Stuck between a rock and a reflection: A tutorial on low-frequency models for seismic inversion

Mark Sams[1] and David Carter[2]

Abstract

Predicting the low-frequency component to be used for seismic inversion to absolute elastic rock properties is often problematic. The most common technique is to interpolate well data within a structural framework. This workflow is very often not appropriate because it is too dependent on the number and distribution of wells and the interpolation algorithm chosen. The inclusion of seismic velocity information can reduce prediction error, but it more often introduces additional uncertainties because seismic velocities are often unreliable and require conditioning, calibration to wells, and conversion to S-velocity and density. Alternative techniques exist that rely on the information from within the seismic bandwidth to predict the variations below the seismic bandwidth; for example, using an interpretation of relative properties to update the low-frequency model. Such methods can provide improved predictions, especially when constrained by a conceptual geologic model and known rock-physics relationships, but they clearly have limitations. On the other hand, interpretation of relative elastic properties can be equally challenging and therefore interpreters may find themselves stuck — unsure how to interpret relative properties and seemingly unable to construct a useful low-frequency model. There is no immediate solution to this dilemma; however, it is clear that low-frequency models should not be a fixed input to seismic inversion, but low-frequency model building should be considered as a means to interpret relative elastic properties from inversion.

Introduction

Inversion of seismic reflection data to elastic impedances can add significant value to seismic reservoir characterization. The benefits of inversion are well-known (e.g., Latimer et al., 2000): increased resolution, conversion from an interface to a layer property, conversion to physical rock properties (impedances), removal of the wavelet, and reduced tuning, all of which lead to an improved interpretability of the seismic data in a qualitative and a quantitative sense. Many of these advantages are dependent on the integration of an accurate low-frequency model with band-limited seismic data to produce the absolute impedances. Most deterministic inversion algorithms operate in the absolute elastic domain and require an estimate of the absolute elastic properties throughout the volume prior to inversion. These volumes of absolute properties are used in various ways depending on the inversion algorithm chosen. They can act as starting models, in which case they may contain a broad range of frequencies, or they can act as low-frequency models that supply data and constraints below the bandwidth of the seismic data. Absolute values are also often used to supply additional inversion constraints based on interrelationships between the different properties being determined (e.g., acoustic impedance to shear impedance and acoustic impedance to density, so called rock-physics constraints), and they can be used as hard constraints, for example, ensuring that Poisson's ratio lies between zero and 0.5. In practice, the results of inversion to absolute properties are supplemented with relative properties, which are commonly the result of applying a band-pass filter to the absolute properties. The seismic reflection data, relative impedances, and absolute impedances should all be considered during interpretation to ensure that conclusions are not being incorrectly influenced by the low-frequency models, although there are many examples in the literature in which this does not seem to be the case. The concern about absolute impedances is due to the difficulty of providing reasonable low-frequency models and the potential that they can negatively bias the absolute estimates and produce incorrect interpretations. Ball et al. (2015) show that with certain assumptions, no low-frequency model is required for inversion to relative impedances when they are appropriately defined (Ball et al., 2014). As such, there is an argument for excluding low-frequency models and relying only on the relative properties for interpretation. On the other hand, the problem with relative impedances is

[1]Ikon Science, Kuala Lumpur, Malaysia. E-mail: msams@ikonscience.com.
[2]KrisEnergy (Gulf of Thailand) Ltd., Bangkok, Thailand. E-mail: david.carter@krisenergy.com.

Manuscript received by the Editor 18 September 2016; revised manuscript received 4 December 2016; published online 21 March 2017. This paper appears in *Interpretation*, Vol. 5, No. 2 (May 2017); p. B17–B27, 9 FIGS.
http://dx.doi.org/10.1190/INT-2016-0150.1.

that with them being relative, their values are dependent on their context. Take an example drawn from a real study. Two wells drilled within 20 km of each other encountered sandstones (sands) at similar depths with very similar absolute elastic properties. The bounding shales, however, are quite different with one well penetrating normally compacted shales and the other well penetrating overpressured shales. In the first case, the sands are softer (they have lower acoustic impedance) than the bounding shales, and in the latter case, the sands are significantly harder. The relative impedances of the sands show considerable variation, whereas the absolute properties are almost constant (Figure 1). Quantitative interpretation, such as porosity estimation from absolute properties in this case would be simpler; however, this begs the question of how to build suitable low-frequency models. In this paper, we review a few techniques for building low-frequency models. There is particular focus on the method of updating a simple initial low-frequency model based on a first-pass inversion through the use of a case study. None of the techniques presented are new: The objective is to highlight the difficulties of low-frequency modeling to expose the limitations and benefits.

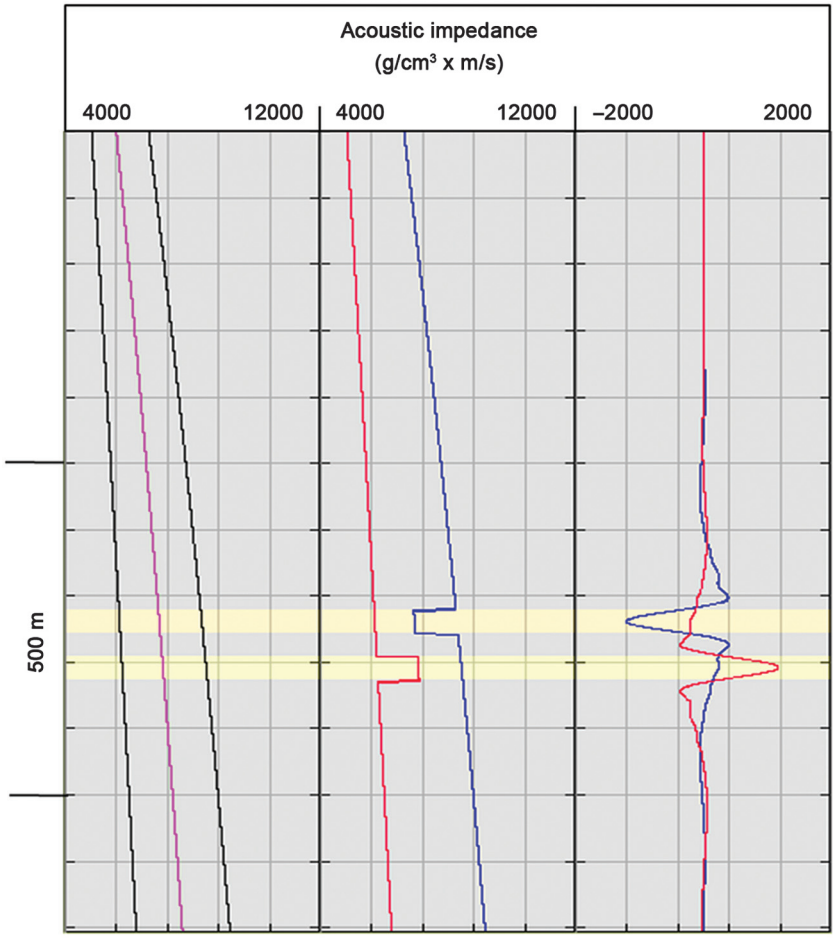

Figure 1. Two wells drilled within 20 km of each other penetrate shales with very different elastic properties but sands with very similar elastic properties. The panel on the left shows the acoustic impedance trend for the sands (pink), the shale normal compaction trend on the right, and the overpressured shale trend on the left. The central panel shows the profiles when considering sands at about the same depth at these two wells. The absolute impedances of the sands are very similar. The third panel shows the relative acoustic impedance profiles derived by applying a band-pass filter. The relative properties of the sands are extremely different.

