Integration of uncertain subsurface information into multiple reservoir simulation models

Michael E. Glinsky and Bruce Asher, BHP Billiton Petroleum, Houston, USA
Robin Hill, Mark Flynn, and Mark Stanley, BHP Billiton Petroleum, Perth, Australia
James Gunning, CSIRO Petroleum, Clayton, Victoria, Australia
Troy Thompson, DownUnder GeoSolutions, Perth, Australia
Jerome Kalifa and Stephane Mallat, Let It Wave, École Polytechnique, Palaiseau, France
Chris White, Louisiana State University, Baton Rouge, USA
Didier Renard, L'École des Mines, Fontainebleau, France

Successful appraisal and development of oil and gas fields requires the integration of uncertain subsurface information into multiple reservoir simulation models. This information includes seismic data, various types of well data, and geologic concepts. Over the past five years, a workflow has been developed by various organizations in conjunction with BHP Billiton Petroleum. This distinctive approach focuses first on building mesoscale reservoir models that can be constrained by seismic data (typically with a resolution up to the stratigraphic seismic loop scale, see Prather et al., 2000), then introducing the finer scale geologic concepts and well data needed for reservoir simulation models (stratigraphic first- and second-order subseismic scale, where each order is about a factor of three in size) via a downscaling step that honors mesoscale model constraints. Uncertainty and correlations of the well and seismic measurements are always taken into account. In fact, they are necessary to be able to combine the various measurements. Bayesian probabilistic techniques are used extensively in this process. The result is an ensemble of reservoir simulation models consistent with all subsurface information.

The application to Stybarrow Field, in the Carnarvon Basin of offshore Western Australia, will be used as an example of this workflow. This workflow starts with a conventional correlated wavelet extraction and sparse spike inversion. The sparse spike inversion gives a preliminary estimate of net rock volume and fluid probabilities. Although it does not consider uncertainty, multiple layer seismic interference or many well constraints, it helps build a layer-based model framework for subsequent steps where these influences are considered.

Because we will be estimating the uncertainty, we need to know what the uncertainty of the seismic data is. To obtain this, the wavelet derivation needs to produce an estimate of the seismic noise—that is, the part of the seismic data that does not correlate with the synthetic of the well log. We used a probabilistic wavelet derivation based on Bayesian concepts that predicts the noise level. Given the wavelet with uncertainty, we then perform a trace-based probabilistic model-based inversion which combines the seismic data with well measurement constraints to give average values, uncertainties, and correlations in important properties such as net reservoir sand, net-to-gross sand, and fluid content.

Several significant technical challenges remain. First, the results must be "massaged" onto the reservoir simulation grid (which is commonly nonuniform in space). Second, transverse spatial correlations that are geologically realistic must be enforced, while eliminating short scale fluctuations that are artifacts of the trace-based inversion. Geostatistical techniques are used to introduce this lateral correlation while respecting the nontrivial correlations produced by the model-based inversion. The final challenge is to "decorate" the model with the stratigraphic first- and second-order subseismic structure needed to allow the reservoir simulation to assess the flow effects of subseismic heterogeneity. The subseismic structure is based on geologic concepts and is constrained by well and seismic information. An "enforcement" step then slightly deforms the decorated model to match the gross thickness and net sand estimates from the inversion. Results of all steps will be shown for Stybarrow Field.

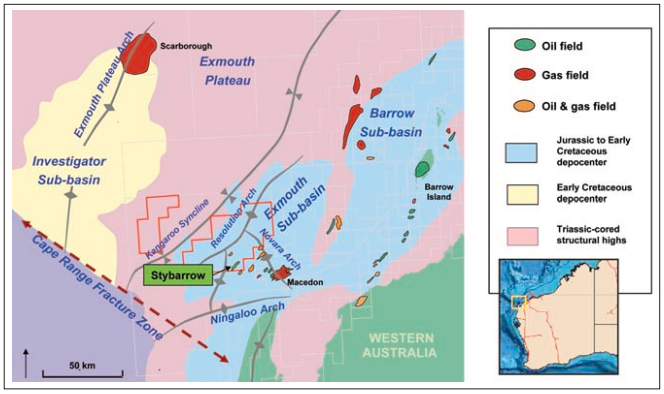

Figure 1. Map showing the location of Stybarrow Field compared to other fields in the region.

