Getting something for nothing—or not: interpolating coherent noise

David C. Henley and Joe Wong

ABSTRACT

Seismic data acquisition techniques continue to evolve towards the increasing use of single sensors and finer receiver spacing, in order to improve both the spatial and vertical resolution of seismic reflection images. Part of the increased resolution comes from the improved ability to estimate and remove coherent source-generated noise from source gathers with reduced receiver spacing. For most practical acquisition situations, however, the receiver spacing cannot be made small enough to completely avoid some aliasing of the slower components of the source-generated noise. Some simple processing tactics which may improve this situation are demonstrated here, using data provided by the physical modelling system at CREWES. These methods involve trace-shift de-aliasing, simple radial trace domain interpolation, and a non-linear radial trace noise attenuation trick. These methods can improve coherent noise removal for a source gather with an existing receiver spacing, but are generally not as effective as halving the receiver spacing during acquisition.

INTRODUCTION

Ongoing development of seismic data acquisition technology has led to a dramatic increase in the number of live channels that can be recorded simultaneously. This, in turn, has encouraged the deployment of more sensors on the ground, as well as the increased use of single sensors, closely spaced (Henley, et al, 2006, 2008). While increased spatial resolution of the resulting reflection image is one of the motivations for these developments, a closely related one is accurate estimation of coherent source-generated noise (Henley, et al, 2006, 2008). The better the noise estimate, the more effectively the noise can be removed from the desired reflections. Effective noise removal then leads to better reflection imaging in terms of both increased signal bandwidth and lateral spatial resolution.

As we showed in 2006 (Henley, et al), the best way to accurately estimate noise is to record it properly, without spatial aliasing. The most straightforward way to accomplish this is to reduce the receiver spacing enough that no significant frequency component of any of the source-generated noise is aliased. We also showed that for some frequency components of ground roll and air blast, the required receiver spacing might be as small as 1m, an unacceptable parameter for most practical field operations (Henley, et al, 2008). We also showed, however, that coarser receiver spacing could be compensated by reduced source interval (Henley, et al, 2009) under some circumstances. When neither of these options can be fully utilized because of time or equipment constraints, is there any way to improve noise estimates without increased field effort?

Simple de-aliasing

Figure 1 illustrates what we mean by aliasing in the context of accurately representing dipping seismic events in a discretely sampled wavefield. As shown in the Figure, dipping events whose analog waveforms are misaligned by as much as a quarter of a wavelength of the predominant frequency no longer represent unique dips. Attempts to reconstruct the analog events from their samples result in ambiguity in event slope and frequency content. As well as illustrating the aliasing problem, Figure 1 also suggests a simple solution to remove aliasing for a single event. If the traces which sample a dipping linear event are incrementally time-shifted with respect to each other (the equivalent of applying linear moveout), the waveforms of the event can be perfectly aligned, removing all aliasing. Unfortunately, if a group of seismic traces contains several events with different dips, only one of the events can be perfectly de-aliased by trace shifting. Also, the act of de-aliasing one event usually increases the aliasing of other events. For example, most seismic surveys are designed to minimize aliasing of the reflected events and any dip they represent; and any incipient reflection aliasing on raw trace gathers is removed by the event alignment provided by NMO correction. Spatial sampling which is adequate to prevent aliasing of reflection events, however, is generally insufficient to prevent aliasing of source-generated coherent noise, except for those noise trains with the highest apparent velocity. Applying trace shifts which align linear noise waveforms, however, badly misaligns the reflection events, hence inducing aliasing. This renders the reflections on such gathers unsuitable for horizontal smoothing or imaging operations. Also, any 2D filtering operation applied to these gathers may be effective for removing noise, but they will simultaneously destroy the reflections.

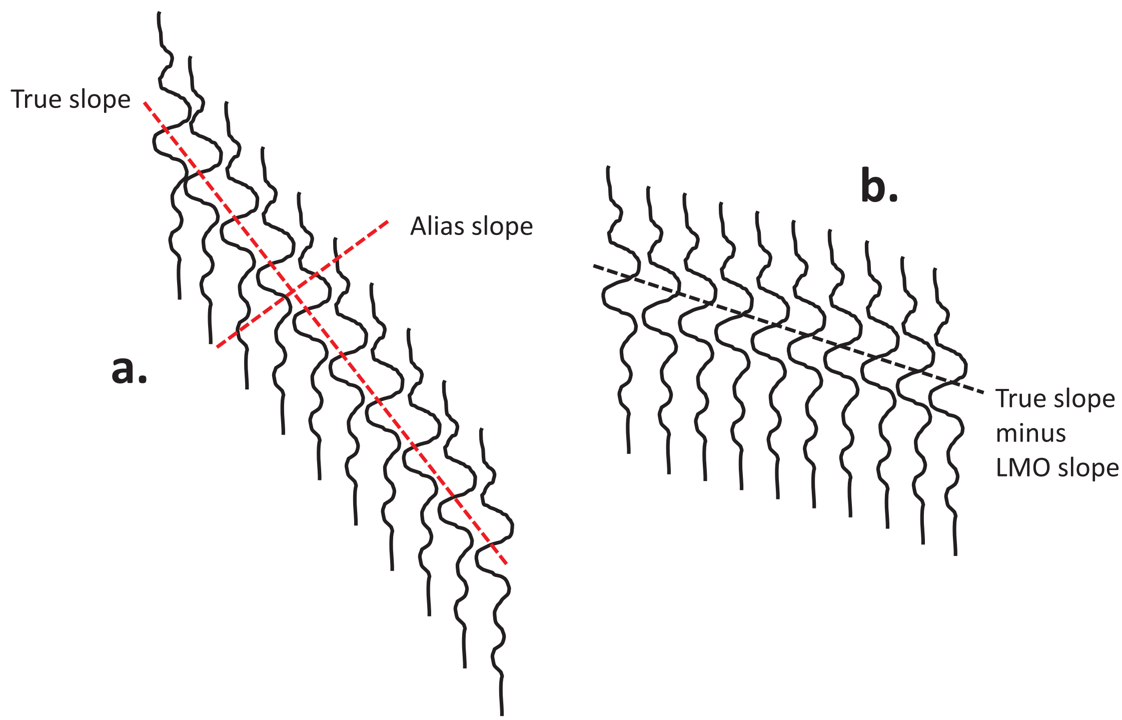

Dipping event which is aliased by spatial undersampling (a.), can be de-aliased by applying a constant velocity LMO, which temporarily reduces the dip, and hence the aliasing (b.)

FIG. 1. This schematic illustrates the concept of spatial aliasing of a dipping event. When an event is too sparsely spatially sampled, event waveforms align along more than one slope (a.). If incremental trace shifts are applied (equivalent to linear moveout,or LMO), the aliasing can be removed.

If, however, we use trace-shifted gathers only to estimate coherent noise, and make no attempt to use the reflection information, we may have a viable approach for modeling noise beyond the physical alias limits. The idea is to apply linear moveout to raw trace gathers to reduce the aliasing of linear events on those gathers, suppress the reflections by mode separation and filtering, remove the linear moveout, and subtract the estimated noise from the original gathers. We normally estimate coherent noise by applying a low-pass filter to the radial trace transform of the original shot gather (Henley, 2003), which preferentially attenuates reflections, leaving only noise in the inverse radial trace transform, to be subtracted from the original source gather.

Source-generated noise can exhibit a wide range of apparent velocities, so a velocity which aligns one noise, removing its aliasing will not properly align all noises. To demonstrate this limitation, Figure 2 shows an example of the notorious ice flexural wave often recorded in the arctic. Because of the relatively sparse spatial sampling, all frequency components of this noise are aliased. If we apply LMO of 1000m/s, which aligns the fastest noise component, we can see in Figure 3 that the slowest noise is still aliased, even though its frequency content is lower. If, on the other hand, we apply LMO of 300m/s to align the slow noise, we have misaligned the higher velocity noise, causing it to be hopelessly aliased (Figure 4).

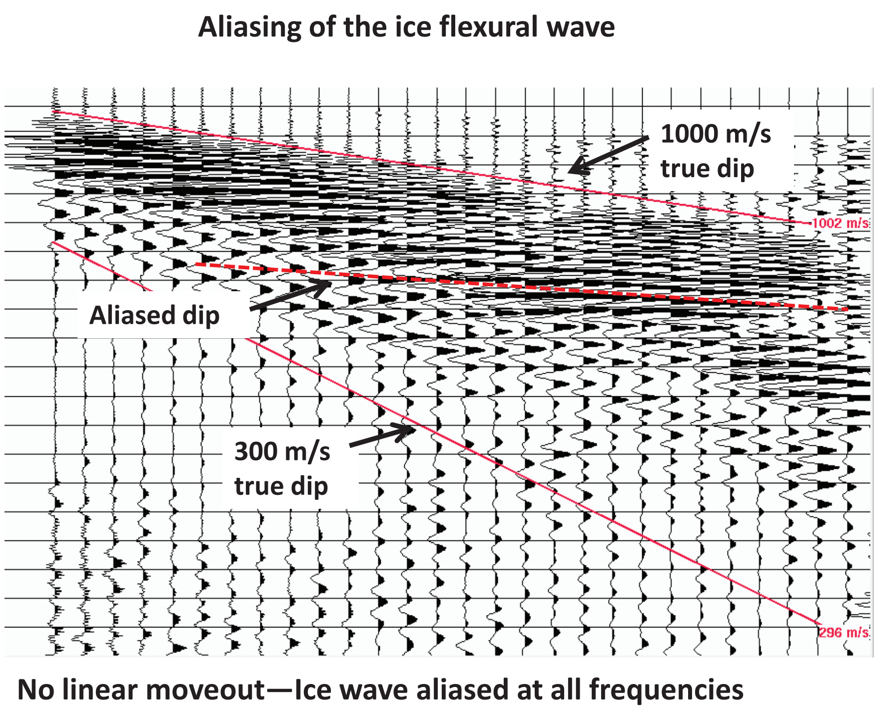

FIG. 2. Example of the ice flexural wave—a coherent noise that is aliased at all frequencies in this example, due to inadequate spatial sampling.

