

Generalized sampling and "beyond Nyquist" imaging

Peter W. Cary

ABSTRACT

We routinely image seismic data beyond the usual Nyquist sampling limits when we stack into limited-offset stacks either before or after migration. Binning or stacking a number of aliased common-offset gathers approximately dealiases the data because of the staggered common-midpoint sampling between common-offset gathers. The interleaved subsurface sampling of individual subsets of the full prestack data allows unaliased poststack images to be obtained from aliased prestack data. The highest recoverable wavenumber in the migrated image is not determined by the smallest sample rate of the input wavefield, as simple sampling theory dictates, but by the average subsurface sample interval, which follows from a more general form of sampling theory. Prestack migration of individual aliased subsets of the data should allow the migration operator to alias in order to recover all frequencies in the stack.

INTRODUCTION

A well-sampled common-offset (CO) section is an ideal dataset for prestack imaging. It has a single live trace at every common-midpoint (CMP) location between one end of the 2-D line and the other. The spatial and temporal sample intervals are less than or equal to the wavefield's Nyquist sample rates so that the continuous wavefield, for that value of offset, could be perfectly reconstructed using Shannon's sampling theorem. Migration of this unaliased, 1-fold dataset would yield a correct, and complete, 1-fold image of the subsurface over the entire 2-D line, except for edge-effects at the ends of the line. Other common-offset images would improve the signal-to-noise ratio, but not the fidelity, of the final image.

Unfortunately, seismic data are rarely acquired in a way that allows CO sections to be sampled correctly. For regular acquisition geometry, the separation between traces with the same offset is the shot interval. Typical marine geometries have a shot interval equal to the receiver interval (or twice the receiver interval), and typical land geometries have a shot interval that is from 1 to 6 times the receiver interval (with a lot of variations). If the wavefield's highest wavenumber requires sampling at the CMP interval, then true CO sections invariably suffer from mild to severe forms of data aliasing.

The standard method of designing seismic surveys (e.g. Liner and Underwood, 1999) chooses the CMP bin size by setting it equal to the largest sample rate that reliably samples the most steeply dipping event on a hypothetical zero-offset stacked section (sampled at the CMP interval), which is assumed to be recoverable during processing. It is not well recognized that this method of deciding on the spatial Nyquist sample rate can cause the high frequencies in the prestack data to be aliased in prestack gathers, unless the shot interval is as small as half the receiver interval. The following example shows that this true.

Fig. 1 shows an example of an unmigrated, constant-velocity zero-offset section (and its F-K amplitude spectrum) with dipping events that are unaliased at the CMP interval. The prestack gathers associated with this zero-offset stack are all aliased for most acquisition geometries. For example, Fig. 2 shows a zero-offset gather that results from an acquisition geometry that has the source interval equal to the receiver interval (and assuming no source or receiver arrays have been used). Since the spatial sample interval is now twice the CMP interval, aliasing of the steeper events occurs. Fig. 3 is a portion of a shot gather from the same dataset, which shows that aliasing also still occurs in this domain, although not as severely. Fig. 4 shows that NMO correction of the shot gather has removed the aliased high frequencies of the steeply-dipping events because of NMO stretch. However, prestack shot migration is performed on data before NMO-correction, so it would still have to confront the aliasing in Fig. 3.

Since we virtually always choose shot intervals to be greater than CMP intervals, Figure 1 to 4 indicate that our prestack gathers could often be aliased. How, then, do we manage to construct unaliased images at the CMP interval? In addition, how do we manage to get unaliased results from prestack migration of shot gathers or CO sections, if these gathers suffer from data aliasing? Are we wrong in assuming that an image that is correctly sampled at the CMP interval can be recovered in processing?

We actually do recover, or at least try to recover, images that are correctly sampled at the CMP interval. However, it is probably fair to say that we are not always aware of how we do it. The main purpose of this paper is to explain how the sampling of combinations of subsets of the data overcomes the sampling limitations of individual subsets of the data. Before explaining how partial stacking approximately dealias seismic data, it is worthwhile to discuss some more general approaches to the dealiasing problem.

DEALIASING BY INVERSION

Clearly, if we are going to produce an image with a Nyquist sample rate that is close to the CMP interval, then we need to somehow dealias our aliased prestack data. Vermeer (1990) discusses methods for dealiasing CO sections that result from “symmetric sampling” acquisition geometry by interpolation of shot and receiver gathers, and Jakubowicz (1994) uses a priori information about the expected slope of the data to interpolate the shot gathers that are not acquired during multisource (“flip/flop”) marine surveys. However, neither of these methods is concerned with obtaining information in the image beyond the wavenumber provided by sampling at the receiver interval.

Ronen (1987) sets up the problem of estimating a zero-offset section that is sampled at the CMP interval (or at any interval, for that matter) from an irregularly sampled nonzero-offset dataset using least-squares DMO. The forward modeling set of equations is of the form

$$\mathbf{d} = \mathbf{Lm},$$

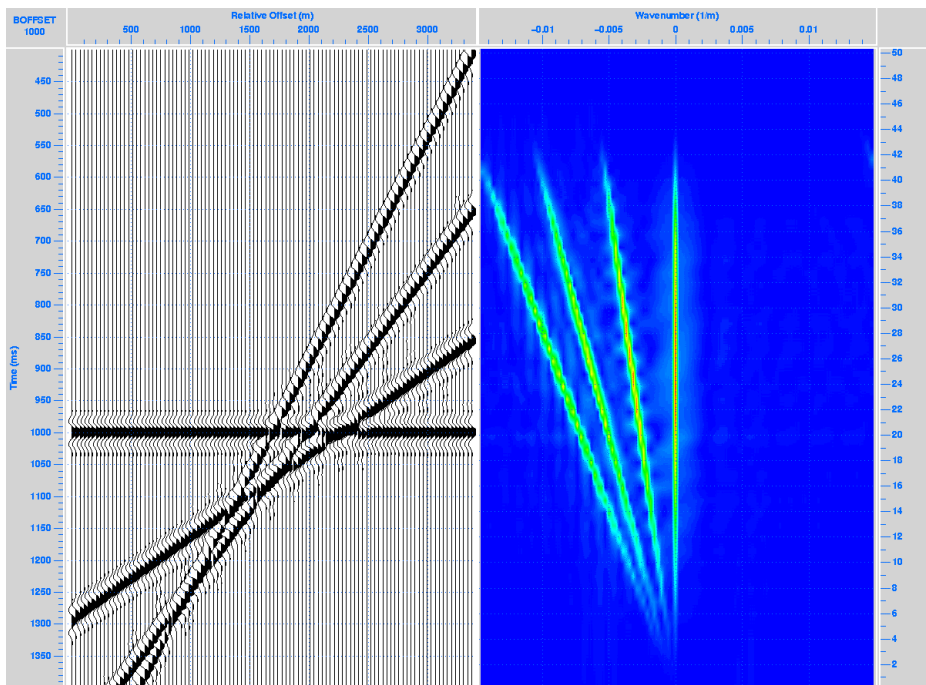


Fig. 1. Dipping events on zero-offset section sampled at the CMP interval. F-K spectrum on the right shows that events are not aliased.

