Multilevel acceleration strategy for the robust estimation of primaries by sparse inversion

Tim T.Y. Lin and Felix J. Herrmann
Department of Earth and Ocean sciences, University of British Columbia, Vancouver, BC, Canada

January 17, 2014

Abstract

We propose a method to substantially reduce the computational costs of the Robust Estimation of Primaries by Sparse Inversion algorithm, based on a multilevel inversion strategy that shifts early iterations of the method to successively coarser spatial sampling grids. This method requires no change in the core implementation of the original algorithm, and additionally only relies on trace decimation, low-pass filtering, and rudimentary interpolation techniques. We furthermore demonstrate with a synthetic seismic line significant computational speedups using this approach.

Motivation

Multiple removal is a crucial aspect of seismic signal processing that constantly face a difficult trilemma between accuracy, robustness, and low computational complexity. Current prediction-subtraction methods such as Surface-Related Multiple Removal (Verschuur, 1992) face limits in accuracy and robustness when confronted with undersampled data of limited quality, prompting recent developments in whole-wavefield inversion/deconvolution techniques to decrease dependence on practitioner guesswork and QC. However, these methods usually require many iterations of multiple prediction to determine an entire unknown discretized wavefield, and thus are far more computationally intensive, posing significant hinderances to practical adaption.

This work proposes a strategy for significantly reducing the costs of one such method, the Robust Estimation of Primaries by Sparse Inversion (REPSI, Lin and Herrmann, 2013) based on a method proposed by van Groenestijn and Verschuur (2009), without fundamentally altering the core implementation of the algorithm. This is achieved by shifting a large part of the computational burden to coarser spatial sampling grids while exploiting the unique way in which REPSI mitigates spatial aliasing, which is the major contributor of multiple prediction errors due to coarsening (Dragoset et al., 2006). The REPSI algorithm looks for a discretized full-band multiple-free seismic impulse response; this full-band structure is enforced by imposing sparsity constraints in the physical time-space samples of the solution wavefield, similar in concept to spiking wavelet deconvolution. Therefore, we can use time-domain low-pass filters to remove spatially-aliased frequencies while expecting the coarse-grid REPSI solution to retain some degree of temporal resolution and accuracy (assuming that time-axis sampling remains unchanged). Figures 1 and 2 demonstrates this effect in practice for two levels of grid decimation.

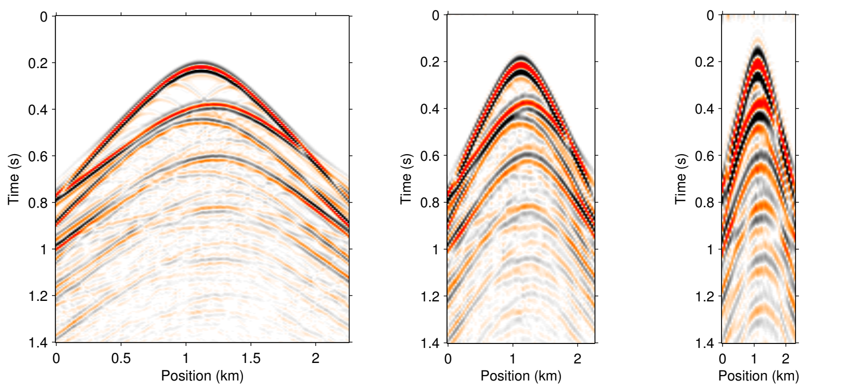

(a) Original sampling (b) 2:1 decimation (c) 4:1 decimation

Figure 1 *Input synthetic data shot gathers, at (a) the original spatial sampling and at (b,c) two levels of spatial decimation, with appropriate time-domain low-pass filters to mitigate artifacts spatial aliasing when computing the surface-related multiple wavefield. The sampling of the time axis remains untouched.*

