EARTH WORKS

Technical Note

Understanding **Stochastic Seismic Inversion**

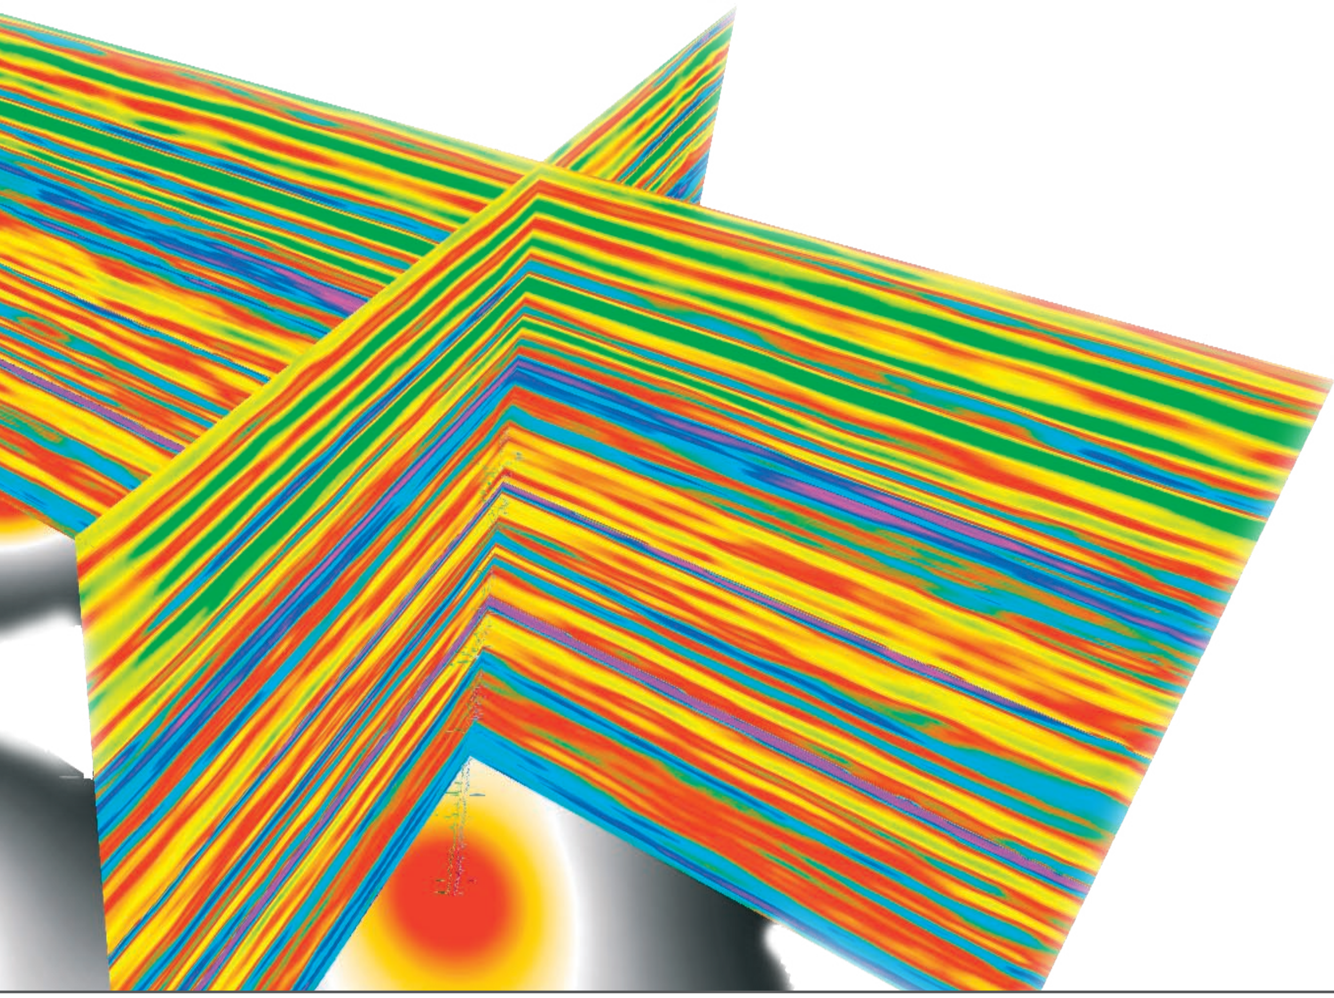

www.sorviodvnvm.co.uk

Earthworks Ltd - Company Overview

Earthworks Environment & Resources Ltd. was formed in June 2001 to provide services and consultancy internationally to the oil and gas industry. Specialising in subsurface geosciences, Earthworks has a broad set of skills in geophysics, geostatistics and risk and uncertainty evaluation.

Earthworks is a specialised company offering the latest technology through staff fully conversant with all technical aspects of the work. Our expertise and understanding in geophysics and geostatistics is second to none and as a consequence we are also a leading trainer to the industry in these disciplines, through courses and technical conferences.

Based in the historic medieval town of Salisbury in the heart of the English countryside, we offer the most modern technology and electronic communications in ancient surroundings. We have ready access to London, the City and its international airports, allowing Earthworks to provides services and consultancy to any locale or operation worldwide.

Earthworks Environment & Resources Ltd
Prestex House
Edison Road
Churchfields
Salisbury
Wiltshire
SP2 7NU

Tel: 01722 343990
Fax: 01722 343991
Web: www.sorviodvnvm.co.uk

Acknowledgement

This technical note was originally published as a tutorial in the November and December 2006 issues in First Break magazine. The original article is reprinted here with permission of the European Association of Geoscientists & Engineers. The original copyright of the EAGE is also acknowledged here.

Introduction

Seismic inversion tools designed to estimate impedance have been available to geophysicists for over twenty years. Most of the available methods are based on forward convolution of a reflectivity model with the estimated wavelet, comparison of the modelled output with the observed seismic trace and then updating the reflectivity model (inverting) to minimise the difference between the modelled and observed traces. Whether generalised linear inversion, sparse spike or simulated annealing, all the algorithms work on this basic principle of minimisation. Methods based on minimisation are commonly referred to as "deterministic". The output of a deterministic inversion is a relatively smooth (or blocky) estimate of the impedance.

Because of its smoothness deterministic inversion is generally unsuited for constraining reservoir models used for volumetric calculations, estimation of connectivity or fluid flow simulation. Stochastic seismic inversion generates a suite of alternative heterogeneous impedance representations that agree with the 3D seismic volume. Taken together, the suite of possible impedance models or realisations capture the uncertainty or non-uniqueness associated with the seismic inversion process. Stochastic seismic inversion is complementary to deterministic inversion. The deterministic seismic inversion is the average of all the possible non-unique stochastic realisations.

Although the principles of stochastic seismic inversion were published over 12 years ago commercial implementation and application has only started to grow in the last five years or so. For many geophysicists, understanding and being able to make a discerning judgement on the possible benefits of this new technique is difficult. A number of misconceptions concerning stochastic seismic inversion have arisen, particularly related to resolution. A commonplace but incorrect statement is that stochastic techniques can "…allow substantially increased resolution, capturing details well beyond seismic bandwidth". It is the purpose of this tutorial to provide a clear theoretical basis for deterministic and stochastic inversion and assist the geophysicist in making an informed decision concerning the application of stochastic seismic inversion to his or her particular reservoir description problem.

In order to understand stochastic seismic inversion we will have to understand some of the limitations of conventional seismic inversion (often referred to as "deterministic" inversion), provide a grounding in some geostatistical concepts and also consider the general problem of estimation at unmeasured locations.

This tutorial will only consider seismic inversion in the sense of estimating an impedance model of the subsurface. "Impedance" will be taken in a very general sense to refer to any rock property estimated from surface seismic data. This could include acoustic or elastic impedance, extended elastic impedance or any other more elaborate pre-stack inversion scheme to estimate Vp, Vs and density. Stochastic inversion may be applied to any of these objectives and the aim of this tutorial is to describe only the general limitations of deterministic inversion and the possible advantages of a stochastic inversion framework and not to consider the specifics of pre- or post-stack inversion objectives.

Deterministic and Best Estimate

Conventional seismic inversion is often referred to as deterministic inversion. A better description would be best estimate inversion. The word deterministic should properly refer to a model from which predictions are determined directly through a functional (physical) relationship. Examples might include Darcy's law for fluid flow or Newton's Laws of Motion.

Best estimate refers to a model in which the objective is to minimise the error in the prediction. Well known examples would include linear regression and geostatistical procedures such as Kriging. Conventional seismic inversion minimises the difference between the forward convolution of a wavelet with a reflectivity model and the seismic trace.

Objective Function

The mathematic objective of an inversion algorithm is to minimize (or maximize) an "objective function". The objective function will always include a quantitative measure of the misfit between the observed data and the data predicted using the inverted model.

The most commonly used measure of data misfit is the L2 norm or least squares difference [the sum of the square of the data residuals, $\sum(d_{obs} - d_{calc})^2$]. "Least squares" inversion is mathematically expedient and stable; however it can be too sensitive to outliers in the data. Other misfit measures are sometimes preferred in order to achieve better noise rejection and/or reduce the non-linearity of the problem. One example is the L1 norm, which is simply the sum of the absolute data residuals, $\sum|d_{obs} - d_{calc}|$.

As well as the data misfit, additional constraints can also be included in the objective function. The most common of which are smoothness/roughness constraints and a priori model constraints. All such constraints are used to reduce the non-uniqueness of the inversion problem by providing additional (a priori) information that is not contained in the data to be inverted.

Seismic inversion is generally understood by geophysicists to be non-unique: there are a large number of alternative models which would have a forward convolution with an acceptable match to the seismic trace. It is not widely appreciated that seismic inversion is non-unique because the seismic trace is band limited. For a given inversion algorithm (with a specific objective function and model parameterisation) the seismic inversion is unique within the seismic bandwidth.

Limitations of Deterministic Seismic Inversion

Deterministic seismic inversion suffers from a number of limitations. These arise because of the limited bandwidth of the seismic data. If we assume a convolutional model for the seismic trace:

$$x(t) = r(t) * w(t)$$

where x(t) is the observed seismic trace, r(t) the true reflectivity and w(t) is the wavelet, then the goal of seismic inversion is to find the inverse operator v(t) which when convolved with the wavelet w(t) gives a Dirac spike (a reflection coefficient). Unfortunately since the wavelet is band limited there cannot exist an operator to convert the wavelet to a Dirac spike, which has a white amplitude spectrum and infinite bandwidth (Oldenburg et al, 1983). Instead, the convolution of our inverse operator with the wavelet is a band limited averaging function which has a width. Both low frequencies below the seismic bandwidth and high frequencies above the seismic bandwidth are missing. This means that deterministic seismic inversion is limited to estimating a sparse reflectivity series which corresponds to a blocky average impedance profile. This limitation is general and applies to all deterministic inversion schemes: we cannot obtain estimates at frequencies which are not present in the seismic signal.