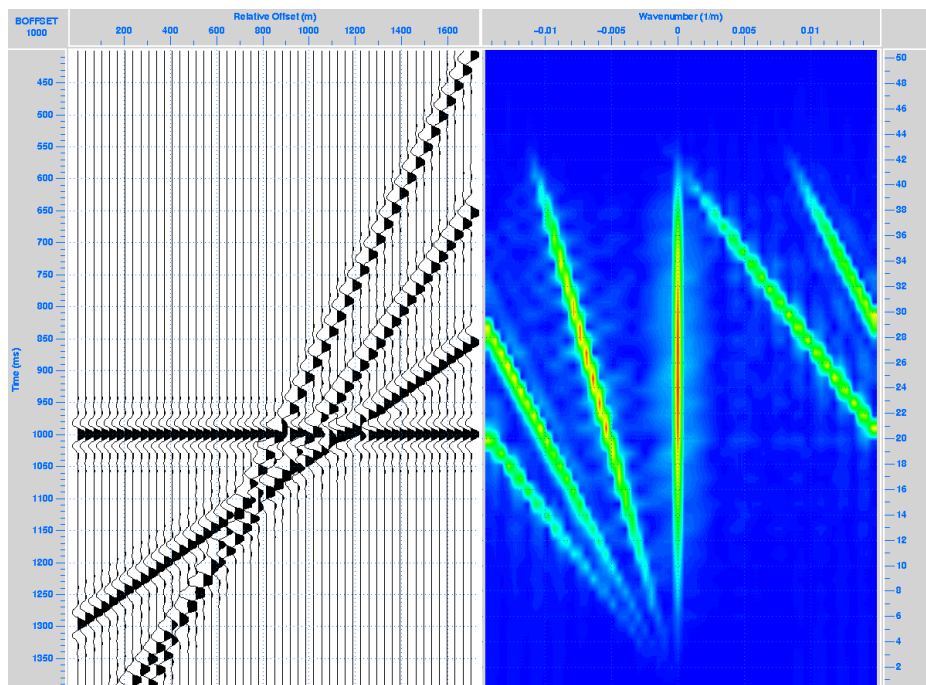


Fig. 2. A zero-offset section as recorded with a shot interval equal to the receiver interval. The steeper events are now aliased since the trace interval is twice the CMP interval.

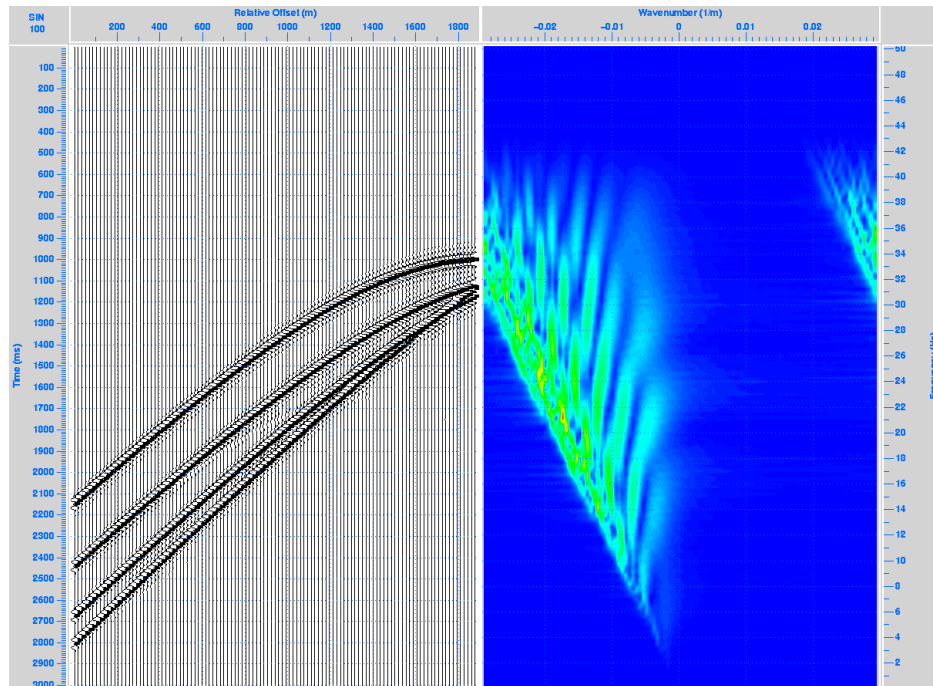


Fig. 3. A portion of a shot gather showing that aliasing also occurs in this domain.

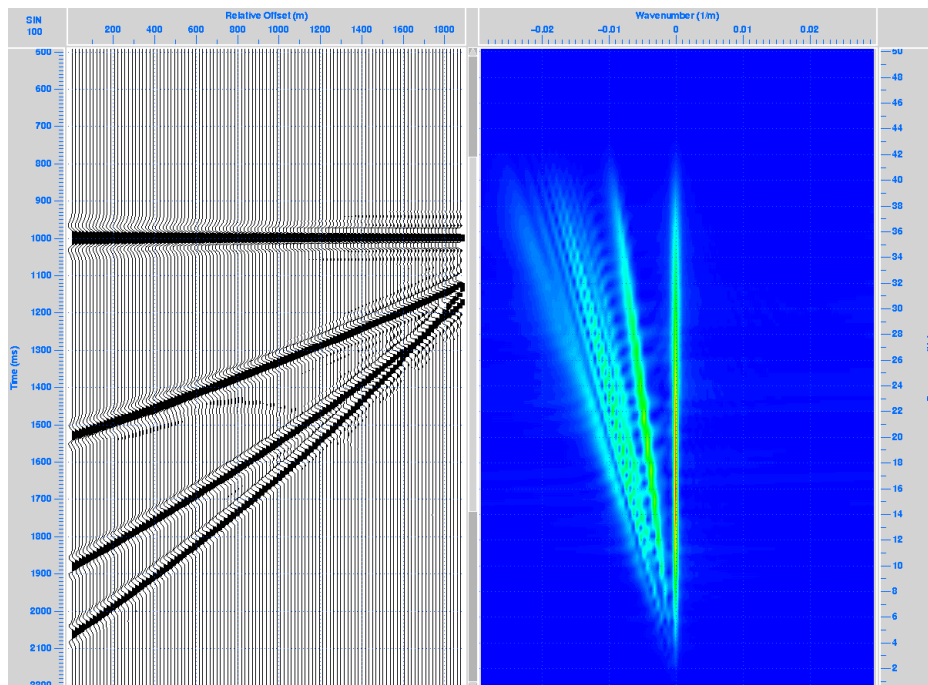


Fig. 4. NMO stretch removes the aliased frequencies from the steeper events of the shot gather.

where \mathbf{d} is the input, irregularly sampled data, \mathbf{m} is the desired zero-offset section, and \mathbf{L} is the inverse-DMO operator. Applying a logarithmic stretch to the time axis makes the DMO operator time-invariant, which allows the problem to be posed in the frequency domain and solved for each frequency individually, which reduces the size of \mathbf{L} considerably. However \mathbf{L} is normally ill-conditioned, so the solution is non-unique.

