Stanford Exploration Project, Report 112, November 11, 2002, pages 151–161

Combined inversion: preconditioning with regularization

Marie L. Clapp[1]

ABSTRACT

Iterative inversion schemes are becoming more common in seismic processing. The high cost of the operators generally used in these inversion schemes makes it very important to minimize the number of iterations needed to obtain a good model. In complex environments, inversion schemes can be improved by styling the model through regularization or preconditioning. At early iterations, regularization provides a result that has a frequency content comparable to that of the "ideal" model. Preconditioning defines a solution at every model point at earlier iterations than regularization. An "improved" model should combine these two characteristics. This paper examines a scheme that uses the result of preconditioned inversion as an initial model for regularized inversion. I show that this scheme allows us to obtain an improved model in fewer iterations than would be needed for preconditioned inversion or regularized inversion alone.

INTRODUCTION

As the search for oil concentrates on ever more complicated areas of the subsurface, we find ourselves needing to balance the benefits of obtaining a better image with the cost of obtaining that better image. To obtain an ideal image, we would have to use an imaging operator that is the inverse of the physical operator propagating our seismic signal into the ground. However, imaging operators such as migration are adjoints rather than inverses (Claerbout, 1995), so in complex areas the resulting image may not be as good as it could be. Unfortunately, finding an operator that is an inverse in complex areas is almost impossible, so we generally approximate the inverse through a process like least-squares inversion (Nemeth et al., 1999; Duquet and Marfurt, 1999; Ronen and Liner, 2000). When using such an iterative technique, the result of iterating to convergence can be thought of as the "ideal" model.

Iterative inversion schemes often have trouble with problems that are unstable or where the mapping operator has a null space (Claerbout, 1991). These issues can be overcome by regularizing the problem (Tikhonov and Arsenin, 1977; Harlan, 1986; Fomel, 1997). However, our regularization operators, which are usually roughening operators, tend to be small. Their influence at any single iteration is limited in range. When our mapping operator has large areas that do not correspond to any data locations this can be especially troublesome. A solution to this problem is to perform a change of variables, turning it into a preconditioned problem

[1]**email:** marie @ sep.Stanford.edu
The author has recently changed her name from Marie Prucha to Marie Clapp.

(Fomel et al., 1997). Using the helix transform (Claerbout, 1998), we can apply the inverse of our small regularization operator, which will be a smoothing operator, whose influence extends a large distance. The advantage of this approach is that we quickly define our solution at all model points. The disadvantage of this approach is that our preconditioning operator dominates early iterations, creating a model that is often too low in frequency. What we ideally would like is a process where the solution is defined everywhere without the reduction in frequency content.

In Prucha and Biondi (2002), we presented a scheme that met these requirements by using the result of the preconditioned inversion as an initial model for a regularized inversion. From that example, and for the purposes of this paper, I will define an "improved" model as one that has a solution defined at every point and has a frequency content comparable to that of the "ideal" model. The combined inversion using preconditioning and regularization allows me to obtain an improved model with fewer iterations than would be needed using preconditioning or regularization alone. In this paper, I will take a closer look at the process of combined inversion.

In order to efficiently examine combined inversion with preconditioning and regularization (CIPR), this paper solves an interpolation problem which is much simpler than the imaging problem in Prucha and Biondi (2002). I will begin by explaining the constructed problem and the operator that is used for interpolation. Then I will present and discuss the results. Finally, I will explain my future plans for this combined inversion scheme.

CONSTRUCTING AN INTERPOLATION PROBLEM

The operators

In order to examine the results of preconditioned inversion, regularized inversion, and my proposed CIPR, I needed a problem that was easier to understand than that shown in Prucha and Biondi (2002). I am concerned with two issues: frequency content and solutions at every model point. To address the first issue, I chose to make my inversion operator a "smoother" that causes the model to have a higher frequency content than the data. This can be expressed as:

$$\mathbf{d} \approx \mathbf{S}\mathbf{m} \tag{1}$$

where $\mathbf{d}$ is the data, $\mathbf{m}$ is the model, and $\mathbf{S}$ is a smoothing operator that maps the average of 5 vertical points in the model to one point in the data. Since the model should be high frequency, the effects of the preconditioned inversion should be quite obvious.