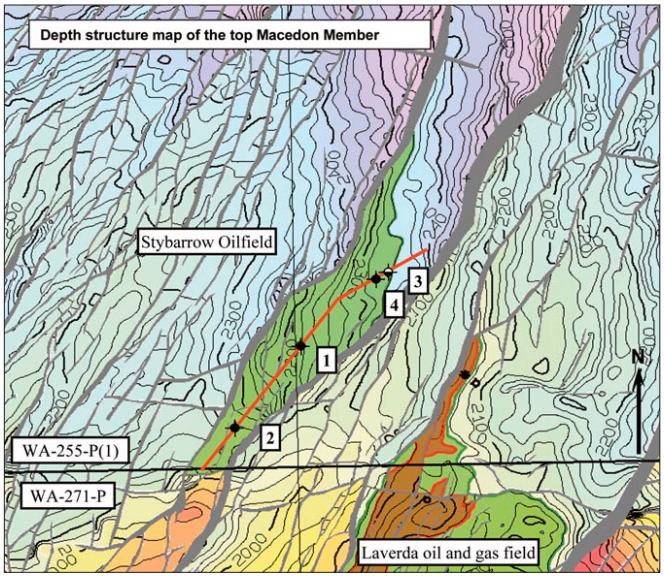

Figure 2. Closeup of Stybarrow Field showing the depth structure map and faults.

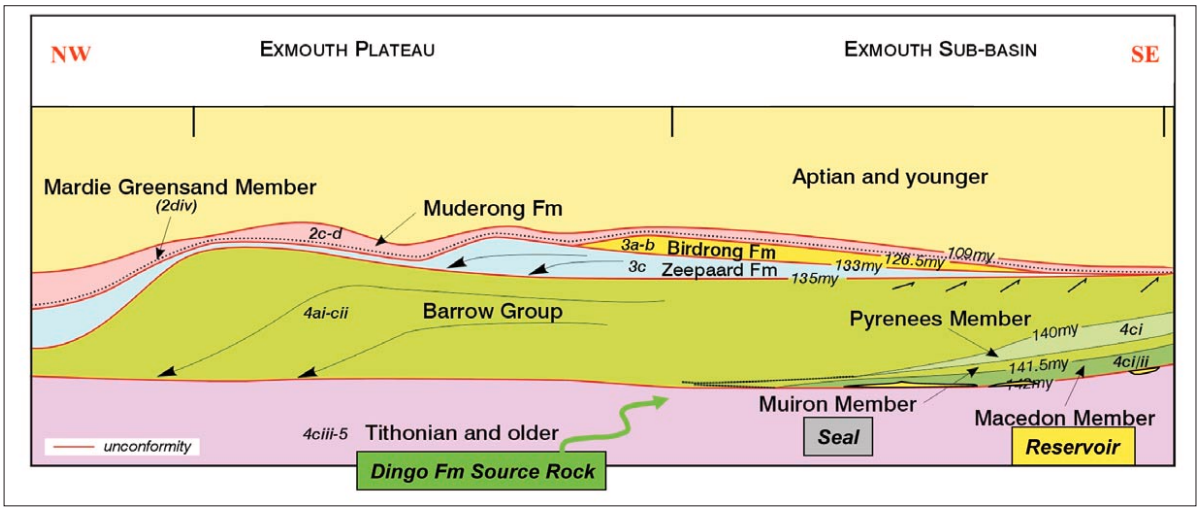

Figure 3. Geologic cross-section showing the context of the Stybarrow reservoir, seal, and source.

Description of Stybarrow Field. Stybarrow Field is in exploration permit WA-255-P(1), some 135 km west of Onslow, offshore Western Australia. The water depth at the location is approximately 800 m. The field lies near the southern margin of the Exmouth sub-basin within the overall Carnarvon Basin (Figure 1).

Oil is trapped in the Early Cretaceous, Berriasian age turbidite and debris flow sandstones deposited on a relatively shallow passive margin slope. The Stybarrow structure comprises a northeast-to-southwest tilted fault block forming a terrace within the westward plunging Ningaloo Arch (Figure 2). The intersection of SW-to-NE and E-to-W trending normal faults establishes an elongate, triangular trap forming structural closure to the southwest. The structure dips from the SW to the NE at about 5°. Top, base, and bounding fault seals are provided by claystones and siltstones of the overlying Muiron Member of the Barrow Group and mudstones of the underlying Dupuy Formation. Oil is sourced from claystones of the Dingo Formation (Figure 3). More information about the field can be found in Ementon et al. (2004).