Non-uniqueness is overcome by imposed *a priori* information about the expected form of the solution into the problem with a model constraint. The solution is obtained by minimizing an error function, E , which is the weighted sum of a data constraint (normally the sum of the squared difference between the actual and modeled data, ϕ_d) and a model constraint, ϕ_m :

$$E = \phi_d + \lambda\phi_m,$$

where λ is a constant. The simplest, and most commonly used, model constraint is to choose the model with the smallest “size” in an l_2 sense, so that

$$\phi_m = \sum_i m_i^2.$$

The result of minimizing E is the familiar damped-least-squares solution:

$$\mathbf{m} = (\mathbf{L}^T\mathbf{L} + \lambda\mathbf{I})^{-1}\mathbf{L}^T\mathbf{d},$$

where \mathbf{I} is the identity matrix.

When Ronen (1987) computed the damped-least-squares solution from poorly sampled input data, he obtained a stack with very poorly balanced amplitudes. A more reasonable solution was obtained with a model constraint that imposed a type of smoothness on the solution in the Fourier domain. The nonuniqueness of this problem was investigated further by Ronen (1994) and Ronen et al.(1995). The question of what model constraint is best to use in all circumstances remains a problem because we do not know beforehand whether the correct model is smooth or rough, or is best described in some other way.

The lesson to learn from Ronen’s results is that aliasing cannot be overcome simply by solving a least-squares problem with wave-equation extrapolators like DMO or prestack migration. The aliasing of prestack gathers is overcome by the combination of the model constraint and by the fact that the sampling of the stack is not as poor as the sampling of each prestack gather. If all the input data taken together are still effectively aliased, then only a severe model constraint can overcome aliasing.

The problem can be broken into smaller pieces by using least-squares offset-continuation or azimuth moveout (Chimengui and Biondi, 1999) to interpolate

irregularly sampled subsets of the data into regularly-sampled subsets. Least-squares migration (Nemeth and Schuster, 1999) is in some ways the most attractive approach to this problem because modelling plus migration can in principle take all the subtleties of wave propagation into account. However, it is also the most expensive approach because the migration operator is space and time-variant. More importantly, the solution will still be nonunique. A solution that overcomes severe sampling limitations of the data can only be obtained by imposing a strong model constraint, such as smoothness of the image. If that is the case, a much less expensive approach that regularizes the data before migration should be capable of obtaining virtually as good a solution.

DEALIASING BY PARTIAL STACKING

The most obvious way to solve the aliasing problem is with partial stacking. It is undoubtedly the most common, and least expensive, method of interpolating CO sections. It is somewhat surprising that the advantages and disadvantages of partial stacking (both before and after DMO or migration) have not been investigated in the geophysical literature more fully since it is used routinely in seismic processing.

Binning together neighboring true CO sections and considering them to be part of one CO section dealiases data by taking advantage of the fact that the CMP sampling is staggered between CO sections. For example, when the shot interval equals the receiver interval, then CO sections that sample the odd CMP numbers (1, 3, 5,...) alternate with CO sections that sample the even CMP numbers (0, 2, 4, ...). By binning neighboring pairs of CO sections together, the odd and even CMP sampling interleaves, so we obtain pseudo-CO sections with full CMP sampling (0, 1, 2, 3, ...). Likewise, if the shot interval is 5 times the receiver interval, such as in the data example that is presented in the next section, then binning together 10 true CO sections yields one trace at each CMP location.

This generalized form of sampling (Papoulis, 1977) would solve the aliasing problem exactly if the neighboring CO sections that are binned together were sampling the same function. However, the neighboring offsets are actually related by an offset-continuation operation (Bolondi et al., 1984). For offsets that are nearly the same, offset continuation is close to being a do-nothing operation. Offset continuation has a bigger effect as the CO sections get farther apart. Offset-continuation to zero offset is DMO.

Partial stacking before migration or DMO is the solution of the dealiasing inverse problem when we assume that the offset continuation operator is a do-nothing operator (the identity operation). We assume that the data can be accurately modeled as a constant within each offset bin (for each value of time and CMP) after partial NMO-correction to the offset at the center of the bin. So the extremely simple forward modeling equations becomes:

$$\mathbf{d} = \mathbf{L}_1 m$$

where $\mathbf{d} = (d_1, d_2, \dots, d_N)^T$ is a vector of the input samples that fall within the offset bin at one particular CMP location and time after partial NMO, m is the single model parameter to solve for at that CMP and time, and $\mathbf{L}_1 = (1, 1, \dots, 1)^T$. The least-squares solution,

$$m = (\mathbf{L}_1^T \mathbf{L}_1)^{-1} \mathbf{L}_1^T \mathbf{d} = \frac{1}{N} \sum_{i=1}^N d_i$$

is simply the average of the data within the offset bin, i.e. the partial stack. The solution is unique because the problem is completely overdetermined. So partial stacking before migration solves the nonuniqueness problem very neatly.

Partial stacking after migration or DMO can be considered to be the 0th iteration of an inverse problem that uses inverse migration or inverse DMO as the forward model. It is just a first guess at a solution to the least-squares problem. The 0th iteration should not be considered as a solution to any least-squares problem since it does not properly take into account the variations in the fold of the input or output data, and therefore can create a form of acquisition footprint.

Since partial stacking corresponds to using crude nearest-neighbor interpolation of true CO sections, it is easy to understand that it must corrupt the image of dipping events to some extent. However, the fact that partial stacking provides an effective method of overcoming the nonuniqueness problem inherent in the dealiasing inverse problem is not widely recognized. For example, we certainly do not usually think of CMP stacking as a method of imaging "beyond Nyquist", but in fact it is.

A quantitative analysis of the effect of partial stacking before or after migration is not easily done, except in the simplest case. If fully-populated common-offset gathers exist for all offsets, then partial stacking simply corresponds to convolving the data with a boxcar filter in the offset direction, where the width of the boxcar is twice the shot interval. The effect of partial stacking will be similar to, but less severe than, the time-dependent, low-pass filtering effect of doing CMP stacking over all offsets without DMO (Rocca and Ronen, 1984). However, since an irregular selection of offsets usually exists at each CMP location, a quasi-periodic, asymmetric distortion of waveforms is more likely to occur on dipping events.

In the next two sections I use simple synthetic data generated with real acquisition geometry to illustrate in a qualitative manner how well partial stacking before and after migration succeeds in solving the aliasing and irregular sampling problems. The results have implications for survey design, operator antialiasing, and noncommon-offset migration methods.

EXAMPLE WITH LARGE SOURCE INTERVAL

The geometry for the synthetic data that is used in this study is taken from a real 2-D land survey in Alberta, Canada. The receiver interval is 34m and the shot interval is 5 times larger (170m). As with most land lines, there are irregularities in the acquisition geometry, such as gaps in the shooting pattern, and skidded shots that attempt to

make up for the gaps. Figure 5 shows the part of the full stacking diagram that is centered on the portion of the line that is imaged in this study.