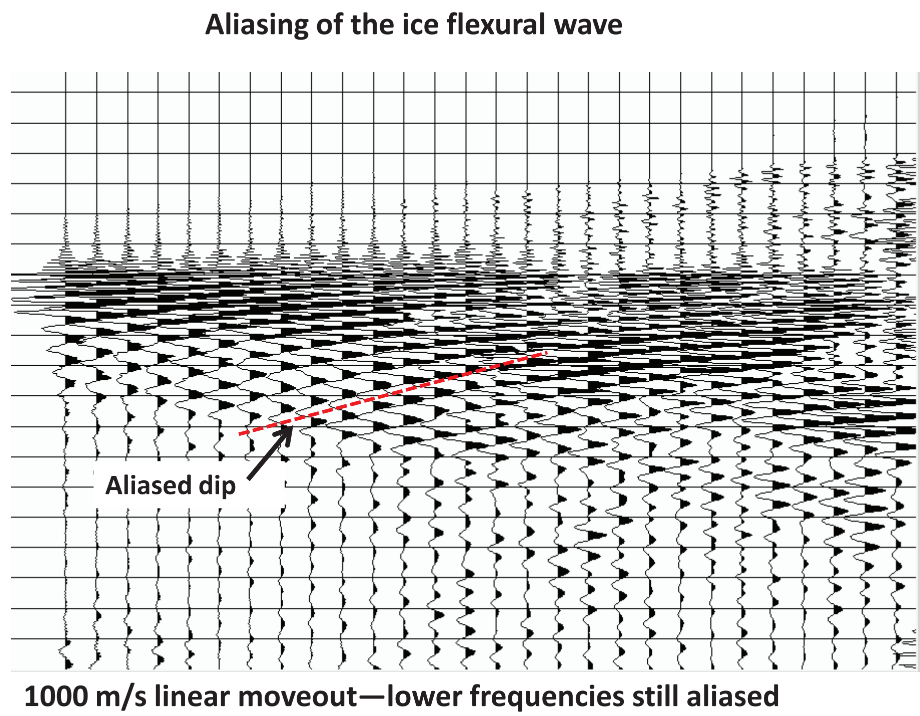

FIG. 3. De-aliasing the high-velocity arrivals is insufficient to de-alias the slower arrivals, even though they are lower in frequency, as well.

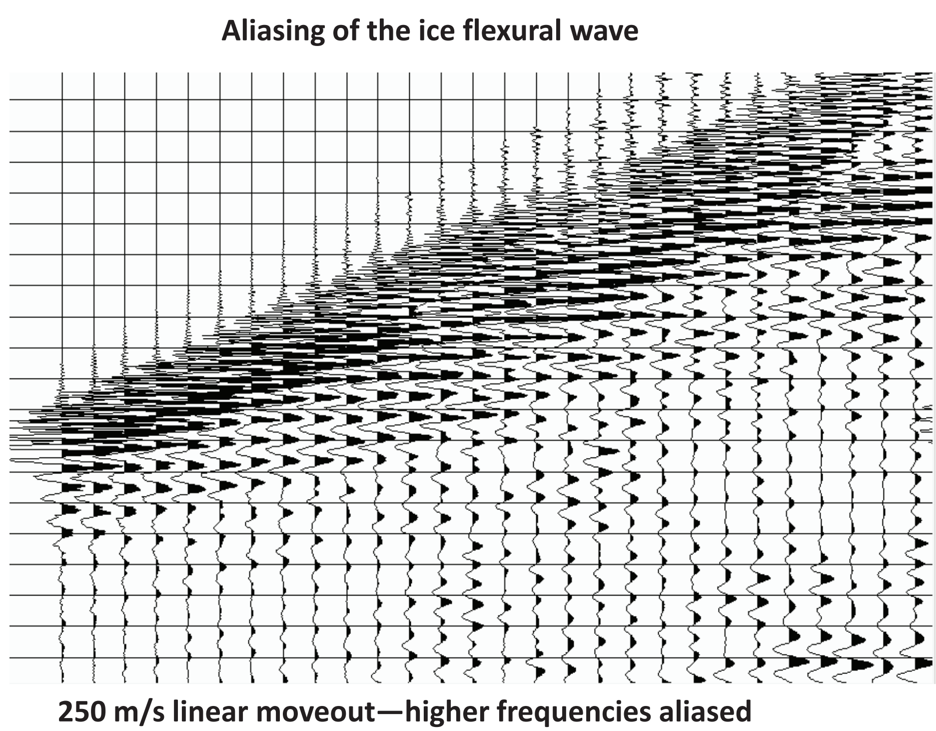

FIG. 4. De-aliasing the low-velocity arrivals leads to even more severe aliasing of the high-velocity arrivals.

For a highly dispersed noise like the ice flexural wave, LMO de-aliasing is not practical. However, if we have a more typical trace gather with a relatively small number of distinct linear noises, each confined to a relatively narrow range of velocities, we can often find a compromise velocity which will partially align the noise, removing the aliasing for at least some of the noise frequencies. In the demonstration that follows, we apply LMO which partially de-aliases a single strong noise component without increasing the aliasing of other noises unacceptably. Since our objective is to model the noise, it is of no concern that we alias reflections, since we filter them out of the noise estimate in any case.

Simple interpolation

Another tool which has proved useful in estimating and removing coherent noise is a simple interpolation scheme whereby we can either interpolate noise estimates to finer spatial sampling, or interpolate entire gathers before making the noise estimate. Henley (2000) showed how to use the radial trace transform to interpolate trace gathers, utilizing a simple two-point l^p-norm interpolation method. It is this simple interpolation that we utilize here.

A nonlinear R-T trick

One further simple processing trick which can be applied in the R-T domain to increase the effectiveness of coherent noise attenuation is to apply AGC in the R-T domain. Ordinarily, we transform an input trace gather from the X-T domain to the R-T domain and low-pass filter the result to obtain an estimate of the coherent noise, in the R-T domain. The inverse R-T transform of this noise estimate can be subtracted from the original trace gather in the X-T domain to remove the noise, a linear process. If we are willing to sacrifice AVO information on a trace gather for more effective noise attenuation, however, we find that AGC, a non-linear process, applied in the R-T domain, can be even more effective than an ordinary R-T domain filter. When we take pains to properly fit the radial trace trajectories of the R-T transform to the noise wavefronts, we find that radial traces just adjacent to a coherent noise wavefront and parallel to it have very low amplitude compared to radial traces which lie on the coherent noise event itself. Simply equalizing the amplitudes of non-noise radial traces with those sampling the noise itself greatly reduces the amplitude of the noise with respect to the background. Although ordinary trace normalization can be used, AGC is even more effective, since it adjusts not only the DC level of an R-T trace but the lower frequency components as well.

PHYSICAL MODEL

The CREWES physical modeling facility has been described by Wong et al (2007, 2008, 2009). It is a convenient tool for generating ‘real’ seismic data which are subject to many of the same limitations as full scale ‘field’ data. Since we can control the physical structure of the model, we can generate data sets that have no statics or elevation differences, and which have real reflection/conversion events from a limited number of interfaces. We also control the acquisition geometry, and can choose to sample almost any event with a pre-determined degree of aliasing. The trace gather that we study here was created with that system; and the spatial sampling was chosen to minimize the aliasing on all visible events, even the slowest. That particular gather and its processed version will be our comparison standard for all other results.

A simple way to test de-aliasing and interpolation schemes for data with aliased events is to start with a perfectly sampled gather, decimate it by discarding every other trace, then test our processing schemes on the decimated gather to see how close we can come to the fully sampled original gather. This simulates a hypothetical field situation in which our receiver spacing was too large to properly sample one or more events, specifically coherent noise, on the source gathers. We produce first the result of our usual coherent noise attenuation on the decimated record. Then we apply one of our proposed de-alias or interpolation methods to the same decimated record for comparison. Figure 5 shows the raw source gather we chose for our testing, and Figure 6 shows the same shot after radial trace (R-T) domain coherent noise attenuation (Henley, 2003) and deconvolution, presumably the best result we can obtain for these data.

TESTING THE METHODS

Figure 6 corresponds to the ideal field situation where we are able to reduce the trace spacing by a factor of two during the acquisition. We show, in comparison, three methods for improving the effectiveness of coherent noise attenuation when we cannot alter the acquisition parameters. While all demonstrate significant improvement, only one method gives results comparable to halving the trace spacing, and that one requires compromising relative reflection amplitudes.

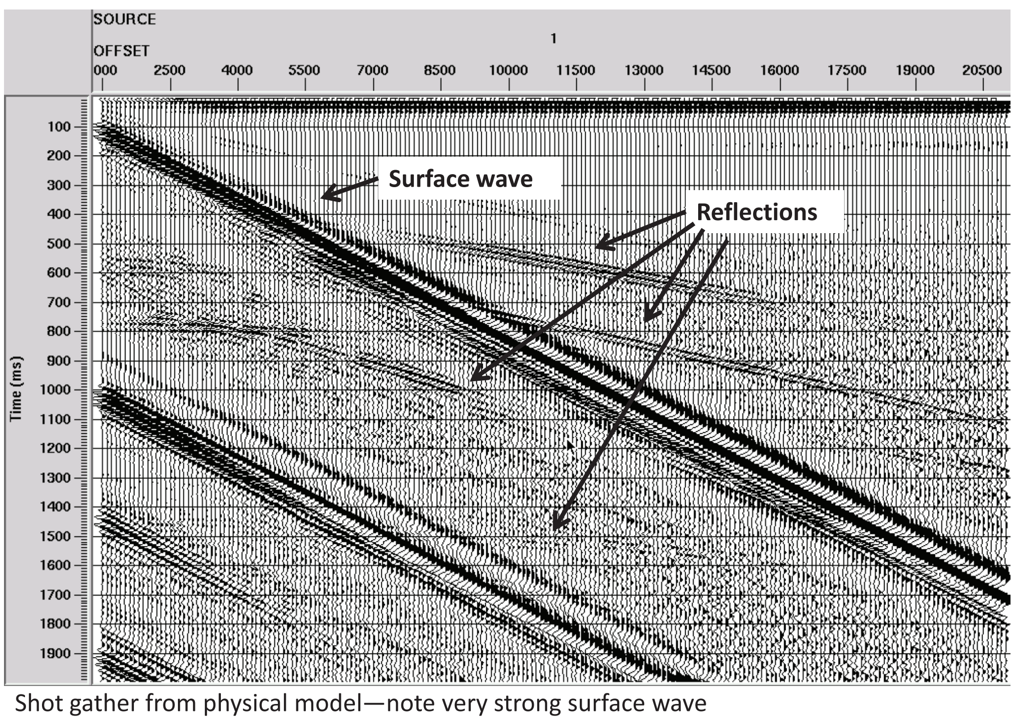

FIG. 5. A shot gather from a physical model which has been adequately spatially sampled.

