Row-action inversion of the Barrick-Weber equations

J. J. Green*, L. R. Wyatt†

2 Jan 2006

Abstract

The Barrick-Weber equations describe the interaction of radar signals with the dynamic ocean surface, and so provide a mathematical basis for oceanic remote sensing. We describe the inversion of these equations with several of the row-action methods commonly used to solve large linear systems with unstructured sparsity. We find that the performance of the methods in inverting both synthetic and measured Doppler spectral data is comparable, with the method of Chahine-Twomey-Wyatt offering a slight advantage in the reliability of the recovery of the full directional wave spectrum and of parameters derived from its integration. We conclude with some remarks and open questions on the ill-posedness of the inversion problem.

1. Introduction

The Doppler spectrum of a high frequency (HF) radar signal backscattered from the ocean's surface contains a wealth of information on the sea-state, and so naturally has attracted interest from the remote sensing community. As was shown in Weber and Barrick (1977) and Barrick and Weber (1977), the interaction of the radar with the sea's surface admits a perturbation analysis on the Doppler spectrum σ,

$$\sigma = \sigma_1 + \sigma_2 + \cdots$$

The first term of this analysis is found to be a linear combination of Dirac delta functions δ:

$$\sigma_1(\omega, \phi) = 2^6 \pi k_0^4 \sum_{m=\pm 1} S(-m\mathbf{k}_0)\delta(\omega - m\omega_\mathrm{b}), \tag{1}$$

where ω is the angular frequency of the Doppler shift, $\mathbf{k}_0$ is the wave vector of the incident radar signal with wavenumber k_0 and direction ϕ, S is the ocean wavenumber spectrum and ω_b is the frequency of the Bragg matched waves

$$\omega_\mathrm{b} = \sqrt{2gk_0}$$

*j.j.green@shef.ac.uk, Department of Applied Mathematics, University of Sheffield, UK.
†l.r.wyatt@shef.ac.uk, Department of Applied Mathematics, University of Sheffield, UK.

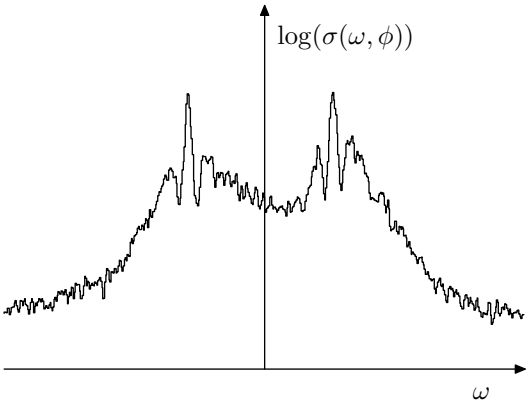

Figure 1: An example of measured Doppler spectra

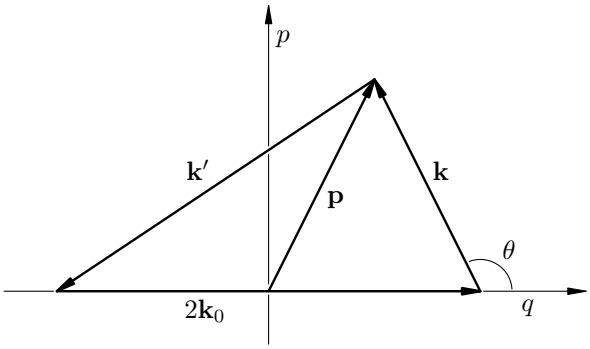

Figure 2: Geometry of the wavevectors at second order

where g is the acceleration due to gravity.

The equation (1) predicts impulses in the Doppler spectrum of the backscatter at shift frequencies $\pm\omega_{\mathrm{b}}$, features which are evident in the measured spectrum, such as is illustrated in Figure 1.

The second order of the perturbation analysis leads to the nonlinear integral equation

$$\sigma_2(\omega,\phi) = 2^6\pi k_0^4 \sum_{m,\,m'=\pm 1} \int_{\mathbf{R}^2} |\Gamma|^2\, S(m\mathbf{k})S(m'\mathbf{k}')\delta(\omega - m\sqrt{gk} - m'\sqrt{gk'})\,\mathrm{d}\mathbf{p}, \tag{2}$$

where the wave vectors $\mathbf{k}$ and $\mathbf{k}'$ satisfy the Bragg resonance condition $\mathbf{k}+\mathbf{k}' = -2\mathbf{k}_0$, and are related to $\mathbf{p}$, the variable of integration, by

$$\mathbf{k} + \mathbf{k}_0 = \mathbf{p}, \quad \mathbf{k}' + \mathbf{k}_0 = -\mathbf{p},$$

as illustrated in Figure 2.

The kernel of (2) is determined by the coupling coefficient $\Gamma = \Gamma(\omega, k, \theta, mm')$ accounting for nonlinear hydrodynamic and electromagnetic effects, and is described in detail in the articles by Weber and Barrick (1977); Barrick and Weber (1977); Barrick and Lipa (1986).