The idealised averaging function will be a spiking deconvolution of the wavelet with its inverse operator. The output will have a strong central peak and relatively low amplitude symmetric side lobes. In sparse spike inversion an iterative procedure is used instead of a spiking deconvolution to place spikes to represent reflection coefficients. This is equivalent to suppressing the side lobes of the spiking deconvolution output. However, for a sparse spike inversion to be stable we have to accept that the seismic data has limited information content due to its bandwidth restriction and not place the spikes too close together, thus violating the resolution limitations of the averaging function.

In deterministic inversion the estimation is a trade-off between resolution and accuracy. We usually include a parameter that allows a compromise between these two requirements. For example, in a sparse spike inversion the number of spikes allowed (the sparseness of the modelled reflectivity) controls this trade-off. Allowing sufficient but few spikes will give a solution which will have thick intervals containing relatively accurate average impedances. Increasing the number of spikes will result in a solution with a higher resolution but also greater uncertainty in the estimated impedances. Proper choice of the sparseness criteria will result in a blocky average impedance between the sparse reflection coefficients which we can describe as *locally* smooth.

A further limitation is the presence of noise in the seismic trace and the estimated wavelet. Noise will be more significant at high and low frequencies close to the bandwidth limits where the seismic signal is weaker. At the high frequency end of the seismic spectrum, attempting to estimate too many reflection coefficients will result in modelling noise as reflectivity, resulting in greater uncertainty in the estimated impedances. This is called “over fitting” the data.

The bandwidth limitation is also important when we consider low frequencies. The (missing) low frequencies contain the critical information concerning the absolute values of impedance. This means that it is impossible to recover the absolute impedance values from a seismic trace. Instead, inversion of the trace can only recover the relative changes in impedance. In order to convert the estimates of relative change in impedance to an absolute impedance additional low frequency information must be added to the inversion. This usually takes the form of a smooth 3D interpolation of the well data impedance values, constrained stratigraphically by seismic horizon picks.

After inversion the smooth low frequency model is embedded in the seismic inversion. Artefacts in the low frequency model manifest themselves as artefacts in the deterministic inversion. We might consider the embedded low frequency model component to be *globally* smooth due to the lack of spatial sampling between the wells.

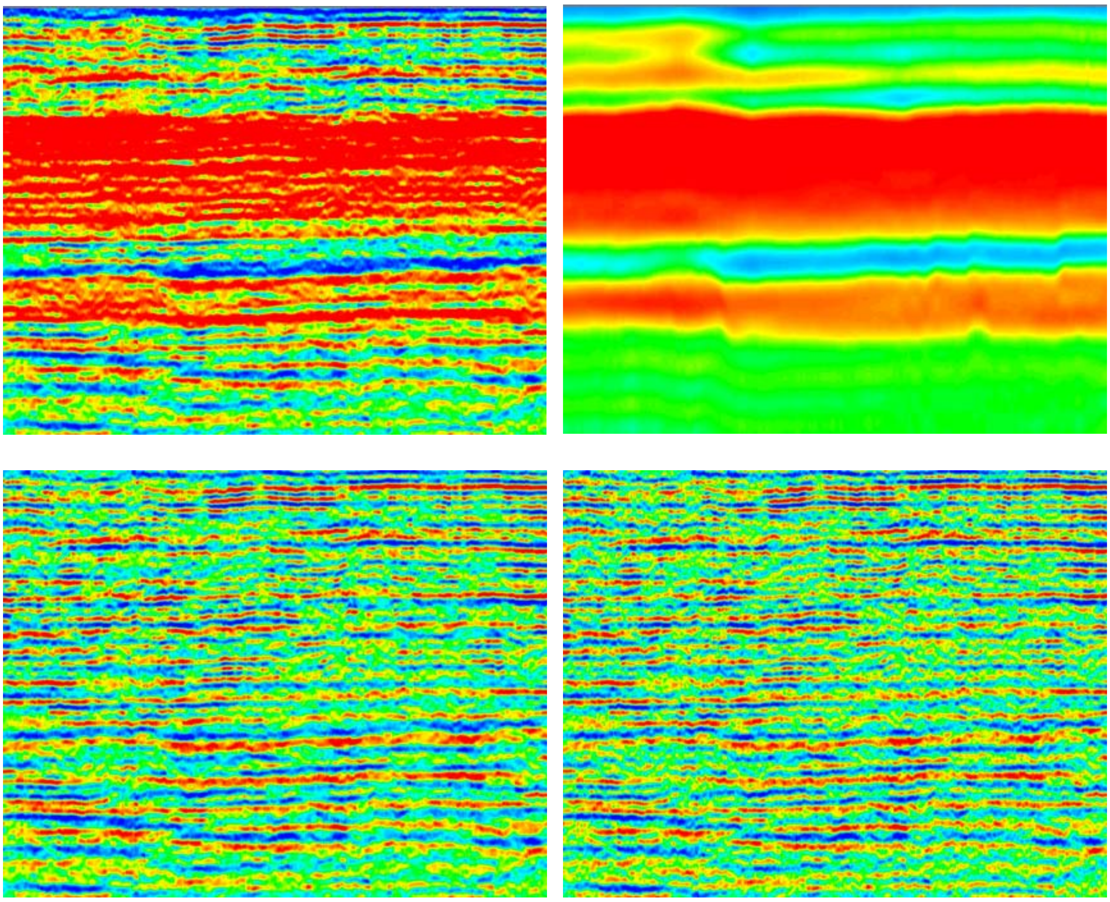

Figure 1 CSSI deterministic inversion: final inversion (top left); low pass filtered (top right); high pass filtered (lower left) and relative impedance (lower right)

Figure 1 shows an example of a constrained sparse spike deterministic inversion (CSSI) and of the limitations described above. The final deterministic inversion (top left) is low pass filtered to reveal the low frequency model (top right). The well locations used in the interpolation can be easily identified along with the "halo" effect around the wells which will result in artefacts in the final inversion.

Of more concern in the example in Figure 1 is the comparison of the lower two panels. The lower left panel is the final inversion after high pass filtering to remove the low frequency model. It is the output of the CSSI inversion algorithm after inverting the seismic but prior to adding the low frequency model The lower right panel is the result of a relative impedance calculation obtained by integrating the original zero phased seismic data. The CSSI is almost indistinguishable from the relative impedance, the only difference being due to a three trace mix applied to the seismic to improve signal:noise ratio prior to CSSI inversion. This suggests that in this case the deterministic inversion offers little value over and above a simple relative impedance calculation. This is because the geological layering in this locale is thin bedded with few resolvable intervals. Any deterministic inversion algorithm is likely to be inappropriate in this geological environment.

Based on these observations we can expect deterministic seismic inversion to perform best when the:

(a) geological system is comprised of relatively thick layers (ie with little vertical variation of impedance within each layer)

(b) stratigraphic layering is conformal, with slow lateral thickness changes
(c) lateral impedance variation within a layer is stationary, meaning that there is no significant trend in the impedance variations.

Criteria (a) ensures that the local smoothing associated with the lack of high frequencies is not important because there is sufficient bandwidth from the seismic frequencies to resolve the layers and our output layer average estimate is therefore a good representation of the impedance of the layer. Geology which meets the criteria (b) and (c) will result in the most reliable estimate of the low frequency model and hence reliable estimates of absolute impedance in the inversion.

Prediction at Unmeasured Locations

The objective of any estimation procedure is to predict the value of an attribute at an unmeasured location. All estimation procedures obtain accuracy by averaging or smoothing. Consider a class room in which some students are seated and some are late, having been detained together elsewhere as a group, so their seats are empty. If we wish to estimate each of the heights of the missing students, our best estimate would be to use as our prediction the average or mean height of the students already present in the room.

Our prediction of the height of each student would then be the same for each empty seat – the mean or average value. If we look at a grid of the seats, where a student is already present we will observe a true value and at the empty seats we will observe our predicted value. The predicted values will be very smooth – in fact they have no variability at all. This is what we mean when we describe estimation as smoothing towards the average. Using the average as a predictor in this way has an important benefit: the square difference between our prediction and the heights of the missing students will be minimised – meeting the so-called least squares criteria.

Now consider that for each seat representing a missing student we want to use our prediction to determine if the missing student is taller than some threshold. If the chosen threshold is equal to or less than the predicted (average) value, all missing students would be predicted as being taller than the threshold. If the threshold were greater than the average value we would predict none of the missing students to be taller than the threshold. This type of threshold or cutoff calculation is integral to volumetric estimation. Our smooth prediction is clearly not suited to this type of prediction and would lead to a bias, either over or under predicting the number of missing students that meet our threshold criteria.

However, for the missing students we can state the probability that each student is taller than the cutoff. Using the data for the heights of students already present, we calculate the number of those students whose height exceeds the cutoff. Dividing this number by the total number of students present gives the probability that any of the missing students are taller than the cutoff.

Spatial Estimation

We can illustrate the use of mapping to predict the value of an attribute at an unmeasured location in the sub-surface using the class room example just described. Each student already in the classroom represents a known data value (analogous to a well) and the empty seats represent the grid nodes at which we are making predictions. The key difference when mapping is that we know there is a special spatial relationship between the sample points, whereas the heights of the students are not spatially related – the students could sit anywhere they choose and the best estimate at the empty seats (unmeasured locations) would still be the average value. The student heights are spatially independent samples.

In mapping the subsurface we can make use of the special spatial relationship between the data points. This special spatial relationship is called spatial continuity or spatial correlation. If we estimate the value at an unmeasured location using a weighted linear combination of the known samples then we can give more weight to samples which are close to (more strongly correlated with) the unmeasured location and less weight to samples far from (less strongly correlated with) the unmeasured location, thus accounting for their relative importance to the estimation. However, because it is a weighted linear combination it is essentially an averaging or smoothing process. In the classroom example, because the known data (the student heights) are spatially independent we combine them with a weighted linear combination (called the average) where each sample is given the same weight.

In geostatistics the function which describes this special spatial relationship between the points is called the variogram. The variogram (or rather its upside equivalent, the covariance) is analogous to the autocorrelation function in time series analysis. The variogram may be anisotropic, representing different spatial correlation behaviour in different directions. Due to geological layering the horizontal variogram of impedance usually identifies spatial correlation over much greater distances than the vertical direction.