Given such a simple inversion operator, creating a need for regularization or preconditioning requires that I cause the model created by inversion (fitting goal (1)) to have points that do not have solutions defined by the inversion operator. I chose to do this by introducing a masking operator $\mathbf{W}$. The combined operator $\mathbf{WS}$ will now have a null space where $\mathbf{W} = \mathbf{0}$. This changes my fitting goal to:

$$\mathbf{d} \approx \mathbf{W}\mathbf{S}\mathbf{m}. \tag{2}$$

To interpolate the model in the areas affected by the null space, I add a second fitting goal to fitting goal (2):

$$\begin{aligned} \mathbf{d} &\approx \mathbf{WSm} \\ \mathbf{0} &\approx \epsilon \mathbf{Am} \end{aligned} \tag{3}$$

where the new operator, **A**, is a regularization operator. I have chosen to make **A** a steering filter (Clapp et al., 1997; Clapp, 2001) generated as described in Prucha et al. (2000, 2001). Briefly, a steering filter consists of dip penalty filters at every model point, meaning that it is a non-stationary roughening operator that acts over short distances. To precondition this problem, I perform a change of variables to replace the model **m** with the preconditioned variable **p**:

$$\mathbf{m} = \mathbf{A}^{-1}\mathbf{p}. \tag{4}$$

Applying this to fitting goals (3) results in a new set of fitting goals:

$$\begin{aligned} \mathbf{d} &\approx \mathbf{WSA}^{-1}\mathbf{p} \\ \mathbf{0} &\approx \epsilon \mathbf{p}. \end{aligned} \tag{5}$$

The inverse of the steering filter ($\mathbf{A}^{-1}$) is applied using the helix transform. The inverse operator will be a smoothing operator that will act over a much larger distance than **A**.

The data

Given the operators I have chosen to use in this experiment, selecting data to test is straightforward. I need data that will result in a model that requires interpolation and will make differences in frequency content of various results obvious. Since the regularization operator is a steering operator, the data can have varying dips. To meet these simple requirements, I chose to take a 2-D slice from the familiar “qdome” model (Claerbout, 1995). The masking operator **W** contains enough zeros to defeat the inversion operator, making the regularization operator necessary. Figure 1 shows the data multiplied by the masking operator (**Wd**) I used for this experiment. I am displaying it this way to make comparison with the inversion results simpler. Figure 1 also shows the “ideal” model that would be obtained if **W** was simply an identity operator.

RESULTS

The first experiments I ran were to simply test the result of the regularized inversion (fitting goals (3)) and the preconditioned inversion (fitting goals (5)). I am concerned with the behavior in early iterations, so I just ran 6 iterations of each. These results are in Figure 2.

The regularized result is high frequency, but it has barely begun to fill in the areas affected by the null space. This is exactly the behavior we expect at early iterations in a regularized inversion. The preconditioned result has completely filled in the areas affected by the null