The Barrick-Weber equations are amenable to linearisation because the signal is dominated by the interaction between long and short ocean waves and the latter can be approximated using wind-wave models (see, for example Wyatt (1986), Lipa and Barrick (1986)). In Green (2003), it is observed that the linearised equation may be viewed, locally, as a weighted projection transform. This leads to a method for discretisation of the linearised equations based on the techniques for that of the general unweighted projection transform, i. e., the discretisation of transmission tomography problems. The discretisation reduces inversion of the linearised Barrick-Weber equations to the solution of a linear system

$$\boldsymbol{\sigma} = A\boldsymbol{\xi} \tag{3}$$

where the vector $\boldsymbol{\sigma}$ represents the normalised measured Doppler spectral values (typically from more than one radar), $\boldsymbol{\xi}$ contains the coefficients of the normalised ocean-wave spectrum representation, and A is the discretisation matrix. The matrix A is large (in practice around 200 rows and 2000 columns) but possesses an unstructured sparsity which can be exploited to obtain rapid solutions to (3), and so to the problem of estimating the ocean-wave directional spectrum from HF backscatter measurements.

In this note we investigate the use of row-action methods in the solution of the discretisation (3). We will compare synthetic wave spectra with those inverted from the deduced backscatter, and compare several ocean parameters as obtained from the inversion of measured backscatter with those obtained from a collocated buoy. In addition to comparing three row-action inversion methods, we seek to optimise a number of parameters controlling these inversions, for the purpose of calibration for forthcoming deployments.

2. Background

In this section we describe the inversion method used for the linearised Barrick-Weber equations. We begin with as full a summary of the functional representation and discretisation as is required to introduce our notation; a more detailed account can be found in Green (2003).

a. Functional representation

Following Weber and Barrick (1977); Barrick and Weber (1977) we consider the normalised directional wave-spectrum $Z = S/(2k_0)^4$, and seek a representation of Z as a linear combination of basis functions b_i

$$Z(\mathbf{y}) = \sum_i \xi_i b_i(\mathbf{y}) \tag{4}$$

where $\mathbf{y}$ is the vector in the direction of $\mathbf{k}$, but with a magnitude $y = \sqrt{k/2k_0}$ (a transformation chosen so that a uniform discretisation of $\mathbf{y}$ scales approximately with the data, see Green (2003, §3) for details). For the basis functions b_i we

take translates of a single function Ψ, radially symmetric with respect to the Euclidean norm,

$$b_i(\mathbf{y}) = \Psi(\mathbf{y} - \mathbf{y}_i) = \psi(\|\mathbf{y} - \mathbf{y}_i\|). \tag{5}$$

The function's window, ψ, is the Lewitt-Kaiser-Bessel function (see Lewitt, 1990)

$$\psi(r) = \begin{cases} (1 - (r/a)^2)^{m/2} I_m(\alpha\sqrt{1 - (r/a)^2})/I_m(\alpha) & (r < a) \\ 0 & \text{otherwise} \end{cases}, \tag{6}$$

where I_m is the modified Bessel function of order m, as in Watson (1944). In (6), a is the radius of support of ψ, α controls the localisation of ψ about the origin, and m determines the smoothness of ψ at $r = a$ and the rate of decay of its Fourier transform.

Functional representation in the form (5) was introduced by Lewitt (1990) in the context of transmission tomography, i. e., in the inversion of the x-ray transform. Such a representation has a number of advantages over the usual pixel-based representation, in particular:

- the smoothness of the basis function, and so of the representation, can be controlled by the adjustment of the parameter m, a fact which is important if the inversion problem is ill-posed;
- the function Ψ is computationally attractive since most interesting quantities associated with it: Fourier transform; gradient; projection or Abel transform, can be calculated explicitly as a result of the recurrence relations enjoyed by the Bessel function (see Lewitt, 1990);
- the choice of parameters a, m and α can be made so as to give the representation beneficial approximation properties (Lewitt, 1990; Green, 2002).

In our implementation of the representation, the centres $\mathbf{y}_i$ are placed on a uniform grid in the $\mathbf{y}$-domain, i. e., on the cylinder. Following the deliberations in Green (2002) we initially choose $m = 2$, $a/g = 1.78$ and $\alpha = 9.2$, where g is the grid spacing: the distance between adjacent centres $\mathbf{y}_i$.

b. Discretisation

The discretisation of the linearised Barrick-Weber equations with the Lewitt-Kaiser-Bessel representation (4) is straightforward. Each Doppler spectral datum

$$\sigma_j = \frac{1}{\Delta\eta} \int_{\eta_j}^{\eta_{j+1}} \sigma(\eta)\, \mathrm{d}\eta$$