The best estimate procedure used in geostatistics is called Kriging. The Kriging estimator calculates the appropriate weights for the estimation by reference to the variogram, thus giving an estimate which is least squares with respect to the spatial correlation function (usually the variogram). Kriging has another advantage: in addition to the best estimate it also outputs the uncertainty we should assign to the estimate at any unmeasured location. Clearly the uncertainty in our estimation procedure will be small when we are close to a measured value and will generally increase the further we make the estimate from the measured data points. This uncertainty usually takes the form of a Gaussian distribution of error (see Figure 2).

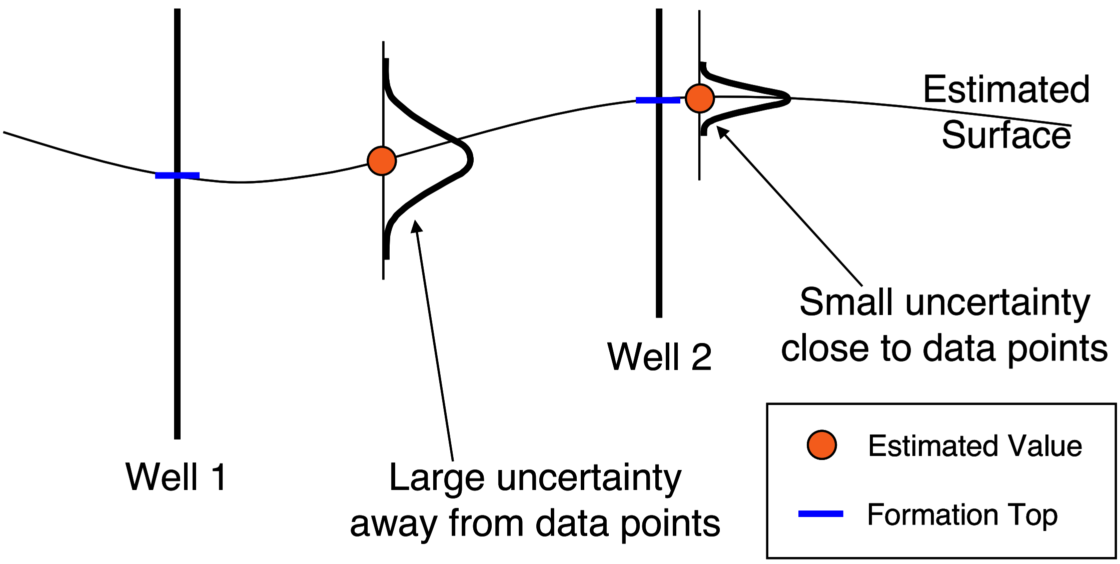

Figure 2 Diagram showing estimation at unmeasured locations on a surface. Far from well control uncertainty is large whereas close to well control the uncertainty will be much smaller

Side Effects of Smooth Estimation

Because all estimators of this type are smoothers, we should be aware of the consequences of smoothness on 2D or 3D estimation problems. Consider smoothing a 2D topographic elevation surface obtained from satellite measurements (see Figure 3). Taking the true measured surface and imposing a false datum and imaginary hydrocarbon contact we can obtain a true gross rock volume (GRV) and area for the structure defined by our imaginary hydrocarbon contact. In the example shown the true volume is 300 M m^3 and the true area 9.84 km^2. Also shown is the main accumulation and its lateral connectivity above the cutoff, coloured in orange in the upper right display of Figure 3. Note that the central structure is not connected to the smaller structure to the east.

Two lower displays of Figure 3 show the effect of smoothing of the original structure using a small moving average filter. In the lower left image the topographic surface (top left) was smoothed using a 3 x 3 moving average element before clipping at the imaginary hydrocarbon contact. The lower right image is with a 5 x 5 moving average. In both cases the GRV and the area of the accumulation are underestimated, the error in the estimates becoming larger as the smoothing increases. A 3 x 3 moving average reduces the GRV by over 7% to 278 M m^3. A 5 x 5 moving average reduces the GRV estimate to 261 M m^3, an underestimate of 13%. This volumetric bias effect is well know in geostatistics and will apply to the result of deterministic inversion. In addition, smoothing also changes the connectivity of the predicted surface. In the image derived after the application of the 5 x 5 moving average, the main structure is erroneously predicted to be connected to the eastern structure.

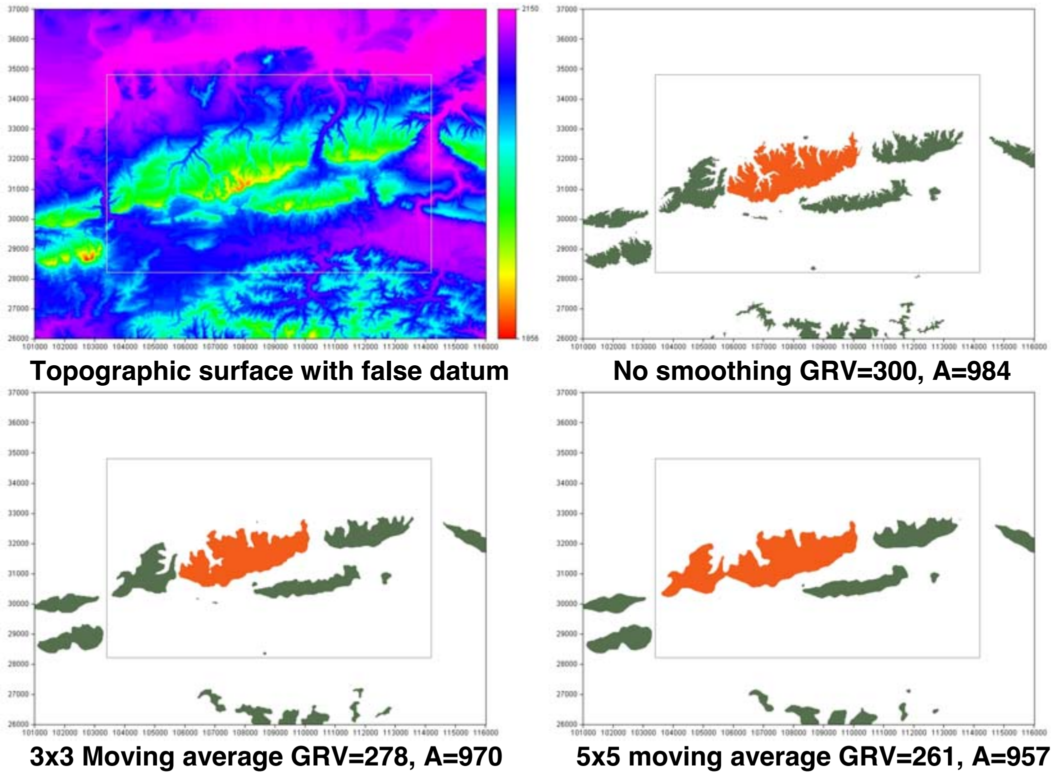

Figure 3 (top left) True topographic surface measured from satellite data; (top right) truncation of true topographic with a pseudo hydrocarbon contact; (lower left) truncation of topographic surface after smoothing with 3 x 3 moving average and (lower right) after smoothing with 5 x 5 moving average.

The smoothing used here could be thought of as analogous to mapping, where we fill in unmeasured nodes with a smoother prediction, or as reduction of the bandwidth from the original image. We can draw two conclusions about smoothing from this example, (a) smoothing causes our estimates to appear more connected than they really are and (b) smoothing gives rise to a bias in volumetric and area calculations.

It is important to note that the moving average (or any analogous mapping procedure) does not result in a bias in the depths, only in the volumes, area and connectivity. The reason can be best understood by comparing the cumulative distribution functions (CDF) of the true surface and the predicted surface.

Figure 4 shows the cumulative distribution function (CDF) for the true data values and the predicted (or in this case, smoothed) data values. Note that whilst the predicted values have been smoothed closer to the mean to give an accurate prediction, the positive errors at low values and the negative errors at high values are balanced and, on average, cancel out. We therefore describe the estimation procedure as *unbiased*.

However, when we attempt a volumetric calculation based on the predicted (estimated) surface, we truncate the CDF and integrate only one end of the distribution so the errors no longer cancel out. For a depth surface used to estimate GRV, the hydrocarbon contact is at the low end of the CDF and we integrate to shallower depths to calculate the volume. It is clear that the number of depth values predicted to be shallower than the contact will be underestimated and also that, because they are smoothed in the estimation procedure and therefore closer to the mean value, their average contribution to column height will also be reduced. This combination results in the underestimate of GRV from the smooth estimate.

Note that in the GRV case the cutoff is usually at the low end of the CDF and the integration to the left. If the integration were to the right (for example, the variable is impedance) we would over predict the volume. Four possible cutoff arrangements and their associated volume prediction bias are illustrated in Figure 5.

The connectivity between cells is increased by smoothing because a smooth estimate is less likely to have abrupt changes between the cells, so the surface is more continuous and therefore more connected.

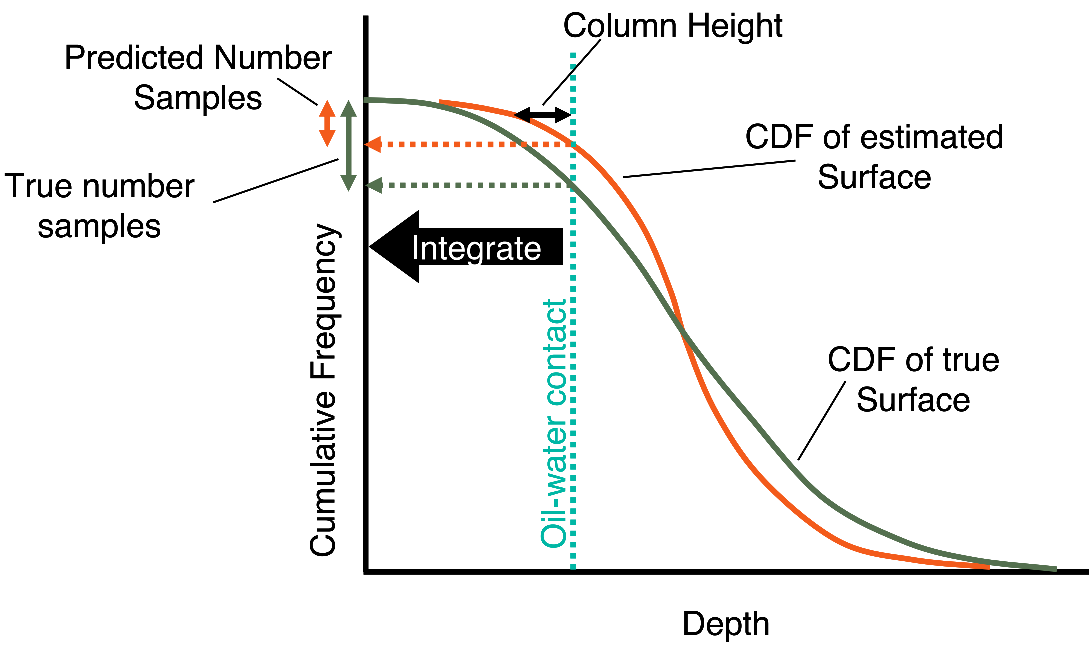

Figure 4 Comparison of CDF of true surface (green) and estimated or smoothed surface (orange) and illustration of mechanism for volumetric bias arising from truncation and integration of the distribution

