Interpolating solutions of the Helmholtz equation with compressed sensing

Tim T.Y. Lin, Evgeniy Lebed, Yogi A. Erlangga, and Felix J. Herrmann, University of British Columbia, EOS*

SUMMARY

We present an algorithm which allows us to model wavefields with frequency-domain methods using a much smaller number of frequencies than that typically required by the classical sampling theory in order to obtain an alias-free result. The foundation of the algorithm is the recent results on the compressed sensing, which state that data can be successfully recovered from an incomplete measurement if the data is sufficiently sparse. Results from numerical experiment show that only 30% of the total frequency spectrum is need to capture the full wavefield information when working in the hard 2D synthetic Marmousi model.

INTRODUCTION

Seismic wavefield modeling is a technique which can be carried out in either the time domain or the frequency domain. When solutions to the Helmholtz equation for all frequencies are available, modeling in the frequency domain is equal to modeling with time-stepping methods (Symes (2007)). However, if we fail to accurately obtain solutions for the entire discretized spectrum of the signal, aliasing artifacts will be visible in the resulting modeled wavefield.

We are motivated by the long-standing observation that seismic imaging algorithms formulated in the frequency domain can yield useful alias-free images by acting only on a very limited number of frequencies; see Sergue and Pratt (2004) and Mulder and Plessix (2004). In this respect, seismic inversion in the frequency domain is considered more serviceable compared to time-stepping, where by definition every time step must be considered in order to obtain results. Sergue and Pratt (2004) showed that the extent to which we can exploit this property of frequency-based inversion is related to the range of available offsets in the seismic data, and furthermore presented a strategy for selecting an optimal subset of frequencies to consider.

Following the spirit of frequency-based seismic inversion, we believe that a similar property of information redundancy exists for the seismic wavefield itself in the frequency domain, which we can apply to save computation in seismic wavefield modeling. As we shall see, our conjecture follows naturally from considering a subset of monochromatic Helmholtz equation solutions as a restricted sampling of the modeled seismic wavefield in the Fourier domain, which has been demonstrated earlier by Lin and Herrmann (2007). In this approach, the mathematical machinery of compressed sensing (Candès et al. (2006); Donoho (2006)) naturally leads to a robust method of anti-aliasing the modeled wavefield in the time domain by solving a de-noising problem, an idea introduced in Hennenfent and Herrmann (2006).

In this paper, we describe an algorithm for computing acoustic wavefields by using only a limited number of frequencies. The algorithm is based on sparsity-promoting recovery formulated in the context of compressed sensing. We also present experimental results showing that this algorithm can compute wavefield using only 30% of the total frequency information even in models with sharp discontinuities, such as the hard 2D Marmousi model.

Computing solutions to the Helmholtz equation

Solving the Helmholtz equation is the standard way to model the propagation of seismic wavefield in the frequency domain through an inhomogeneous media. By applying a Fourier transform in the time axis, the standard acoustic wave equation can be converted to the Helmholtz equation:

$$\mathscr{H}(\omega)u(\omega,x_s;x) := \left(-\nabla^2 + \frac{\omega^2}{c^2(x)}\right) u(\omega,x_s;x) = b(\omega,x_s), \tag{1}$$

with $\omega = 2\pi f$, f being the frequency in Hz, $u(\omega,x_s;x)$ the solution in terms of scalar pressure field in the spatial domain. The initial term $b(\omega,x)$ is the frequency component of the source wavefield corresponding to ω in $\mathscr{H}(\omega)$, and $c(x)$ is the background velocity model. To obtain the propagation behavior of the initial source wavefield $f(\omega,x)$ we need to solve Equation 1 for all frequencies. In matrix notations, this computation can be expressed as, for one shot position x_s,

$$\begin{bmatrix} u(\omega_1) \\ \vdots \\ u(\omega_{n_f}) \end{bmatrix} = \begin{bmatrix} \mathscr{H}^{-1}(\omega_1) & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & \mathscr{H}^{-1}(\omega_{n_f}) \end{bmatrix} \begin{bmatrix} b(\omega_1) \\ \vdots \\ b(\omega_{n_f}) \end{bmatrix}, \tag{2}$$

with n_f the number of frequencies. Note that we have dropped x and x_s to simplify our notations. We then obtain the modeled wavefield in the time domain with an inverse Fourier transform $\mathscr{F}^{-1}$ in the time axis, defined as $u(t) = \mathscr{F}^{-1}u(\omega)$.