(where $\Delta\eta$ is the sample width of the normalised Doppler shift $\eta = \omega/\omega_\mathrm{b}$) corresponds to a strip integral over the normalised directional spectrum $Z = (2k_0)^4 S$, and making a suitable approximation we find that

$$\sigma_j = \frac{1}{\Delta\eta} \sum_i \Lambda_i B_{i,j} \xi_i$$

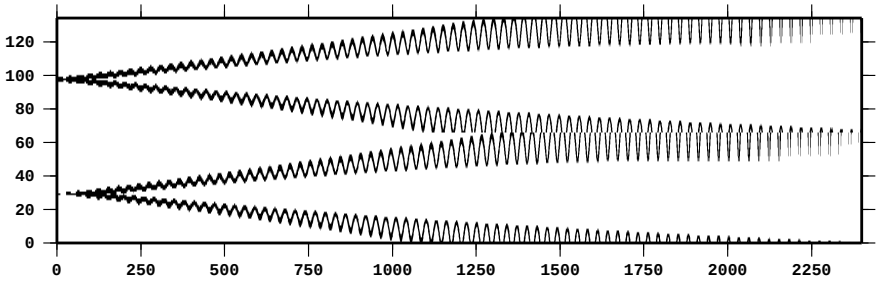

Figure 3: Sparsity structure of the discretisation matrix A, nozero entries shown black. Note the exaggerated vertical scale of this 128×2400 matrix.

where Λ is an expression accounting for the kernel of the Barrick-Weber equations, a linearisation term and Jacobians of domain transformations (see Green (2003, Eq. 2.6)), $\Lambda_i = \Lambda(\mathbf{y}_i)$ and $B_{i,j}$ is the jth strip integral over the ith basis function. Thus the elements, $A_{i,j}$, of the matrix A in (3) are given by

$$A_{i,j} = \frac{1}{\Delta\eta}\Lambda_i B_{i,j}.$$

Note that the matrix $B = [B_{i,j}]$ is sparse provided that the supports of the basis functions are small, for then each integration strips intersects only a few of them. Naturally the discretisation matrix, A, inherits this sparsity, as seen in Figure 3.

In our implementation the number of basis functions used is dependant on the wavenumber limits, which are constrained by the radar frequency and noise level, and the choice of the parameters controlling the functional representation as mentioned above. For the examples in this note a typical discretisation has 30 centres in angle and 60-80 in wavenumber, so 1,800–2,400 basis functions are used.

c. Inversion by row-action method

Large sparse linear systems are common in many applications, and numerous techniques are available for solving them. Systems, such as (3), having an (essentially) unstructured sparsity, are an important subclass as they arise in the discretisation of tomographic and similar problems. A popular class of solution methods are the *row-action* methods — iterative techniques, where the next iterate $\boldsymbol{\xi}^{(k+1)}$ is found using only the data $\boldsymbol{\xi}^{(k)}$, $\boldsymbol{\sigma}$ and the $i(k)$th row $A_{i(k)}$ of the matrix A, where the sequence $(i(k))$ cycles through the rows of A in some fashion.

The obvious computational advantages of row-action methods has prompted a substantial research effort into their convergence properties, techniques which improve their speed and stability, and so on. We refer the interested reader to Censor's (1981) excellent review.

It is notable that the algorithm of Wyatt for the inversion of the linearised Barrick-Weber equations (Wyatt, 1990), has many of the characteristics of a row-action method. The success of Wyatt's algorithm in practical real-time oceanic remote-sensing applications (Wyatt et al., 2003), provided the motivation for the discretisation described above, and one finds that the natural reimplementation of the iterative step in Wyatt's algorithm produces an *explicit* row-action method. In the remainder of this note we investigate and optimise this and two other row-action methods for the solution of the linearised Barrick-Weber equations.

1) Algebraic reconstruction technique

The algebraic reconstruction technique (ART) is the oldest of the row-action methods, being proposed as a solution method for the convex feasibility problem by Kaczmarz in 1937. The algorithm was later rediscovered by Gordon et al. (1970) and implemented in the EMI computerised tomography scanner (Hounsfield, 1973). For ART the iterative step is

$$\xi_j^{(k+1)} = \xi_j^{(k)} + \lambda \frac{\sigma_{i(k)} - A_{i(k)} \cdot \boldsymbol{\xi}^{(k)}}{\left\| A_{i(k)} \right\|_2^2} A_{i(k),j} \qquad (j = 1, \ldots, n),$$

where $0 < \lambda < 2$ is the *relaxation* parameter. The ART has an attractive geometric interpretation when $\lambda = 1$, for then the iterative step gives the orthogonal projection of $\boldsymbol{\xi}^{(k)}$ onto the hyperplane defined by the $i(k)$th row of the system (3).

A number of theoretical results are available for ART, particularly in the underrelaxed ($0 < \lambda < 1$) case. For example, when applied to a consistent but underdetermined system, underrelaxed ART will converge to the solution of minimum norm (Censor, 1981, §4.4). It has also been shown (Fleming, 1990) that the early termination of the ART iteration is equivalent to Tikhonov regularisation, an important consideration when the problem discretised is ill-posed, as is the case for the inversion of the x-ray transform.