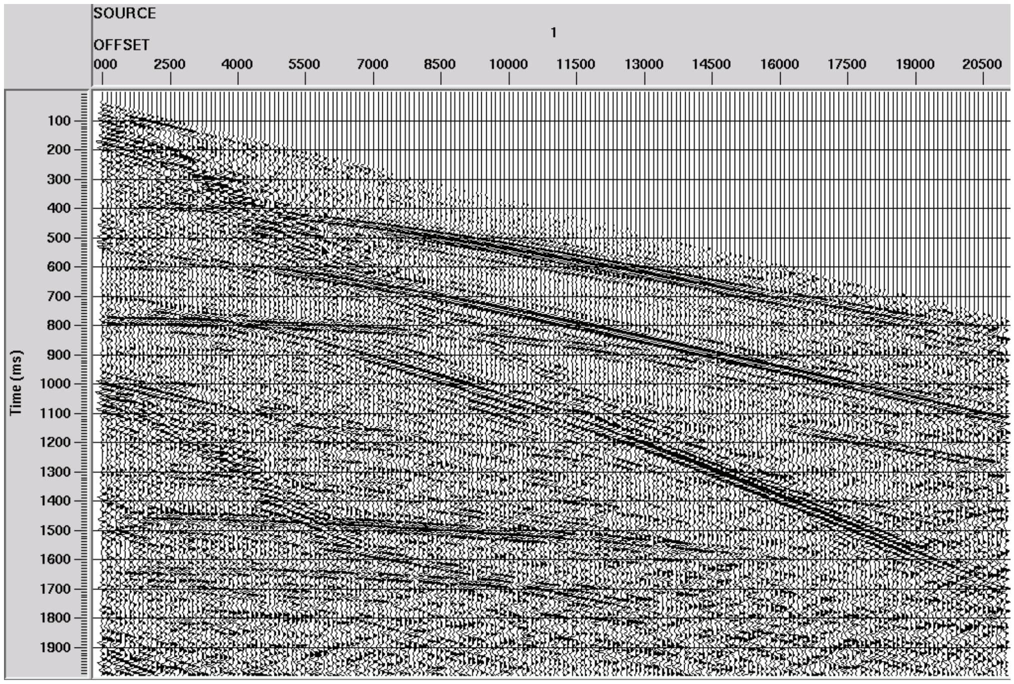

Physical model shot gather after radial trace coherent noise attenuation

FIG. 6. Shot gather from the physical model after R-T domain coherent noise attenuation.

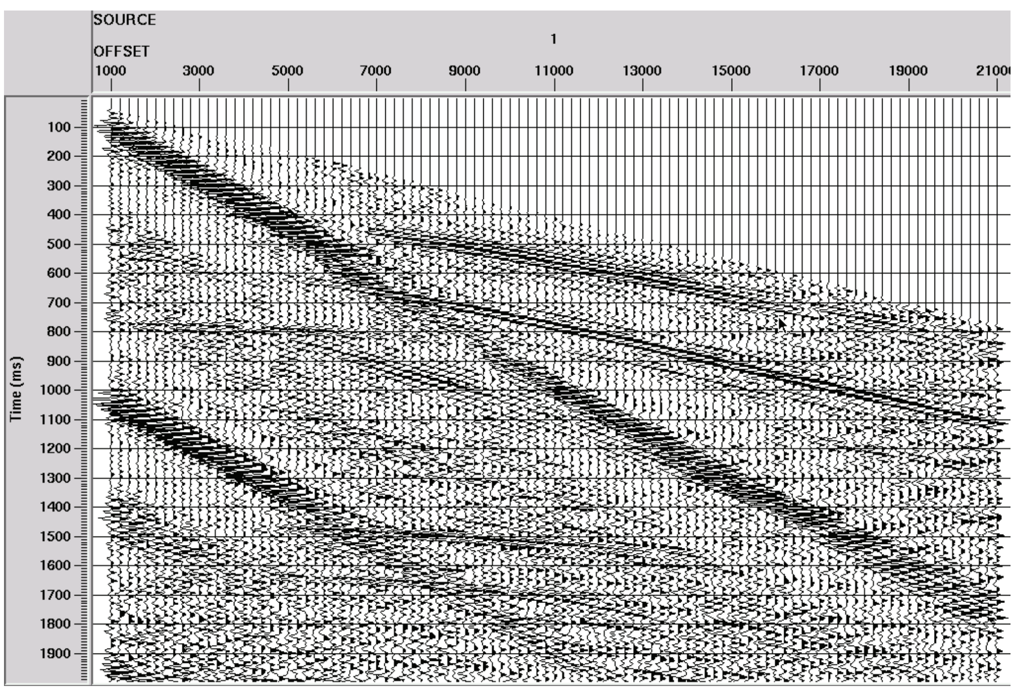

Physical model shot gather after decimation, radial trace coherent noise attenuation

FIG. 7. The physical model shot gather after discarding every other trace (decimation). This simulates a gather which has been spatially undersampled. The R-T domain coherent noise attenuation which was effective in Figure 6 is much less effective on this gather, due to aliasing.

De-aliasing the noise

We decimated the trace gather shown in Figure 5, and applied the same coherent noise attenuation and deconvolution as in Figure 6, to obtain the result in Figure 7, much inferior to the image in Figure 6. The remnant surface wave is much stronger, and the reflections in the background correspondingly weaker. It is obvious that the finer spatial sample interval of the original source gather allows much more effective attenuation of the very strong surface wave.

To test noise de-aliasing, we first applied linear moveout (LMO) to the decimated source gather, sufficient to nearly align the waveforms of the surface wave arrivals, as in Figure 8. Note that the faint reflections are now badly aliased...but that doesn't matter, since we filtered them out in the next step (a low-pass filter in the R-T domain: Figure 9). Removing the LMO yields the noise estimate in Figure 10. The noise estimate produced by the same procedure, but *without* the LMO de-aliasing step, is shown in Figure 11. Note that the waveform of surface wave arrivals here is much lower in frequency than in Figure 10. This is due to the aliasing of the higher frequencies in the decimated gather. The residual noise in Figure 7 is also higher frequency, indicating that the R-T domain noise estimate was deficient in these frequencies. Figure 12 shows the result of subtracting the de-aliased estimate from the decimated gather. While still not as good as the original properly sampled result, there is a significant improvement over the result with no de-aliasing applied. The overall strength of the remnant noise relative to the reflections is significantly less, particularly in the lower right of the trace gather where the surface wave interferes with an underlying reflection, completely masking it in Figure 7.

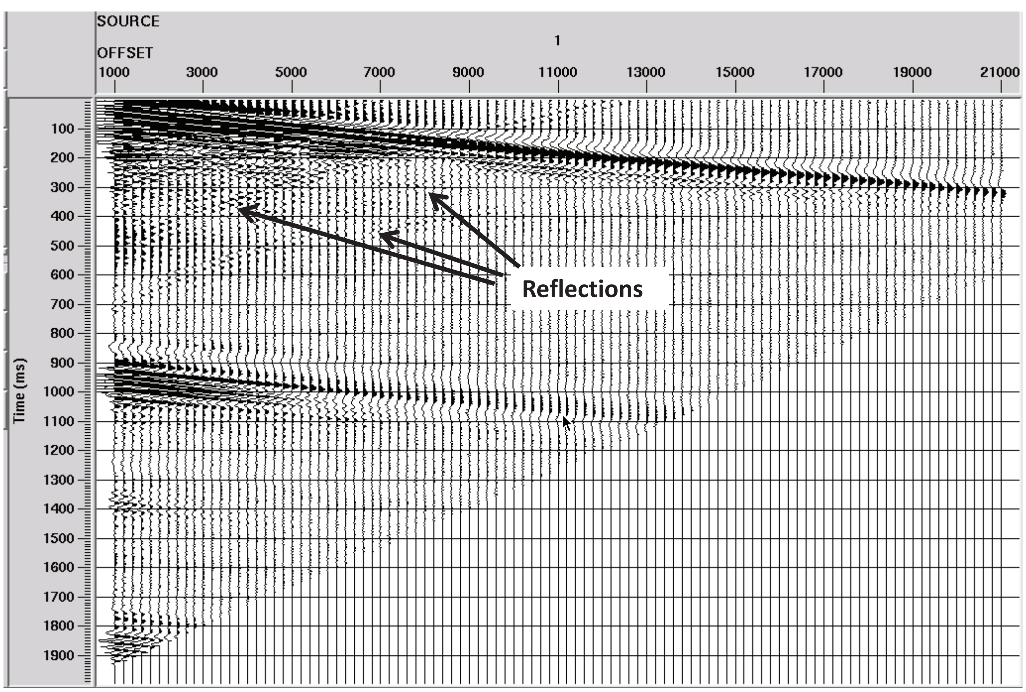

Decimated physical model shot gather after LMO correction to de-alias surface wave noise

FIG. 8. Linear moveout (LMO) removed from decimated physical model shot gather to de-alias the surface wave. Reflections are now visibly aliased.

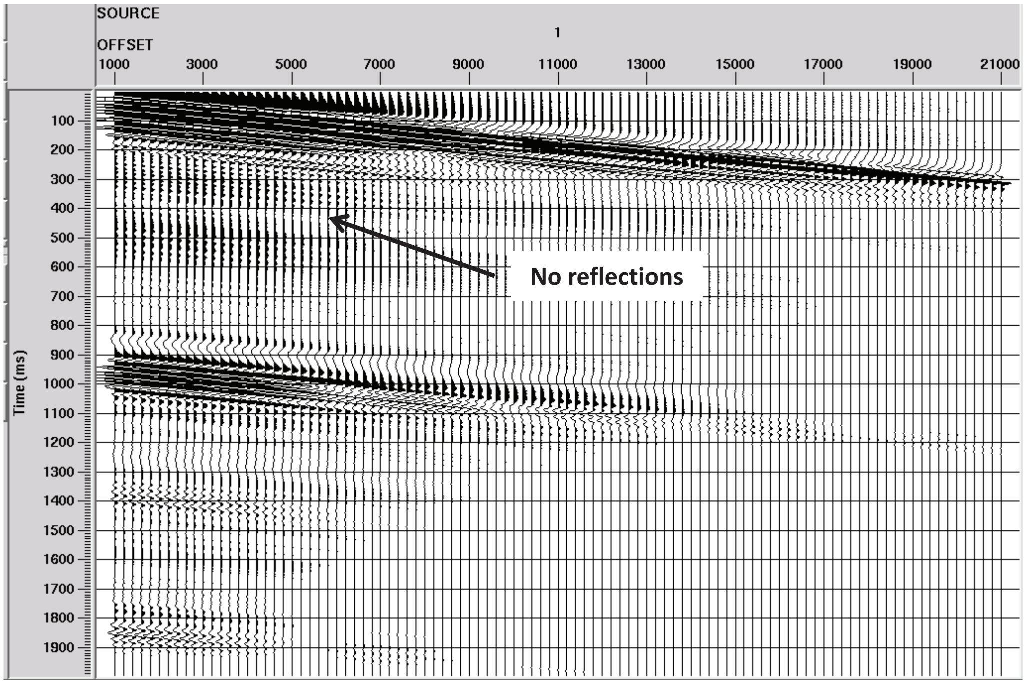

Decimated physical model shot gather after LMO, low-pass filter in R-T domain

FIG. 9. Low-pass filter in the R-T domain effectively removes reflections from the surface wave estimate.

