Dec. 2007.

[14] C. Berger, S. Zhou, J. Preisig, and P. Willett, "Sparse channel estimation for multicarrier underwater acoustic communication: From subspace methods to compressed sensing," *IEEE Trans. on Signal Processing*, vol. 58, no. 3, pp. 1708–1721, March 2010.

[15] W. Bajwa, A. Sayeed, and R. Nowak, "Learning sparse doubly-selective channels," in *Proc. 46th Annual Allerton Conf. on Communication, Control, and Computing*, Urbana-Champaign, IL, USA, Sept. 2008, pp. 575–582.

[16] S. Mallat and Z. Zhang, "Matching pursuits with time-frequency dictionaries," *IEEE Trans. on Signal Processing*, vol. 41, no. 12, pp. 3397–3415, Dec. 1993.

[17] M. Figueiredo and R. Nowak, "An EM algorithm for wavelet-based image restoration," *IEEE Trans. on Image Processing*, vol. 12, no. 8, pp. 906–916, Aug. 2003.