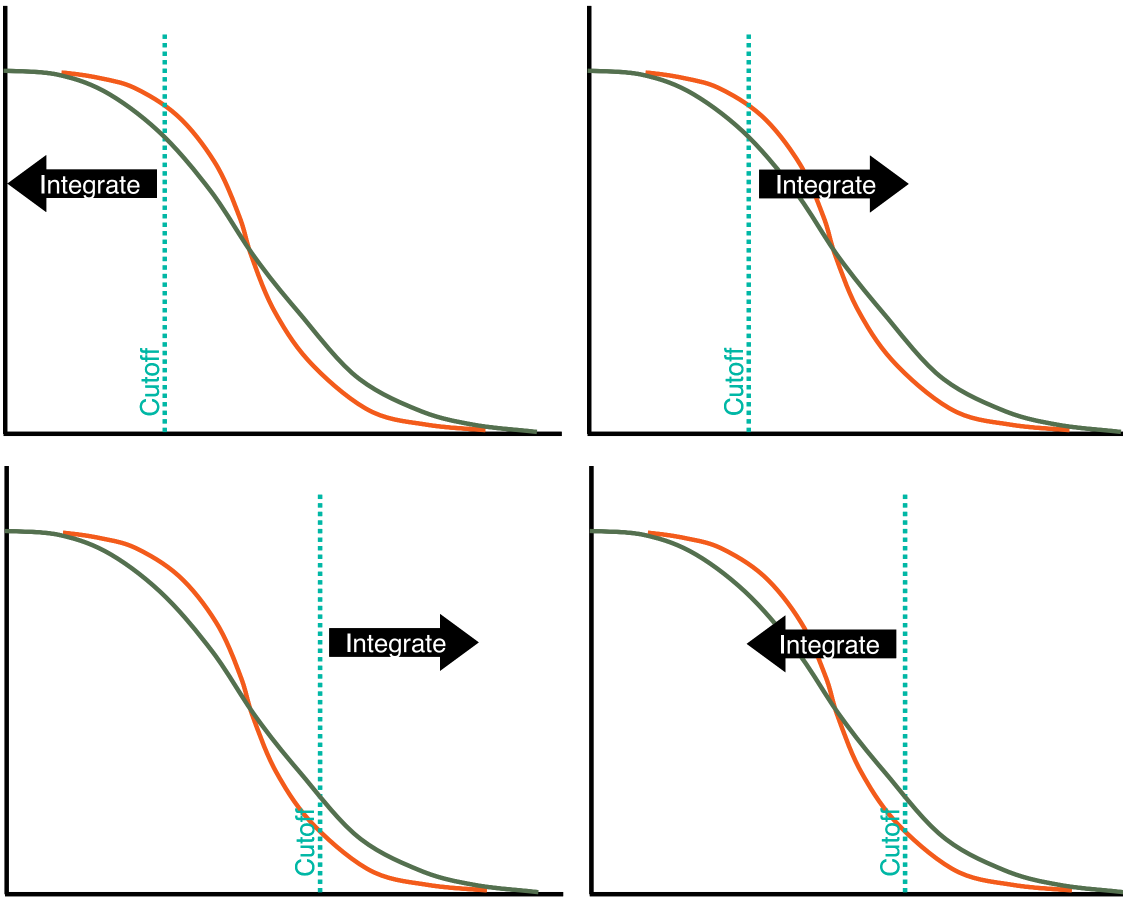

Figure 5 Four possible cutoff arrangements for different types of measurements. Those on the left will result in under prediction of volume and those on the right over prediction of volume.

Geostatistical Simulation

Both of these problems are widely understood in geostatistics and addressing these problems is the purpose of geostatistical simulation methods. To address the smoothing we need to introduce some additional variation to our estimates which corrects the CDF. It is not possible to do this uniquely so, like the classroom example, we introduce it as a probability using a stochastic method. The missing information which causes us to smooth in estimation is the uncertainty in the estimation. The uncertainty arises due to the limited bandwidth of our measurements (for example, seismic data) and insufficient sampling (for example, of the inter-well space).

The missing variability is simulated using a Monte Carlo method based on a random number generator. However, simply adding random noise to our smooth estimates would only correct the volumetrics globally. In reservoir problems we also need to consider the connectivity as this determines whether we have single or multiple compartments. Therefore the missing variability must be added systematically, consistent with the spatial correlation function of the data. Finally we should note that we cannot just add a random component to our smooth estimate, even if it were spatially constrained, because then we may generate a surface which no longer passes through (is *conditional to*) our measured data points such as wells.

A valid geostatistical simulation therefore must be (a) conditional to our measured data points; (b) reproduce the histogram of the data and (c) honour a spatial correlation function, such as a variogram.

Because geostatistical simulation is not unique, there are many possible surfaces (or in 3D seismic inversion, impedance volumes) which could meet these criteria. To explore the uncertainty and correctly estimate volume we must look at many simulated surfaces or volumes. Each possible simulation is referred to as a *realisation*. A simple way to understand the simulation process is to consider throwing a six-sided die for which the faces are numbered 1 through 6. The *expected value* of the outcome of rolling the die is 3.5, the average of all the possible outcomes. If we have exact knowledge of the distribution of all possible outcomes then this would also be our *best estimate* or prediction at an unmeasured location (throw of the die), since this minimises the prediction error. However, the numbers 1 to 6 represent the possible *realisations* of throwing the die. Note that our smooth best estimate surface is in fact the average of all the possible realisations we could generate for the process. The best estimate is not and can never be a realisation as it fails two of our criteria for a valid realisation: it does not reproduce the histogram of the data and it does not honour the spatial correlation function.

Having generated a suite of realisations the best estimate of volume is obtained by performing our volume calculation on each realisation to obtain a suite of volume estimates. The average of the volume estimates is then our best estimate of the true volume. Note that the average of the suite of realisations will be the same as our smooth best estimate, so averaging the realisations and then calculating the volume is not correct and will give the same biased volume estimate as the smooth surface.

A further application of the suite of realisations is to estimate the probability of occurrence of a criteria such as exceeding an impedance threshold. In the example of the students in the class room the probability can be estimated in a straightforward manner from the histogram heights already measured. For the seismic inversion problem, the uncertainty associated with estimating impedance away from well control will be reduced by including the seismic trace constraint, but it is difficult to estimate the uncertainty directly without going through the simulation process. If we use stochastic seismic inversion to generate a suite of possible impedance realisations, for each sample we can count the number of realisations which meet our threshold criteria. If the impedance threshold is exceeded for 70 out of 100 realsiations and the threshold is likely to indicate the presence of sand then our estimate of the probability of sand occurrence at that sample would be 0.70 or a chance of 70%. By checking all samples in a 3D volume, a probability volume can be constructed for a given impedance test.

Because the stochastic inversion impedance realisations reproduce the spatial behaviour described by a spatial correlation model the realisations can also be used to estimate connected probabilities and their associated volumes. For example, each realisation can be seeded at a well location and the samples connected to this location identified by a connectivity algorithm. The number of samples connected to the well then gives an estimate of the volume of reservoir in direct communication with the well. By repeating the calculation for each realisation, the uncertainty in this connected volume estimate can be established and shown as a CDF along with the associated connected probability seismic cube.

In order to generate reliable statistics for estimating volumes, probability cubes and probability of connectivity we must produce sufficient impedance realisations. The question of how many realisations is “sufficient” has been considered in the context of geostatistical depth conversion (Samson et al, 1996). The general criteria for judging if sufficient realisations have been generated is to check the resulting volumetric CDF: a smooth function would be acceptable whereas a CDF with discontinuities or gaps would suggest a greater number of realisations should be computed. Samson et al note that for their geostatistical depth conversion example 100 realisations were probably insufficient but tests of 250 and 500 realisations showed little difference in the estimated CDF of GRV, so 250 realisations were considered adequate to capture the volumetric uncertainty.

<u>Sequential Gaussian Simulation</u>

There are many different algorithms which can be used to generate geostatistical realisations, satisfying our three criteria of conditioning to measured samples, histogram and spatial correlation. Some algorithms, such as the spectral (frequency domain) methods, are very fast. The methods usually include some form of non-conditional spatial generation mechanism and a Kriging step to provide the conditioning to the measured well data.

The easiest geostatistical simulation algorithm to describe is the sequential Gaussian simulation or SGS. SGS elegantly combines the simulation and Kriging (well data conditioning) steps. The method follows a basic recipe of:

(a) Randomly select an unmeasured grid node at which a value has not yet been simulated
(b) Estimate the value and the Gaussian uncertainty at the unmeasured location by Kriging using the measured data points (for example, the wells)
(c) Draw a random number from the distribution defined by the estimate and uncertainty in step (b) and assign this simulated value to the grid node
(d) Include the newly simulated value in the set of conditioning data as though it were a real data point
(e) Repeat steps (a) – (d) until all unmeasured grid nodes have a simulated value.

Depending on the exact implementation there may be some variation from this basic procedure. There are also mathematical requirements that the measured data must follow a Gaussian distribution. If it does not then pre-conditioning of the data through a transformation to a Gaussian distribution will be required. This can be achieved through straightforward techniques such as the *normal score transform*, similar to the manner in which logarithmic distributed data can be converted to a Gaussian distribution through a lognormal transform. SGS and the normal score transform are standard procedures in reservoir modelling packages for generating realisations of porosity and permeability fields.

SGS for Seismic Inversion

The SGS method is simple to implement for problems such as 3D porosity simulation within a reservoir model. However, to perform stochastic seismic inversion to generate realisations of impedance is a more difficult simulation task. This is because in seismic inversion we have to include an additional constraint in the geostatistical simulation process: the seismic amplitudes represented by the seismic trace.

Including additional constraints in a geostatistical simulation (or a Kriging) using methods such as cokriging or external drift is usually straightforward. For example, in depth conversion the combination of well depths and a corresponding picked seismic horizon (perhaps through a velocity model) using a geostatistical method such as collocated cokriging or Kriging with external drift (or their simulation equivalents) has been widely practised by geophyscists (Galli and Meunier, 1987; Scola and Ruffo, 1992). Similar applications using horizon based seismic attributes and well data such as porosity have also been successful and popular (Doyen, 1998).