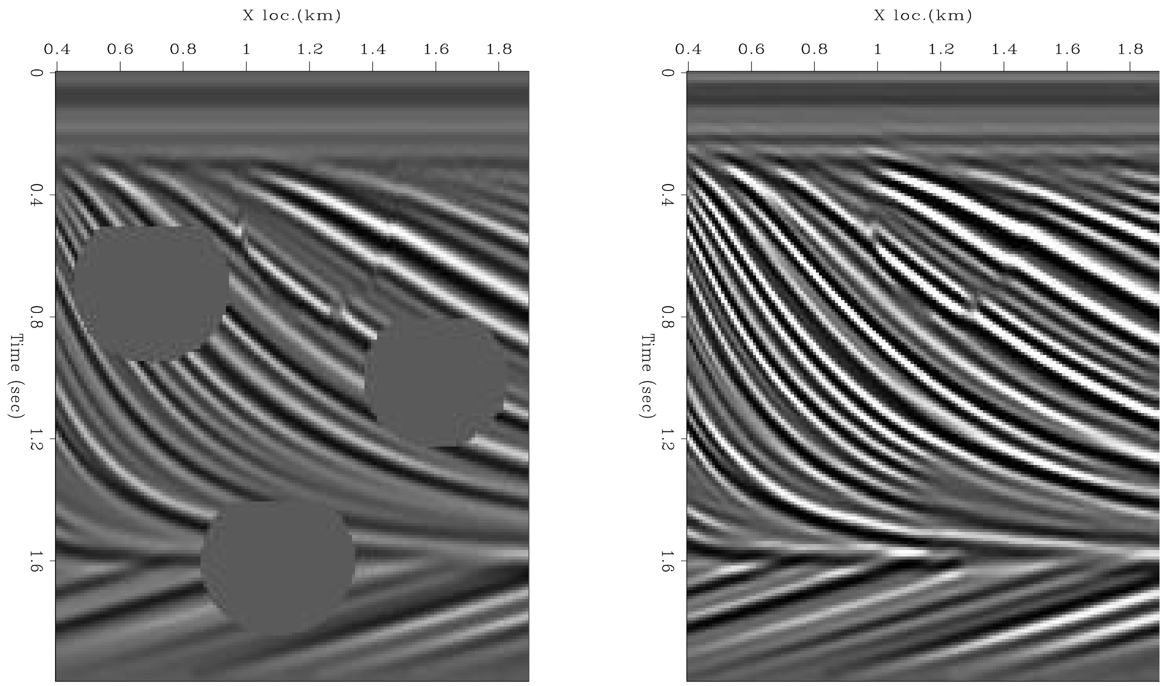

Figure 1: Left panel is the data weighted by the masking operator used for the inversion problems, right panel is the ideal model we get when the masking operator is replaced with an identity operator. marie1-datmod [ER,M]

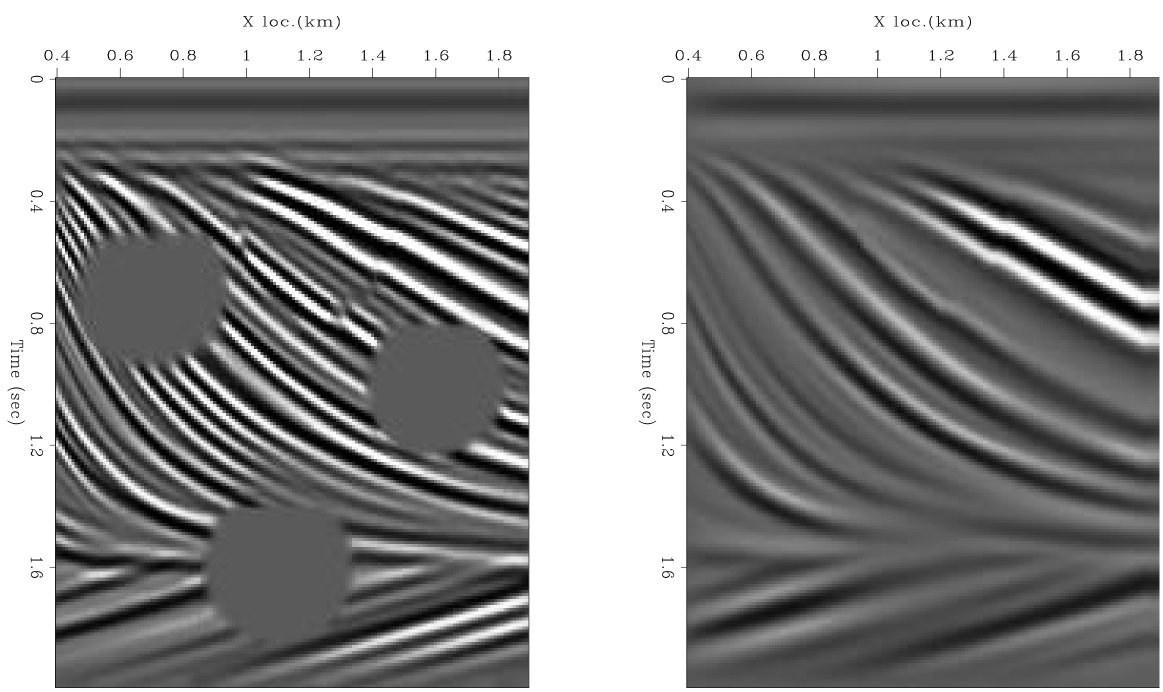

Figure 2: Left panel is the result of 6 iterations of just regularized inversion, right panel is the result of just 6 iterations of preconditioned inversion. marie1-regprec [ER,M]