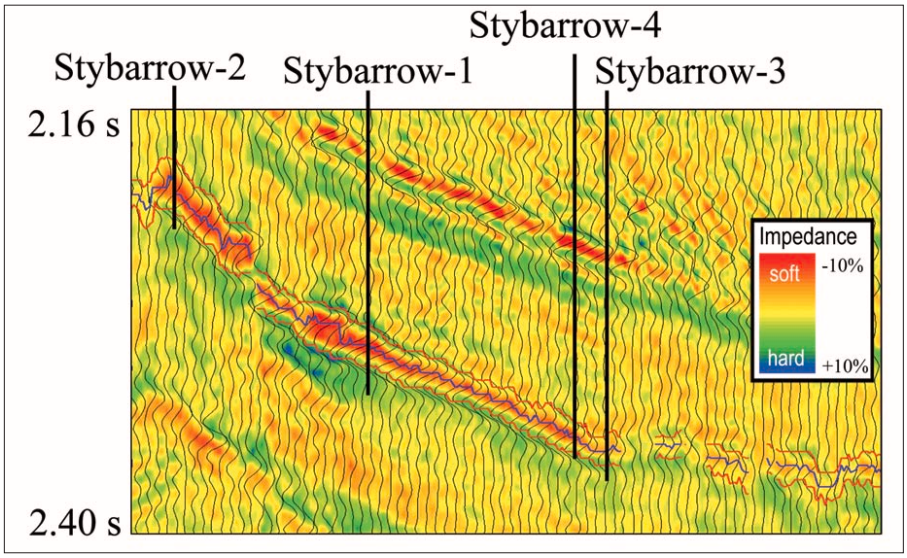

Figure 4. Cross-section through the four Stybarrow wells. Seismic data are shown as wiggle traces (reflection coefficient phase, right kick is a soft reflection). Acoustic impedance of the sparse spike inversion is shown in color. The units are those of reflection coefficients. Three events are shown. The two in red are the secant points used as a baseline for both amplitude and area calculations. They are shown as P1 and P2 in Figure 5. The third, shown in blue, is the peak used in the amplitude calculation, shown as P3 in Figure 5.

Use of sparse spike inversion and the correlation wavelet. We started by derivation of a wavelet from well log information using traditional correlation techniques. The wavelet was the basis for a sparse spike inversion. The net sand was estimated by extraction of the secant area (Figures 4-6) of the inversion normalized by the factor $V_{sand}/2R_{sand}$, where V_{sand} is the interval velocity of the reservoir sand-shale mixture and R_{sand} is the modeled average reflection coefficient of pure sand. Our sparse spike inversion outperforms commonly used sparse spike inversions in estimating net sand for a benchmark wedge model with a shale channel (Figure 7). We also estimated the hydrocarbon probability by a Bayesian update of a constant probability map using the observed reflection coefficients derived from a secant amplitude extraction and the estimate of the range of possible reflection coefficients for the target sand (Glinsky et al., 2004). The results are shown in Figure 8.

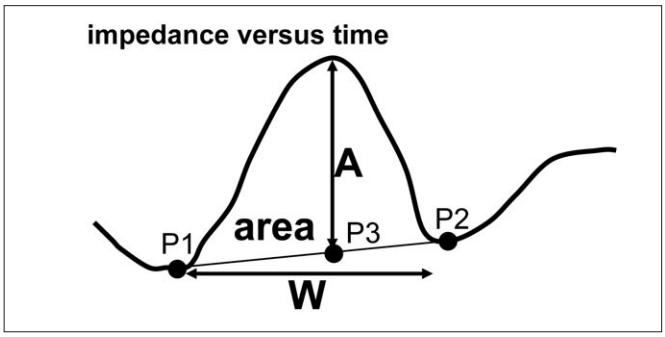

Figure 5. Diagram showing definition of secant amplitude and secant area.

What we do not have from this analysis is uncertainty, inclusion of the well constraints, and removal of the effects of multiple layer tuning. Although this analysis is sufficient in many cases for exploration screening, it is not good enough for making appraisal and development decisions. Because of this, a more detailed analysis was done, a probabilistic model-based inversion. The starting point for this is a structural

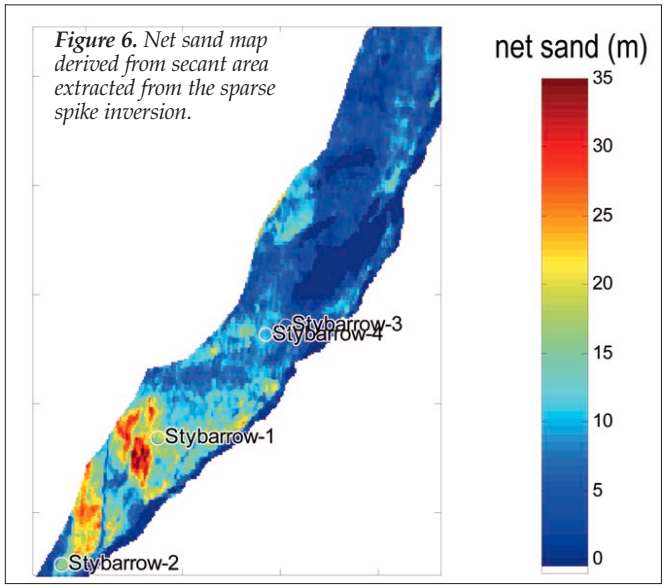

Figure 6. Net sand map derived from secant area extracted from the sparse spike inversion.