Building low-frequency models

Low-frequency models need to be constructed for all properties, for which inversion is to be carried out. The properties depend on the seismic data that are being inverted, but often they are some form of the elastic property or combination of elastic properties. In the case of seismic inversion of a single stack, the property inverted for is often acoustic or elastic impedance depending on the nature of the stack. Currently, it is common to invert multiple partial angle stacks simultaneously for multiple properties, and therefore multiple low-frequency models are required (e.g., acoustic impedance, shear impedance, and density).

There are two main sources of low-frequency data that can be used to help construct low-frequency models: well-log data and seismic velocity data. Well-log measurements of elastic properties contain a full range of frequencies from zero to well above the highest seismic frequencies, and may have measurements for all required elastic properties (P-sonic, S-sonic, and density). Seismic velocities contain a very narrow range of reliable frequencies from zero up to approximately 2–3 Hz (a point that will be discussed later), and they normally provide compressional velocity data only. There are potentially other sources of data, such as electromagnetic data, which contain low frequencies that might be considered as an aid to building a low-frequency model (e.g., Mukerji et al., 2009), but these are not discussed here.

The most common method to construct a low-frequency model is a simple interpolation of well data within a structural and stratigraphic framework (Pendrel and van Riel, 2000). Interpreted horizons and faults define a structure, and stratigraphic relationships are assigned between these horizons, such as conformal, or parallel to the top or base. The well data are interpolated along the stratigraphic microlayers according to this model. Current inversion approaches contain many interpolation methods that can be used, such as the inverse distance, triangulation, or kriging (Doyen, 2007; Pedersen-Tatalovic et al., 2008). An example of a well interpolation using kriging is shown in Figure 2. The map in this figure shows a time slice

through an acoustic impedance model constructed by interpolation of data from three wells using ordinary kriging with an isotropic, exponential variogram with a range of 1500 m. The appearance of bull's eye effects around the wells indicates that kriging in this simple form is not an appropriate interpolation algorithm. In this case, the rock types are fluvial sands and shales, with sand having a relatively low impedance compared with the background shales. Clearly, data from where a well penetrates a sand body should not be extrapolated beyond the sand-body margin. A variogram might be appropriate to distribute continuous properties within a lithology if the underlying assumption of a slowly varying (or constant) mean is geologically valid, but it is not suitable to propagate properties across relatively abrupt changes in lithology that are not captured by the variogram. Other simple interpolation algorithms perform equally as poorly, although it should be noted that those that extrapolate toward a local or global mean can give vastly different results from those that project a trend from poorly sampled well data. The main weakness of interpolation is that it is a mathematical process with no geologic insight: The results are strongly dependent on the distribution of wells and not on local geologic variations. It is also common practice to interpolate well data cokriged to seismic velocities that are used as a secondary data set. Ideally, the seismic velocities incorporate information from local geologic variations. Even if this is the case, the attempt to predict the frequencies above those in the seismic velocities but below those in the seismic reflection data is normally forlorn: The results are still dependent on the distribution of the wells with the intermediate frequencies forming bull's eyes around the wells. If these are the commonest methods for building low-frequency models, it can be no surprise that the interpretation of relative properties is often preferred. With these interpolation methods, the predictions based on absolute properties away from well locations will be poor and may indeed impact negatively on interpretation.

Attempts have been made to overcome the limitations of well distribution by the inclusion of pseudowells that have been created with geologic interpretation (Glenn et al., 2005). In that study, a simple well interpolation was made to generate an initial model at high detail. Synthetic seismic stacks were generated and compared with the recorded seismic data. Pseudowells were positioned where the differences were large. The lithology and fluid distribution in the pseudowells were interpreted based on the relative impedances from the initial inversion and a conceptual geologic model. The lithologies were populated with properties consistent with local trends, rock-physics models, and the interpreted fluids. Adjustments were made until a match was found between synthetic and seismic data. Pseudowells were added iteratively to the database, and the 3D model was rebuilt. The process was repeated until sufficient convergence. One of the significant disadvantages of this method is the considerable amount of effort required. A very large number of wells are needed to match the geologic variation because the construction of the final low-frequency model still relies on interpolation between wells. It might be possible to reduce the number of pseudowells required by using image guided interpolation (Naeini and Hale, 2015), so that, for example, only one well that penetrates a channel sand might be sufficient rather than many wells positioned along the channel sand's length.

In a similar, unpublished study, several geologists were employed to create and modify the pseudowells. After a first attempt, several pseudowells were incorporated into the model and severe bull's eye effects were seen in the low-frequency model. The reason was that some geologists had interpreted a particular seismic signature as two separate thin sands and others had interpreted it as single thick sand. This suggests a link between problems in building low-frequency models and problems interpreting relative properties from inversion, as will be discussed later.

A more efficient and potentially more accurate process attempts to use volume based interpretation rather than interpretation at discrete pseudowell locations. If the presence of geologic bodies can be captured based on relative impedances from an inversion, the information can be used to build a more realistic low-frequency model that can be used to update the inversion. Capturing bodies on the updated absolute impedance can enhance the interpretation, and so the process can be iterated. A simple example for a channel sand with constant impedance in a constant impedance background highlights some of the benefits and limitations. Figure 3 shows the results of applying a sequence of inversions and low-frequency model updates to predict the low-impedance channel starting with the assumption of

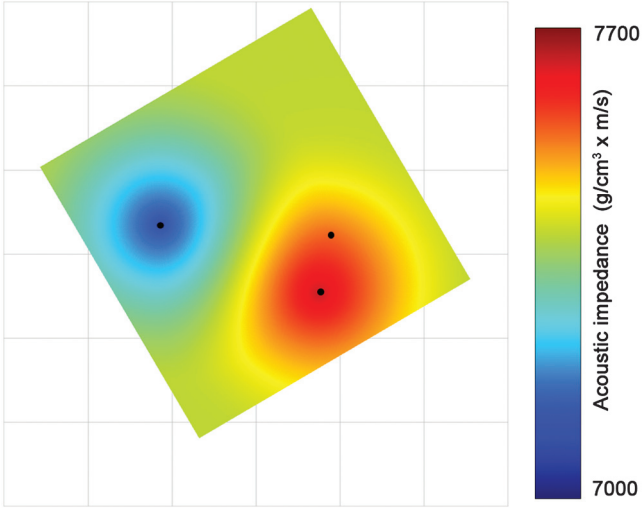

Figure 2. A slice through a model of acoustic impedance constructed by interpolating three wells within a structural framework and filtering with a high-cut filter (6–12 Hz). Grid spacing is 1000 m. Note that the range of values varies by 10% even though the wells (black dots) are only approximately 2000 m apart.