space, meaning that it has defined solutions at every point (although not the ideal solution), but it is very low frequency. Once again, this is expected and has been seen in earlier works with imaging operators (Prucha et al., 2000; Prucha and Biondi, 2000, 2002).

The previous example helps to demonstrate two important points made by Claerbout (1999). First, both regularized inversion and preconditioned inversion take a great many iterations to converge. While this is not a problem for a toy problem like the one presented in this paper, it is impossible for a real geophysical problem like imaging in complex areas. The operators used in such a problem are infinitely more complex than those used in this simple interpolation problem, so it is vital that we minimize the number of iterations needed (Biondi and Vlad, 2001). Secondly, when we limit ourselves to a small number of iterations, we encounter several problems with both regularization and preconditioning. These problems include:

- A regularized inversion using a small roughening operator will not fill the null space.
- The result of a preconditioned inversion will not contain high frequencies.

Clearly, in order to obtain a high frequency result with defined solutions at every point in a small number of iterations, we need some combination of the regularized and preconditioned inversions. I chose to run a small number of preconditioned iterations then use that result as an initial model for a small number of regularized iterations. I chose to test two different combinations, one with 3 iterations of preconditioned inversion and 3 iterations of regularized inversion and one with 5 iterations of preconditioned inversion and 1 iteration of regularized inversion. These results are in Figure 3.

Both of the CIPR results contain higher frequencies than the purely preconditioned result (right panel Figure 2) and fill the areas affected by the null space better than the purely regularized result (left panel Figure 2). Determining which CIPR result is "better" is fairly subjective, but I chose to compare them by looking at their frequency spectrums. This can be seen in Figure 4. The frequencies shown in this figure are the average over all of the traces.

Figure 4 shows the frequency spectra of the results in Figure 2 and Figure 3 along with the frequency spectrum of the "ideal" model in Figure 1. As expected, the frequency content of the regularized inversion is close to that of the ideal model and the frequency content of the preconditioned inversion is much lower than the ideal model. It is more interesting to compare the frequency contents of the two different CIPR results. This shows us that the inversion using 3 iterations of preconditioning with 3 iterations of regularization has a frequency content closer to the ideal model than that of the inversion using 5 preconditioned iterations and 1 regularized iteration. This is particularly interesting because it indicates that both preconditioning and regularization are important to get the most improvement.

In this paper, I will consider the CIPR result using 3 iterations of preconditioned inversion and 3 iterations of regularized inversion to be my "best" result. Given this result, I felt it would be instructional to see how many iterations of just preconditioned inversion (fitting goals (5)) it would take to get an equivalent frequency content. It took 30 iterations of preconditioned inversion to get the same frequency content as the "best" result. The frequency content of the