2) Multiplicative ART

The multiplicative algebraic reconstruction technique (MART) is the row-action method with iterative step

$$\xi_j^{(k+1)} = \left(\frac{\sigma_{i(k)}}{A_{i(k)} \cdot \boldsymbol{\xi}^{(k)}} \right)^{\lambda A_{i(k),j}} \xi_j^{(k)} \qquad (j = 1, \ldots, n).$$

It can be shown that, if the system $\boldsymbol{\sigma} = A\boldsymbol{\xi}$ is consistent then this iteration converges to the solution $\boldsymbol{\xi}$ which minimises the *entropy*

$$\sum_j \xi_j \log \xi_j,$$

provided that the elements of A satisfy $0 \leq A_{i,j} \leq 1$. To ensure that this condition is satisfied, Byrne (2000) recommends a rescaling of the system to obtain the iteration

$$\xi_j^{(k+1)} = \left(\frac{\sigma_{i(k)}}{A_{i(k)} \cdot \boldsymbol{\xi}^{(k)}} \right)^{\lambda A_{i(k),j}/m} \xi_j^{(k)} \qquad (j = 1, \ldots, n),$$

where $m_i = \max\{A_{i,j} : j = 1, \ldots, n\}$. This rescaling is included in our implementation of MART.

3) CHAHINE-TWOMEY-WYATT

Our final algorithm is a reimplementation of that proposed by Wyatt (1990) specifically for the inversion of the linearised Barrick-Weber equations. The original specification of the algorithm included a discretisation which has been replaced by that described in Section 2.2. Wyatt's iteration, a two-dimensional version of that of Chahine (1968) later modified by Twomey (1996), has the iterative step

$$\xi_j^{(k+1)} = \left(1 + \lambda \frac{A_{i(k),j}}{\left\| A_{i(k)} \right\|_\infty} \left(\frac{\sigma_{i(k)}}{A_{i(k)} \cdot \boldsymbol{\xi}^{(k)}} - 1 \right) \right) \xi_j^{(k)} \qquad (j = 1, \ldots, n).$$

Here $\|A_i\|_\infty$ is the maximum value of the $A_{i,j}$ for $j = 0, \ldots, n$ and, again, $0 < \lambda \leq 1$ is the relaxation parameter. The motivation for the form of the iteration is detailed by Twomey (1996) who, interestingly, notes his impression that the iteration is particularly suited to the recovery of functions with a large dynamic range — a typical property of ocean-wave directional spectra.

3. Synthetic spectra

To investigate the behaviour of the iterative schemes in inversion of the wave spectrum we have generated synthetic Doppler spectral test data. Note that there is no question of an inverse crime (Colton and Kress, 1998, pp. 133, 304) being committed here: our synthesis of Doppler spectra is by a direct discretisation of the nonlinear equation (2) along the contours defined by the Dirac constraint therein; our inversion uses the radial basis discretisation of the linearised equation described in Section 2.2.

Synthetic Doppler spectra were produced assuming directional ocean spectra with wind-waves of Pierson-Moskowitz type, in some cases with added swell, and using the Barrick-Weber equation (2). Two such Doppler spectra were then inverted using the various methods and these compared with the original model spectrum. The likeness for the various inversion was broadly similar, but with a definite superiority in recovery for CTW over ART, and for ART over MART.

A typical recovery for CTW with unit relaxation and 200 iterations is shown in Figure 4, where we note that the inversion has the gross features of the original wave-spectrum. However, as can be seen in Figures 4(a) and 4(b),

there is some "splitting" of the spectral peak of the recovered spectrum at lower wave frequencies, and the emergence of spurious modes (this splitting is more pronounced for the ART and MART inversions, not shown). We interpret this as indicating that the smoothing imposed by the discretisation alone is not sufficient to regularise the inversion problem. Moreover, there is a limit on how much smoothing can be imposed by our discretisation: the radial basis functions must have small support if the approximations underpinning the discretisation are to remain valid.

Further, the regularising effects of taking a small relaxation parameter or of the use a small number of iterations are difficult to exploit here; such inversions have problems recovering swell components of the directional wave spectrum which typically need scores of iterations to emerge. We would suggest that the underlying problem is that row-action methods not address the anisotropy of the inversion problem.

Our solution is to apply a mild smoothing on candidate solutions between iterative steps, implemented (for simplicity) in a 3-point smoothing kernel applied to the representation coefficients in both wavenumber and direction. We find that good results are obtained if less smoothing is applied to higher frequencies than to lower, and so apply with a kernel $[0.1, 0.8, 0.1]$ for large wavenumbers decreasing linearly to $[0.2, 0.6, 0.2]$ for small wavenumbers. This seems to provide sufficient regularisation to produce inverted wave spectra without peak-splitting, as shown Figure 4(c). Moreover, this smoothed algorithm is able to distinguish between the different modes in a genuinely bimodal spectrum, as in Figure 4(d).

4. Wavebuoy comparison

EuroROSE (Wyatt et al., 2003) was an EU-funded project to demonstrate the use of radar sensors for vessel traffic service applications. The WERA HF radar, developed by the University of Hamburg (Gurgel et al., 1999) was deployed in two experiments in 2000. One of these was at the Norwegian coastal islands north-west of Bergen, the site of the Vessel Traffic Service centre guiding large oil tankers into mainland ports. The radar was deployed for a period of about 6 weeks providing measurements of the wave and current field from the coast to up to 40km offshore with 1km resolution every ten minutes. A microwave X-band radar, a directional wavebuoy, an acoustic-Doppler current profiler (ADCP) and wave and current models were also used and the radar and model data was made available in near real time to the staff in the centre.