no channel but with the correct background impedance. The geometry has been chosen to include variations in sand thickness and a horizontal surface. At each stage, the channel delineation is made based on the current inversion results and a selected rock-property cutoff. Where the channel is predicted after each inversion, the low-frequency model is updated. In this example, the correct channel properties are used where delineated by the appropriate cutoff, and the model is then low-pass filtered prior to the next inversion step. The choice of cutoff is critical. If it is too high, it will allow more data at the thin margins of the channel and overestimate the thickness and therefore the corresponding low-frequency contribution. If the threshold is too low, then the center of the channel will never get filled. That is, the optimal cutoff value varies with thickness. The cutoff value chosen in this example attempts to balance these concerns. However, it is clear that at the channel edges, the thinnest parts have no low-frequency update (because no channel facies is predicted), and there is a sudden jump to where the low frequency is being overestimated (too much channel facies predicted) due to the limitations of seismic resolution. Clearly, the final low-frequency model is wrong, although the final inversion only varies by at most 4% from an inversion result using the correct low-frequency model.

A semiautomated process based on the approach just described has some appeal because it does not require as heavy a human intervention compared with the previous method of Glenn et al. (2005). Although the channel example is the application to a single geologic body, it can be applied to multiple bodies if they are resolved. The process removes the requirement for interpolation. However, it assumes that the band-limited seismic data are able to drive the low frequencies. The information within the seismic bandwidth is being extrapolated back to a very low frequency. This is achieved by inserting sharp contrasts at the boundaries between the background and the captured bodies. The sharp contrasts contribute to frequencies below the seismic bandwidth, and they effectively assume a series of vertical step functions, in which the amplitude and phase of the low frequencies are defined. This may be satisfactory on some occasions, but it will not always be appropriate. When adjacent rocks have significantly different impedances, low frequencies are present due to the sharp contrast at the boundary, which can be detected by seismic and therefore has the possibility to influence the low-frequency model update. On the other hand, low frequencies are often present due to the accumulation of small contrasts (trends), which may not be detected by seismic and therefore may be missed by such a process.

There are several other points of interest that we take from this example. The first is that the use of a constant background (the first iteration) provides a result that requires further interpretation. Note that the top of the channel does not appear to be flat after the first inversion, whereas the channel top appears to become more so after each iteration and is perfectly flat in the actual model. We can see that the simple updating process provides us with an absolute image that is in many respects much simpler than the relative image and requires less interpretation (this is because the construction of the low-frequency component is an interpretation in itself). This

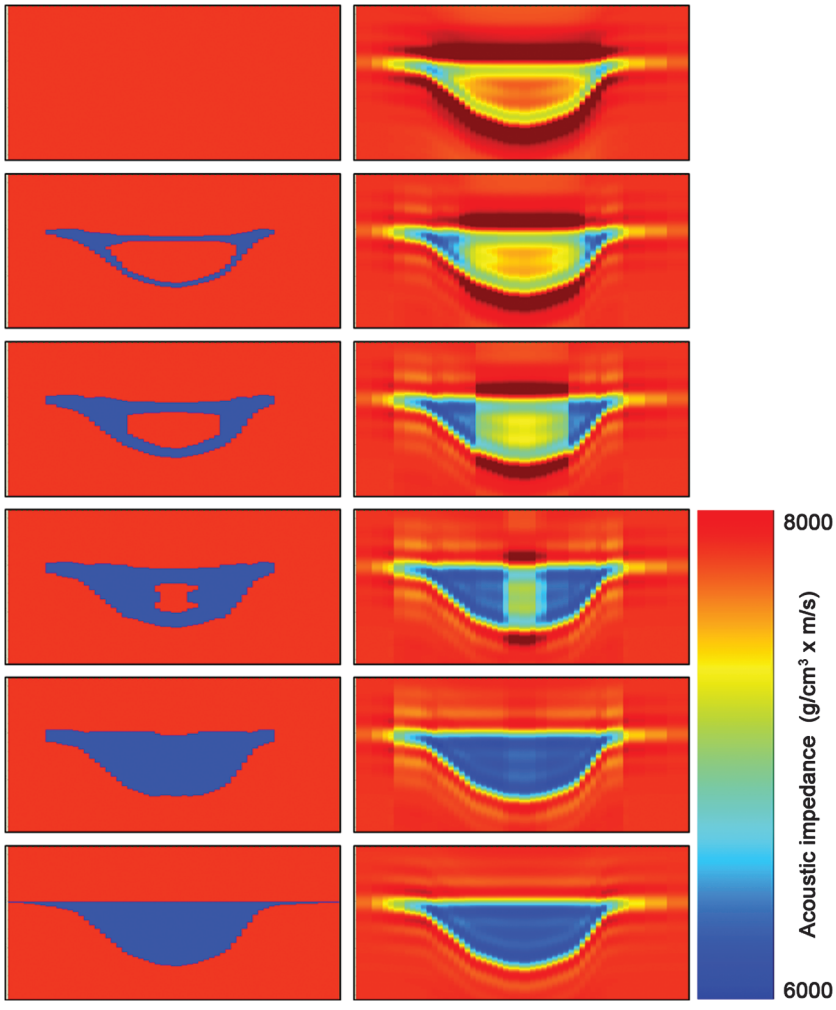

Figure 3. A sequence of pictures representing the progress by row of iteratively updated low-frequency modeling for inversion. The column on the left represents a model that is the result of updating the previous model by using that previous model for the low-frequency component of an inversion. The column on the right represents the absolute inversion using the low-frequency component of the model on its left. The final row shows the exact model (not a prediction) on the left and an inversion based on that correct low-frequency model.

suggests that despite it being wrong, the final low-frequency model is worth pursuing even if only to provide a better representation of the body. The final inversion is clearly wrong for the thin edges, and therefore will not provide accurate quantitative estimates of, say, net sand. The reason is because we either do not capture the body because the cutoff is too harsh, or we capture too much due to tuning effects. Note that although inversion improves the resolution (accurate quantification of the top and base of such a body) over seismic reflection data, it can only do this if the low-frequency model is correct. For the thin margins, at the first iteration, the low frequency is not correct and the bandwidth has not been increased. Because this causes the incorrect estimation of the thickness and therefore the low-frequency model, the resolution either does not improve or gets worse for the thin margins.