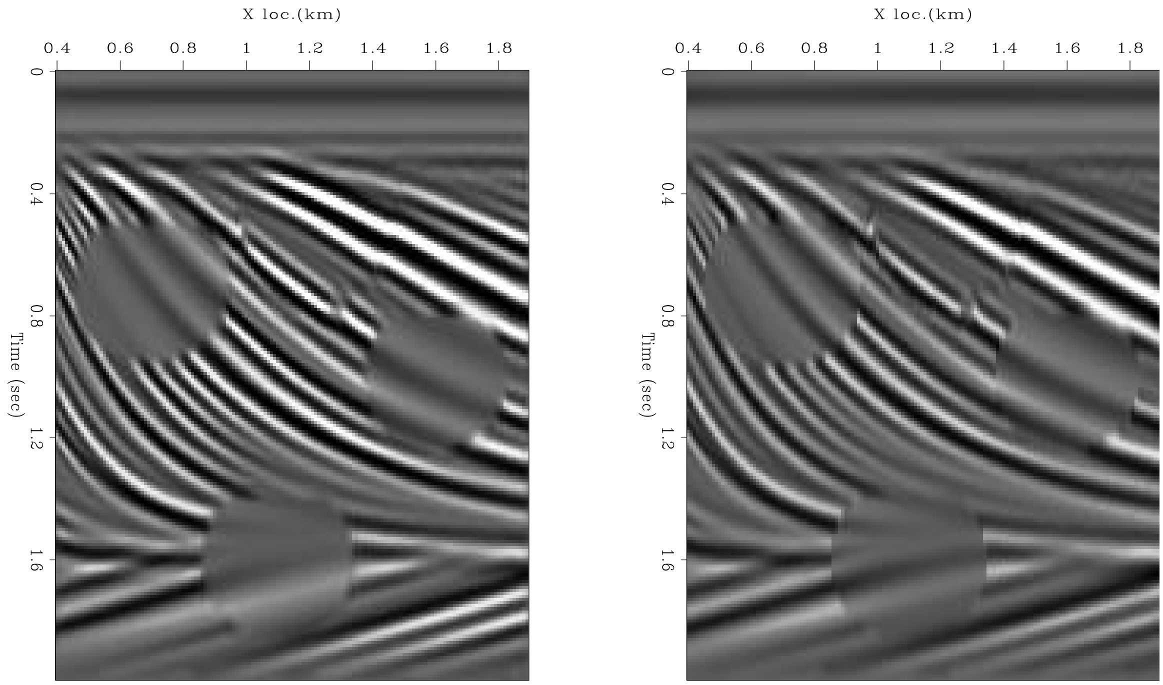

Figure 3: Left panel is the result of 3 iterations of preconditioned inversion followed by 3 iterations of regularized inversion, right panel is the result of 5 iterations of preconditioned inversion followed by only 1 iteration of regularized inversion. marie1-precreg [ER,M]

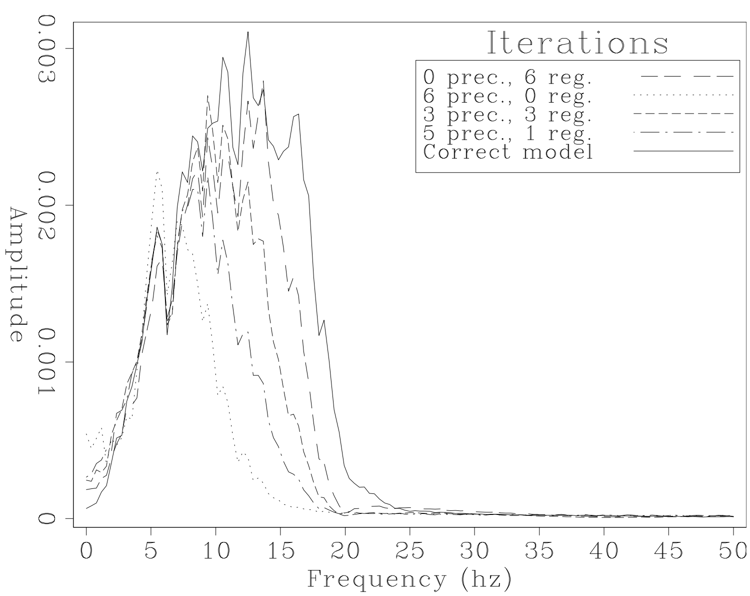

Figure 4: Comparison of the frequency content of the resulting models seen in Figures 2 and 3 along with the frequency content of the correct model (right panel of Figure 1). marie1-spectrum [ER]

result can be seen in Figure 5. One again, the frequencies shown here are the average over all of the traces.