Four reflectors, with dips of 0° , 15° , 30° and 45° , are included in the model, and velocity is constant at 2000m/s. A 20 Hz Ricker wavelet was used. Modeling was done with a Kirchhoff approximation that sums the contributions of tiny segments that make up each reflector. The time and amplitude of each contribution is determined by ray tracing (Deregowski and Brown, 1983) with an analytic formula. Figure 6 shows the Kirchhoff migration of a pseudo-CO section where the offsets of all traces have been artificially set to zero, in order to indicate the quality of image that we would like to get from each common-offset migration. The image in Figure 6 still has some artifacts because of a few dead traces in the input CO section.

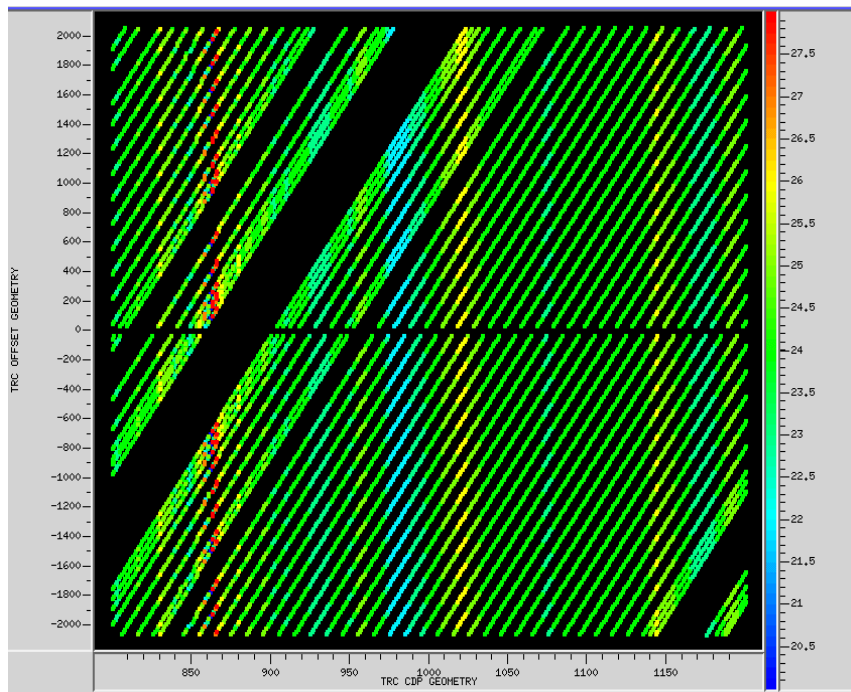


Fig. 5. Stacking diagram of a real 2-D line from Alberta with one shot every five receiver stations. The synthetic data in this study use this geometry. Offsets of $1000 \pm 170\text{m}$ are used since this provides nominal 1-fold coverage for one pseudo-common-offset section.




Fig. 6. Kirchhoff migration of a well-sampled zero-offset section.

The spacing between traces within true common-offset sections is the shot interval (170m), which guarantees that most events are severely aliased since only about every 1 trace in 10 has data, as in the true 1000m CO section in Figure 7.




Fig. 7. A severely aliased, true 1000m common-offset section.

Migrating each true CO section produces a result that is dominated by artifacts, as indicated by Figure 8.




Fig. 8. Migration of the true 1000m common-offset section in Fig.3.

The offset bin-width of a pseudo-CO section needs to be twice the shot interval in order to populate each output CMP location. Fig. 9 is a partial stack of all traces with absolute values of offset equal to $1000 \pm 170\text{m}$. Positive and negative offsets have been stacked together since it tends to fill gaps in the sampling of the partial stack. In Fig. 9, full NMO to zero offset has been applied to all traces, but before prestack migration, partial NMO to the offset at the center of each offset bin is applied. Live traces now occupy most, but not all, CMP locations. Obviously, partial stacking before migration severely distorts the image of the 45° reflector. However, the flat reflector is perfectly preserved since that is what NMO stack is designed to do. The distortion of waveforms that occurs with a period of roughly 10 traces on the 30° and 45° events is due to the fact that 10 true CO sections have been binned together under the assumption that all reflectors are flat. This is evidence of the breakdown of the assumption that the offset-continuation operator is a do-nothing operator. The 45° event appears to be made up of two interleaved events, which is due to both positive and negative offsets being include in the pseudo-CO section.




Fig. 9. Partial stack (pseudo-common-offset section) of traces with offsets of $1000 \pm 170\text{m}$.

Figure 10 is the prestack Kirchhoff time migration of the partial stack in Figure 9. As expected, the 45° reflector is very poorly imaged. However, the other reflectors are well imaged, especially the 0° and the 15° events.




Fig. 10. Prestack migration image by doing partial stack before migration.

For comparison, Figure 11 is a poststack Kirchhoff time migration of the partial stack in Figure 9. Notice that the 30° reflector is slightly mispositioned compared to Figures 6 and 10. However, the differences in the 15° and 0° reflectors are very small, so poststack migration is adequate when dips are small, as expected.

Figure 12 is the result of partial stacking after prestack migration. The 45° reflector is obviously imaged much better by using the correct offsets to migrate each input trace, and then stacking into the pseudo-CO section. The other events are also imaged in a coherent fashion. Notice that there is background migration noise (wavefronting) in all three migrated images (Figs. 10, 11, 12).




Fig. 11. Poststack migration of the partial stack in Fig. 9.




Fig. 12. Prestack migration image by doing partial stack after migration.

A cursory comparison of Figures 10 and 12 would favor the image generated by partial stacking after migration because of the better imaging of the steeper events. However, a close look at the waveforms on the low-dipping reflectors favors the partial stack before migration image. Figures 13 and 14 are close-up views of the flat-lying reflector. The partial stack before migration image (Fig. 13) is close to being perfect (the imperfections are mostly due to missing input traces in the partial stack). The partial stack after migration image (Fig. 14) has large variations in the waveforms that cause apparent 5 to 10 ms traveltimes variations, and peak-amplitudes that go up and down by about a factor of two.

Therefore, if your data consists of largely flat-lying reflectors, and you are interpreting subtle changes in the waveforms, you are better off using either partial stack before prestack migration or poststack migration than partial stack after prestack migration.

The partial stack of the real data (offset = $1000 \pm 170\text{m}$) that was recorded on this line is shown in Figure 15. It is typical Alberta plains data, with predominantly flat-lying events and steeply dipping coherent noise from near-surface scatterers.

Figures 16, 17 and 18 show the results on the real data of poststack migration, partial stack before prestack migration and partial stack after prestack migration, respectively.




Fig. 13. Close-up of flat-lying reflector with partial stack before prestack migration.