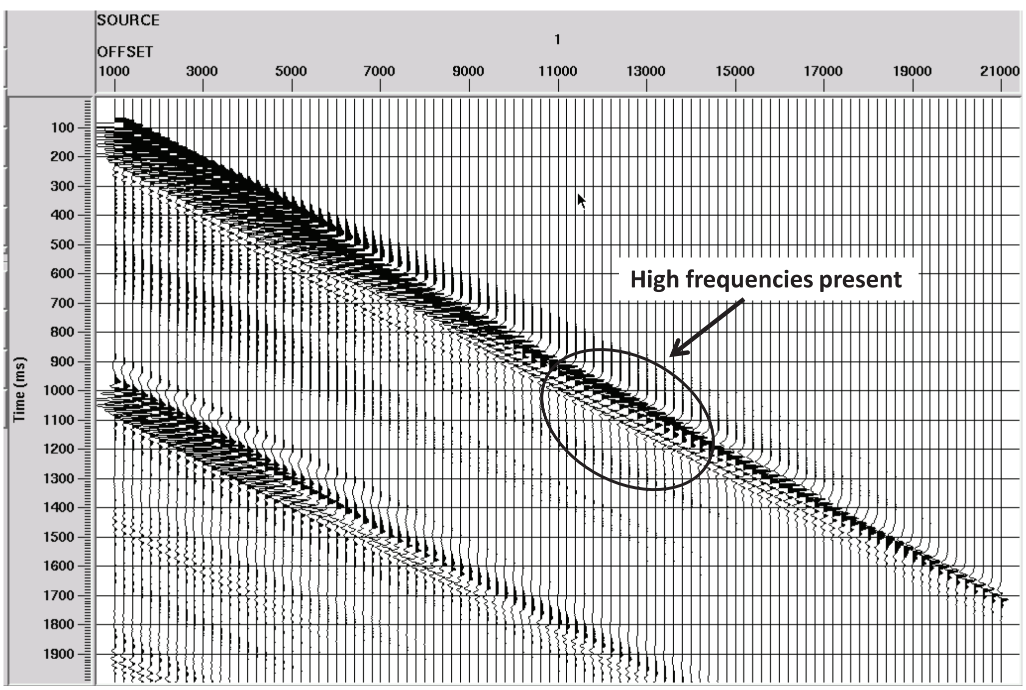

Decimated physical model shot gather after LMO, low-pass filter in R-T domain, inverse LMO

FIG. 10. Restoring LMO completes the surface wave estimate, which can now be subtracted from the raw decimated physical model shot gather.

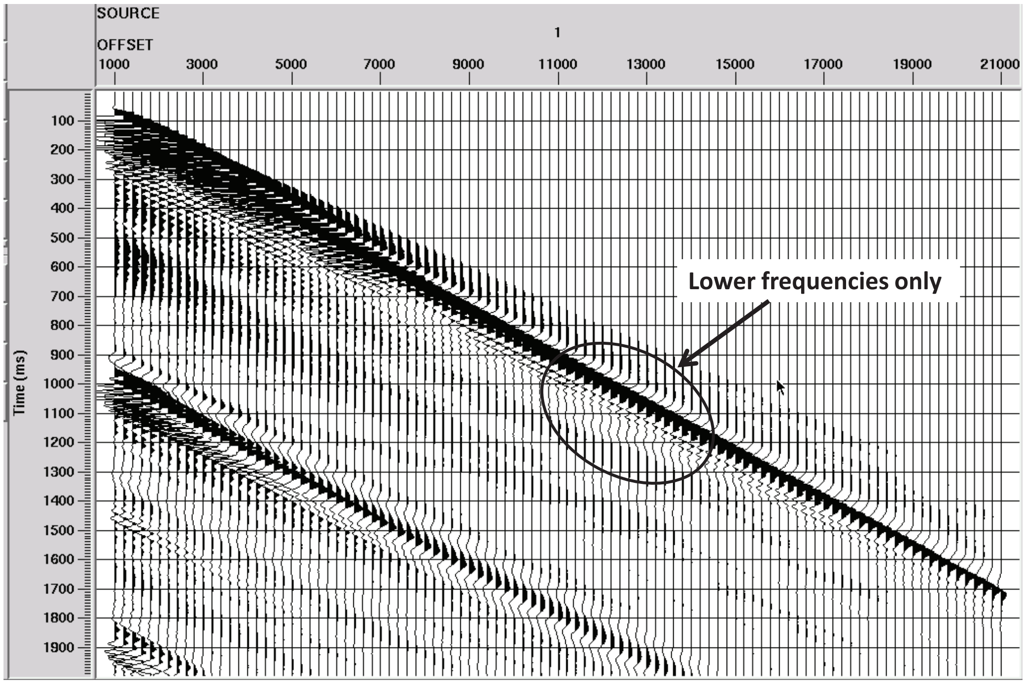

Decimated physical model shot gather, no LMO de-alias, low-pass filter in R-T domain only

FIG. 11. The surface wave estimate that would be created from the decimated shot gather without applying LMO de-aliasing. Only the lower frequencies are adequately represented.

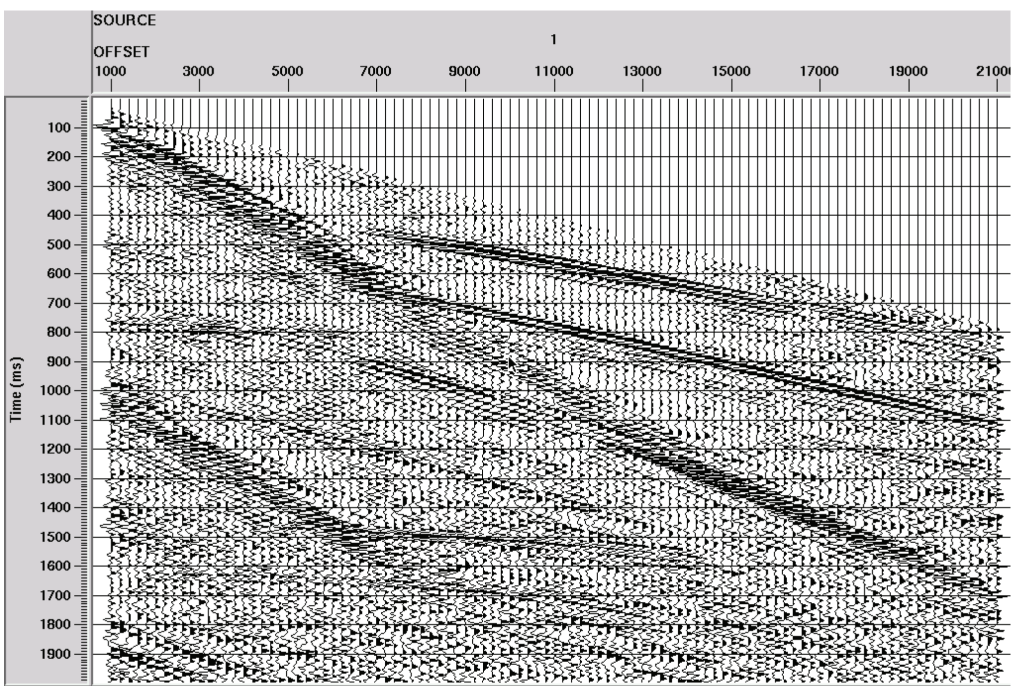

Decimated physical model shot gather after LMO de-aliased noise subtraction

FIG. 12. The decimated physical model shot gather after subtraction of the de-aliased surface wave estimate in Figure 10. Compare with Figure 7, where no de-aliasing was applied.

Interpolating the noise

Another technique that we tested for attenuating coherent noise on an undersampled and aliased trace gather was the following: We applied a partial LMO correction to the trace gather, where the LMO velocity was such that the coherent surface wave was partially de-aliased, but the reflections were *not* aliased. Figure 13 shows the gather after this LMO correction. To see what effect this partial LMO has on the aliasing of the noise, we look at the radial trace transform of the original shot gather in Figure 14, and at the R-T transform of the decimated gather in Figure 15, on which the aliasing of the surface wave is obvious. The R-T transform of the decimated gather after partial LMO is shown in Figure 16, compared with the R-T transform of the original gather, also after LMO, in Figure 17. The aliasing of the surface wave is significantly decreased. After this partial de-aliasing, we interpolate the decimated shot gather in the R-T domain to obtain the gather shown in Figure 18, and remove the LMO to obtain the gather in Figure 19, our estimate of the original gather in Figure 5, which is our goal for these tests. To see how well the surface wave has been interpolated, we apply the usual noise attenuation and deconvolution sequence to the gather in Figure 19 to obtain the filtered gather in Figure 20. For comparison, the same noise attenuation and deconvolution applied to the decimated gather is shown in Figure 21. The original shot gather with the same processing appears in Figure 22. Comparing the three figures, we see that the partial de-aliasing and interpolation method gives better results than nothing at all, but still not as good as filtering the original shot gather.

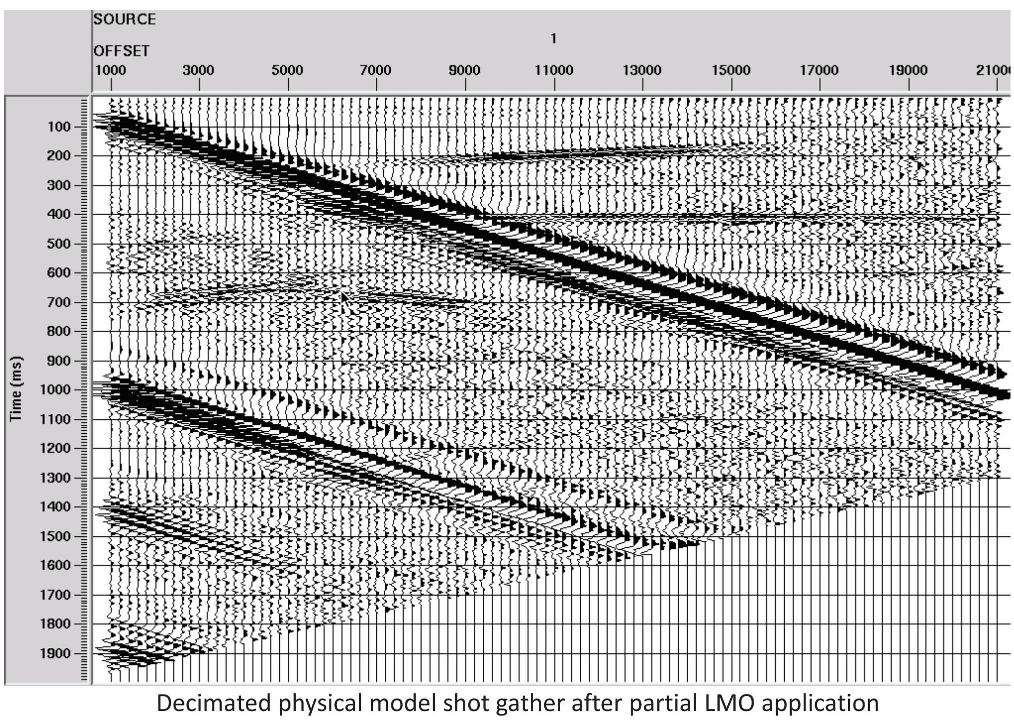

FIG. 13. Decimated physical model shot gather after partial LMO applied. Note that the reflections have not been aliased by the partial LMO.