Detailed comparisons between the wave-measuring instruments and model are presented in Wyatt et al. (2003). These confirmed previous work (Wyatt et al., 1999) showing an overestimation in HF radar measured waveheight and underestimation in wave period in high sea conditions. This is thought to be due to limitations in the scattering model used by Weber and Barrick (1977) and is the subject of current research. For the assessment of the accuracy of the inversion method discussed here we have therefore selected a period of relatively low seas. Since the experiment took place in the winter on an exposed

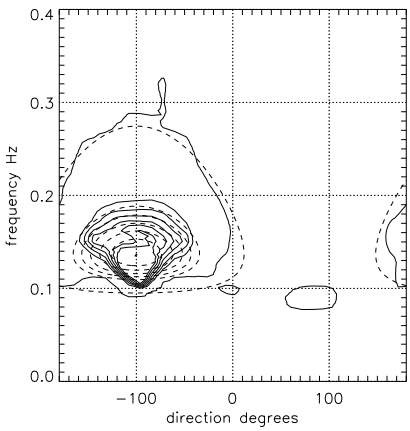

(a) Monochromatic 10ms^{-1} wind-sea spectrum, inverted without smoothing

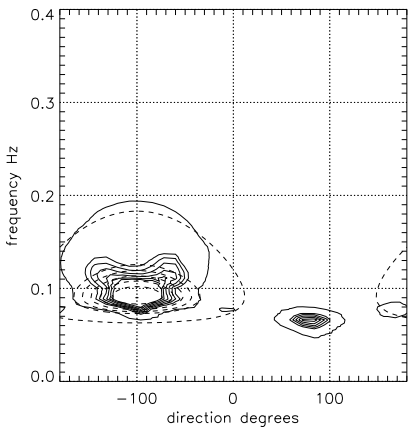

(b) Monochromatic 15ms^{-1} wind-sea spectrum, inverted without smoothing

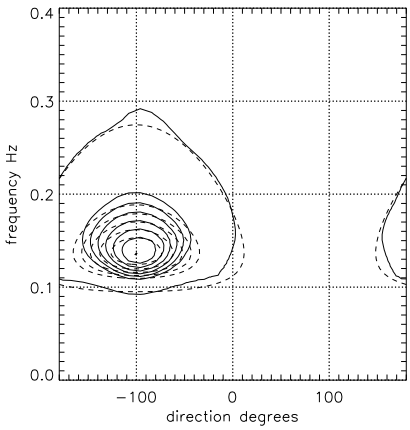

(c) As (a), but inverted with smoothing

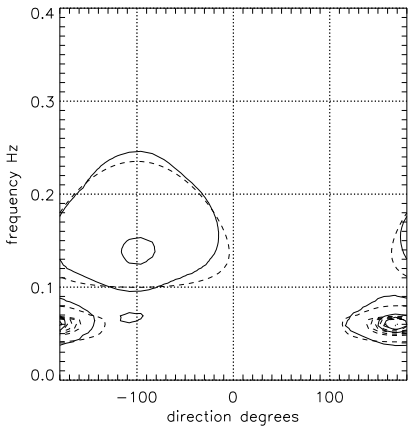

(d) Bimodal spectrum, with smoothing

Figure 4: Reconstruction of ocean wave spectrum with CTW (solid) from Doppler spectra calculated from a synthetic Pierson-Moskowitz ocean spectrum (dotted).

North-Atlantic coast and was therefore dominated by storms, the period selected is quite short: 17–21 Feb. 2000, but does include enough variation in waveheight and period (see Figure 5) and also in direction to give us confidence in the generality of our results as far as the numerical methodology is concerned. The comparison presented here is with a directional waverider buoy which is currently the accepted standard for wave measurements. This buoy provides measurements of five Fourier coefficients of the directional distribution, a_0, a_1, a_2, b_1, b_2 all functions of wave frequency (Tucker, 1991). This is thus a more limited measurement than can be obtained with the radar which provides the full directional spectrum, $S(\mathbf{k})$. Here we will focus on just three integrated parameters that are commonly used to describe the wavefield, significant waveheight, mean period and mean direction. We refer the reader to Tucker (1991) for the definitions of and formulae for these quantities.

The data available for the comparison were recorded at different times: inverted spectra every 10 minutes, buoy data every 30. Consequently we paired each buoy observation with the nearest inversion (with a maximum difference of 10 minutes) yielding 334 such pairs in the test period. Summary statistics on the differences between the integrated parameters were then calculated for a range of values of parameters affecting the row-action inversion; the number of iterations and the relaxation parameter λ amongst them.