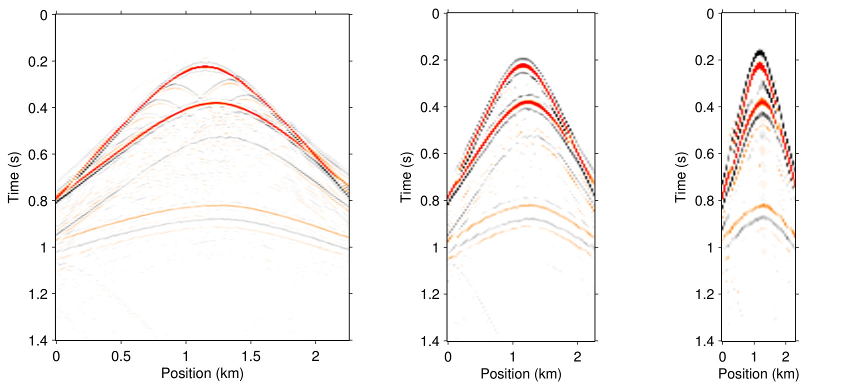

(a) Original sampling (b) 2:1 decimation (c) 4:1 decimation

Figure 2 *Computed multiple-free seismic impulse response shot gathers using a straightforward application of the REPSI algorithm, from the original data and from its decimated, low-pass filtered versions shown in Figure 1. Even though its input data was low-pass filtered at 15Hz, the solution shown in (c) retains a wide-band, "deconvolved" appearance with good resolution of the two separate events at t=0.9s.*

Multilevel inversion strategy for REPSI

The REPSI algorithm solves, for a given regularized observed wavefield $\mathbf{p}$ containing multiples (hereby understood to be surface-related multiples), the following optimization problem in the space-time domain:

$$\min_{\mathbf{g},\,\mathbf{q}} \|\mathbf{g}\|_1 \quad \text{subject to} \quad \|\mathbf{p} - M(\mathbf{g},\mathbf{q})\|_2 \leq \sigma, \tag{1}$$

with unknown variable $\mathbf{g}$ representing the discretized (identically with $\mathbf{p}$) multiple-free impulse response of the subsurface, and $\mathbf{q}$ the source-side signature function. The modeling function $M(\mathbf{g},\mathbf{q})$ attempts to reconstruct the observed data by separately modeling the multiple wavefield (by wavefield convolution between $\mathbf{p}$ and $\mathbf{g}$ at the multiple-generating surface, similar to the SRME prediction step) and the primary wavefield (convolution between $\mathbf{g}$ and $\mathbf{q}$). The problem essentially asks for the sparsest (via minimizing the ℓ_1-norm) multiple-free impulse response that explains $\mathbf{p}$, along with a source signature $\mathbf{q}$ that matches $\mathbf{g}$ to the wavefronts of the recorded data, while ignoring some amount of noise as determined by σ. The details of how the REPSI algorithm accomplishes this can be found in Lin and Herrmann (2013).

Roughly speaking, REPSI converges through carefully-controlled projected gradient iterations, and by far the most time-consuming part of each iteration is the computation of $M(\mathbf{g},\mathbf{q})$ and its adjoint operation to form the gradient updates on $\mathbf{g}$. Both of these involve a complete wavefield convolution, which can be computed in the time-frequency domain by matrix multiplication between mono-frequency slices of the two wavefields. Excluding the Fourier transform cost, each iteration requires two $\mathscr{O}(n_f n_x^3)$ operations, where n_x denotes the number of grid points per spatial axis, and n_f the number of discrete frequencies, to compute the multiple wavefield for a single seismic line. For a well-sampled 3D dataset, this complexity approaches $\mathscr{O}(n_f n_x^6)$. We can see that by halving both the spatial sampling and upper frequency limit, we can disproportionally reduce the computational operation by an order of magnitude.

Consequently, solving the complete REPSI problem on this coarser grid (typically requiring less than 80 iterations) is comparable in computation to just a few much more computationally expensive gradient iterations on the finer-scale problem. Recalling our observation in the previous section that these coarse-scale solutions closely resemble decimated versions of the original solution, it is thus reasonable to postulate that the computation budget would be more efficiently spent on updates at the coarser level first until all coarse-scale information is exhausted, assuming that we also have an effective mapping (i.e., interpolation) from the 2:1 decimated grid to the original grid. This degree of interpolation is often performed in multiple removal and can be achieved reasonably well by, e.g., NMO-corrected trace averaging. Recursion of this approach leads to a multilevel strategy for efficiently solving the REPSI problem, which we outline in Algorithm 1. This method should outperform unmodified REPSI provided that the interpolated coarse-scale solution brings $\mathbf{g}$ closer to the true finer-grid solution than the computationally equivalent number of iterations on the finer-grid problem. Note that, because the REPSI problem involves inherently solving an integral equation of the second kind (Frijlink et al., 2011), our proposed method bears some theoretical connection to the general class of multigrid algorithms.