In actual computation, the operator $\mathscr{H}(\omega_i)$, $i = 1,\ldots,n_f$, is usually discretized with finite difference or finite element methods into a matrix representation $\mathbf{H}(\omega_i)$. This results in relatively large linear systems

$$\mathbf{H}(\omega_i)\mathbf{u}(\omega_i) = \mathbf{b}(\omega_i). \tag{3}$$

For example, if finite difference is used then $\mathbf{H}(\omega_i)$ is represented by a large sparse matrix consisting of diagonal bands. Although relatively trivial to implement, this method is highly susceptible to numerical dispersion. Fine grid spacings must be used to combat the numerical errors introduced by such dispersion effects. A common rule of thumb among numerical analysts is that a typical 5-pt stencil finite difference implementation of $\mathbf{H}(\omega_i)$ will require a minimum of 12 samples per wavelength for acceptable accuracy. For 3D, or even large 2D systems, where the signal band exceeds 50Hz, this would quickly lead to an astronomically large $\mathbf{H}(\omega_i)$. The computational cost and storage associated with explicit matrix inverse methods such as LU decomposition make them practically prohibitive for this purpose.

For systems too large for explicit methods, iterative methods are usually called upon to solve for a solution to Equation 3. The strategy is to start with a reasonable guess for $\mathbf{u}$ and then iteratively converge it to the correct solution by updates via matrix-vector multiplications with $\mathbf{H}(\omega_i)$. The matrix $\mathbf{H}(\omega_i)$ is typically very sparse and quite storable in memory, so in terms of computation costs multiplication with $\mathbf{H}(\omega_i)$ is relatively cheap. However, iterative methods are quite unstable compared to explicit methods, and had traditionally been plagued with exponentially increasing convergence difficulties for frequencies over 5Hz.

A recent work by Erlangga and Nabben (2007) show that it is indeed possible for the Helmholtz equation to achieve a convergence rate for iterative methods which is practically independent of frequency. They accomplish this by utilizing the multigrid method within a multilevel Krylov method setting introduced earlier by Erlangga and Nabben (2008). In this case, multigrid is applied to a damped, or complex shifted, Helmholtz operator, used as preconditioner for $\mathbf{H}(\omega_i)$. This method, termed as the MKMG method, can reach satisfactory convergence for frequencies as high as 50Hz in an acceptable number of iterations; see for details another contribution by one of the authors to the proceedings of this meeting. This is a drastic improvement over previously available schemes for iteratively inverting the Helmholtz operator, which would fail to converge at frequencies as low as 5 Hz. For the first time, iterative Helmholtz methods are approaching the threshold of practicality for those interested in the fine details of the inverted image.

INTERPOLATION IN THE FOURIER DOMAIN

From equation 2 it is clear that solutions to the Helmholtz are completely independent in the frequency domain. Following the discretization of $\mathscr{H}(\omega)$ described above, Equation 2 becomes:

$$\mathbf{u}(\omega) := \begin{bmatrix} \mathbf{u}(\omega_1) \\ \vdots \\ \mathbf{u}(\omega_{n_f}) \end{bmatrix} = \begin{bmatrix} \mathbf{H}^{-1}(\omega_1) & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & \mathbf{H}^{-1}(\omega_{n_f}) \end{bmatrix} \begin{bmatrix} \mathbf{f}(\omega_1) \\ \vdots \\ \mathbf{f}(\omega_{n_f}) \end{bmatrix}. \tag{4}$$