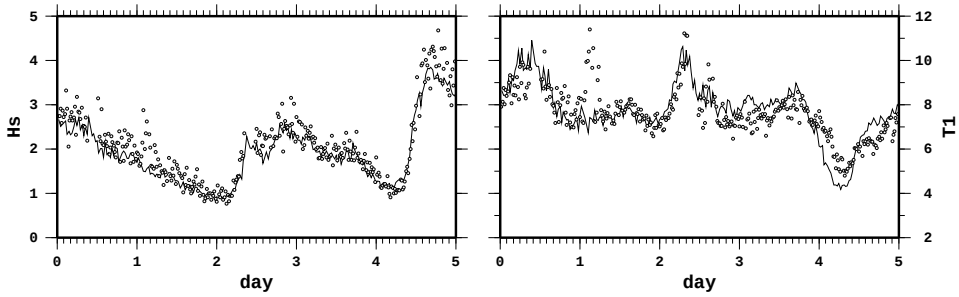

Figure 5: Significant waveheight (Hs, left) and mean period (T1, right) measured by buoy (solid) and CTW inversion with 200 iterations and smoothing (circles) over the test period

Sample results from the comparison for the number of iterations (in the range 1–500) with unit relaxation are illustrated in Figure 6. As can be seen, each row-action method has rather similar behaviour, with the mean difference between buoy and inversion decreasing to a stable value with increasing iterations. The variance in the difference behave similarly, albeit with more iterations needed to achieve stability, particularly the differences in mean direction (bottom row of Figure 6).

Some differences between the row-action methods are apparent from these comparisons: the MART method seems to have a small negative bias in mean period, an effect which is not so pronounced in ART and CTW, and clearly ART

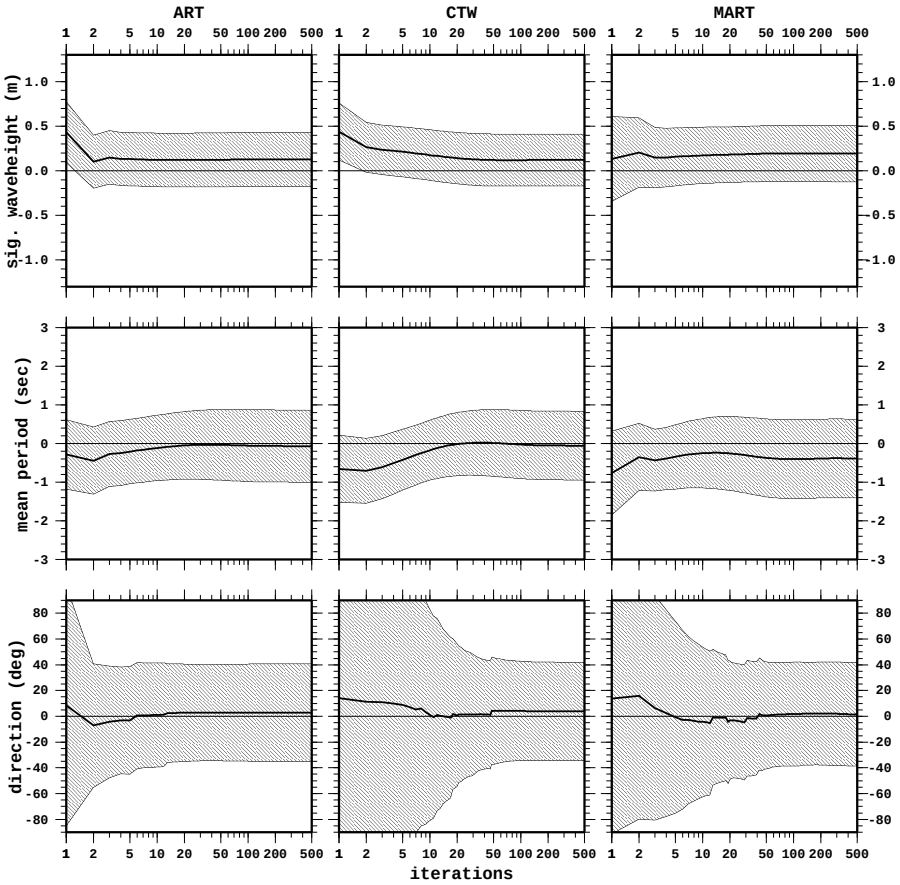

Figure 6: Mean absolute disrepancy between buoy and row-action inversion: various parameters against the number of iterations. The shaded region covers a standard deviation and so (assuming Gaussian statistics) would be expected to contain the inner 38% of the sample.

outperforms both CTW and MART in the number of inversions needed to achieve convergence to the mean direction value.

The comparisons for various values of the relaxation parameter λ (not shown for reasons of space) for a fixed number of iterations reveal a similar picture, with the difference between buoy and inversion decreasing to a stable value as λ increases from a small value. Indeed, the results suggest that the inversion behaves as if a superposition were occurring: a row action inversion with n iterations and relaxation of λ is very similar to one with $2n$ iterations and relaxation of $\lambda/2$.