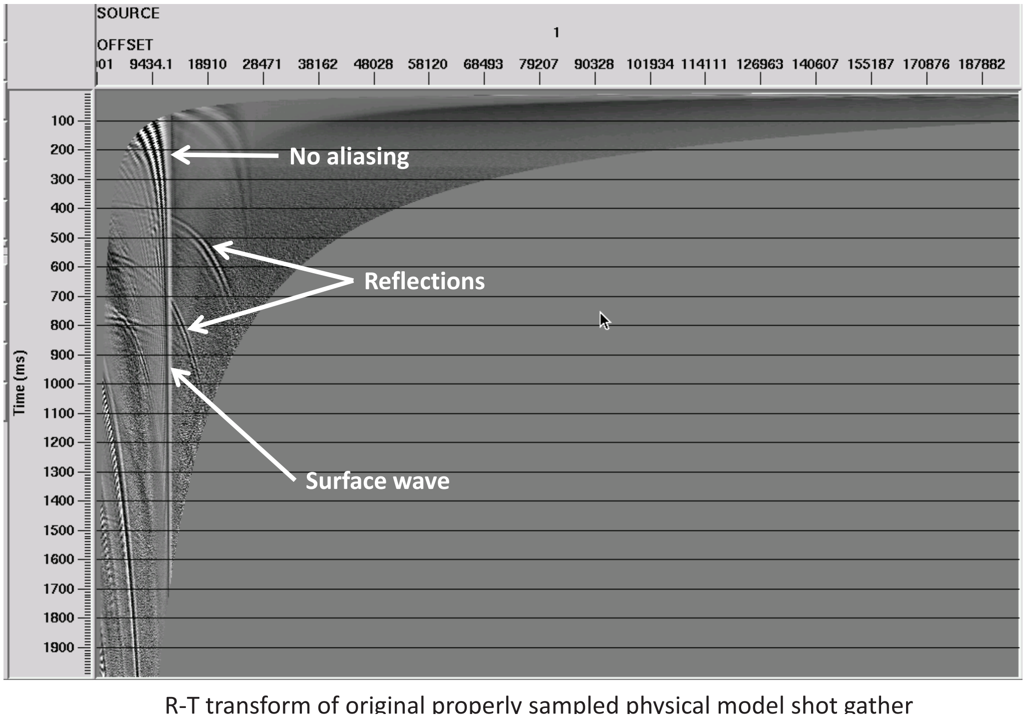

FIG. 14. R-T transform of the original properly spatially sampled physical model shot gather. No events show evidence of aliasing.

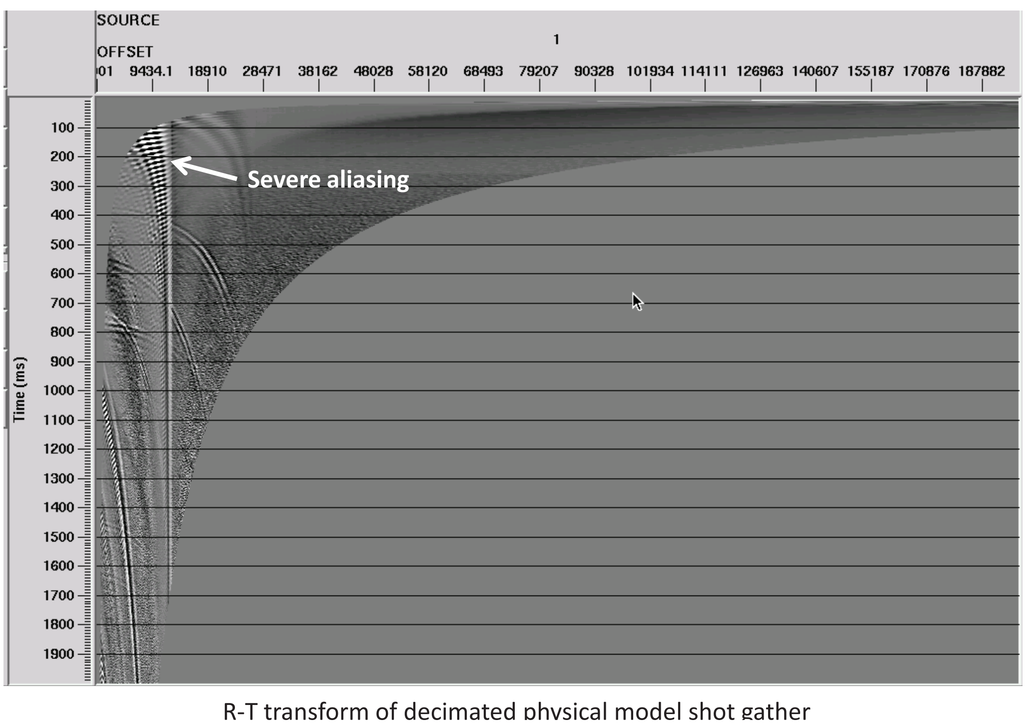

FIG. 15. R-T transform of the decimated physical model shot gather. While the reflections show no signs of aliasing, the surface wave is severely aliased and cannot be properly estimated by a low-pass filter in this domain.

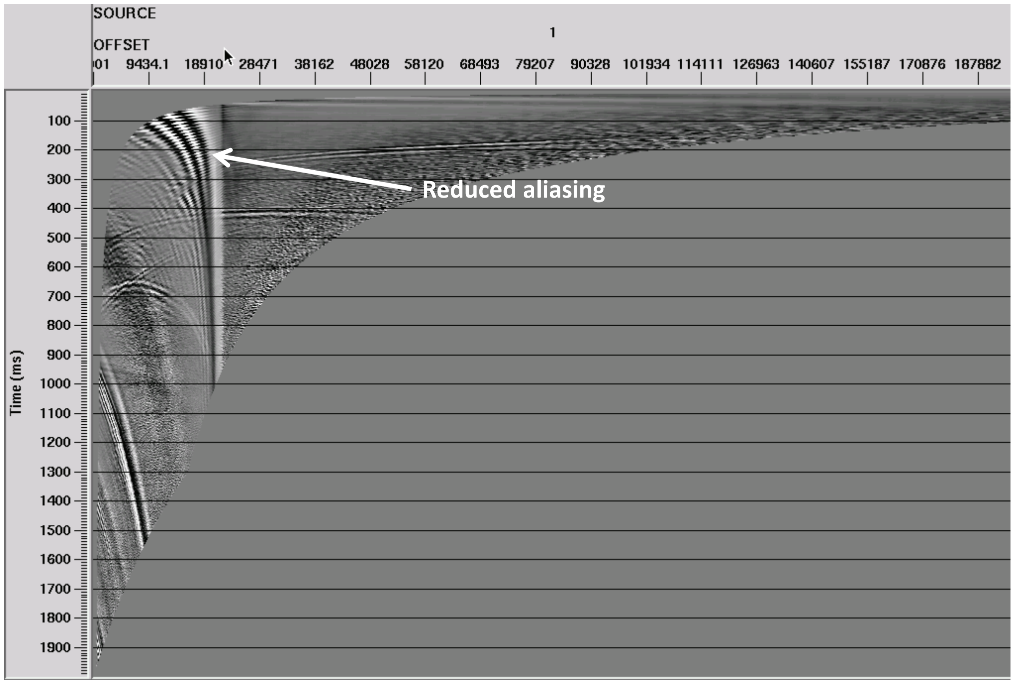

R-T transform of decimated physical model shot gather after partial LMO

FIG. 16. R-T transform of the decimated physical model shot gather after partial LMO. The aliasing of the surface wave has been greatly reduced.

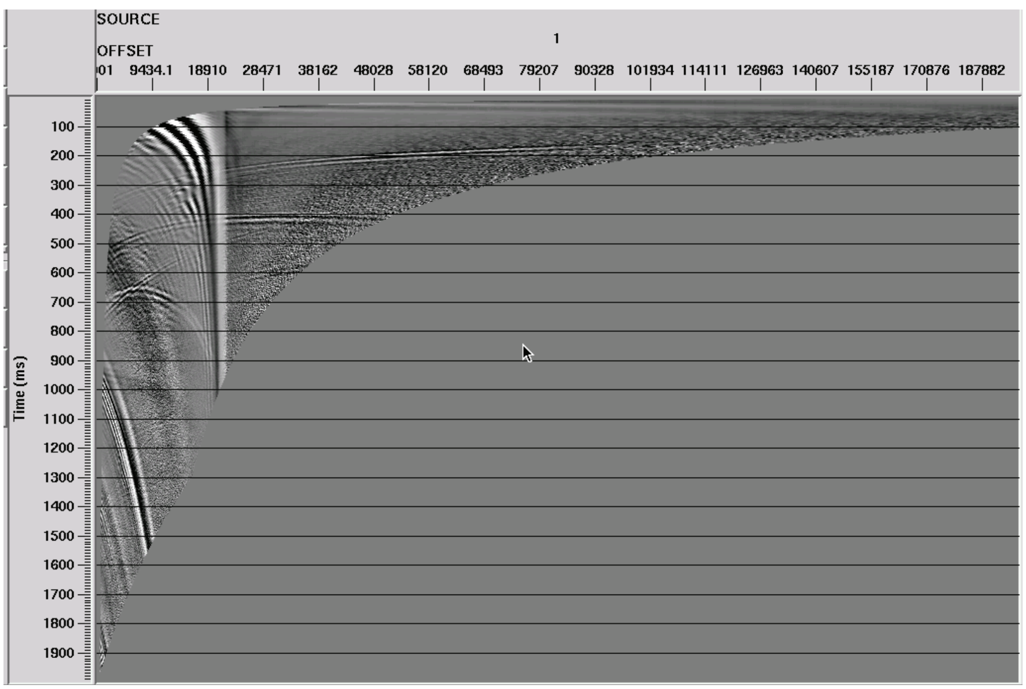

R-T transform of original shot gather after partial LMO

FIG. 17. R-T transform of the properly sampled physical model shot gather after partial LMO, for comparison with Figure 16.