Now consider hypothetically the situation where solutions to the Helmholtz equation are not available for all frequencies. This would be the case, e.g., if iterative methods failed to converge for certain frequencies, or if the computer exhausted its storage while explicitly computing the inverse for all frequencies. We can mathematically express this by defining a restriction operator $\mathbf{R}$ which removes from its operand a predefined set of frequency components:

$$\hat{\mathbf{u}}(\omega) = \mathbf{R}\mathbf{u}(\omega), \tag{5}$$

with the hat symbol denoting that $\hat{\mathbf{u}}(\omega)$ is incomplete in the Fourier domain. The purpose of this section is to develop a method which, when provided only with $\hat{\mathbf{u}}(\omega)$ and knowledge of which frequencies are missing from it, we can use to fully recover the full set of solutions $\mathbf{u}(\omega)$.

Stable recovery of sparse signals

The problem of stable signal recovery (SSR) is to recover a vector $\mathbf{x}_0 \in \mathbb{R}^n$ from an incomplete set of linear measurements $\mathbf{y} = \mathbf{A}\mathbf{x}_0$, where $\mathbf{A}$ is an n by m, highly underdetermined matrix with $n \ll m$. This problem is resolved by solving the following ℓ_1-regularization minimization:

$$\widetilde{\mathbf{x}} = \arg\min_{\mathbf{x}} \|\mathbf{x}\|_1 \quad \text{s.t.} \quad \mathbf{A}\mathbf{x} = \mathbf{y}. \tag{6}$$

Recent results show that if the matrix $\mathbf{A}$ obeys the uniform uncertainty principle, and that $\mathbf{x}_0$ is sufficiently sparse, then the solution $\widetilde{\mathbf{x}}$ to Equation 6 is going to be exactly $\mathbf{x}_0$, Candès et al. (2006); Donoho (2006).

For data which are not sparse, it is often necessary to employ a modified formulation of Equation 6 to instead solve for a sparse representation of $\mathbf{x}_0$:

$$\begin{cases} \check{\mathbf{x}} = \arg\min_{\check{\mathbf{x}}} \|\check{\mathbf{x}}\|_1 \quad \text{s.t.} \quad \mathbf{A}\mathbf{S}^H\check{\mathbf{x}} = \mathbf{y}, \\ \widetilde{\mathbf{x}} = \mathbf{S}^H\check{\mathbf{x}}. \end{cases} \tag{7}$$

Here $\mathbf{S}$ is a linear operator, which maps a vector of data $\mathbf{x}_0$ to a sparse vector. Hence, $\mathbf{S}$ is a sparsity transform. Different types of data will have sparse representation under different transforms–i.e., there is no single transform that will sparsely represent all types of data. For example, piecewise constant images can be sparsely represented by spatial finite differences. Real-life images are known to have a sparse representation in the discrete cosine transform (DCT), wavelet transform or curvelets domain.

The intricacies of choosing a suitable sparsity basis for successful recovery will be discussed in the next section.

Application of stable signal recovery to anti-aliasing

We now return to the problem of recovering the full set of Helmholtz solutions from frequency-incomplete Helmholtz solutions $\hat{\mathbf{u}}(\omega)$. To connect with the theory of stable signal recovery, we note that $\hat{\mathbf{u}}(\omega)$ can be written as

$$\hat{\mathbf{u}}(\omega) = \mathbf{R}\mathbf{F}\mathbf{u}(t), \tag{8}$$

where $\mathbf{F}$ is a discrete implementation of the Fourier transform. In other words, $\hat{\mathbf{u}}$ can be described as an incomplete Fourier measurement of the full Helmholtz solution in the time domain. To see the similarity of Equation 8 to incomplete measurements in stable signal recovery, we consider the restricted Fourier transform $\mathbf{RF}$ as the matrix $\mathbf{A}$ in Equation 6, and $\mathbf{u}(t)$ the full signal $\mathbf{x_0}$ we wish to recover. Here we see that $\hat{\mathbf{u}}(\omega)$ plays the role of the incompletely measured signal $\mathbf{y}$. We can now propose to recover $\mathbf{u}(\omega)$ from $\hat{\mathbf{u}}(\omega)$ by writing it as a stable signal recovery problem:

$$\begin{cases} \check{\mathbf{u}}(t) = \arg\min_{\check{\mathbf{u}}} \|\check{\mathbf{u}}\|_1 \quad \text{s.t.} \quad \mathbf{R}\mathbf{F}\mathbf{S}^H\widetilde{\mathbf{u}}(t) = \hat{\mathbf{u}}(\omega), \\ \mathbf{u}(t) = \mathbf{S}^H\check{\mathbf{u}}(t). \end{cases} \tag{9}$$

The results here are significant: if there exists a suitable sparsity transform $\mathbf{S}^H$, we can then recover the solution of the acoustic wave equation in the time domain by solving the Helmholtz equation for only a subset of frequencies. The important thing to note is that the success of the recovery of a signal depends on the mutual coherence between the measurement basis and the sparsifying basis. Numerically, the coherence $\mu(\mathbf{M},\mathbf{S})$ of

a measurement basis $\mathbf{M}$ and a sparsifying transform $\mathbf{S}$ is computed by $\mu(\mathbf{M},\mathbf{S}) = \sqrt{m}\cdot\max_{k,l}|\mathbf{m}_k(\mathbf{s}_l)^H|$ with $\mathbf{m}_k$ the k^{th} row of $\mathbf{M}$ and $\mathbf{s}_l$ the l^{th} row of $\mathbf{S}$. Results from compressed sensing tell us that these factors directly influence the number of frequencies we can restrict while still obtaining correct solutions. Hennenfent and Herrmann (2006) solved a similar problem where the measurement basis was the Dirac basis and the sparsifying transform was the curvelet transform. They showed that curvelets are incoherent with the spike measurement basis, but have compact support in Fourier. However, for our purposes we take the restricted Fourier transform as our measurement basis. Since curvelets are highly coherent in Fourier, curvelets are not a satisfactory transform to use in this case. Recent work by Lebed and Herrmann which are presented in this conference have experimentally shown that if the measurement basis is a restricted Fourier transform, then a particularly good sparsity basis is the shift-invariant wavelet basis. Although it does not promote sparsity in seismic signals as effectively as curvelets, shift-invariant wavelets are extremely incoherent in the Fourier domain, which in turn leads to better performance when performing stable signal recovery from a restricted Fourier basis.

ALGORITHM FOR INTERPOLATING HELMHOLTZ SOLUTIONS

Recently an algorithm has been proposed by Herrmann and Hennenfent (2007) to solve problems which can be described by the form stated in Equation 7. The solver is based both on cooling method optimization and an iterative thresholding algorithm (see Daubechies et al. (2004)). The cooling method aims at finding the optimal multiplier λ^* for $\mathscr{L}(\mathbf{x},\lambda) := \lambda\|\mathbf{x}\|_1 + \|\mathbf{A}\mathbf{S}^H\mathbf{x}-\mathbf{y}\|_2^2 - \varepsilon^2$, the Lagrangian function of Equation 7, such that the residual $\mathbf{r} := \|\mathbf{A}\mathbf{S}^H\mathbf{x}-\mathbf{y}\|_2 \le \varepsilon$. For the sake of legibility in presenting the algorithm, in this section we overload the symbol $\mathbf{x}$ with the definition of $\check{\mathbf{x}}$ in Equation 7. The value ε is set by the desired accuracy of the returned vector $\mathbf{x}$ as a solution to the system $\mathbf{y} = \mathbf{A}\mathbf{S}^H\mathbf{x}$. The algorithm is as follows:

$\mathbf{x}_0 :=$ initial guess
$\lambda_0 :=$ initial Lagrange multiplier
while $\mathbf{r} > \varepsilon$
    $\min_{\mathbf{x}} \mathscr{L}(\mathbf{x},\lambda_k)$
    $\lambda_{k+1} = \alpha_k\lambda_k$ with $0 < \alpha_k < 1$
end while.