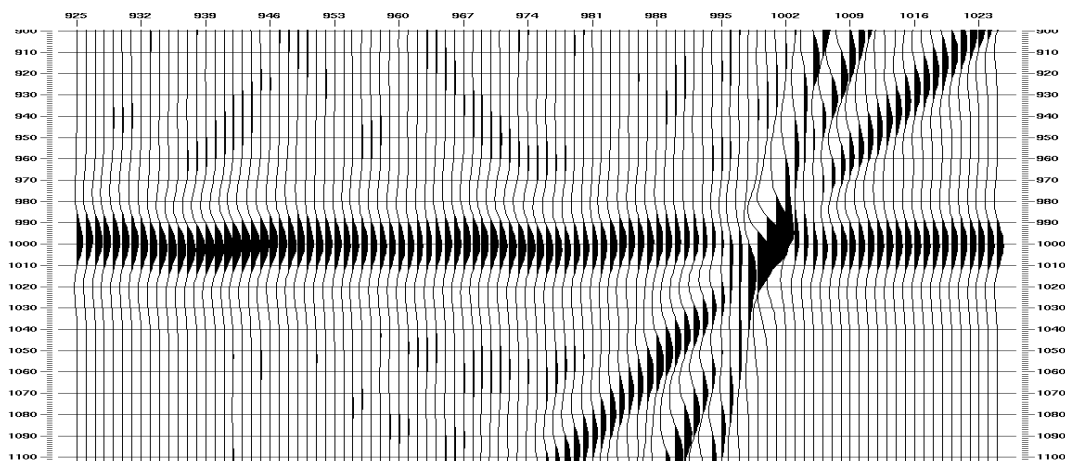


Fig. 14. Close-up of flat-lying reflector with partial stack after prestack migration.




Fig. 15. Partial stack of $1000 \pm 170\text{m}$ real data




Fig. 16. Poststack migration of 1000m pseudo-common-offset section.




Fig. 17. Prestack migration of real data (partial stack before migration).




Fig. 18. Prestack migration of real data (partial stack after migration). Notice the waveform perturbations due to migration artifacts.

The partial stack before poststack migration and the partial stack after prestack migration images (Figs. 17 and 18) are quite similar at first glance. However, there are several places where the waveforms on the partial stack after prestack migration image (Fig. 18) differ substantially from the waveforms in the other two images. The boxes identify one place where a substantial difference in waveforms exists between Figure 18 and Figures 16 and 17. An unwary interpreter who is given just Figures 16 and 18 to compare could easily misinterpret the waveform variation on the strong event in Fig. 18 as a true geologic feature that prestack migration reveals, but poststack migration fails to reveal. In fact it is just a migration artifact.

EXAMPLE WITH SMALL SOURCE INTERVAL

The example in the previous section showed that it is possible to overcome the sampling restrictions imposed by the coarse sampling within true CO sections, which could be considered to be an example of “beyond Nyquist” imaging. However, it does not show clearly that it is possible to produce an image that is resolved down to the CMP interval, even when the smallest sample interval of the continuous wavefield is never less than twice the CMP interval.

In order to illustrate that point clearly I have used the same model as above, but I have now simulated a dataset that is acquired with the shot interval equal to the receiver interval. This was done by taking the same set of pseudo-CO traces as in the previous example (offsets = $1000 \pm 170\text{m}$), where the shot move-up is five stations, and making the offsets alternate between 983m and 1017m for positive offsets, and -983 and -1017 for negative offsets. The result is pseudo-CO sections with absolute offsets of $1000 \pm 17\text{m}$. True CO sections are still aliased with this type of sampling. For example, the 45° reflector in the true CO section in Figure 19 clearly is aliased at its dominant frequency (the center frequency of the Ricker wavelet has been increased to 25Hz to ensure that this is true).

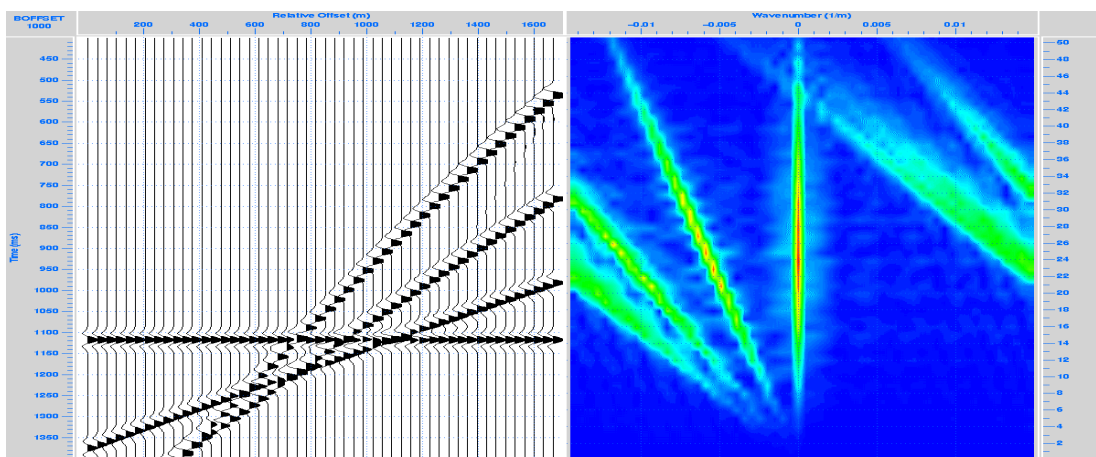


Fig. 19. True common-offset section (sampled at twice the CMP interval) and its F-K spectrum showing aliasing of the steeper events.

If we migrate this true CO section with partial stacking before Kirchhoff prestack migration, and choose the degree of operator antialiasing by assuming that the input traces are correctly sampled at the CMP interval, then the migration tries to image the high, aliased frequencies in the input traces, and the result is that the aliasing artifacts in Figure 20 are generated.

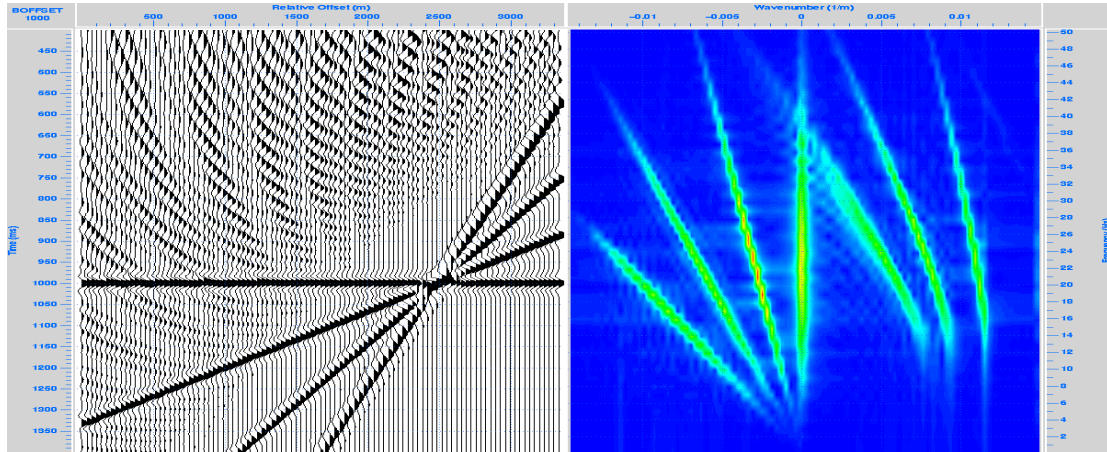


Fig. 20. Prestack migration (with partial stack beforehand) with operator antialiasing that assumes correct sampling at the CMP interval. The migration was output at the CMP interval.