Unfortunately the form of the seismic constraint for stochastic seismic inversion is less tractable than the depth conversion or seismic attribute applications just mentioned. There are two considerations. Firstly, the seismic trace represents the convolution of the reflectivity series with the wavelet and is therefore non-linearly related to the 3D impedance realisations we seek to generate. Because it is a convolution we cannot incorporate seismic information on a sample by sample basis: to make the simulated impedances conditional to the seismic we have to consider a whole trace at a time. Secondly, for the successful applications using additional constraints previously described the secondary constraint such as the picked time structure surface from a 3D survey informs the geostatistical procedure of lateral changes in the mean value of the surface, commonly referred to as trends. A seismic trace cannot do this because it does not contain low frequency information, so the seismic constraint in stochastic inversion is only within the seismic bandwidth and the stochastic inversion is constrained only by the well data in the low and high frequency ranges.

Haas and Dubrule published their SGS stochastic seismic inversion algorithm in First Break in 1994. Because of the difficulty of simulating impedance and conditioning to a seismic trace, stochastic seismic inversion using SGS includes two modifications compared to the grid node by grid node recipe described above. To constrain to the seismic trace we have to simulate the problem a whole trace at a time, so the method described by Haas and Dubrule selects a random trace position and then uses the well data to simulate an impedance profile for that trace. The impedance trace is converted to reflectivity, convolved with the wavelet and the modelled result compared to the actual seismic trace at the same location.

Because the simulation generated in this way is not constrained by the seismic trace, the simulated trace may not be at all similar to the seismic trace. If this is the case then we would reject the simulated trace and simulate another trace at the same position. We would continue simulating traces until we generated (by chance) a simulated trace for which the forward convolution is a good match to the seismic trace at that location. When such a trace is found we would accept it and insert the resulting impedance profile into our data set as though it were a real profile. The usual practise for the accept/reject approach is to simulate a large number (perhaps several hundred) possible profiles, convolve and compare all of them to the seismic trace and select the trace with the best fit where the fit exceeds some minimum criteria such as a correlation coefficient threshold. A typical choice of minimum correlation coefficient would be the observed correlation between the wells and the seismic. The accepted trace is then included as conditioning data along with the measured data (wells), as in step (d) of the SGS recipe given above.

Simulation at the next randomly selected trace location would then include both our real (well) and simulated impedance profiles. The procedure is repeated until all of the seismic trace locations have been simulated to give a 3D impedance realisation conditional to the well impedances, the histogram, the spatial correlation function and, by forward convolution, the observed seismic volume.

The SGS method is straightforward to describe, mathematically rigourous and elegant apart from the accept/reject strategy. However, the SGS approach is relatively slow. All SGS methods are generally slow because they require a Kriging step for every unmeasured sample and every realisation. Other geostatistical algorithms for generating stochastic realisations are computationally more efficient, requiring just one Kriging per unmeasured location irrespective of the number of realisations generated. The accept/reject requirement for conditioning to the seismic trace is an additional factor that further slows down the algorithm. Because stochastic seismic inversion involves large 3D seismic data volumes and many realisations, computational efficiency is an important consideration in the choice of algorithm. Because of this modern methods often use a combination of SGS and 1D Markov Chain Monte Carlo, (MCMC) to speed up the local trace simulation. There is also a direct method for stochastic seismic inversion which does not require an accept/reject stage and uses a single Kriging and inversion for all realisations. The stochastic spatial simulation is based on a spectral (frequency domain) method which is one of the fastest geostatistical simulation methods available.

Finally we should note that a stochastic seismic inversion can be run at any required output sample rate. Note this does not imply stochastic has a higher *resolution* because the resolution is controlled by the frequency content and bandwidth of the seismic conditioning data. Rather a stochastic seismic inversion can *simulate* the broad band impedance and so properly represent the uncertainty in the seismic inversion. However, a stochastic seismic inversion will usually extract more detail from the seismic for a given frequency content and bandwidth because it does not locally smooth like a deterministic inversion and this changes the trade-off between resolution and accuracy at the low and high frequencies at the ends of the seismic spectrum where the signal:noise ratio is poor.

Stratton example

The data-set used here to compare deterministic and stochastic seismic inversion is the Stratton Field 3D seismic and well log data package, prepared by the Bureau of Economic Geology, Austin, Texas, USA (Levey et al, 1994).

Stratton Field is an on-shore gas field producing from the Oligocene Frio Formation in the NW Gulf Coast Basin, Gulf of Mexico. There is little faulting in this interval and the formations relatively undeformed and flat lying. Reservoir facies of the middle Frio are interpreted as multiple amalgamated fluvial channel-fill and splay sandstones. The composite channel fill deposits range from 10 to 30 ft thickness and show either an upward fining or a blocky log profile. The composite channel deposits can be up to 2,500 ft in width. Splay deposits show typical thicknesses of 5 to 20 ft and are proximal to the channel systems.

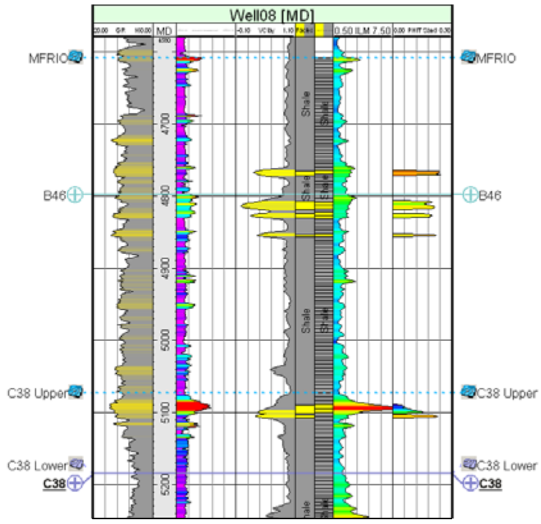

Figure 6 Log display for Stratton Field well 8. (L-R) logs are: gamma ray acoustic impedance; Vclay; resistivity and total porosity.

Sands are generally indicated by high impedances and have typical velocities of 3,650 m s^{-1}, a 30 ft sand having a two-way time thickness of about 5 ms thick in time. In the wells sand facies have been identified with a combination of ILM (resistivity), Vclay and acoustic impedance cutoffs (see Figure 6). An impedance cutoff of about 8,150 m s^{-1} * g cm^{-3} gives a good seismic discrimination between sands and shales, as illustrated by the acoustic impedance histograms of sand and shale in Figure 7.

A deterministic inversion has been performed over the 3D seismic cube and a detail of the resulting impedance volume showing the B46 sand interval through well 8 is shown in Figure 8. High impedances (blue/purple) will approximately correspond to the sands. In this example a model-based deterministic inversion has been used. A model-based method will give good results when the initial model is a good representation of the geology, which, for a good well tie, is always the case close to the well location. Away from well control model-based inversion will performs well if the geological layering is conformal and laterally relatively invariant, as is the case for this level in the Stratton Field.

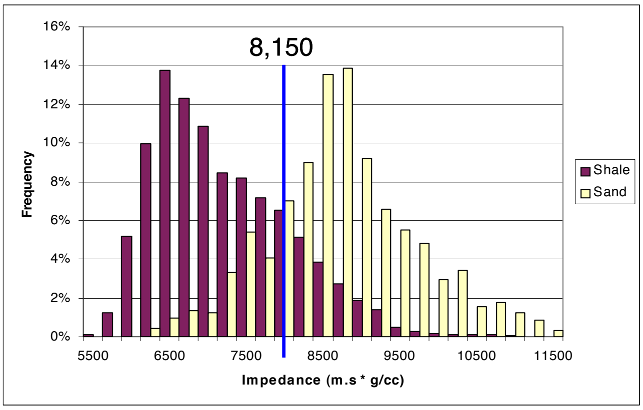

Figure 7 Histograms of acoustic impedance for sand and shales in Stratton Field with simple impedance cutoff to identify sands

A total of 100 stochastic seismic inversion realisations have been generated using a spectral domain direct stochastic seismic inversion. The spatial constraint utilises 3D anisotropic variograms. The mean of the 100 stochastic impedance realisations can be compared to the deterministic inversion in Figure 8. The same detail interval over the B46 sand interval through well 8 for four of these realisations are shown in Figure 9. Note the significant variation in possible sand geometries and also the increased dynamic range of the stochastic inversion realisations.

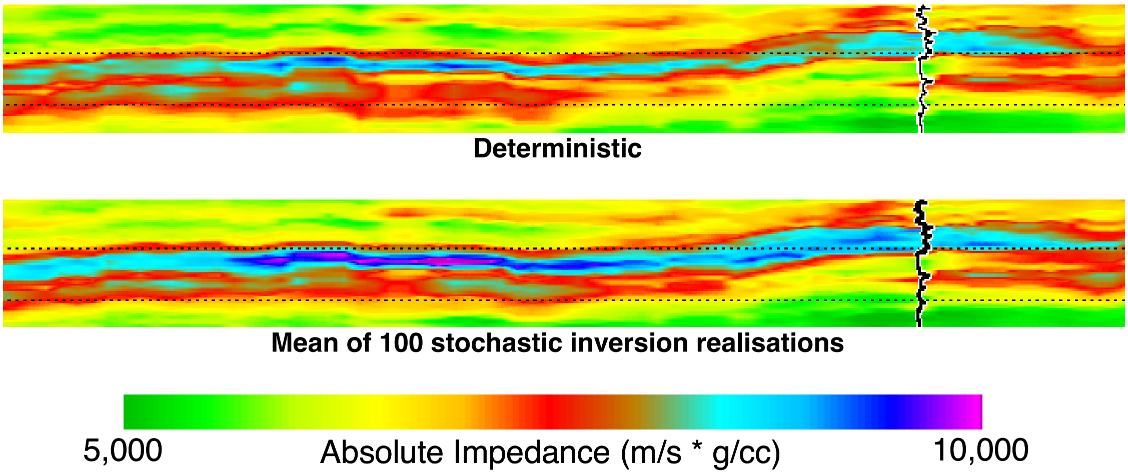

Figure 8 Detail of cross-line through well 8 at B46 sand level showing deterministic inversion result (top) and mean of 100 stochastic impedance realizations (lower). Sands indicated by blue/purple colours.