Another point of interest is that when the low-frequency model is wrong, for example, after the first inversion with a constant low-frequency model, there are strong side lobes. Such artifacts can be used to identify potential errors in low-frequency models. If the inversion is run when the low-impedance channel is populated with too high or too low impedances, there are either high-impedance sidelobes or low-impedance halos. Although this is easy to see for this simple model, it can be quite hard to judge when there are a lot of seismic reflections that interfere and when there is noise in the data, as is most often the case for real seismic data. It should also be noted that if there were low- and high-impedance events, such as a low-impedance channel in the presence of tight sands, that needed to be included in the low-frequency model, the updates must be done carefully. For example, an automatic body capturing interpretation might capture side lobes rather than actual bodies. That is, the process is not always simple and additional user interaction and interpretation might be required. Processes such as the application of a threshold combined with connectivity criteria and lateral smoothing can be used.

A real data example of this process is taken from the Gulf of Thailand. Here, the logged strata are Tertiary in age and typical of many parts of the Gulf of Thailand and portions of other Tertiary sedimentary basins in Southeast Asia. In this example, a large portion of the Miocene succession several thousand feet thick is comprised of an alternating sequence of red-bedded sandstone and shale, which forms the target zone for the seismic inversion. Comparison of well and seismic data shows that many of the sandstone bodies have a sinuous form with varying degrees of width and sinuosity. The wider sandstone bodies are interpreted to result from deposition and lateral accretion during the lateral migration of meandering fluvial channels. Their preserved geometries have highly contorted and complex forms with relatively narrow margins. Log data show that the thicker sandstones often have a blocky log character with varying degrees of a fining-upward profile in their uppermost portions. The sandstone bodies are contained within a shale-dominated section that is inferred to have varying degrees of marginal marine influence. The sequence of low-frequency model construction for one well is shown in Figure 4. Only the acoustic-impedance logs are shown in the figure for the sake of simplicity. The rocks of interest are low-impedance channel sands saturated with brine, oil, or gas. The interbedded shales follow a reasonably simple and consistent depth/time trend as determined by the well-log data from three wells. A model of the elastic properties of shales was built in the time domain for acoustic and shear impedance and density and used

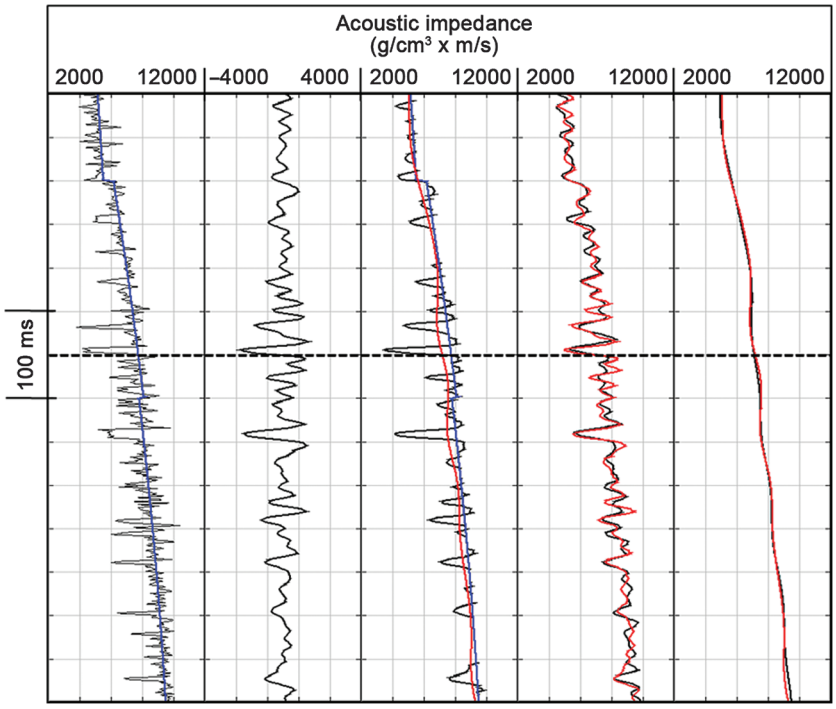

Figure 4. Panels showing a sequence through the development of a low-frequency model. Panel 1: measured acoustic impedance at well-log scale (black) and interpreted shale trend (blue). Panel 2: relative acoustic impedance from a simultaneous inversion using a shale trend for the low-frequency model. Panel 3: the shale trend (blue), the updated trend based on cutoffs and scaling applied to the relative acoustic impedance (black), and filtered back to the low-frequency content required (high-cut filtered at 6–12 Hz) in red. Panel 4: The updated trend combined with the relative acoustic impedance to provide absolute acoustic impedance (red) and the measured acoustic impedance from the well high cut filtered to the same upper limit as the inversion (60–70 Hz). Panel 5: The predicted low-frequency model (red) and the equivalent frequency content from the measured acoustic impedance from the well (black). The dashed line represents the level of the time slice shown in Figure 5.

as a low-frequency model in a simultaneous inversion. The seismic angle stack data (five stacks ranging from 5° to 40°) used in the inversion were low-cut filtered (6–12 Hz linear filter) to avoid any possibility for the seismic reflection data to directly contribute to the low frequencies in the final model. A standard simultaneous inversion was run using the shale trends as a low-frequency model to produce the absolute elastic properties. The relative elastic properties were calculated through the application of a band-pass filter and for acoustic impedance produced a high correlation (>0.8) with the well data within the equivalent bandwidth. The next step was to update the initial shale-based low-frequency model. First, the distribution of sand and shale was determined by applying cutoffs to the relative elastic property results. The shale trends were then modified where sands were interpreted to be. There are different possibilities for populating these interpreted sands with elastic properties. One possibility is to use depth trends derived from the well data. This would ideally include different trends for different fluids, which would require delineation of fluids from the first-pass inversion. Another approach, which was used for this study, is to set the relative elastic properties to zero within the interpreted shales prior to adding these modified relative properties to the shale trends. The sharp changes in the modified relative properties at the upper and lower boundaries of the interpreted sands will then contribute to the low frequencies. This approach avoids the need to interpret the fluid content of the sands prior to the next iteration of inversion. It was found to be necessary to apply a global scaling to the modified relative elastic properties prior to adding to the shale trend. The scaling was adjusted to optimize the match of the predicted low frequency of the model to the low-frequency component of the well-log data. This updated low-frequency model was then used for a new inversion.