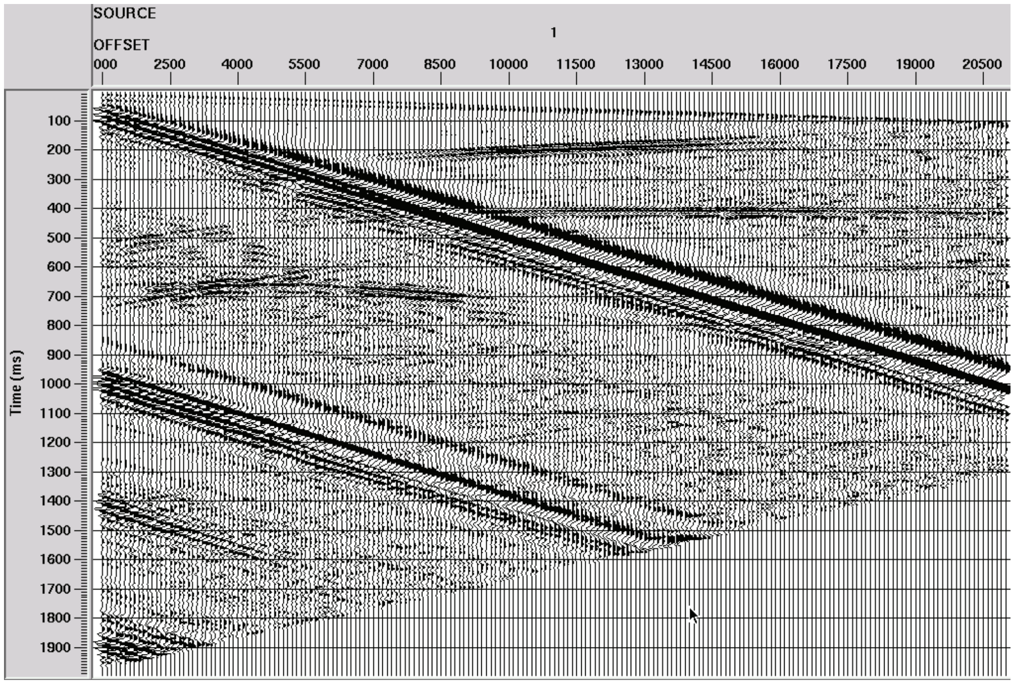

Decimated physical model shot gather after partial LMO, R-T domain interpolation

FIG. 18. The decimated physical model shot gather, with partial LMO applied, has been interpolated in the R-T domain, to the same spatial sampling as the original physical model shot gather.

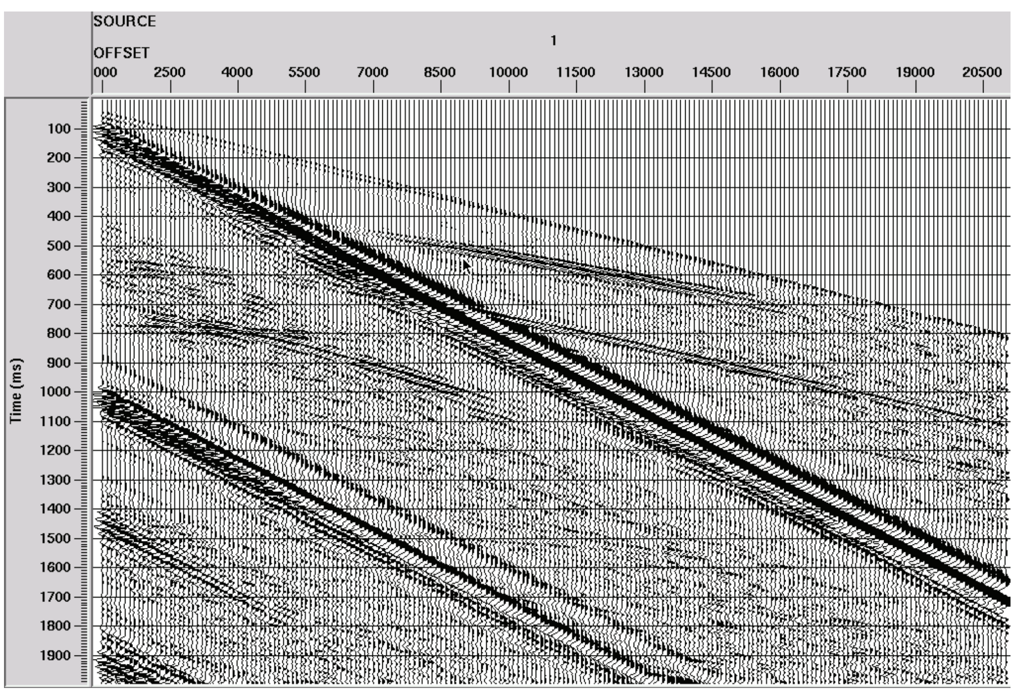

Decimated physical model shot gather after partial LMO, R-T domain interpolation, inverse LMO

FIG. 19. The LMO has been removed from the gather shown in Figure 18. This is an estimate of the original properly spatially sampled gather, in which the surface wave has hopefully been interpolated without aliasing.

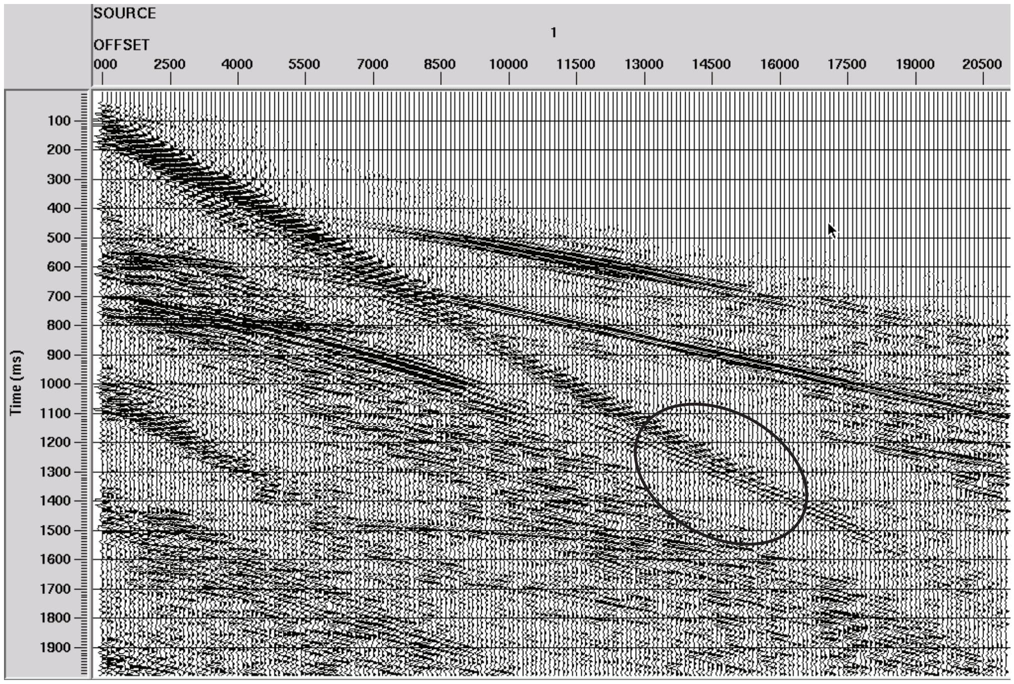

Decimated physical model shot gather after noise interpolation, R-T noise attenuation

FIG. 20. The interpolated shot gather from Figure 19 after R-T domain coherent noise attenuation. Residual noise is still present, but underlying reflections can now be seen (elipse). This is still not as good as proper spatial sampling during the acquisition!

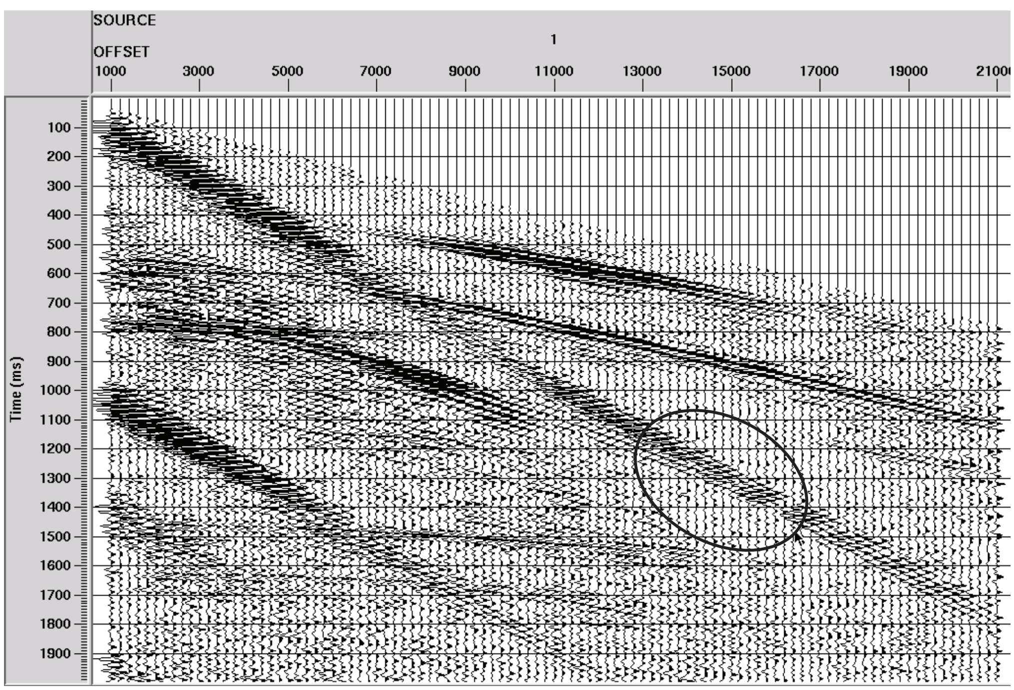

Decimated physical model shot gather with no noise interpolation, R-T noise attenuation

FIG. 21. The decimated shot gather with R-T domain noise attenuation. Residual noise obscures underlying reflections.