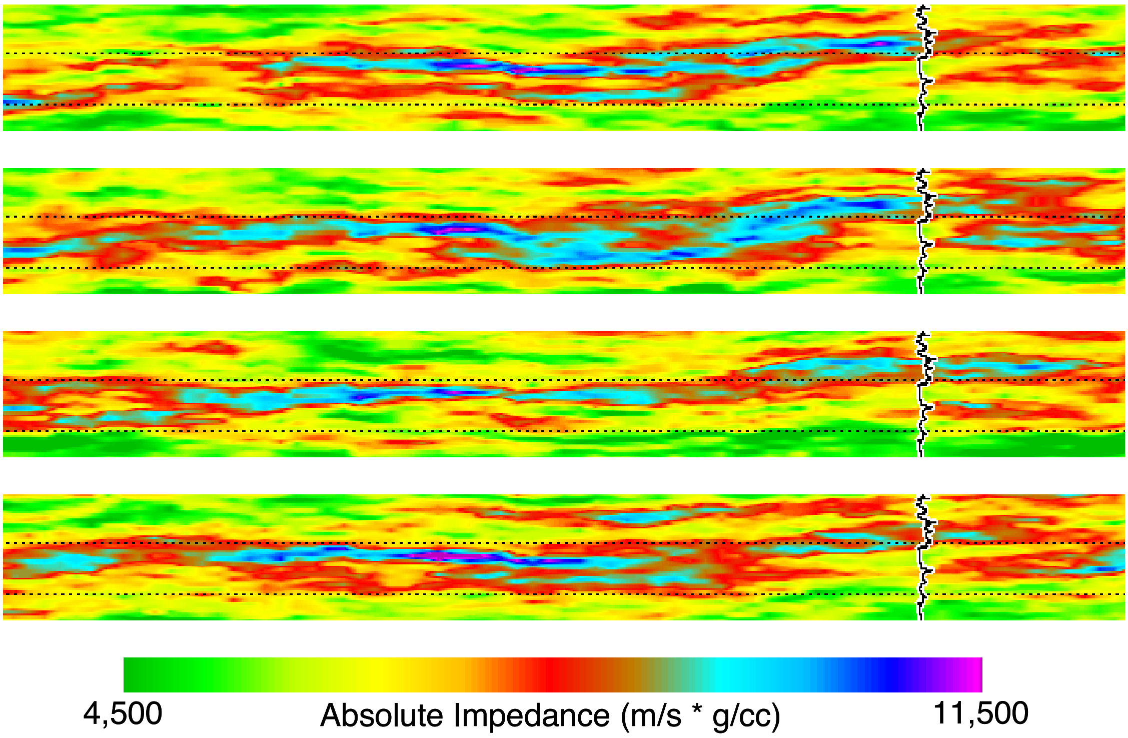

Figure 9 Four stochastic inversion impedance realisations through well 8 at the B46 sand level. Note the significant variation in possible sand geometries between the realisations. Sands indicated by blue/purple colours.

Figure 10 compares the forward convolution of each of the realisations of Figure 9 to the observed seismic traces. Despite the significant variation between the realisations shown in Figure 9, their forward convolutions are the same and a good match to the seismic. By comparing the realisations of Figure 9 to the deterministic inversion result shown in Figure 8 we can appreciate how much uncertainty is smoothed out by a deterministic inversion.

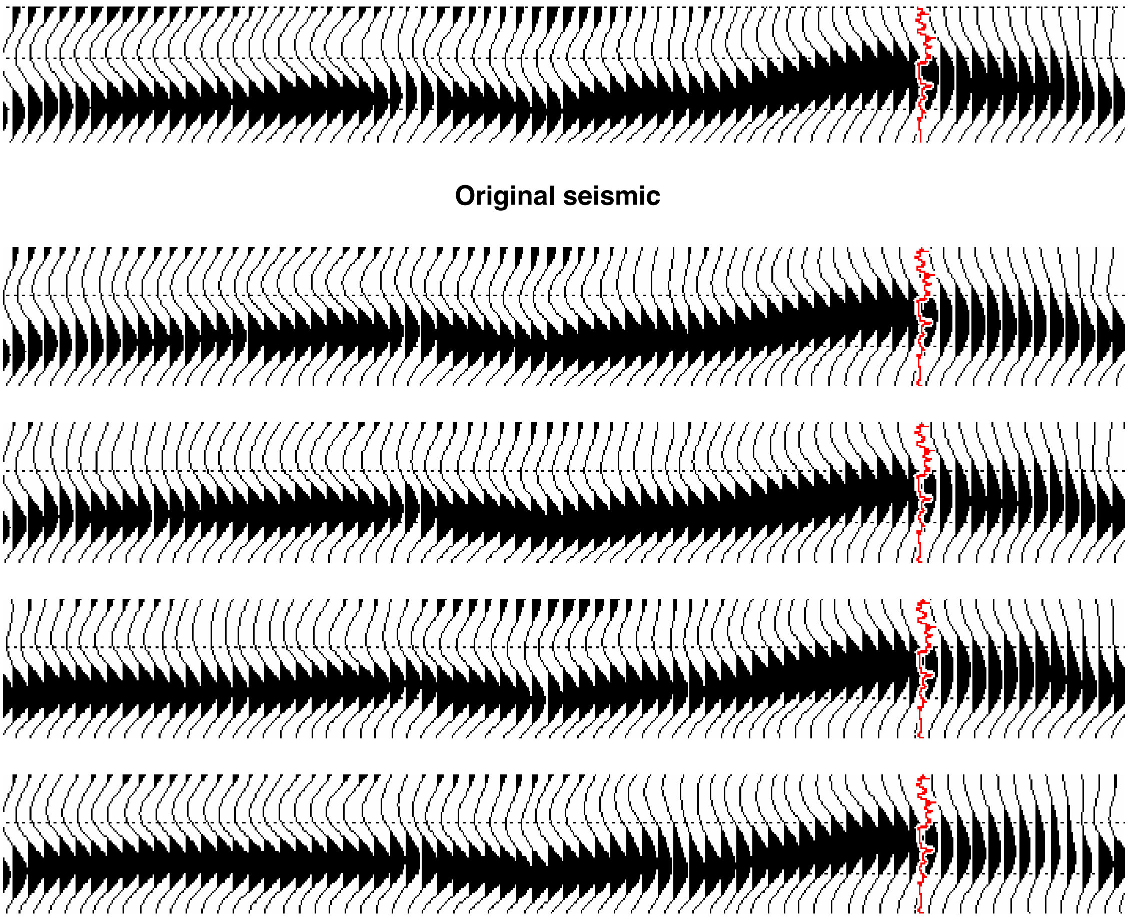

Figure 10 Lower four panels show forward convolution with wavelet of the stochastic impedance realisations of Figure 9. Comparison with original seismic (top panel) shows all realisations are conditional to the seismic response.

Using an impedance threshold of 8,150 we can compare the sand prediction and volume of sand for the deterministic and stochastic inversions. The upper panel of Figure 11 shows the B46 sand predicted for the deterministic inversion. The colour table has been adjusted to show only impedances greater than 8,150 m s^{-1} * g cm^{-3}, corresponding to our simple sand indicator identified in Figure 7. The lower panel of Figure 11 shows the sand probability for the B46 sand obtained by classifying each of the 100 stochastic impedance realisations and counting for each sample the number of realisations for the which the impedance is greater than the sand impedance threshold of 8,000 m s^{-1} * g cm^{-3}. The colour table for the sand probability display is clipped to only show greater than 50% sand probability. A simple visual comparison of the two panels suggests more sand is predicted by the stochastic inversion.

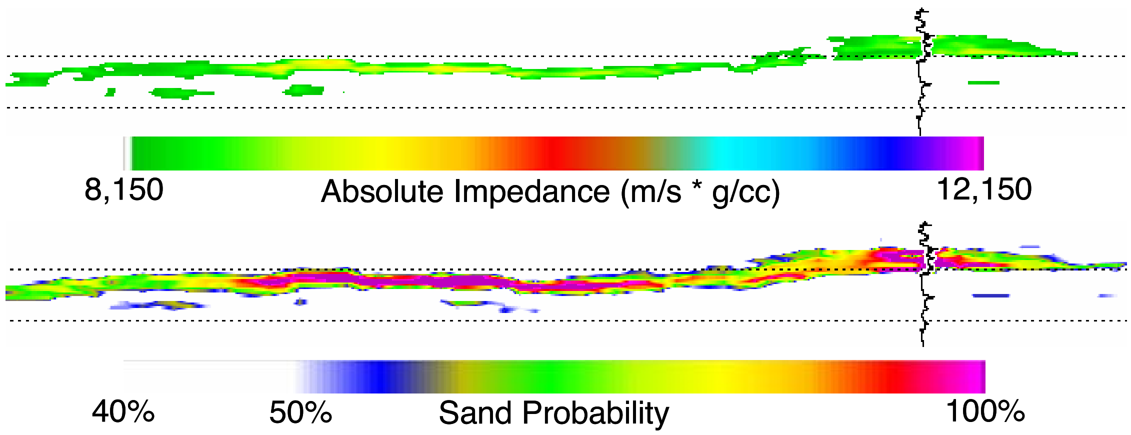

Figure 11 Predicted B46 sand through well 8 from (top) deterministic inversion and (lower) at P50 probability from stochastic inversion

Using the impedance criteria of > 8,150 m s^{-1} * g cm^{-3} to indicate sands, the net sand in the entire impedance cube has been computed for the model-based deterministic inversion and for each of the 100 stochastic impedance realisations. The deterministic inversion gives an estimate of net sand fraction of 8.5 %. The wells show an average net sand of 13.2%: the under-estimate of net sand from the deterministic inversion shows significant biased compared to the wells. In this case sands are indicated by high impedances so the volume bias when the CDF is truncated is that shown in the lower left example of Figure 5.

The cumulative distribution function (CDF) for net sand derived from the 100 stochastic impedance realisations is shown in Figure 12, the net sand estimate from each of the 100 realisations is represented by the a red dot (after sorting by their magnitude). The net sand fraction calculated from the 100 realisations varies from a minimum of 11.6% to a maximum of 15.5% with a mean value of around 13.5% net sand. This is in close agreement with the net sand fraction in the wells. The error (bias) in the net sand estimated using deterministic seismic inversion is readily apparent when overlaid on Figure 12. Note that *any* realisation gives a more accurate estimate of the net sand fraction than the deterministic inversion.