The prediction of the low-frequency component shown in the final panel of Figure 4 is very good. A comparison is made between the measured well-log acoustic impedance and the predicted acoustic impedance with a band-pass filter with corner points of 1–2–6–12 Hz applied. The lowest frequencies were removed because those data come predominantly from the basic shale trend and not the update. The correlation of this band-pass filtered prediction with the equivalent band-pass filtered well-log data is 0.94 with an rms error of 60 $g/cm^3 \times m/s$. By comparison, if the low-frequency model is constructed based on the extrapolation of the first well drilled, then the equivalent correlation at this well is only 0.75 with an rms error of 288 $g/cm^3 \times m/s$. Even when the three wells are considered together, an interpolation does not honor the variations within the low-frequency model caused by the complex distribution of the channels in time and space. Figure 5 compares a time slice through the acoustic impedance low-frequency models derived from interpolation of the three wells (left) and through the updated shale trend model (right). The differences are obvious. It should be noted that simple interpolation provides accurate results at the wells, but the updating method is not perfect. The updating process will never be perfect for a variety of reasons that include seismic data quality and resolution limits, but the errors in this case are small compared with the errors from simple interpolation away from well control. The low-frequency models for shear impedance and density have been updated in a similar fashion (Figure 6). The good correlation for density suggests that it is being accurately predicted from the inversion and low-frequency model update. There is no implication, however, that density can be predicted directly from reflection data through simultaneous inversion. The simultaneous inversion algorithm used here includes a constraint between acoustic impedance and density, which is dominating the density solution for the settings chosen. Note that if the relative properties from the first-pass inversion are not correct, the errors will impact the low-frequency model update, and therefore the errors in the relative inversion and low-frequency model will be compounded in the absolute properties.

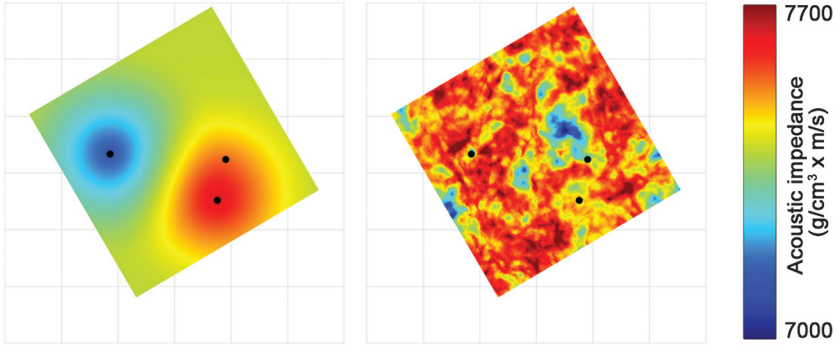

Figure 5. The left map shows a time slice through an acoustic impedance low-frequency model constructed by interpolation of three wells. The map on the right shows the same slice through a low-frequency model constructed through an update of a shale trend using the relative impedance results from a first-pass inversion. It is shown in Appendix A that it would require a well spacing of approximately 150 m to achieve an appropriate result through well interpolation.

A full automation of this process can be achieved (Kemper and Gunning, 2014), in which the simultaneous inversion solves for the distribution of facies and elastic properties. The low-frequency model is a product of the inversion, and it is constrained by the elastic property depth trends associated with each facies and the inverted distribution of the facies. This fully automated process has the potential to improve on the manually updated process in particular with regard to very thin beds because the limits of resolution can be overcome when conditions are favorable. When applied to the simple channel case shown earlier (Figure 3) with noise-free seismic data and

nonoverlapping facies trends, a perfect image of the channel can be obtained for all thicknesses. When applied to the real data of the above example, a similar lateral variation in the low-frequency component as shown in the right-hand plot of Figure 5 is obtained. It must be noted that any method (manual or automatic) that relies on the seismic data and depth trends will be limited by seismic quality, bandwidth, and angle range, as well as by the size and variation of the elastic property contrasts between layers/rocks/fluids of interest.

In some cases, the updates for a low-frequency model may vary laterally (or vertically) within a single body. Mesdag et al. (2010) provide an example of this, in which the strength of the relative impedance contrast at the top of a salt body was used to design a low-frequency update to the high-impedance salt layer. Use of a constant value to populate the salt layer had caused the predicted porosity in the underlying sand reservoir to be too high. When the reflectivity strength at the top of the layer was taken into account, an update was made with a laterally varying property, and an improved porosity model was produced. The assumption in the applied method being that the contrast at the top of the body could determine the properties of the salt layer relative to a constant overburden, and that the contrast at the base of the layer was a combination of the layer property and the porosity of the reservoir. That is, it was necessary to remove the effects of the varying salt layer to reveal the actual variations in the porosity. In such a case, the derivation of an accurate porosity from relative impedance would be extremely difficult because the relative impedance of the reservoir would be contaminated by the side lobe of the base salt layer reflection. Note that if the vertical variation within a body is slowly varying, then the exact nature of the variation might not be readily determined from the seismic data. To emphasize: The extrapolation of information from within the seismic bandwidth to lower frequencies will not always be successful.

Other examples of low-frequency model updating include an update to a porosity model, which is transformed to an elastic property update through a rock-physics model (Sams et al., 2012), and a 4D update to time lapse low-frequency models (Mesdag et al., 2007).

Using seismic velocities

Seismic velocities contain frequencies from zero hertz upward. Therefore, they represent a potential source of very low-frequency information. The upper limit of the

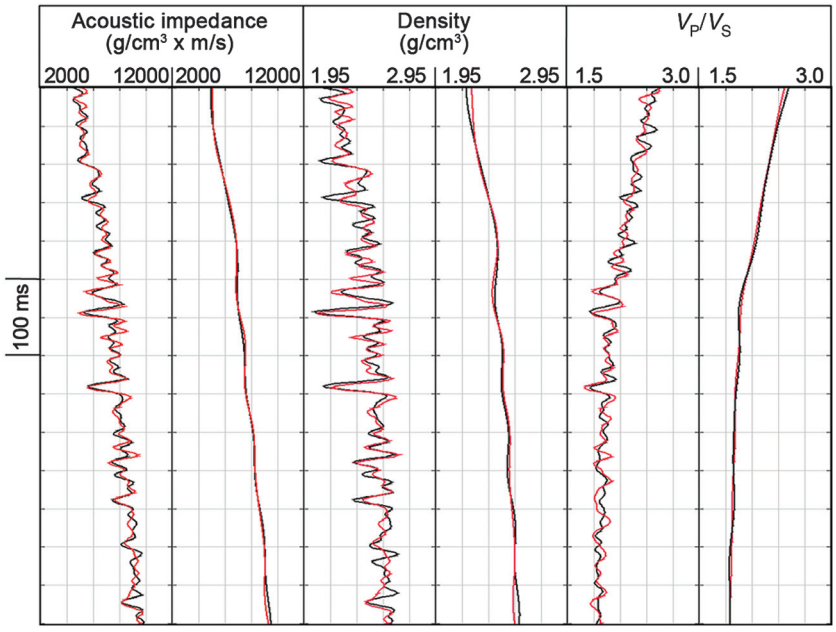

Figure 6. Panels showing the comparison of the predicted (red) and measured (black) acoustic impedance, density, and for frequency ranges of zero to 60–70 Hz (left) and zero to 6–12 Hz (right).