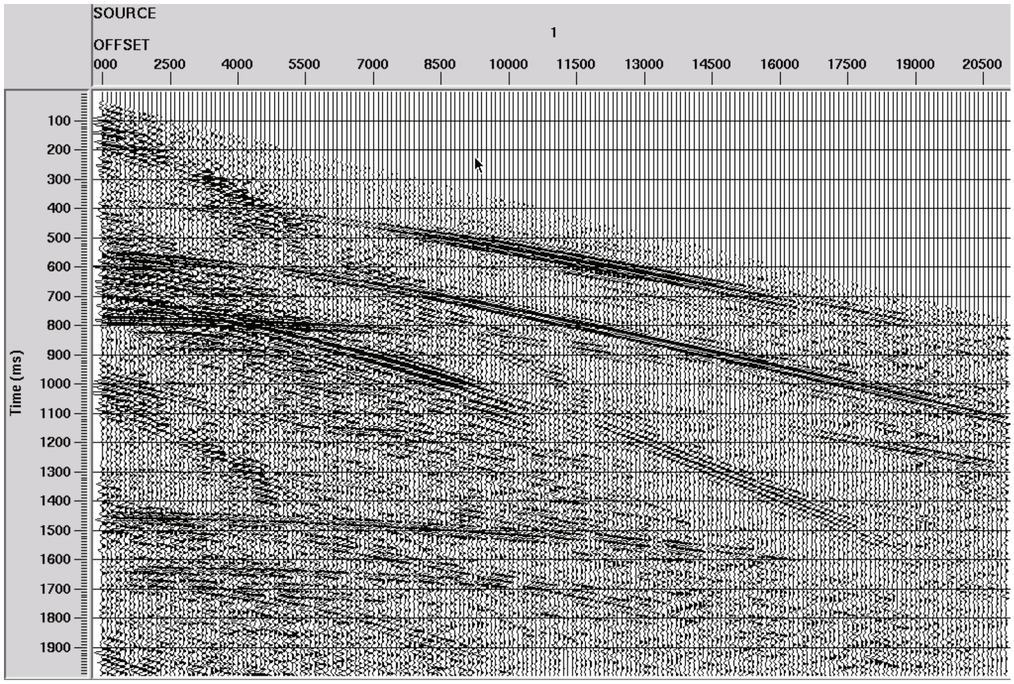

Fully sampled physical model shot gather after R-T coherent noise attenuation

FIG. 22. The original properly spatially sampled shot gather after R-T domain noise attenuation, for comparison.

Going nonlinear

A final tactic tested for coherent noise removal was to sacrifice relative trace amplitude information and apply AGC in the R-T domain. The AGC operation, when applied in the context of an R-T dip transform whose velocity is that of the surface wave, dramatically reduces the amplitude of the surface wave relative to the background reflections. Nothing is free, however, and in using this operation, we sacrifice any AVO relationships within the original trace gather (Henley, 2011).

Two noise attenuation sequences were tested. In one case, the decimated gather was subjected to the R-T domain AGC, then a set of R-T domain filters, followed by R-T domain interpolation. In the other, the decimated gather was first interpolated, then subjected to the R-T domain AGC and R-T filtering. The order of the interpolation was reversed between the two sequences. For comparison, the original fully sampled gather was subjected to the same R-T domain AGC and filters. In all three cases, Gabor deconvolution and F-X deconvolution were applied, as well, to sharpen events and reduce random noise. Figure 23 shows the fully sampled gather after this processing; little can be seen of the original surface wave except shadow zones where it formerly interfered with reflections. Some residual dipping noise at small offsets and early times can still be seen, as well as random noise; but this gather is a dramatic improvement over the original in Figure 5. Figure 24 shows the result of starting with the decimated gather, interpolating it in the R-T domain, and applying the R-T domain AGC and filter sequence. The results show less random noise than Figure 23, but the deeper reflections are weaker than in the original. Somewhat surprisingly, if the decimated gather is subjected to the noise attenuation sequence before being interpolated in the R-T domain, the results, as seen in

Figure 25, are actually better than when the gather is interpolated first. Furthermore, the dipping noise at shallow times and near offsets is actually less than on the original gather. The deeper events are still not as strong as on the original, however. Hence, if we are willing to sacrifice AVO in the interest of imaging, the nonlinear operation of AGC in the R-T domain can be used to attenuate isolated linear coherent noises nearly as effectively as decreasing the original trace spacing by a factor of two.

To further illustrate why R-T AGC is so effective, we show, in Figure 26, the R-T dip transform of the unfiltered shot in Figure 5. Note that because the dip velocity of the transform closely matches the velocity of the linear surface wave, all the linear noise has been rotated into vertical traces whose lateral amplitude variations essentially mirror the waveform of the noise. By applying a short-window AGC to the radial traces, these lateral amplitude variations are greatly diminished (Figure 27). The much weaker linear noises then may be almost totally eliminated by the usual R-T filters (Figure 23). This technique works best for isolated linear noises whose velocity can be carefully matched by radial trace velocity parameters.

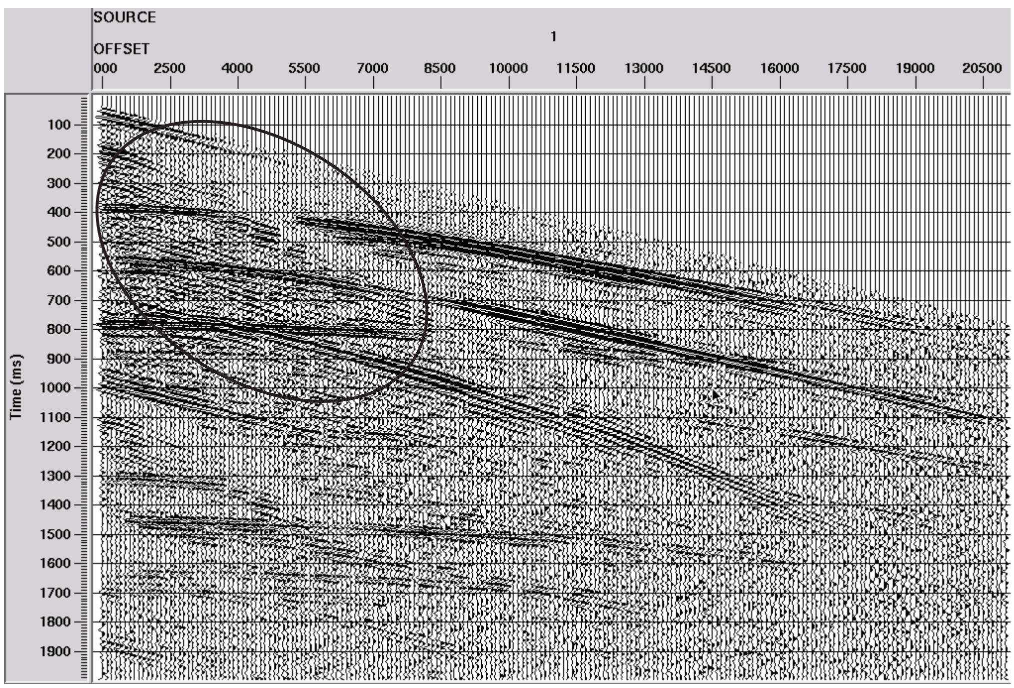

Fully sampled physical model shot gather after R-T domain AGC, R-T noise attenuation

FIG. 23. Properly sampled physical model shot gather after R-T domain AGC, R-T domain noise attenuation. Some residual linear noise still visible.

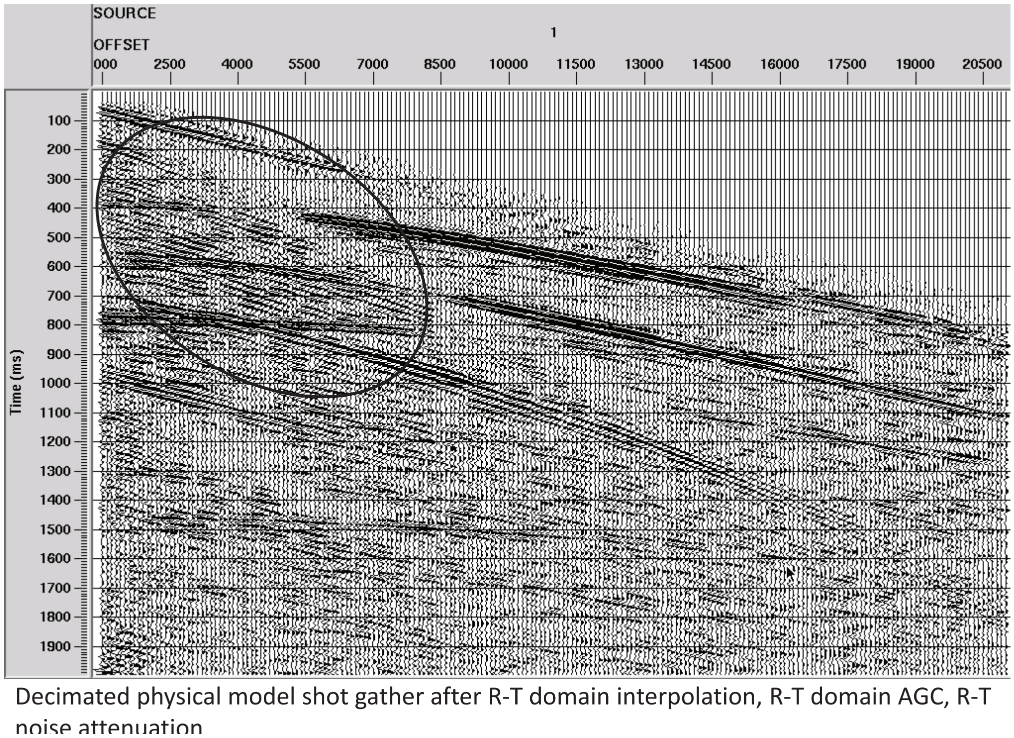

FIG. 24. Decimated physical model shot gather after R-T domain interpolation, R-T domain AGC, and R-T domain noise attenuation. Deeper events are difficult to see; residual linear noise still present.