We also mention that the comparisons seem rather insensitive to variation in a number of other details of the row-action method, and even the discretisation. There seems to be little effect on the buoy intercomparison when varying the the amount of smoothing between iterations (in contrast to the results of Section 3, but perhaps not surprising given that we are comparing integrated parameters). We also find that reordering the rows in the discretisation matrix prior to inversion (a common strategy for improving convergence of row-action inversion of problems in transmission tomography) has little effect on the statistics. Finally, we find that varying the number of basis functions used in the discretisation has little effect above a certain threshold; this threshold presumably determined by breakdown in validity of a number of approximations used in the derivation of the discretisation.

5. Ill-posedness

Many problems of mathematical physics can be expressed in the form $y = Ax$, where A is a (possibly nonlinear) operator between function spaces, y is some measured data and x is to be found. Haddamard considered such problems and declared *well-posed* those for which a solution x exists, is unique, and depends continuously on the data y. Ill-posed problems, those which are not well-posed, are by no means unphysical; the equation $y = Ax$ is ill-posed whenever the operator A acts to smooth its argument, as is the case for the integral operators arising in the practical problems of metrology, geophysics and so on. Ill-posedness was, for many years, thought to be mathematically intractable; as late as 1961 Courant was to write, in *Methods of Mathematical Physics II*

> So far, unfortunately, little mathematical progress has been made in the important task of solving or even identifying such problems which are not "properly posed" but still are important and motivated by realistic situations.

The same decade was to see a revolution in the approach to ill-posedness, that of *regularisation*: approximation of an ill-posed problem by a *family* of well-posed problems, each, in a precise sense, near to the original. Tikhonov, Lavrentiev and Morozov in the East and John and Twomey in the West reduced the intractability to a problem of a careful choice of from the family of approximants.

As has been mentioned earlier in this note, an expectation that the inversion of the Barrick-Weber equations is ill-posed has been carried into the design of

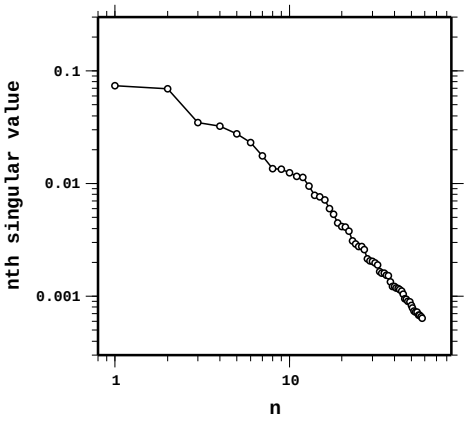

Figure 7: The first 60 singular values, σ_n, of a typical discretisation matrix

the discretisation and the choice of inversion methods, although there are, as far as we know, no analytic results in this direction. In the remainder of this section we present *circumstantial* evidence that the inversion *is* ill-posed, albeit mildly so.

The ill-posedness of the problem $y = Ax$ for linear A is intimately related to the singular-value expansion (SVE) of the operator A, as discussed widely in the literature (see Hansen 1998, §1.2.2,Ch 2.2 and Engl et al. 2000, §2.2, for example). For some problems, for example when A represents the Radon transform, the SVE can be obtained exactly (Natterer and Wübbeling, 2001, Th. 2.9). Even when this is not the case we can draw some conclusions from the singular value decomposition (SVD) of the discretisation matrix of the problem; the singular values of the operator approximate those of the matrix under mild conditions (Hansen, 1998, §2.1). With this motivation we have found the first 60 singular values of our discretisation of the linearised Barrick-Weber equations using the SVDPACK codes (Berry, 1992), the results shown in Figure 7. A regression on these values gives an estimate of $\sigma_n \sim n^{-1.9}$ and repeating the calculations for a range of 237 locations (and so discretisation geometries) as would occur in practical inversion we find a decay estimate of $\sigma_n \sim n^{\alpha}$, where α has mean -1.88 and standard deviation 0.105.

This apparent subquadratic decay of the singular values would suggest that the inversion is, to use the informal classification of Engl et al. (2000, §2.2), mildly ill-posed and so requiring only moderate regularisation to achieve stability. It is worth noting that the problem limited-angle tomography (inversion of the Radon transform when measured data is available in a restricted set of directions) is found to have exponential decay of the singular values, and so is severely ill-posed (Natterer and Wübbeling, 2001, §6.2), in marked contrast with the Barrick-Weber inversion (all inversion presented here used exactly two Doppler spectra and so there are four "directions"). A notable difference between the two problems is the highly directional kernel $|\Gamma|^2$ present in the Barrick-Weber equations.

6. Conclusions

Our investigations of the Barrick-Weber problem suggests that the smoothing imposed by our discretisation is not sufficient to fully stabilise the inversion by row-action method. However, the addition of a mild smoothing between row-action iterations does provide this regularisation giving reasonable results in inverting both synthetic and real Doppler spectral data. Consistent with the evidence of mild ill-posedness, we find that all of the row-action methods give similar results but with ART and CTW having a slight advantage over MART in the speed and accuracy of the recovery of integrated parameters. With the added smoothing and unit relaxation around 50–100 iterations seem sufficient to achieve convergence.