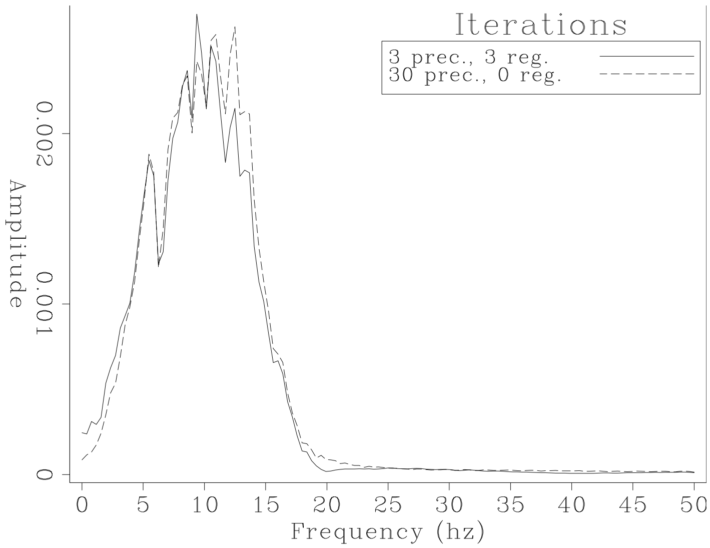

Figure 5: Comparison of the frequency content of the results of 3 iterations of preconditioned inversion with 3 iterations of regularized inversion and 30 iterations of just preconditioned inversion. marie1-speccomp [ER]

Figure 6 displays the models resulting from the "best" solution and the solution using 30 iterations of preconditioned inversion. The model resulting from 30 iterations has done a better job of filling the areas affected by the null space, as we would expect for an inversion process that used 5 times as many iterations. I have also included a model that has filled the areas affected by the null space equally well as that used only regularized inversion (fitting goals (3)). This result took 50 iterations.

CONCLUSIONS

The simple experiment conducted in this paper has compared two familiar inversion schemes, preconditioned and regularized, with a new combined inversion scheme that uses the result of a small number of preconditioned iterations as an initial model for a small number of regularized iterations (CIPR). I used a simple interpolation problem to test CIPR's ability to reduce the number of iterations needed to get an "improved" model. This "improved" model has a solution defined at every point and has a frequency content close to that of an "ideal" model. I have shown that to obtain an "improved" model, CIPR takes far fewer iterations than either of the other schemes. This makes CIPR an interesting option for many types of seismic inversion problems.

FUTURE PLANS

The simple problem presented in this paper will allow me to more thoroughly understand CIPR. One issue I plan to examine is its effect on amplitudes. An extension of this issue is the possibility of using the model normalization described by Rickett (2001a,b) to normalize