The aliasing artifacts may be largely removed by modifying the operator antialiasing so that it assumes the input data are correctly sampled at the receiver interval. The result is that the high frequencies are selectively filtered out of the steeply dipping events, and the lower frequencies are correctly imaged, as in Fig. 21. Aliasing artifacts that remain in Fig. 21 are due to imperfections in the operator antialiasing algorithm.

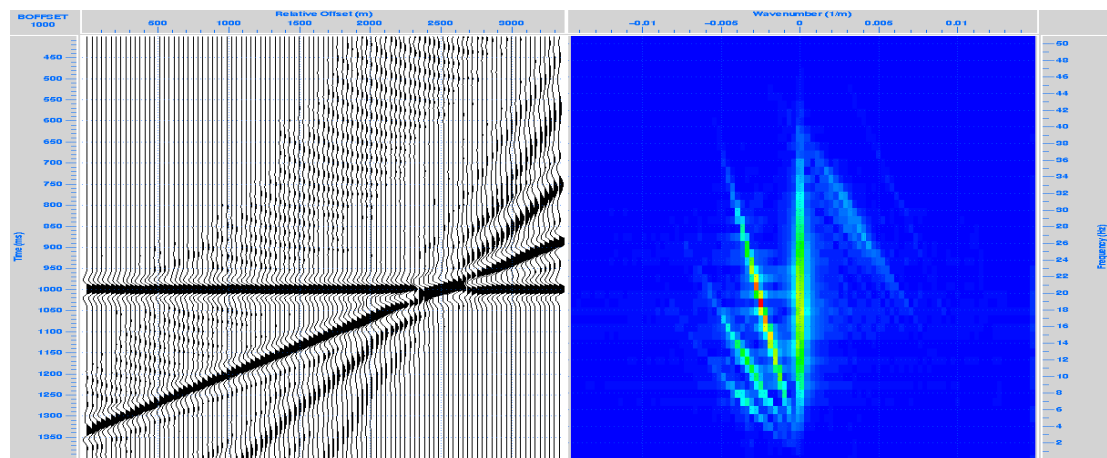


Fig. 21. Prestack migration (with partial stack beforehand) with operator antialiasing that assumes correct sampling at the receiver interval.

The high frequencies on the steeply dipping events can be imaged very well by partial stacking of two true CO sections together, as in Figure 22. It is clear from a

comparison Figures 19 and 22 that partial stacking has largely succeeded in dealiasing the main frequencies of the steepest event.

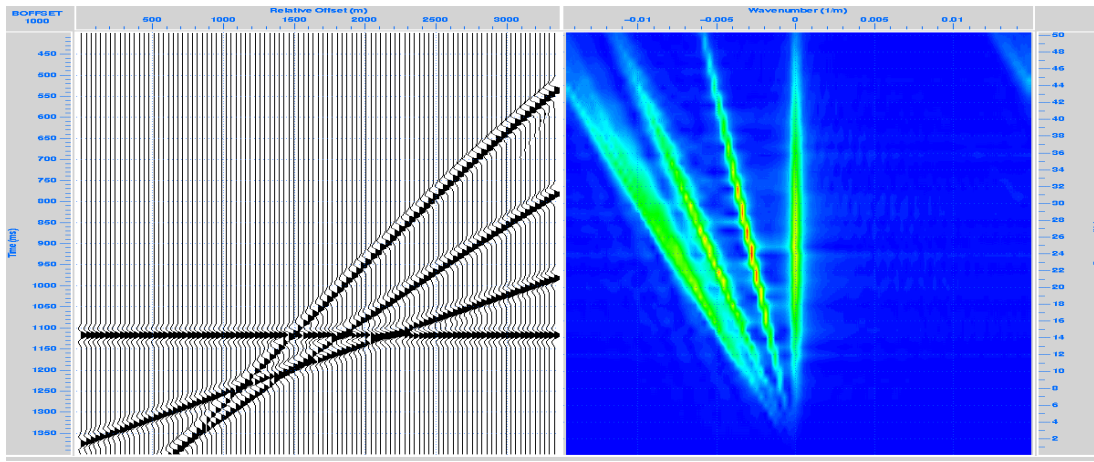


Fig. 22. Partial stack of two true common-offset sections. Notice that the steeper event are now mostly unaliased.

Partial stacking before migration in Fig. 23 and partial stacking after migration in Fig. 24 now both succeed in imaging all events very well. Operator antialiasing that assumes correct sampling at the CMP interval was used for both migrations.

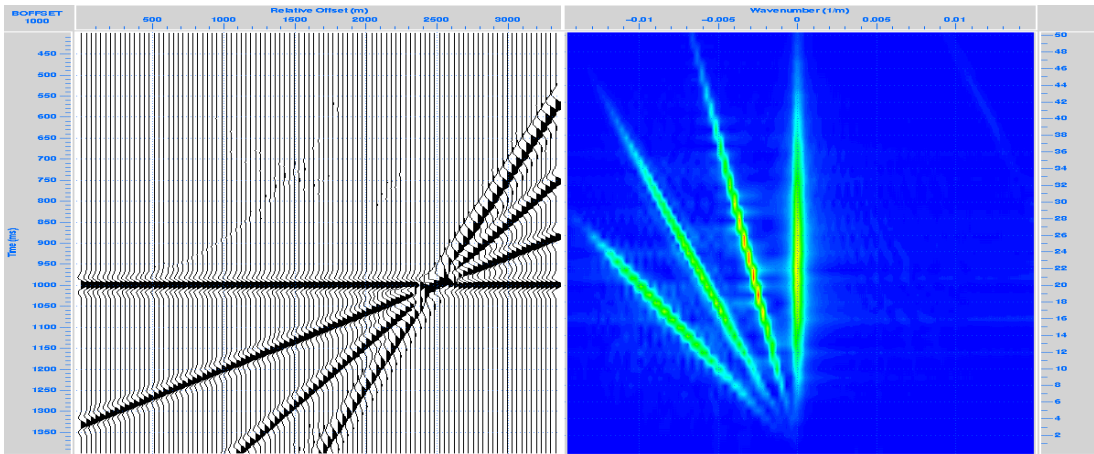


Fig. 23. Partial stack before migration. All dips are imaged with reliable waveforms.