The critical part of this algorithm is the minimization of $\mathscr{L}(\mathbf{x},\lambda_k)$ done by the iterative thresholding algorithm presented in Daubechies et al. (2004). At each sub-iteration, evaluation of

$$\mathbf{x}_{i+1} = \mathscr{S}_{\lambda_k}\left(\mathbf{x}_i + \mathbf{S}\mathbf{A}^H(\mathbf{y}-\mathbf{A}\mathbf{S}^H\mathbf{x}_i)\right) \tag{10}$$

with

$$\mathscr{S}_{\lambda_k}(x) := \text{sign}(x)\cdot\max(|x|-\lambda_k,0) \tag{11}$$

yields an approximate estimate for $\mathbf{x}$ which converges to the solution of the sub problem for a large enough number of iterations. The next algorithm outlines this procedure.

Result: Estimate for $\mathbf{x}$
initialization;
$\mathbf{x}_0 \longleftarrow$ initial guess for $\mathbf{x}$;
$\lambda_0 \longleftarrow$ initial Lagrange multiplier;
while $\|\mathbf{A}\mathbf{x}-\mathbf{y}\|_2 \ge \varepsilon$ **do**
    while $\|\mathbf{x}_{i+1}-\mathbf{x}_i\|_2 \ge \hat{\varepsilon}$ **do**
        $i \longleftarrow i+1$;
        $\mathbf{x}_{i+1} \leftarrow \mathscr{S}_{\lambda_k}\left(\mathbf{x}_i + \mathbf{S}\mathbf{A}^H(\mathbf{y}-\mathbf{A}\mathbf{S}^H\mathbf{x}_i)\right)$;
    end
    $\lambda_{k+1} \longleftarrow \alpha_k\lambda_k$ with $0 < \alpha_k < 1$;
    $k \longleftarrow k+1$;
end

Algorithm 1: Iterative soft thresholding

In the above algorithm $\hat{\varepsilon}$ refers to the tolerance (or accuracy) to which we want to the solution of $\min_{\mathbf{x}}\mathscr{L}(\mathbf{x},\lambda_k)$ and in practice, one only needs to approximately solve each sub problem, which significantly accelerates the overall procedure. In the examples in the next section we run the algorithm for a fixed number of inner and outer iterations.

NUMERICAL EXPERIMENT

We perform recovery experiments from forward-modeling in the full 2D *hard* (unsmoothed) Marmousi velocity model of size 762-by-2502 with spatial sampling $\Delta x = \Delta z = 4.167$m, shown in Figure 1. A single Ricker wavelet shot centered around 15Hz is fired at $t = 0.2$s, $(x,z) = (5200\text{m}, 200\text{m})$ as indicated by the orange arrow. Receivers are placed at depth $z =$ 2500m and is depicted by the gray line. They are activated for a total of 5 seconds with $\Delta t = 20$ms for a total of 250 time samples. The corresponding discretized frequency domain consists of 125 frequency samples from 0 to 50Hz at $\Delta f = 0.4$Hz. Using the MKMG method as described above, we calculated to convergence a full set of solution to the Helmholtz equation for this problem in the frequency domain.

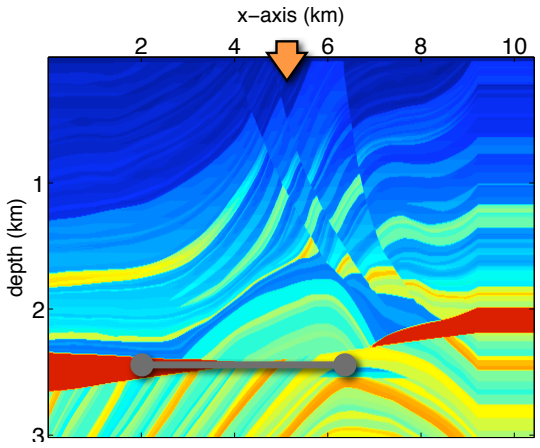

Figure 1: The hard marmousi model used in numerical experiment. Orange arrow indicates shot position and the grey line represents receiver positions.