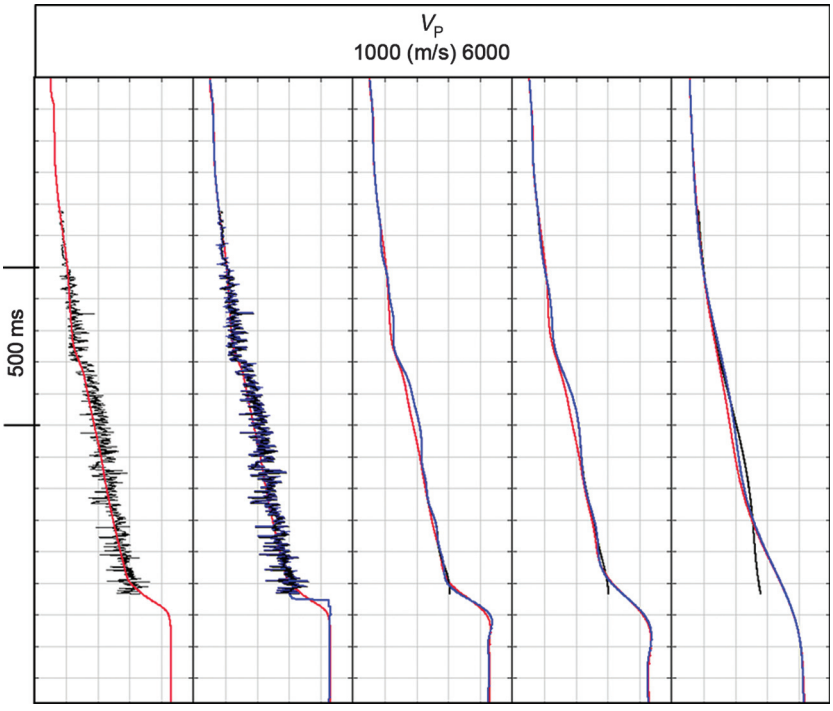

Figure 7. Comparisons of well-log and seismic velocities at different frequencies. The seismic velocities extracted along the well are shown in red. The well velocity from P-sonic is shown in black and extended to cover a larger vertical interval shown in blue. In the first two panels, the data are unfiltered. In the last three panels, the data are high-cut filtered at 5–10, 3–6, and 1–2 Hz.

frequency content depends on how the velocities were picked or generated. In some cases in which automatic high-density velocity analysis or full-waveform inversion has been performed, the frequency content can be quite high, perhaps 4 or 5 Hz. However, this should not be confused with the upper limit of useful frequency content. Extraction of the seismic velocity along well tracks and comparison with good quality sonic logs or checkshot velocities often indicates that the upper frequency limit of reliable information for use in constructing low-frequency models for inversion is no higher than 2 Hz and sometimes much less. Even then, there are usually systematic variations between well-log velocities and seismic velocities, for example, as discussed in Appendix B. It is common, therefore, to condition the seismic velocities through vertical and lateral filtering and to apply a calibration process to reconcile the seismic velocities to the well velocities. Several problems are encountered during this conditioning and calibration. First, how to choose the optimal frequency filter to apply; second, how to accurately calculate calibration factors when the well data are usually of limited extent, and therefore have high uncertainty in the very low-frequency (0–2 Hz) content; third, whether there is a good reason for calibration or whether the data are of insufficient quality; and fourth, how to interpolate calibration factors between wells.

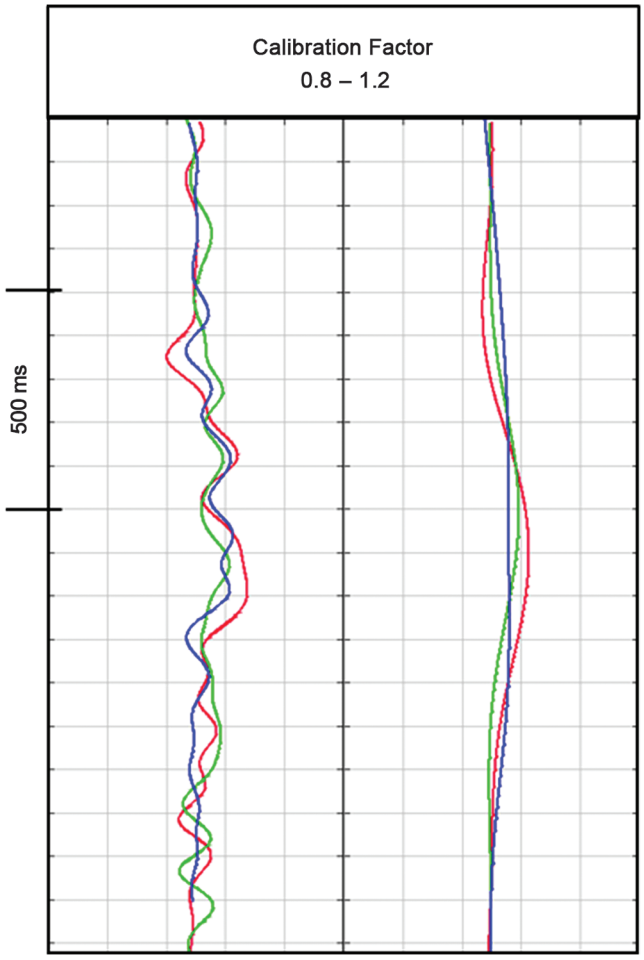

Figure 8. A comparison of calibration factors that are required to adjust the seismic velocities to the well velocities for three wells (each color is a different well). The first panel shows the calibration factors for a smoothing of 5–10 Hz, and the second panel shows the calibration factors for a smoothing of 1–2 Hz.

An example can be drawn from the same Gulf of Thailand data used above. The seismic velocity that was used for prestack depth migration (PSDM) was extracted at each of the wells. The highest frequencies in the seismic velocities are less than 5 Hz. A comparison at one of the wells is shown in Figure 7. The first panel of the figure shows that the seismic velocities have a very good match to the well-log data, although some systematic differences exist. It can also be observed that the presence of the basement is indicated in the seismic velocities just below the base of the well-log data. To make a fair comparison between the seismic and well-log velocities, filtering is applied. It is obvious that without the well-log velocities being extended in depth, a fair comparison cannot be made. With an extension to the surface and to the depth, the well-log velocities can be filtered to make a comparison and highlight the areas of systematic variation. At this stage, these differences need to be understood before deciding whether and how to compensate across the entire survey. A comparison of the calibration factors (well-log velocity/seismic velocity) measured at the three wells for two frequency scenarios is shown in Figure 8. The calibration factors show some similarities at the wells, but it is still not clear what approach should be taken to calibration. Even in such a simple case in which the calibration factors are small and reasonably consistent between the wells, there are errors whatever solution is chosen.