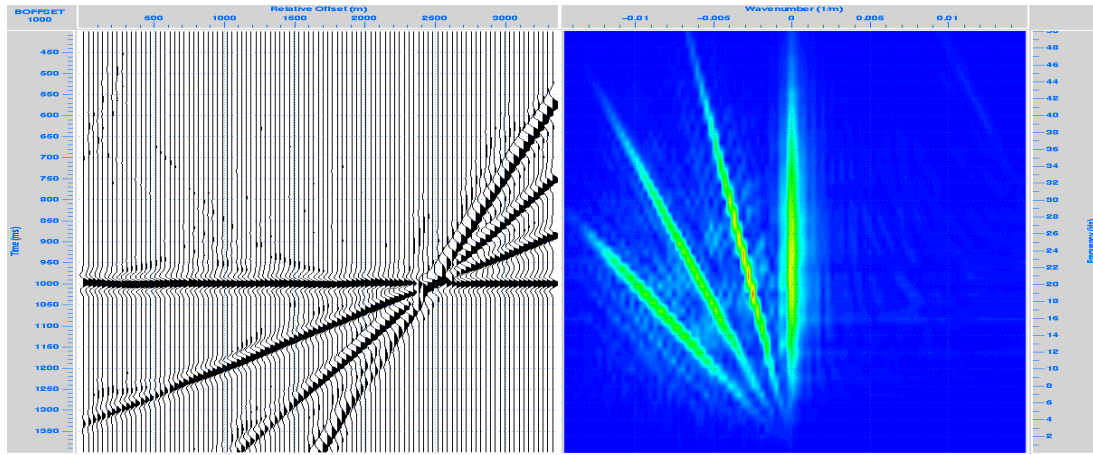


Fig. 20. Partial stack after migration. Waveforms are less reliable than in Fig.19.

Notice that, in terms of fidelity of the waveforms, the result of partial stacking before migration is now clearly superior to stacking after migration for all events, not just the shallow dips. The reason for this is that stacking before migration has the advantage of having input data with regularized geometry. In other words, the dealiasing inverse problem has been solved very well by partial stacking. Stacking after migration does not properly account for irregularities in the fold of the data, and therefore leaves the acquisition footprint that perturbs the waveforms in Fig. 20.

OBSERVATIONS

In summary, the following observations can be made from the examples:

With severe aliasing (shot interval = 5 receiver intervals):

- 1) Stacking after migration yields an image that has correctly-positioned reflections, but all events have unreliable waveforms.
- 2) Stacking before migration yields an image with very poor imaging of steeply dipping events, but excellent imaging of events with small dip (better than stacking after migration).

With mild aliasing (shot interval = receiver interval):

- 3) Stacking before migration and stacking after migration both yield good images, but the waveforms on all reflections are much better when stacking before migration.

REMARKS AND CONCLUSIONS

Based on the examples presented in this paper, the following conclusions about the impact of acquisition geometry on the fidelity of the imaged signal can be made. It is important to realize that final decisions about survey design should be based on the

ability to image signal reliably *and* reject noise. The impact of noise has not been considered in these tests.

Images that have correctly positioned reflections from flat and steeply-dipping events can be obtained from data acquired with large shot intervals by stacking enough aliased CO sections after migration. However, the waveforms on all events will be unreliable due to acquisition footprint.

Excellent prestack migrated images of low-dipping events (less than 30 degrees) can be obtained, with reliable waveforms, from the migration of data acquired with large shot intervals by stacking enough true CO sections into pseudo-CO sections before migration. For data with dips this small, poststack migration also provides good results.

If both correct positioning *and* reliable waveforms on steeply dipping reflections are to be obtained, then a small shot interval is required. The receiver interval can still be twice the Nyquist sample rate.

In areas with low geologic dips (e.g. Western Canadian Sedimentary Basin) signal can be very well imaged with large shot intervals. For 3-D surveys in these areas, large shot line and receiver line intervals should be adequate for imaging signal. The present standard method of doing 3-D prestack migration (stacking after migration) yields unreliable waveforms. Partial stacking of 3-D data before migration into pseudo-common-offset-vector volumes (Cary, 1999), which are the 2-D analog of pseudo-CO sections, would yield more reliable waveforms.

The largest wavenumber in the recorded wavefield that can be reliably imaged is not determined by the smallest spatial interval at which the continuous wavefield is sampled (normally the receiver interval), as basic sampling theory dictates. Although the subject requires further study, it is clear that an upper limit on the largest recoverable wavenumber is determined in an imperfect sense by the average spatial sample interval of a group of CO sections taken together. Of course, we have to keep in mind that the correct Nyquist sample rate for the continuous wavefield is not determined by sampling. It is determined by the highest wavenumber of signal in the continuous wavefield (and we may also choose to determine it from the highest wavenumber of the noise). We can only hope that the smallest sample rate that we can reliably image is at least as small as the true Nyquist sample rate.

In most cases (P-P data, average dip = 0), the CMP interval is probably a good practical estimate of the effective Nyquist sample rate of the image, which is what we have always assumed is true during survey design. For cases where the average dip is significantly different from zero (e.g. flat-lying sediments abutting against the flank of a salt dome), the apparent separation between subsurface reflection points of traces within a pseudo-CO section will be less than the CMP interval, so it seems that the imaging procedure could be modified to not only reliably image temporal frequencies of steeply dipping events beyond normal Nyquist criteria, as Biondi (1999) has demonstrated, but also to reliably image spatial wavelengths that are beyond the usual Nyquist limit of twice the CMP interval.

The situation with P-S converted-wave data is interesting because the analysis of the fold within common-conversion-point bins (Eaton and Lawton, 1992) indicates that regular subsurface fold on flat reflectors is obtained with a “natural” bin size that is larger than the CMP bin size. This might lead one to conclude that the highest recoverable wavenumber from migration of P-S data is lower than for P-P data, but this would be wrong because it ignores the bandlimited character of the data. If we look at the problem in terms of the average subsurface coverage of pseudo-common-offset gathers, we see that the resolution limit of P-S data is basically the same as for P-P data. The separation between traces in true P-S CO sections is equal to the source interval, just as for P-P data, and also the separation between subsurface reflection points for a true P-S CO section is half the source interval. The only difference with P-P data is that the reflection points are shifted towards the receiver. The average separation between subsurface reflection points for a fully-populated pseudo-CO P-S gather, which has been constructed the same way as a pseudo-CO P-P gather, is also equal to the CMP interval, just like P-P data, although there will be some variation about that average value along the flat reflector. So the highest recoverable wavenumber for P-S data is no different than for P-P data. The need for a larger “natural” bin size for P-S data evaporates once the bandlimited character of the wavefield is taken into account.

The way that we normally decide on the bin size for our seismic surveys, which assumes that we can resolve images down to the CMP interval, is basically correct. From the point of view of imaging signal, Vermeer’s prescription for survey design (symmetric sampling) therefore seems to be overly restrictive. However, it is quite likely that symmetric sampling may be optimum in some sense for noise rejection. The subject of the impact of noise requires further study.

The results of this study have implications for migration operator antialiasing and for migration methods that do not work with CO sections (e.g. shot migration).