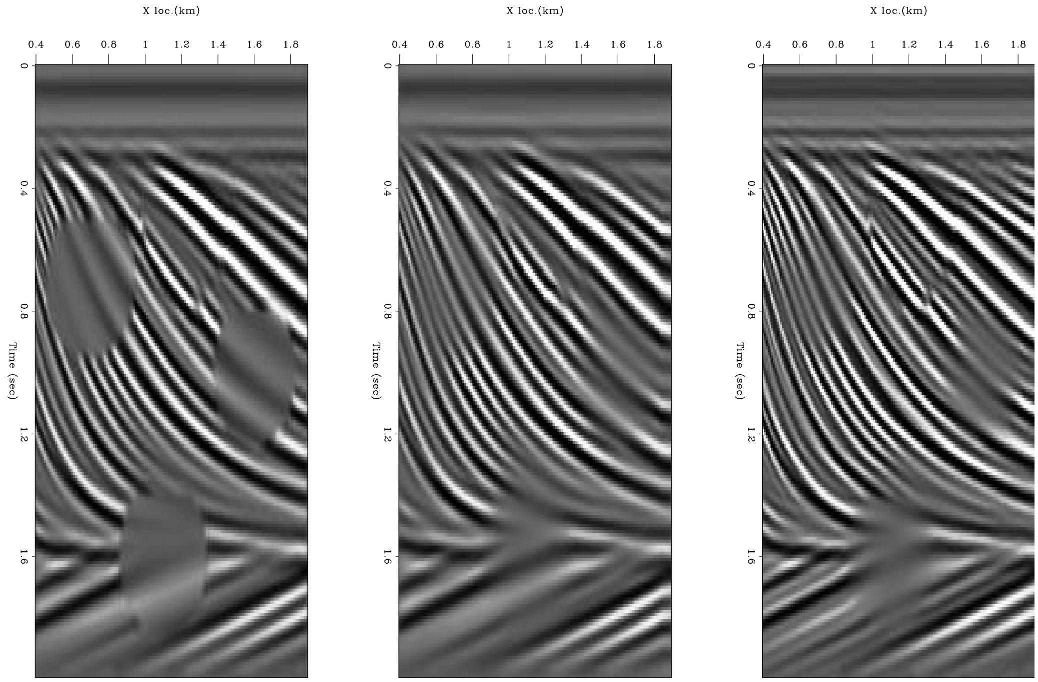

Figure 6: Comparison of the models resulting from 3 iterations of preconditioned inversion with 3 iterations of regularized inversion (left panel), 30 iterations of preconditioned inversion (center panel), and 50 iterations of regularized inversion (right panel). marie1-compits [ER,M]

the result of the preconditioned iterations before it is sent to the regularized inversion as an initial model. Another issue is the possibility of applying a mask that will only allow the preconditioning and regularization to occur within a specified area. Also, I plan to find some way to determine the optimal ratio of preconditioned iterations to regularized iterations. This may be related to another concern I intend to scrutinize, which is the effect of CIPR on the final residuals.

REFERENCES

Biondi, B., and Vlad, I., 2001, Amplitude preserving prestack imaging of irregularly sampled 3-D data: SEP–**110**, 1–18.

Claerbout, J. F., 1991, Design of inverse Kirchhoff-style filters by LSCG: SEP–**72**, 33–38.

Claerbout, J. F., 1995, Basic Earth Imaging: Stanford Exploration Project, `http://sepwww.stanford.edu/sep/prof/`.

Claerbout, J., 1998, Multidimensional recursive filters via a helix: Geophysics, **63**, no. 5, 1532–1541.

Claerbout, J., 1999, Geophysical estimation by example: Environmental soundings image enhancement: Stanford Exploration Project, `http://sepwww.stanford.edu/sep/prof/`.

Clapp, R. G., Fomel, S., and Claerbout, J., 1997, Solution steering with space-variant filters: SEP–**95**, 27–42.

Clapp, R. G., 2001, Geologically constrained migration velocity analysis: Ph.D. thesis, Stanford University.

Duquet, B., and Marfurt, K. J., 1999, Filtering coherent noise during prestack depth migration: Geophysics, **64**, no. 4, 1054–1066.

Fomel, S., Clapp, R., and Claerbout, J., 1997, Missing data interpolation by recursive filter preconditioning: SEP–**95**, 15–25.

Fomel, S., 1997, On model-space and data-space regularization: A tutorial: SEP–**94**, 141–164.

Harlan, W. S., 1986, Signal-noise separation and seismic inversion: SEP–**47**.

Nemeth, T., Wu, C., and Schuster, G. T., 1999, Least-squares migration of incomplete reflection data: Geophysics, **64**, no. 1, 208–221.

Prucha, M., and Biondi, B., 2000, Amplitudes and inversion in the reflection angle domain: SEP–**105**, 203–208.

Prucha, M. L., and Biondi, B. L., 2002, Subsalt event regularization with steering filters: SEP–**111**, 1–17.

Prucha, M. L., Clapp, R. G., and Biondi, B., 2000, Seismic image regularization in the reflection angle domain: SEP–**103**, 109–119.

Prucha, M. L., Clapp, R. G., and Biondi, B. L., 2001, Imaging under salt edges: A regularized least-squares inversion scheme: SEP–**108**, 91–104.

Rickett, J., 2001a, Model-space vs data-space normalization for finite-frequency depth migration: SEP–**108**, 81–90.

Rickett, J., 2001b, Spectral factorization of wavefields and wave operators: Ph.D. thesis, Stanford University.

Ronen, S., and Liner, C. L., 2000, Least-squares DMO and migration: Geophysics, **65**, no. 5, 1364–1371.

Tikhonov, A. N., and Arsenin, V. Y., 1977, Solution of ill-posed problems: John Wiley and Sons.