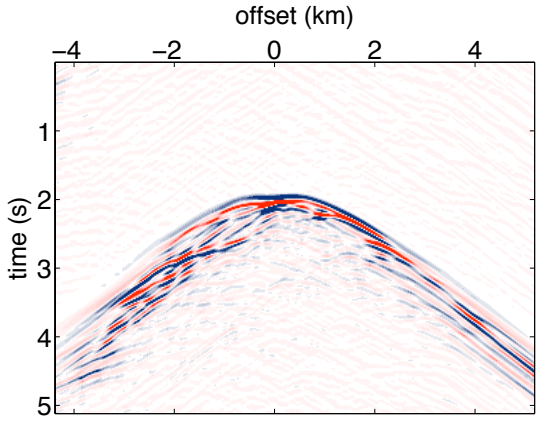

Figure 2: The full time-domain wavefield solution at depth $z = 2500$m

Figure 2 shows the time-domain wavefield observed by the full range of receivers at $z = 2500$m using a complete set of the solution $\mathbf{u}(\omega)$. To better present the quality of the interpolation, we windowed the traces to a range of 2000 to 6350m, indicated by the ball ends of the grey line in Figure 1. To obtain $\hat{\mathbf{u}}$, we selected subsets of frequencies at certain percentages of the whole domain discretized frequency domain and used it to define a frequency-domain masking operator $\mathbf{R}$. These frequencies are chosen from a weighted random sampling scheme over the whole domain, using the power spectrum of the shot source signature as weights. We then attempted to recover the complete solution $\mathbf{u}(\omega)$ using the iterative soft thresholding algorithm described in the previous section with $\mathbf{A} := \mathbf{RF}$. The sparsity transform $\mathbf{S}$ is chosen to be the shift-invariant wavelet transform with 7 scales in both axis.

The interpolation results are shown in Figure 3. Figure 3(a) shows the full solution using all frequencies of the domain. Figures 3(b)-3(f) shows the frequency-domain interpolated result from stable signal recovery, such that $\hat{\mathbf{u}}(\omega)$ is the subset of solutions from 50%, 40%, 30%, 20%, 10% of all available frequencies, respectively. The signal-to-noise ratios of the interpolated results are in Table 1. The results shown here indicates that we can limit ourselves to only a random subset of 30% of the full set of frequencies when computing for the Helmholtz solutions. The stable signal recovery algorithm shown in this abstract can then be used to recover the full time-domain wavefield.

% freq. used	50	40	30	20	10
SNR	24.55	17.36	17.02	12.50	7.81

Table 1: SNR values for the frequency-domain interpolated wavefields

CONCLUDING REMARKS

The results from our numerical experiments are acquired under conditions which are known to behave poorly for numerical computations, using the unsmoothed Marmousi model and setting receivers at a large depth for significant scattering. Even then, we see that not all frequencies are required to obtain the full result from frequency-domain modeling. With the model of stable signal recovery formulated under the theory of compressed sensing, we showed that even under such unfavorable condition a 30% subset of frequencies is enough to contain most of the wavefield information in time domain. Together with the newly developed MKMG algorithm as a practical means of obtaining solutions at high frequencies, exploration seismologists now have some powerful tools for practically working in the frequency domain.