In addition, seismic velocities are compressional velocities, and for any inversion, more information is required, for example, for simultaneous inversion: density and shear velocity. Conversion of P-velocity to density or S-velocity is not straightforward. The relationships vary vertically through the section and laterally across the area. In the real-data example, the seismic velocity does not represent just the shale trend, but a combination of the shales and sands. If a low-frequency model is being generated through the updating approach, then care must be taken not to account for sands twice within the lowest frequency range. One approach would be to not to attempt to integrate the seismic velocities into the low-frequency model initially but to predict the low-frequency model as described previously and then compare this prediction with the seismic velocities within a common bandwidth. Systematic differences can be analyzed and interpreted. In the real-data example, a comparison between the updated low-frequency model result and the seismic velocities at a given depth reveals that the seismic velocities are not sufficiently detailed laterally to be used directly in building the low-frequency model. However, there are systematic lateral variations in the seismic velocity that

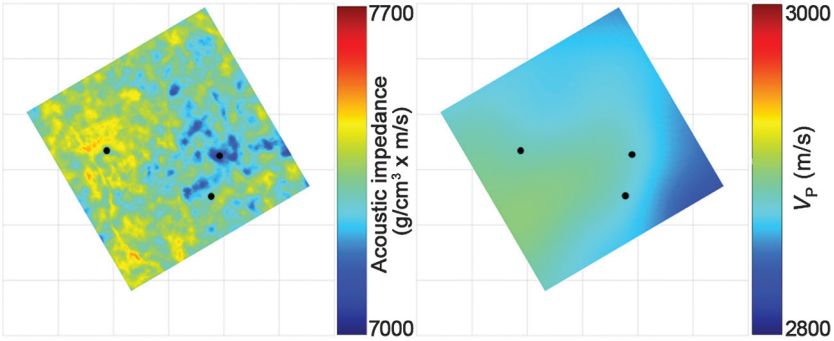

Figure 9. A slice through the predicted acoustic impedance high-cut filtered at 1–2 Hz (left) and the seismic velocity high-cut filtered at 1–2 Hz (right). A constant density is used to transform the color bar of acoustic impedance to velocity, so that the scales are comparable.

might need to be incorporated in the low-frequency model (Figure 9). There is evidence that the variations in the seismic velocities correlate to the depth of the basement and that the current assumption of a flat datum for shale trends in the inversion is insufficient away from the well control. Therefore, an approach here would be to use the lateral variations in the seismic velocities to establish a variable shale trend for use as a starting model for the inversion and then update.

Conclusion

In a deterministic framework, uncertainties are ignored and choices are made to produce the most likely result given the current data and information. For seismic inversion, choices are made that affect the relative and the absolute elastic properties. The examples shown in this paper strongly suggest that the most likely low-frequency model cannot be produced prior to an inversion having been performed. The simple interpolation of well-log data, even though it may be within a geologic framework and somehow conditioned to seismic velocities, will usually produce a low-frequency model that is meaningfully different from the most likely. Unfortunately, there is strong evidence, as provided by many publications and presentations, that in practice most low-frequency models are constructed prior to inversion and not revisited before interpretation of the absolute elastic properties. Under such circumstances, it is often better to only interpret the relative impedances. Indeed, many project objectives might be achieved to a sufficient degree through relative impedance interpretation. On the other hand, interpretation of relative impedances is not necessarily straightforward. The major issue being that relative impedances are only relative and can be misinterpreted without constructing either explicitly or within the interpreter's mind a possible low-frequency model. The example cited in the Introduction of normally pressured sands in an environment, in which the bounding shales are variously pressured, is best explored in the absolute domain, in which a low-frequency model has been developed that possibly includes information from the seismic velocities, assuming that they are of sufficient quality.

The methods for improving the construction of low-frequency models, outlined in this paper, imply that low-frequency model construction is an interpretation and prediction process that attempts to produce broad bandwidth absolute impedance models that are consistent with the relative impedances, well data, and seismic velocities as well as any rock physics and geologic constraints that are known. The aim is to produce the most likely result for all frequencies from zero up to the highest frequencies in the seismic data without constraining any particular frequency range to strictly honor the well data. The methods rely in part on extrapolating information from within the seismic bandwidth to lower frequencies. This process will not always be appropriate; for example, errors in the manual or automatic interpretation of relative impedances would cause the low-frequency model to be wrong. If a single thick body is misinterpreted as two separate bodies from relative impedance, this error might never be corrected through an iterative update of the low-frequency model. What is important is that there should be a shift in focus from concern about if and how a low-frequency model is biasing the interpretation of an inversion result to a concern about how to interpret relative impedances from an inversion, in the absence of low frequencies. Building a low-frequency model as part of the inversion process is one method of interpreting relative impedances and takes us one step further from seismic reflection data and closer to an accurate model of the rocks.

Acknowledgments

The authors would like to thank Kris Energy and partners, Mubadala Petroleum, and Northern Gulf for permission to present their data.

Appendix A

Problems with well data interpolation

Laterally, fluvial sandstone bodies often show clearly defined margins by lateral discontinuities of seismic amplitude or impedance, which result from their deposition in association within a relatively narrow fluvial channel. In many cases, a fluvial sandstone body is the result of sandbar formation and lateral accretion during its deposition, which in turn results in variations in the preserved sandstone body width and thickness. These characteristics are typical of many fluvial reservoirs that are exploited for oil and gas worldwide. In terms of 3D seismic data, a sandstone body's laterally confined margin often approaches a step function of impedance in horizontal (XY) space (on the scale of the seismic resolution); all spatial frequencies are needed

to describe a lateral step function accurately. Sampling theorem dictates that the lateral sample spacing of regular sampled data must be less than or equal to a half cycle of a spatial frequency component for it to be represented correctly. In a 3D seismic data set, the highest spatial frequency of the temporal signal present is limited by its associated first Fresnel zone width, to varying degrees by structural dip, or by the seismic acquisition geometry and process, for example, the bin size that defines the spatial Nyquist frequency of the seismic data. In contrast, low spatial frequencies are generally oversampled.

If a low-frequency model for seismic amplitude inversion is derived by an interpolation of well data alone, then the local well spacing itself defines a spatial Nyquist frequency for the interpolated low-frequency components. As such, the local average well spacing must be less than or equal to a half-cycle of the maximum spatial frequency component being interpolated for it to be represented appropriately; the azimuthal distribution of wells are also an important contributing factor.