Algorithm 1: Multilevel continuation strategy for accelerated REPSI

Data: wavefield data $\mathbf{p}$ containing surface multiples regularized to a fully-sampled grid
Result: approximate solution of the (surface) multiple-free subsurface impulse response $\widetilde{\mathbf{g}}$ to the REPSI problem (Equation 1)
Choose maximum data decimation factor 2^S, set scale variable $s = S$
Set initial solution $\mathbf{g}_{s=S}$ to zero vector
repeat
- $\mathbf{p}_s \leftarrow$ decimate or re-regularize original data in all spatial coordinates by factor of 2^S
- Low-pass filter in the time domain on $\mathbf{p}_s$ as needed, to remove spatial aliasing
- Solve the REPSI problem on decimated data $\mathbf{p}_s$ starting from initial solution $\mathbf{g}_s$
- $\mathbf{g}_{s-1} \leftarrow$ interpolate solution $\mathbf{g}_s$ in all spatial coordinates by factor of 2
- $s \leftarrow s - 1$ (go to finer scale)

until *scale variable* $s = 0$ *(reached original spatial sampling)*
Solve the REPSI problem on original data $\mathbf{p}$, starting from initial solution $\mathbf{g}_0$

Numerical example

Using the fully-sampled, fixed-spread 2D synthetic seismic line data shown in Figure 1(a), we illustrate the performance uplift possible with our proposed multilevel strategy. This dataset served extensively as examples in early works on REPSI (van Groenestijn and Verschuur, 2009; Lin and Herrmann, 2013) and was shown to be amenable to the REPSI process. A 15m spatial sampling grid was used in both the source and receiver coordinates. The data was modeled using a source function of zero-phase Ricker wavelet with peak frequency at 30Hz, and then subsequently low-pass filtered at 60Hz to eliminate spatial aliasing as a factor in multiple prediction.

As shown in Figure 2, directly applying the unmodified REPSI algorithm to decimated, smoothed versions of the original data yields a decimated, *non-smooth* version of the original solution due to the sparse inversion. The impulse response shown in Figure 2(b) is computed from the data shown in 1(b), which is a 2:1 decimated version of the original data (spatial sampling at 30m, dropping every other trace and every other shot record from the original data) followed by low-pass filter with one-sided Hanning window at 30Hz. Similarly, 2(c) follows from 1(c), with 60m spatial sampling and low-pass at 15Hz. The three different solutions in Figure 2 all took roughly 75 iterations of the REPSI algorithm. Due to the cubic scaling of compute operations with the number of grid points, obtaining the coarsest-scale solution 2(c) is equal in runtime to just 2 iterations of the original undecimated problem.

We next further iterate on our coarse-scale solutions (Figure 2b,c) at the finer sampling grids, which allowed us to obtain the finer-scale solution much faster. The coarse-scale solutions are first mapped to a finer grid via linear interpolation as moveout-corrected common-midpoint gathers (using constant velocity 1600m/s), then subsequently used as initial solutions for the REPSI algorithm applied to finer-sampled data shown in Figures 1(a,b), respectively. Compare the results in Figure 3, which only required roughly 30 iterations, with their counterpart in Figure 2 (roughly 75 iterations).

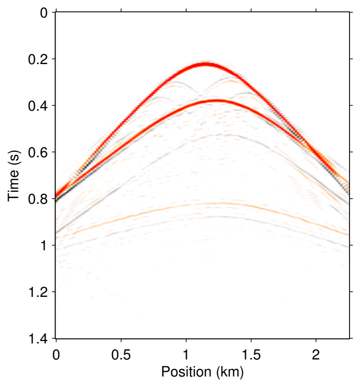

(a) Original sampling

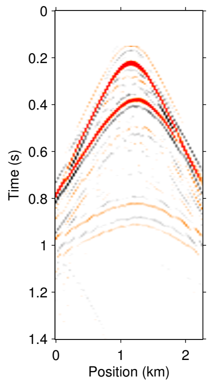

(b) 2:1 decimation

Figure 3 *Computed impulse response shot gathers, using as initial solutions (after NMO-based interpolation) the coarser-scale solutions shown in Figures 2(b,c) respectively. The results are comparable to 2(a,b) while requiring much less iterations compared to starting from the empty solution.*

Following the multilevel continuation strategy completely, we take the 2:1 scale solution (Figure 3b) initialized from the 4:1 scale solution, and use it after interpolation as the initial solution in a REPSI problem at the original scale. We obtained a reasonably accurate solution after just 15 iterations. Figure 4 compares the final demultipled wavefields as computed through this method with a reference computed completely at the original scale. The results are comparable despite the multilevel approach requiring only a quarter of the runtime of the original method, as detailed in Figure 5.

Application to physically undersampled data

Although we utilized a *deliberate* decimation of data to improve computation complexity, it is possible to modify the coarsest grids to coincide with the poorest sampled directions in field acquisition. This would potentially provide more control in mitigating regularization errors due to the typical asymmetry of sampling density in seismic data and its influence on the multiple removal process.