Further research is underway to determine methods whereby the smoothing in the inversion is incorporated more closely into the discretisation, so as to give a more fine-grained control over the regularisation.

Acknowledgements

The first author was partially supported in this work by the UK Natural Environment Research Council (NERC) under the grant NER/A/S/2001/00453.

References

Barrick, D. E. and Lipa, B. J. (1986). The second-order shallow-water hydrodynamic coupling coefficient in interpretation of HF radar sea echo. *IEEE J. Oceanic Eng.*, OE-11(2):310–315.

Barrick, D. E. and Weber, B. L. (1977). On the nonlinear theory for gravity waves on the ocean's surface. Part II: Interpretation and applications. *J. Phys. Oceanog.*, 7:11–21.

Berry, M. (1992). Large scale singular value computations. *Int. J. Supercomput. Appl.*, 6(1):13–49.

Byrne, C. (2000). Block-iterative interior point optimization methods for reconstruction from limited data. *Inverse Problems*, 16:1405–1419.

Censor, Y. (1981). Row-action methods for huge and sparse systems and their applications. *SIAM Rev.*, 23(4):444–466.

Chahine, M. T. (1968). Determination of the temperature profile in an atmosphere from its outgoing radiance. *J. Opt. Soc. Am.*, 58:1634–1637.

Colton, D. and Kress, R. (1998). *Inverse Acoustic and Electromagnetic Scattering Theory*, volume 93 of *Applied Mathematical Sciences*. Springer, Berlin, 2nd edition.

Engl, H. W., Hanke, M., and Neubauer, A. (2000). *Regularization of Inverse Problems*, volume 374 of *Mathematics and its Applications*. Klewer Academic Publishers, Dordrecht. Paperback edition.

Fleming, H. E. (1990). Equivalence of regularization and truncated iteration in the solution of ill-posed image reconstruction problems. *Linear Algebra Appl.*, 130:133–150.

Gordon, R., Bender, R., and Hermann, G. T. (1970). Algebraic reconstruction techniques (ART) for three-dimensional electron microscopy and x-ray photography. *J. Theoret. Biol.*, 29:471–481.

Green, J. J. (2002). Approximation with the radial basis functions of Lewitt. In Leversley, J., Anderson, I., and Mason, J. C., editors, *Algorithms for Approximation IV*, pages 212–219. University of Huddersfield.

Green, J. J. (2003). Discretising Barrick's equations. In Sajjadi, S. G. and Hunt, J. C. R., editors, *Proceedings of Wind over Waves II: Forecasting and Fundamentals of Applications*, pages 219–232, Chichester. IMA and Horwood.

Gurgel, K.-W., Antonischki, G., Essen, H.-H., and Schlick, T. (1999). Wellen Radar (WERA), a new ground-wave based HF radar for ocean remote sensing. *Coast. Eng.*, 37(3–4):219–234.

Hansen, P. C. (1998). *Rank-Deficient and Discrete Ill-Posed Problems*. SIAM Monographs on Mathematical Modeling and Computation. SIAM, Philadelphia.

Hounsfield, G. N. (1973). Computerized transverse axial scanning (tomography), I: Description of system. *Brit. J. Radiol.*, 46:1016–1022.

Lewitt, R. M. (1990). Multidimensional digital image representations using generalized Kaiser-Bessel window functions. *J. Opt. Soc. Am. A*, 7(10):1834–1846.

Lipa, B. J. and Barrick, D. E. (1986). Extraction of sea state from HF radar sea echo: Mathematical theory and modeling. *Radio Sci.*, 21(1):81–100.

Natterer, F. and Wübbeling, F. (2001). *Mathematical Methods in Image Reconstruction*. SIAM Monographs on Mathematical Modeling and Computation. SIAM.

Tucker, M. J. (1991). *Waves in ocean engineering: measurement, analysis, interpretation*. Ellis Horwood.

Twomey, S. (1996). *Introduction to the Mathematics of Inversion in Remote Sensing and Indirect Measurements*. Dover.

Watson, G. N. (1944). *A Treatise on the Theory of Bessel Functions*. Cambridge, second edition.

Weber, B. L. and Barrick, D. E. (1977). On the nonlinear theory for gravity waves on the ocean's surface. Part I: Derivations. *J. Phys. Oceanog.*, 7:3–10.

Wyatt, L. R. (1986). The measurement of the ocean wave directional spectrum from H.F. radar Doppler spectra. *Radio Sci.*, 21:473–485.

Wyatt, L. R. (1990). A relaxation method for integral inversion applied to HF radar measurement of the ocean wave directional spectrum. *Int. J. Remote Sens.*, 11:1481–1494.

Wyatt, L. R., Green, J. J., Gurgel, K.-W., Borge, J. C. N., Reichert, K., Günther, K. H. H., Rosenthal, W., Saetra, O., and Reistad, M. (2003). Validation and intercomparisons of wave measurements and models during the EuroROSE experiments. *Coast. Eng.*, 48:1–28.

Wyatt, L. R., Thompson, S. P., and Burton, R. R. (1999). Evaluation of HF radar wave measurement. *Coast. Eng.*, 37:259–282.