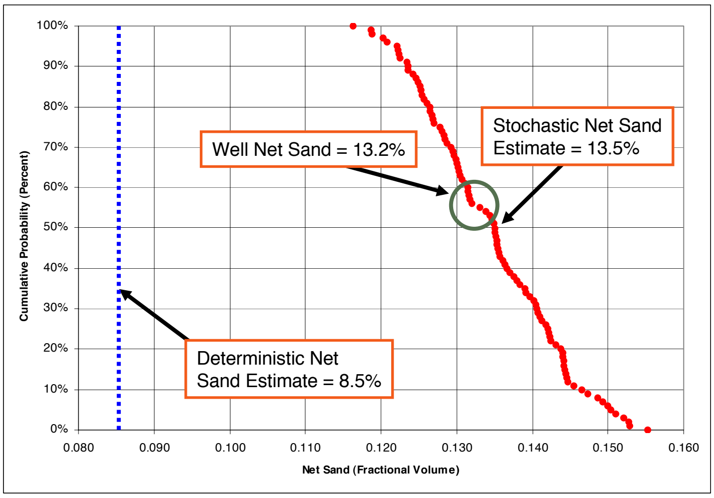

Figure 12 Cumulative distribution function of net sand fraction estimated by analysis of 100 stochastic impedance realisations (red) is consistent with well net sand fraction. Net sand fraction estimated from deterministic inversion (blue) is clearly biased and shows significant underestimate of net sand.

A further point to note is that the CDF of net sand derived from 100 realisations is not particularly smooth and this would suggest that 100 realisations is barely adequate in this case to define the range of uncertainty. This observation is consistent with the recommendations of Samson et al described earlier.

The upper part of Figure 13 shows the P50 sand thickness map for the B46 reservoir interval derived from the stochastic seismic inversion. The north-south seismic cross-line (coloured red) indicates the location of the previous shown seismic displays intersecting well 8. An interpretation suggests an east-west fairway of amalgamated channels in the southern half of the map, perhaps with evidence of crevasse splay. A single thinner channel, again oriented east-west, is identified in the north part of the display.

In the lower part of Figure 13 the interpretation made using the P50 sand thickness probability map from the stochastic seismic inversion has been overlaid on the B46 sand thickness map derived from the deterministic inversion. Note (a) thick, resolvable sand is generally correctly predicted, (b) thinner sands are under-predicted and (c) the thin northern channel is not predicted.

A 3D voxel view (see Figure 14) representing a different sand in the 3D volume clearly shows the difference in the proportion of net sand predicted from the deterministic inversion and at the P50 or greater probability from the stochastic inversion.

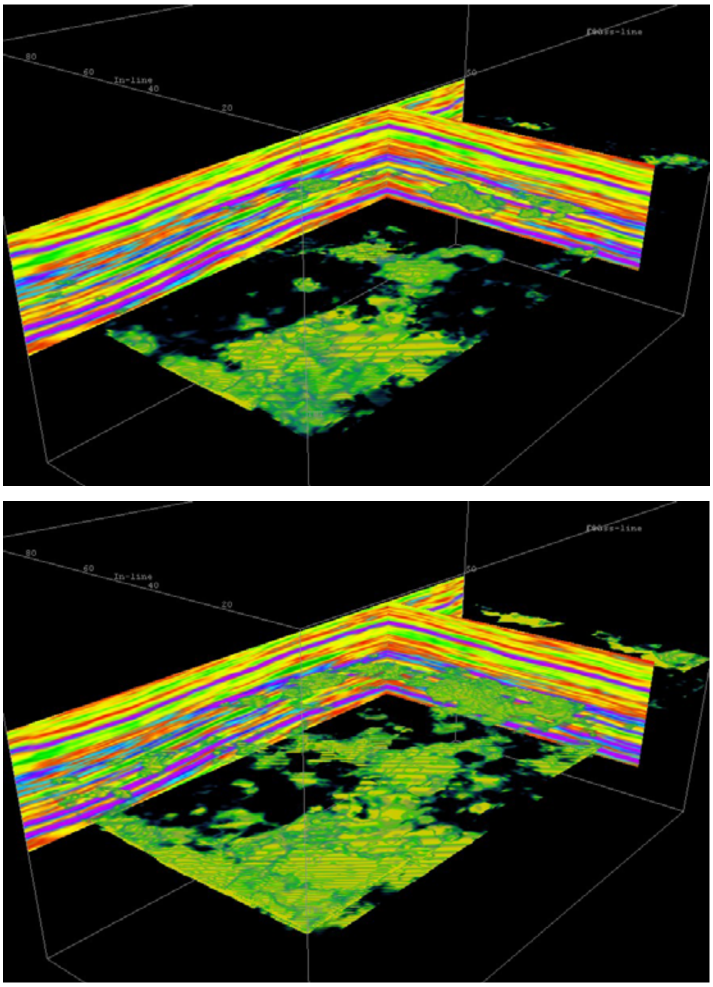

Figure 14 3D voxel representation showing sand predicted from deterministic inversion (top) and at P50 probability from stochastic inversion (below). Deterministic inversion clearly underestimates net sand in this example

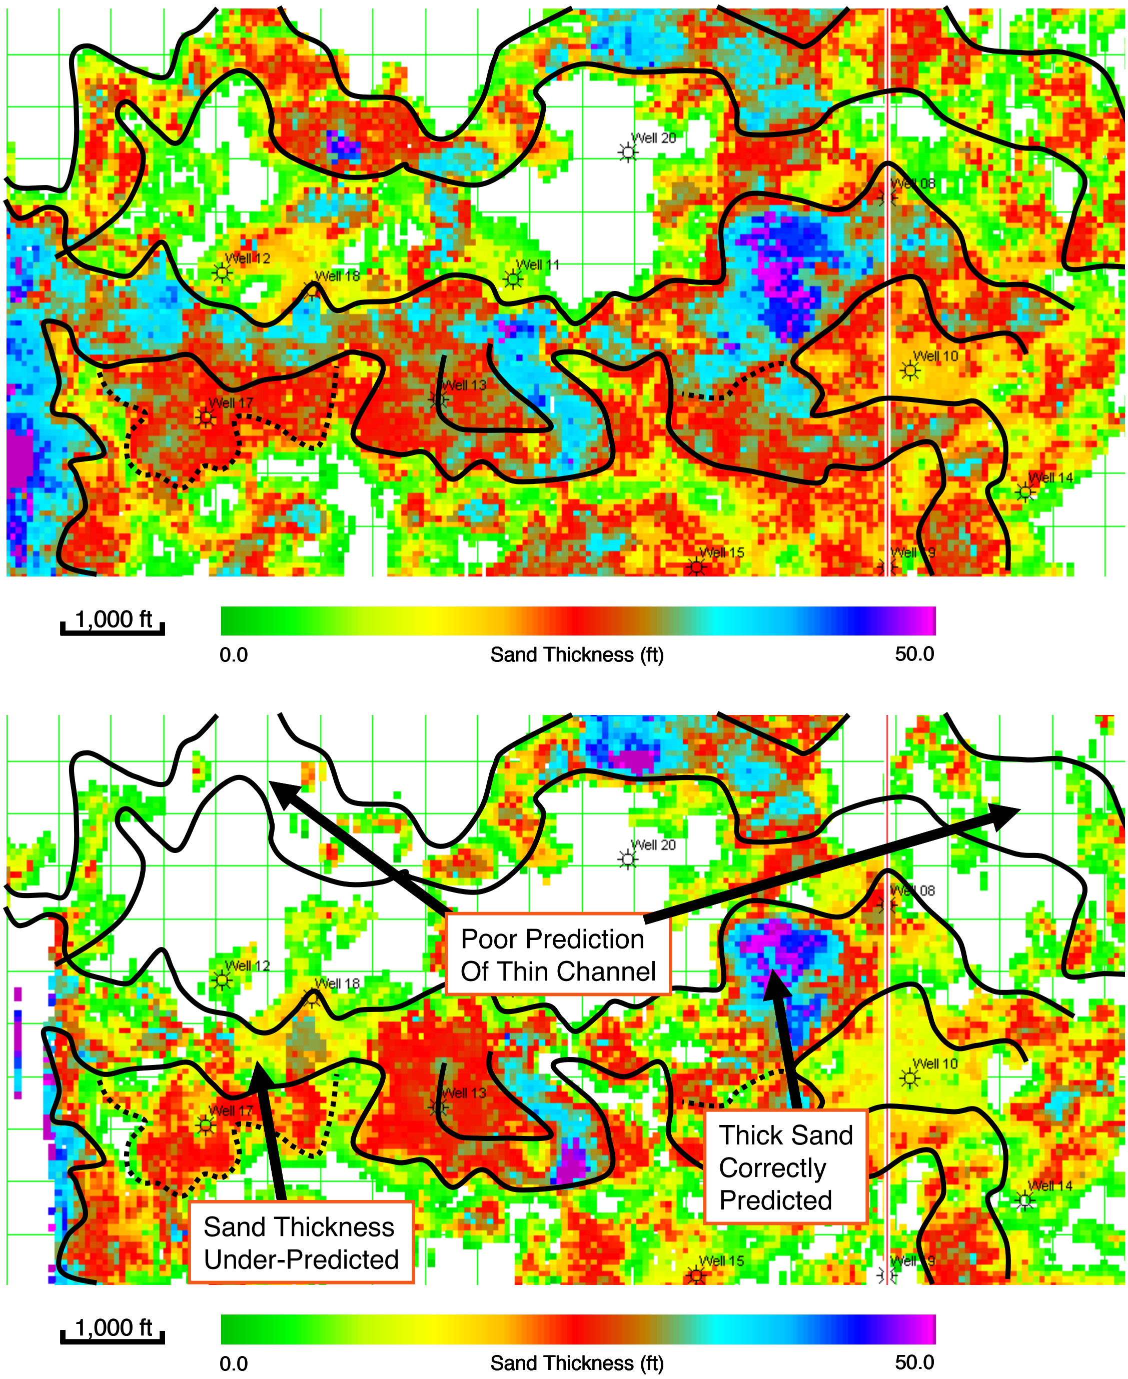

Figure 13 Upper display shows B46 sand thickness map for a P50 sand probability from 100 stochastic seismic inversion realisations. Note thin channel in running east-west across northern area of map. Lower display shows sand thickness for same interval estimated using deterministic inversion. Note overall less sand with poor prediction of thin channel.

Reservoir modelling

The use of deterministic seismic inversion to generate quantitative seismic impedance estimates is becoming common as an input to reservoir modelling. Determinisitic seismic inversion combines data from wells and seismic to create a broad bandwidth impedance model of the earth. However, we must always remember that the seismic resolution is at a coarser scale than the cell size of most reservoir models and so the seismic is usually informing the model of variations in the average properties over a zone, zones or part zone, depending on the actual zone thickness and seismic resolution.

In the vertical direction, the typical averaging from seismic inversion is over a thickness of perhaps 20 m (10 – 50m, depending on resolution). In the horizontal direction, seismic measurements (again depending on resolution) average over a typical 50 x 50 m area. It is important to note that the seismic time sample rate and trace spacing do NOT define the resolution. Resolution is determined by frequency content and bandwidth.