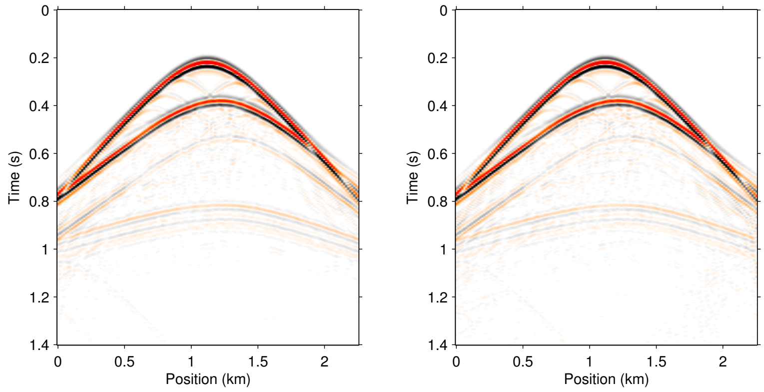

(a) Reference primary (b) Multilevel REPSI primary

Figure 4 *Comparison shot-gathers of the final demultipled primary wavefield between the REPSI algorithm run at the original spatial sampling, and the proposed multilevel strategy beginning at a 4:1 decimation of the data.*

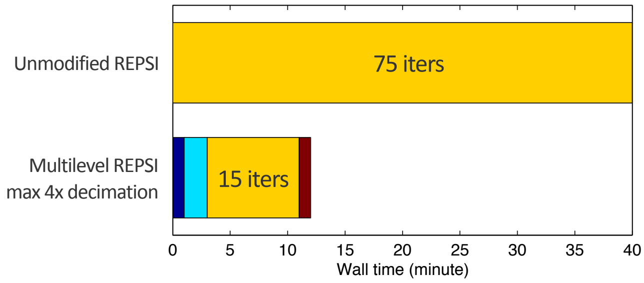

Figure 5 *Runtime comparison between the REPSI algorithm run at the original spatial sampling **(40 minutes)**, and the proposed multilevel strategy beginning at a 4:1 decimation of the data **(12 minutes)**. Different colors denote REPSI computation at different decimation levels of the data, with yellow the original undecimated sampling, light blue the 2:1 decimation, and dark blue the 4:1 decimation. Brown color denotes incidental computations such as interpolation and moveout correction.*

Summary

The main motivation behind this work is the observation that intermediate solutions to the REPSI problem can be computed on a coarser gird and interpolated to a finer grid orders-of-magnitude faster than the time required to reach comparable accuracy working directly on the finer grid. To exploit this discrepancy, we proposed a method based on a multilevel inversion strategy to substantially reduce the computational requirements of REPSI. Demonstrations on a synthetic seismic line show that, without appreciable loss in solution quality, significant speedups up to four times are possible using this approach. Further speedups can be expected for 3D geometry due to the higher computation complexity order as a function of grid point density.

Acknowledgements

This work was in part financially supported by the NSERC Discovery Grant (22R81254) and the CRD Grant DNOISE II (375142-08). This research was carried out as part of the SINBAD II project with support from the following organizations: BG Group, BGP, BP, CGG, Chevron, ConocoPhillips, ION, Petrobras, PGS, Statoil, Total SA, WesternGeco, and Woodside.

References

Dragoset, B., Moore, I. and Kostov, C. [2006] The impact of field-survey characteristics on surface-relatedmultiple attenuation. *Geophysical Prospecting*, **54**(6), 781–791, doi:10.1111/j.1365-2478.2006.00581.x.

Frijlink, M.O., van Borselen, R.G. and Söllner, W. [2011] The free surface assumption for marine data-driven demultiple methods. *Geophysical Prospecting*, **59**(2), 269–278, doi:10.1111/j.1365-2478.2010.00914.x.

Lin, T.T.Y. and Herrmann, F.J. [2013] Robust estimation of primaries by sparse inversion via one-norm minimization. *Geophysics*, **78**(3), R133–R150, doi:10.1190/geo2012-0097.1.

van Groenestijn, G.J.A. and Verschuur, D.J. [2009] Estimating primaries by sparse inversion and application to near-offset data reconstruction. *Geophysics*, **74**(3), A23–A28, doi:10.1190/1.3111115.

Verschuur, D.J. [1992] Adaptive surface-related multiple elimination. *Geophysics*, **57**(9), 1166, doi:10.1190/1.1443330.