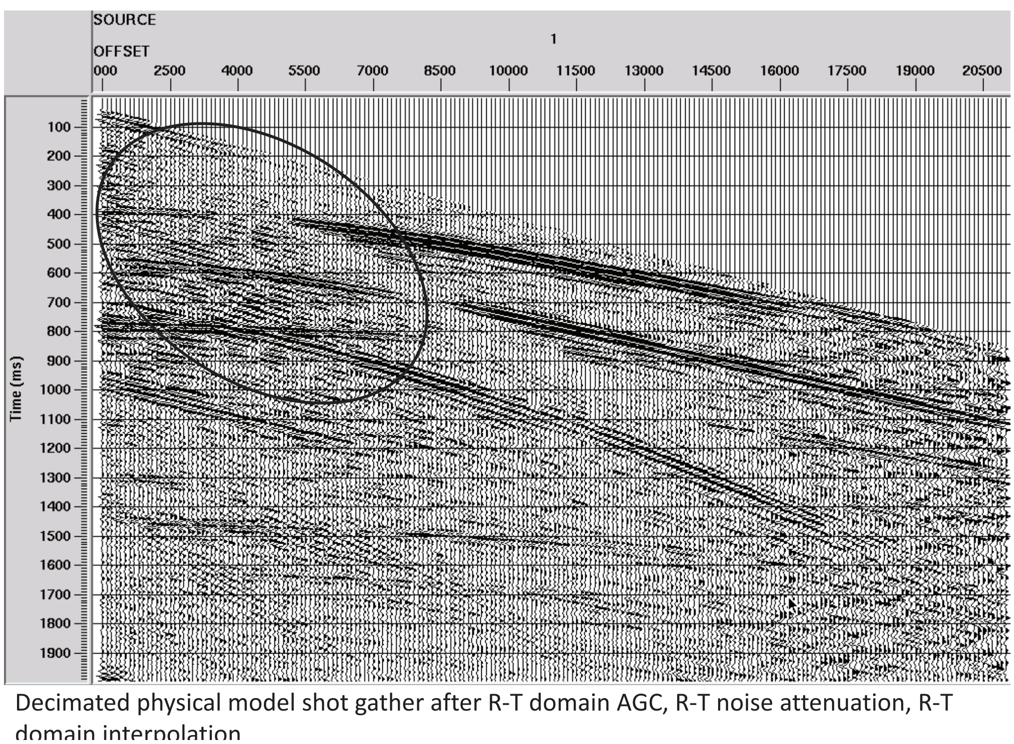

FIG. 25. Decimated physical model shot gather after R-T domain AGC, R-T domain noise attenuation, and R-T domain interpolation. Level of residual linear noise is lower, events at longer offset are stronger.

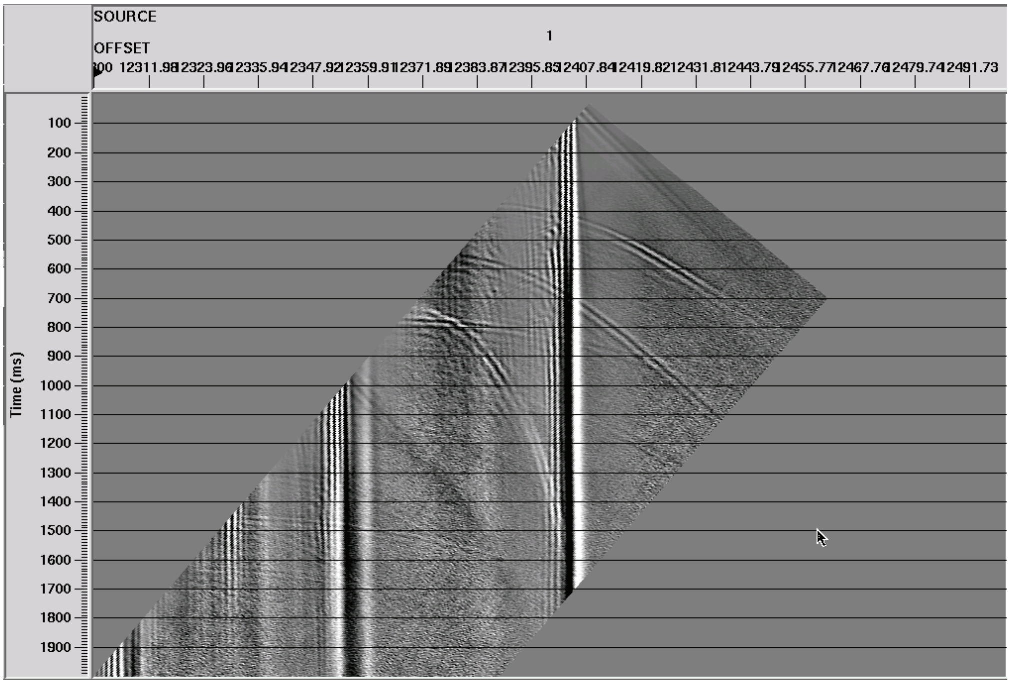

R-T dip transform of raw shot record in Figure 5—dip velocity=12,400m/s

FIG. 26. R-T dip transform of the shot gather in Figure 5. The very strong surface wave is manifested as very large trace-to-trace amplitude fluctuations because the coherent noise wavefront velocity has been carefully matched by the R-T dip velocity.

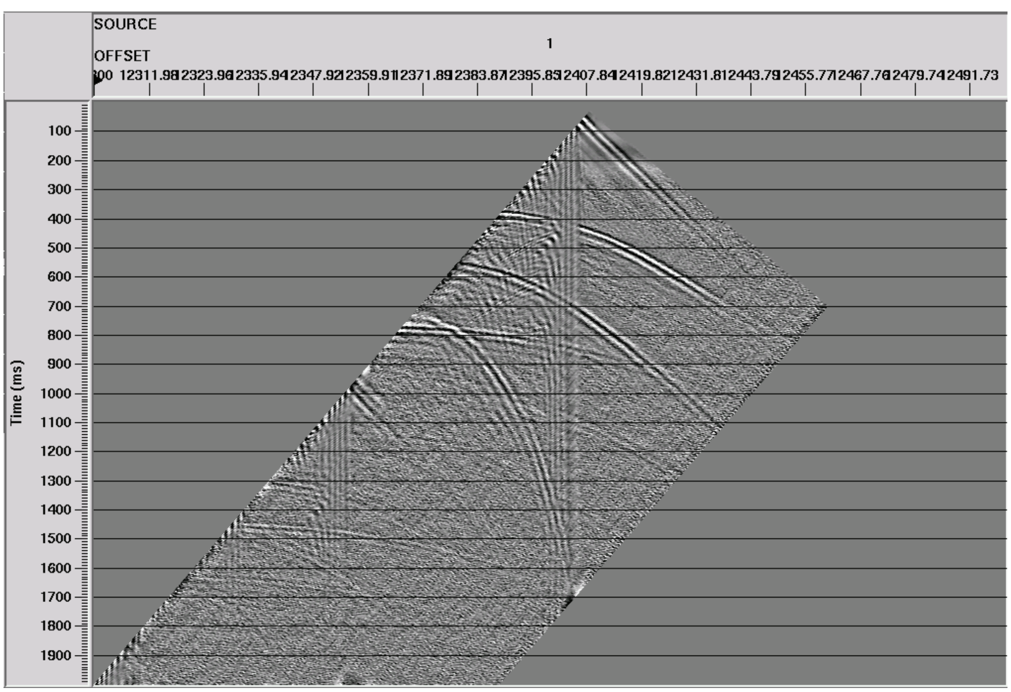

R-T dip transform of raw shot record in Figure 5—dip velocity=12,400m/s—AGC applied.

FIG. 27. R-T dip transform of the shot gather in Figure 5 after application of AGC. Since the large trace-to-trace amplitude fluctuations are gone, the surface waves are also largely gone.

DISCUSSION AND CONCLUSIONS

We have argued in the past that effective attenuation of linear coherent noise requires seismic data recorded with proper spatial sampling to prevent aliasing of the slowest coherent noise (Henley et al, 2006, 2008, 2009). In this report we have tested three different processing schemes aimed at increasing the effectiveness of noise attenuation when we are prevented by acquisition circumstances from sampling at the optimum anti-alias spacing. Our relatively modest goal was to emulate spatial sampling of ½ the actual recorded sampling by using various simple processing manipulations to reduce the inevitable noise aliasing of the original record. The availability of a well-sampled physical model shot record meant that we could decimate the record by discarding every other trace, and try to attenuate coherent noise as effectively as we could on the original record. While we were successful in reducing the overt aliasing of coherent noise, the results were never as good as having sampled at the actual anti-alias spacing in the first place, unless we adopted a non-linear method which compromised relative trace amplitudes so that AVO analysis could not be performed on the filtered record.

The simplest method we tried was to de-alias the coherent noise by applying linear moveout during the estimation of the coherent noise for subtraction from the original data. While the method was effective for the lower frequencies in the coherent noise, it was not as good for the frequencies most severely aliased by the undersampling in the test record. We showed that a combination of partial de-aliasing and R-T domain interpolation produced a reasonable result, though still not as good as a properly sampled record. The most effective method we tested actually involved a nonlinear operation in the R-T domain which disturbs relative trace amplitudes and leaves them unsuitable for AVO analysis. In this case, noise attenuation on the decimated test record was actually slightly better on some parts of the record than that on the original fully sampled record.

While simple processing measures can be taken to improve coherent noise attenuation on records which have been inadequately spatially sampled, they are no substitute for proper sampling during initial acquisition. Even under the most favourable circumstances, their effectiveness is not as good as reducing the spatial sampling by a factor of 2.

ACKNOWLEDGEMENTS

The authors thank CREWES sponsors for continuing support.

REFERENCES

Henley, David C. 2003, Coherent noise attenuation in the radial trace domain, Geophysics, **68**, No. 4, pp 1408-1416.

Henley, D.C., Bertram, M.B., Hall, K.W., Bland, H.C., Gallant, E.V., and Margrave, G.F., 2006, The power of high effort seismic acquisition: the Longview experiment: CREWES 2006 Research Report, **18**.

Henley, D.C., Hall, K.W., Bertram, M.B., and Gallant, E.V., 2008, Pushing the resolution limits: the Priddis 3C-2D survey: CREWES Research Report, **20**.

Henley, D.C, Hall, K.W., Bertram, M.B., Gallant, E.V., Lu, H-X., and Maier, R., 2009, Shaken, not stirred: Priddis 2009 3C-2D hi-res acquisition: CREWES Research Report, **21**.

Henley, D.C., 2011, How processing affects relative reflection amplitudes: CREWES Research Report, **23**.

Wong, J., Hall, K.W., Bland, H.C., Gallant, E.V., and Bertram, M.B., 2007, Some 2D results from the U of C Seismic Physical Modelling Facility, CREWES Research Report, **19**.

Wong, J., Hall, K.W., Gallant, E.V., and Bertram, M.B., 2008, Mechanical and electronic design for the U of C Seismic Physical Modelling Facility, CREWES Research Report, **20**.

Wong, J., Hall, K.H., Gallant, E.V., Maier, R., Bertram, M.B., and Lawton, D.C., 2009. Seismic physical modeling at the University of Calgary, CSEG Recorder, **34**, 36-43