Seismic data is bandlimited. In particular, it does not contain low frequencies and therefore absolute impedances cannot be recovered directly from the seismic trace. All inversion schemes with an absolute impedance output require a low frequency model or constraint. The low frequency scalar is usually obtained from interpolation of well data, stacking velocities or a combination of these. After inversion the low frequency model is embedded in the deterministic inversion. Artefacts in the low frequency model manifest themselves as equivalent artefacts in the deterministic inversion. Conditioning a reservoir model using this type of data is equivalent to conditioning to an impedance map of the wells. For this reason deterministic inversions should NOT be used to condition reservoir models.

The output from a deterministic inversion may be used to constrain a reservoir model if the low frequency component is filtered out to remove the influence of the wells from deterministic inversion. The influence of the wells will be re-introduced in the reservoir modelling algorithm. Alternatively, a seismic attribute or relative impedance (such as coloured inversion) could be used.

The ideal inversion to constrain a reservoir model would be impedance realisations from a stochastic inversion. However, care should be exercised as some stochastic inversion schemes do not output realisations of impedance but instead probability surfaces at P15 or P85 of properties such as porosity. These should not be used to constrain reservoir models as they greatly exaggerate the range of uncertainty: the probability value refers to the point probability, not the whole surface. To have all values on a surface at probability of P15 or P85 is a vanishingly small probability. A P50 surface is equivalent to the output from a deterministic inversion.

Our ideal for constraining a reservoir model with 3D seismic data would be to have the model fully consistent with pre-stack seismic amplitude information. Currently there are several possible approaches considered for solving this problem including 1D probability transforms, stochastic seismic inversion to generate impedance realisations and stochastic inversion schemes based on facies and petrophysicial properties where the seismic constrain is incorporated through a rock physics model.

So far we have only considered simulation of impedance. Some stochastic seismic inversion algorithms are more ambitious, simulating the rock or facies type as well as impedance properties. Such methods are attempting to move the stochastic seismic inversion process closer to the reservoir modelling process, by constructing a facies and reservoir property model which is conditional to the seismic. This is an even more demanding implemention and requires an additional facies simulation step as well as a rock physics transform to convert reservoir properties such as porosity and saturation to impedance estimates prior to conditioning to the seismic.

One option to improve computational efficiency is to ignore the spatial constraint and generate only 1D simulations. This allows the rock property uncertainty to be examined in more detail and speeds up the generation of realisations, but the realisations from this approach do not contain spatial information, so the results must be used in the same manner as deterministic inversion, albeit with an improved uncertainty estimate attached.

The proprietary method Promise, developed by Shell, is often used in this manner, although it also has a 3D spatial constraint capability. An OpenSource algorithm based on the same kind of approach is available called Delivery (see Gunning and Glinsky, 2004)

The idealised goal of stochastic seismic inversion, partially implemented by the facies and rock property approaches such as Promise or Delivery, is to construct a stochastic reservoir model for which the forward convolution via a rock physics model and wavelet is consistent with the 3D seismic volume. The ideal workflow is illustrated in Figure 15. A number of attempts have been made to implement this type of workflow with stochastic seismic inversion. The forward model aspect is straightforward, but the updating of the reservoir model is not. Facies based stochastic seismic inversion schemes are attempting to replace the current reservoir modelling steps and this can give rise to compatibility problems when trying to include these outputs within existing reservoir modelling schemes.

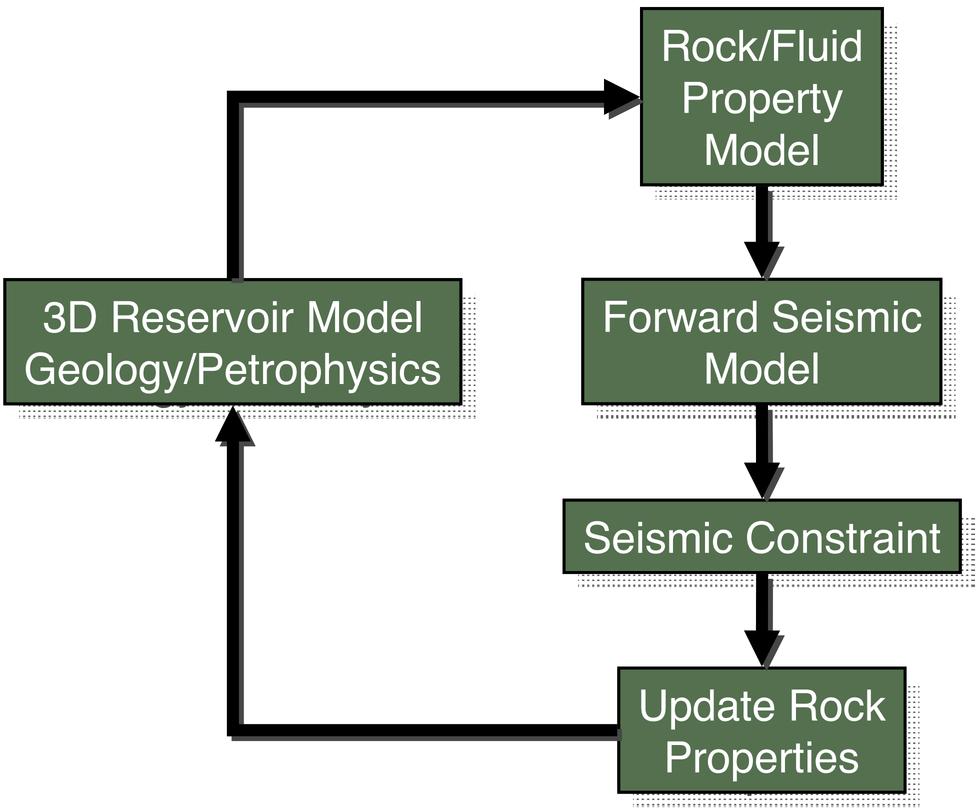

Figure 15 Idealised reservoir modeling workflow to include full seismic constraint

A further difficulty is the incorporation of different spatial constraints on the facies geometries. A spatial relation defined by a variogram can be implemented but the variogram is a two-point statistical measure. As such it cannot describe complex geometries associated with geological facies such as channels and bars. This limitation is why reservoir models containing these geometries are often built using object models. Object models are very difficult to constrain to seismic volumes.

Current reservoir modelling algorithms for incorporating seismic impedance data are generally poor and fail to account for the scale change between the reservoir model cell dimensions and the resolution limitations of seismic inversion schemes. This is especially a problem for deterministic inversion where the smooth impedance output may extend over multiple layers or zones within the model. Conditioning directly to this type of impedance data without accounting for the scale change (known as the *support effect* in geostatistics) is a significant limitation of current reservoir modelling schemes. Stochastic seismic inversion to impedance using a high sample rate close to the cell size in the reservoir model and then conditioning successive reservoir model realisations with different stochastic inversion realisations is currently the most practical approach.

Summary

The objective of any estimation procedure is to predict the value of an attribute at an unmeasured location. Mapping and seismic inversion are estimation procedures. All estimation procedures obtain accuracy by averaging or smoothing based on an error minimisation criteria such as least squares. Methods based on minimisation are commonly referred to as “deterministic”.

The output of a deterministic seismic inversion is a relatively smooth (or blocky) estimate of the average impedance. Deterministic seismic inversion suffers from a number of limitations. These arise because of the limited bandwidth of the seismic data. Missing low frequencies contain the critical information concerning the absolute values of impedance. This means that it is impossible to recover the absolute impedance values from a seismic trace. Missing high frequencies mean the best we can estimate is a local (blocky or smooth) average of the impedance.

Because of its smoothness deterministic inversion is generally unsuited for constraining reservoir models used for volumetric calculations, estimation of connectivity or fluid flow simulation. Deterministic seismic inversion will perform best when the:

(a) geological system is comprised of relatively thick layers (ie with little vertical variation of impedance within each layer)
(b) stratigraphic layering is conformal, with slow lateral thickness changes
(c) lateral impedance variation within a layer is stationary, meaning that there is no significant trend in the impedance variations.

The missing information which causes us to smooth in estimation is the uncertainty in the estimation. The uncertainty arises due to the limited bandwidth of our measurements (for example, seismic data) and insufficient sampling (for example, of the inter-well space).

In geostatistical simulation the missing variability is simulated using a Monte Carlo method based on a random number generator. For a spatial problem such as stochastic seismic inversion the missing variability must be added systematically, consistent with the spatial correlation function of the data.

The average of all the possible non-unique stochastic impedance realisations is the deterministic seismic inversion. A stochastic seismic inversion produces a suite of possible impedance realisations which are valid geostatistical simulations. A valid geostatistical simulation must be (a) conditional to our measured data points; (b) reproduce the histogram of the data and (c) honour a spatial correlation function, such as a variogram.

Analysis of the multiple impedance realisations generated from a stochastic seismic inversion allows more accurate estimation of volume, connectivity and uncertainty than is possible with a deterministic inversion. A stochastic seismic inversion can be run at any required output sample rate. Note this does not imply stochastic has a higher *resolution* because the resolution is controlled by the frequency content and bandwidth of the seismic conditioning data.

Finally it should be noted that current reservoir modelling algorithms are for incorporating seismic impedance data are generally poor and fail to account for the scale change between the reservoir model cell dimensions and the resolution limitations of seismic inversion schemes.

References

Doyen, P.M., 1988, Porosity from seismic data: a geostatistical approach. Geophysics 53(10) pp 1263 – 1275

Galli, A. and Menuier, G, 1987, Study of a gas reservoir using the external drift method. In Geostatistical Case Studies edited by Matheron, G and Armstrong, M, Reidel Publishing Company

Gunning, J. and Glinsky, M.E., 2004, Delivery: an open-source model-based Bayesian seismic inversion program. Computers and Geosciences 30 (6) pp 619 - 636

Haas, A. and Dubrule, O., 1994, Geostatistical inversion – a sequential method of stochastic reservoir modelling constrained by seismic data. First Break 12 (11) pp 561 – 569

Levey, R.A., Hardage, R.A., Edson, R. and Pedleton, V., 1994, 3D seismic and well log data set: Fluvial Reservoir Systems, Stratton Field, South Texas. Published by Bureau of Economic Geology, University of Texas, Austin, Texas, 78713-7509, USA.

Oldenburg, D.W., Scheuer, T. and Levy, S., 1983, Recovery of the acoustic impedance from reflection seismograms. Geophysics 48 (10) pp 1318 – 1337

Samson, P., Dubrule, O. and Euler, N., 1996, Quantifying the impact of structural uncertainties on gross-rock volume estimates. SPE 35535

Scola, V. and Ruffo, P., 1992, An application of geostatistical simulation to probability estimation of prospect geometry. Abstract B031 in proceedings of 54th EAEG, Paris.