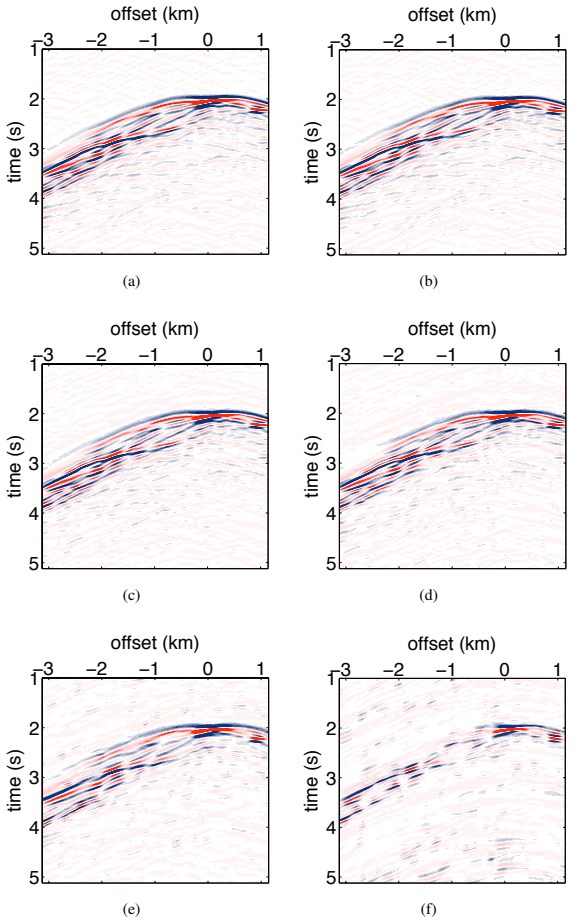

Figure 3: Frequency domain interpolation results. **(a)** the original full solution windowed in the x-axis from 2000m to 6350m, **(b)** wavefield interpolated with iterative soft thresholding using 5 inner loop, 30 outer loop with 50% frequencies represented in $\hat{\mathbf{u}}(\omega)$, **(c)** same as (a) but with 40% frequencies used, **(d)** 30%, **(e)** 20%, **(f)** 10%

ACKNOWLEDGEMENTS

This work was in part financially supported by the NSERC Discovery (22R81254) and CRD Grants DNOISE (334810-05) and was carried out as part of the SINBAD project with support, secured through ITF, from BG Group, BP, Chevron, ExxonMobil and Shell.

REFERENCES

Candès, E., J. Romberg, and T. Tao, 2006, Stable signal recovery from incomplete and inaccurate measurements: Comm. Pure Appl. Math., **59**, 1207–1223.

Daubechies, I., M. Defrise, and C. de Mol, 2004, An iterative thresholding algorithm for linear inverse problems with a sparsity constraint: Comm. Pure Appl. Math., 1413–1457.

Donoho, D. L., 2006, Compressed sensing: IEEE Trans. Inform. Theory, **52**, 1289–1306.

Erlangga, Y. A. and R. Nabben, 2007, On multilevel projection krylov method for the preconditioned helmholtz system: Elec. Trans. Numer. Anal., submitted.

———, 2008, Multilevel projection-based nested Krylov iteration for boundary value problems: SIAM J. Sci. Comput., to appear.

Hennenfent, G. and F. J. Herrmann, 2006, Seismic denoising with non-uniformly sampled curvelets: Comp. in Sci. and Eng., **8**, 16–25.

Herrmann, F. J. and G. Hennenfent, 2007, Non-parametric seismic data recovery with curvelet frames: Technical report, UBC Earth & Ocean Sciences Department. (TR-2007-3).

Lin, T. T. Y. and F. J. Herrmann, 2007, Compressed wavefield extrapolation: Geophysics, **72**, SM77–SM93.

Mulder, W. and R. Plessix, 2004, How to choose a subset of frequencies in frequency-domain finite-difference migration: Geoph. J. Int., **158**, 801–812.

Sergue, L. and R. G. Pratt, 2004, Efcient waveform inversion and imaging: A strategy for selecting temporal frequencies: Geophysics, **69**, 231–248.

Symes, W. W., 2007, Reverse time migration with optimal checkpointing: Geophysics, **72**, SM213–SM221.