Migration operators need to be antialiased in order to prevent artifacts from being generated in the image when the slope of the operator exceeds the “Nyquist slope” (the ratio of temporal to spatial sample intervals). Operator antialiasing criteria can be derived by converting migration integrals to summations with the assumption that the integrand (the data) is bandlimited and the sampling is at or above the Nyquist rates in time and space (Cary, 1998). The smallest true spatial sample rate of the data in this study is the receiver interval, so you would initially think that the operator antialiasing should be determined by the receiver interval. However, the effective sample rate after partial stacking is the CMP interval. Therefore, the correct spatial sample rate to use in determining the degree of operator antialiasing is, in most cases, the CMP interval. Using values greater than the CMP interval will filter out the higher frequencies of the dipping events unnecessarily. In some cases, the use of a large amount of antialiasing may be considered to be beneficial for noise attenuation purposes since operator antialiasing is essentially a frequency-dependent dip filter.

Finally, we are left with the somewhat disturbing knowledge that, if we can image events in the final stack that require a spatial sample rate that is smaller than the receiver interval, then *all* the prestack datasets might suffer from data aliasing. We have investigated two simple ways of dealiasing CO sections, but what about

processes like prestack migration or DMO of shot records, or 3-D prestack migration or DMO of cross-spreads? The usual argument is that each shot record can be imaged separately (Berkhout, 1984), or that each cross-spread can be DMO'd separately (Vermeer, 1998) and the final image can be "tiled" together from each small, but complete, sub-image. This is the argument from the "minimal dataset" approach (Padhi and Holey, 1997). We now see that these arguments assume that the spatial Nyquist sample rate is given by the receiver interval (as well as the shot interval for cross-spreads).

Let us assume that a 2-D seismic survey has been designed in the usual fashion so that the spatial Nyquist sample rate is the CMP interval, not the receiver interval, so the shot gathers are spatially aliased. In that case, we can still reliably image data up to the CMP sample rate by shot-record migration. We simply migrate each aliased shot gather, and stack together all the aliased images so that the aliasing noise partially stacks out. This corresponds to partial stacking after migration of CO sections. Notice that, regardless of which prestack migration method is used, the spacing of the output grid (the CMP interval) has to be smaller than the input grid (the receiver interval) in order to prevent aliasing of the output data. In addition, operator antialiasing has to be controlled by the output grid spacing, not the input grid spacing, in order to prevent the loss of high frequencies on steeply dipping events.

For 2-D DMO of shot gathers, or 3-D DMO of cross-spreads, the same type of argument can be applied. For example, each aliased cross-spread is sampled at the receiver spacing in the inline direction and the shot spacing in the crossline direction, but the aliased DMO output should be sampled at the inline and crossline CMP intervals. The period of the 2-D spatial sinc function that is used for interpolation of the output traces (Beasley and Mobley, 1997) should be determined by the inline and crossline CMP intervals. A single cross-spread does not, by itself, image the high frequencies of the final stack. The high frequencies are imaged when all the cross-spreads are added together. Notice that this approach is contrary to the "tiling" approach of Vermeer (1998).

It is important to realize that we never really image data beyond the true aliasing limits imposed by the "effective" spatial sample rate, so we are not violating any sampling theory when we say that we can image our data "beyond Nyquist". All of our results can be explained by recognizing that, although individual subsets of the data are spatially aliased, the data can be approximately dealiased when the related subsets are considered together. This point has been realized for a long time (Papoulis, 1977). Previous geophysical examples of "beyond Nyquist" imaging that are based on this basic idea have been provided by Ronen (1987), Wisecup (1998) and Cary (1997). It is equally important to realize that the two simple dealiasing procedures investigated here, as well as the more elaborate inverse theory approaches, do not provide perfect solutions to the dealiasing problem, even in the noise-free case.

REFERENCES

- Beasley, C.J., and Mobley, E., 1997, Spatial dealiasing of 3-D DMO: Extended Abstracts, 67th Annual Internat. Mtg., Soc. Expl. Geophys., 1119-1122.
- Berkhout, A. J., 1984, Seismic Migration: imaging of acoustic energy by wavefield extrapolation B. Practical Aspects, Elsevier, Amsterdam, 274pp.

- Biondi, B., 1999, Kirchhoff imaging beyond aliasing: submitted to Geophysics.
- Bolondi, G., Loinger, E., and Rocca, F., 1984, Offset continuation in theory and practice: *Geophys. Prosp.*, **32**, 1045-1073.
- Cary, P.W., 1997, 3-D stacking of irregularly sampled data by wavefield reconstruction: *Extended Abstracts, 67th Ann. Internat. Mtg., Soc. Explr. Geophys.*, 1104-1107.
- Cary, P.W., 1998, The simplest discrete Radon transform: *68th Annual Internat. Mtg., Soc. Expl. Geophys.*, 1999-2002.
- Cary, P.W., 1999, Common-offset-vector gathers: 11th Annual CREWES Report.
- Chemingui, N., and Biondi, B., 1999, Data regularization by inversion to common offset (ICO): *Extended Abstracts, 69th Annual Internat. Mtg., Soc. Expl. Geophys.*
- Deregowski, S.M., and Brown, S.M., 1983, A theory of acoustic diffractors applied to 2-D models: *Geophys. Prosp.*, **31**, 293-333.
- Eaton, D.W.S., and D.C.Lawton, 1992, P-SV stacking charts and binning periodicity: *Geophysics*, **57**, 745-748.
- Jakubowicz, H., 1994, Wavefield reconstruction: 56th Mtg. and Tech. Exhibition., EAGE, H55.
- Liner, C.L., and Underwood, W.D., 1999, 3-D seismic survey design for linear $v(z)$ media: *Geophysics*, **64**, 486-493.
- Nemeth, T., Wu, C. and Schuster, G.T., 1999, Least-squares migration of incomplete reflection data: *Geophysics*, **64**, 208-221.
- Padhi, T. and Holley, T.K., 1997, Wide azimuths—why not?: *The Leading Edge*, **16**, 175-177.
- Papoulis, A., 1977, Generalized sampling expansion: *Inst. Electr. Electron. Eng., Trans. Circuit Syst.*, **CAS 24**, 652-654.
- Rocca, F., and Ronen J., 1984, Improving resolution by dip moveout: *Extended Abstracts, 54th Annual Internat. Mtg. Soc. Expl. Geophys.*, 611-614.
- Ronen, J., 1987, Wave-equation trace interpolation: *Geophysics*, **52**, 973-984.
- Ronen, S., 1994, Handling irregular geometry: equalized DMO and beyond: *Extended Abstracts, 64th Annual Internat. Mtg. Soc. Expl. Geophys.*, 1545-1548.
- Ronen, S., Nichols, D., Bale, R. and Ferber R., 1995, Dealiasing DMO: good-pass, bad-pass and unconstrained: *Extended Abstracts, 65th Annual Internat. Mtg. Soc. Expl. Geophys.*, 743-746.
- Vermeer, G.J.O., 1990, *Seismic Wavefield Sampling: a wavenumber approach to acquisition fundamentals*, Society of Exploration Geophysicists, Tulsa, 120pp.
- Vermeer, G.J.O., 1998, 3-D symmetric sampling: *Geophysics*, **63**, 1629-1647.
- Wisecup, R.D., 1998, Unambiguous recovery above the Nyquist using random sample interval imaging: *Geophysics*, **63**, 763-771.