In traditional inversion methods, the seismic temporal frequency bandwidth often dictates the choice of an upper frequency limit in a well-based low-frequency model, for example, the lowest effective signal frequency component present in the seismic data. This choice in turn limits the minimum lateral sample spacing necessary to interpolate the low-frequency model spatially to an equivalent degree as the seismic data. For example, the typical lowest effective temporal frequency in a seismic data set might be approximately 10 Hz (although this is being reduced via improved seismic acquisition methods in some areas). The Fresnel zone of a 10 Hz temporal frequency component collapses to a 150 m diameter, assuming a velocity of 3000 m/s and perfect 3D migration; the process ideally collapsing the Fresnel zone diameter to a half wavelength of the associated temporal frequency (Bacon et al., 2003). Clearly, in this 10 Hz example, a well spacing of 150 m that may provide adequate sampling is rarely achieved in oil and gas fields. Such laterally undersampled well data often lead to aliasing in the spatial frequency domain that causes bull's eye effects in traditional low-frequency models and their associated seismic amplitude inversion results. In laterally undersampled situations, a particular spatial frequency component being interpolated, but undersampled, doubles back around the spatial Nyquist frequency, and it is represented by an artificially lower spatial frequency component (given that lateral well distribution). This is the principal physical reason why the simple spatial interpolation of well data is often inappropriate.

In the oil and gas fields of the Gulf of Thailand, from where the main example in the paper originates, well spacing varies dramatically on the scale of several hundred meters to several kilometers. Here, the average well spacing may be as low as 400 m (Harr et al., 2011), although this is often along a locally preferred azimuth, for example, being parallel to a nearby fault. In this latter situation, spatial aliasing of a 10 Hz frequency band may still occur during simple well-based data interpolation, but the azimuth-dependent spatial Nyquist frequency often results in a preferred azimuthal bias in the distortion caused by the spatial aliasing process (seen to some extent in Figure 5). Clearly, if the tight laterally sampled seismic data could be used reliably to map the low-spatial-frequency components, as an alternative to traditional methods, then the detrimental effects of spatial aliasing in a low-frequency model could be minimized.

Appendix B

Reliability of seismic velocities

The process of seismic velocity estimation incorporates a large range of factors that may influence the quality and reliability of velocity estimates. Such factors arise from seismic acquisition and its surface environment, the geologic succession, seismic processing, and velocity interpretation.

Velocity estimates from prestack time migrated (PSTM) seismic data suffer from the fact that they often preserve a signature from a relatively large area of the overburden. For example, a lateral step function of interval velocity in the overburden results in a lateral wavelet-like distortion in the PSTM velocity estimates at a target level below (Toldi, 1984). This results from the associated raypaths of individual source-receiver pairs that comprise a prestack gather. The lateral dimension of such a distortion can approach several kilometers, and it is a function of (1) the depth difference between the target and the lateral velocity change, (2) the velocity gradient, and (3) the acquisition geometry, for example, the effective streamer length of a marine seismic survey. In reality, many such lateral velocity variations often occur in the overburden, particularly in a channel sand geologic succession, and result in a complex composite distortion. In many cases, such distortions are biased by the preferred source-receiver azimuth distribution during seismic acquisition. Therefore, in some cases, this may form an overriding factor in the reliability of the PSTM velocity estimates.

In contrast, the iterative velocity model building processes used in PSDM attempts to collapse this type of lateral wavelet-like distortion and isolate a lateral step function of interval velocity in the overburden velocity model. However, PSDM velocities are still a product of seismic imaging, and the velocity model of real data is never perfect. Lateral smoothing is also often performed that may degrade the velocity estimates. It follows that, in addition to data quality, a choice of whether seismic velocities are to be used in a low-frequency model for inversion may depend on the type of velocities available, their lateral sampling and processing, and in some cases the severity of the overburden velocity variations. Low-frequency seismic acquisition and advanced processing methods, such as those associated with full-waveform inversion, may

help to provide more reliable velocity estimates for low-frequency models.

References

Bacon, M., R. Simm, and T. Redshaw, 2003, 3-D seismic interpretation: Cambridge University Press.

Ball, V., J. P. Blangy, C. Schiott, and A. Chaveste, 2014, Relative rock physics: The Leading Edge, **33**, 276–286, doi: 10.1190/tle33030276.1.

Ball, V., L. Tenorio, C. Schiott, J. P. Blangy, and M. Thomas, 2015, Uncertainty in inverted elastic properties resulting from uncertainty in the low-frequency model: The Leading Edge, **34**, 1028–1035, doi: 10.1190/tle34091028.1.

Doyen, P., 2007, Seismic reservoir characterization: An earth modeling perspective: EAGE Publications bv.

Glenn, D., F. Hariyannurgraha, R. Schneider, K Kirschner, S. Walden, S. Smith, E. Berendson, C. Skelt, L. Lisapaly, R. van Eykenhof, and M. Sams, 2005, Enhancing consistency between geological modeling and seismic prestack amplitude inversion with pseudo-wells: An example from the Gendalo Field: 30th Annual Convention and Exhibition, Indonesian Petroleum Association, Proceedings, IPA05–G029.

Harr, M. S., S. Chokasut, C. Bhuripanyo, P. Viriyasittigun, and R. Harun, 2011, Evaluation of factors in horizontal well recovery in the Pattani Basin in the Gulf of Thailand: International Petroleum Technology Conference, IPTC, Paper 14981–MS.

Kemper, M., and J. Gunning, 2014, Joint impedance and facies inversion — Seismic inversion redefined: First Break, **32**, 89–95.

Latimer, R. B., R. Davison, and P. van Riel, 2000, An interpreter's guide to understanding and working with seismic-derived acoustic impedance data: The Leading Edge, **19**, 242–256, doi: 10.1190/1.1438580.

Mesdag, P. R., M. Feroci, L. Barens, P. H. Prat, and W. Pillet, 2007, Full bandwidth inversion for time lapse reservoir characterization on the Girassol Field: 69th Annual International Conference and Exhibition, EAGE, Extended Abstracts, A032.

Mesdag, P. R., D. Marquez, L. de Groot, and V. Aubin, 2010, Updating low frequency model: 72nd Annual International Conference and Exhibition, EAGE, Extended Abstracts, F027.

Mukerji, T., G. Mavko, and C. Gomez, 2009, Cross-property rock physics relations for estimating low-frequency seismic impedance trends from electromagnetic resistivity data: The Leading Edge, **28**, 94–97, doi: 10.1190/1.3064153.

Naeini, E. Z., and D. Hale, 2015, Image- and horizon-guided interpolation: Geophysics, **80**, no. 3, V47–V56, doi: 10.1190/geo2014-0279.1.

Pedersen-Tatalovic, R., A. Uldall, N. L. Jacobsen, T. M. Hansen, and K. Mosegaard, 2008, Event-based low-frequency modeling using well logs and seismic attributes: The Leading Edge, **27**, 592–603, doi: 10.1190/1.2919576.

Pendrel, J., and P. van Riel, 2000, Effect of well control on constrained sparse spike inversion: CSEG Recorder, **25**, 18–26.

Sams, M. S., T. J. Focht, and J. Ting, 2012, Porosity estimation from deterministic inversion: 74th Annual International Conference and Exhibition, EAGE, Extended Abstracts, X015.

Toldi, J., 1984, Estimation of a near-surface velocity from stacking velocities: SEP, **41**, 99–120.

Biographies and photographs of the authors are not available.