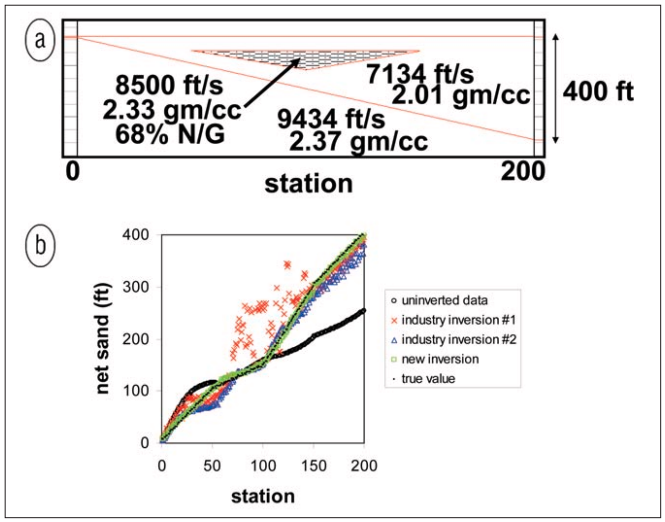

Figure 7. *Benchmark of net sand derived from secant area extracted from sparse spike inversion. (a) Model used to benchmark sparse spike inversions. (b) Comparison of net sand estimates calculated from the uninverted data, two industry standard inversion programs, the new inversion method that we used, and the correct answer.*

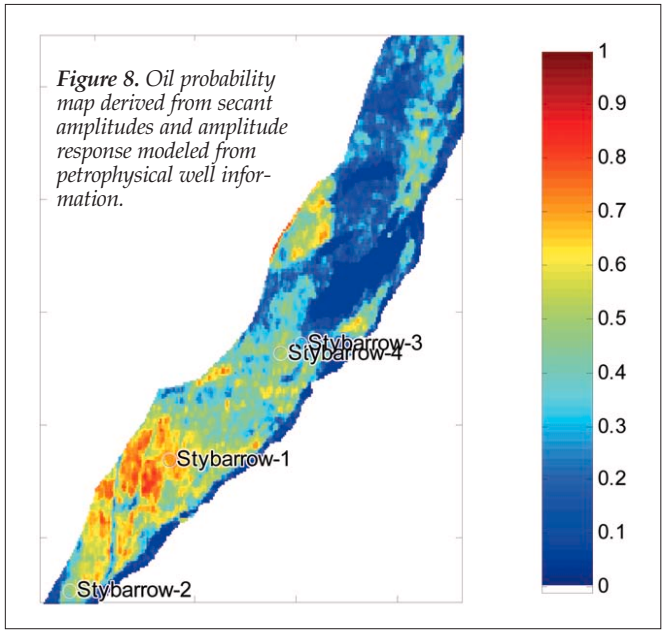

Figure 8. Oil probability map derived from secant amplitudes and amplitude response modeled from petrophysical well information.

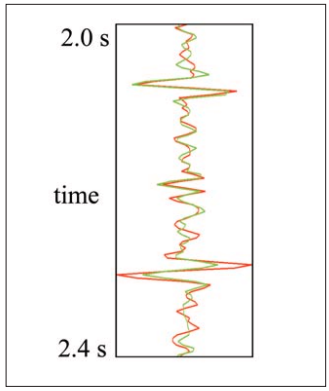

Figure 9. *Synthetic seismic with most likely wavelet (red) compared to the seismic data (green) at the Stybarrow-1 well.*

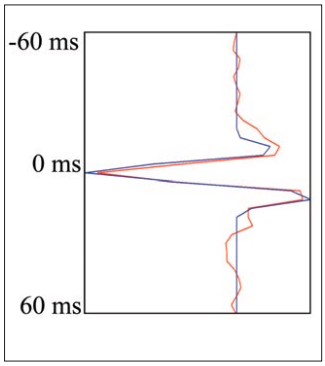

Figure 10. *Wavelet derived using standard correlation techniques (red) compared to the wavelet derived with the new probabilistic process (blue).*

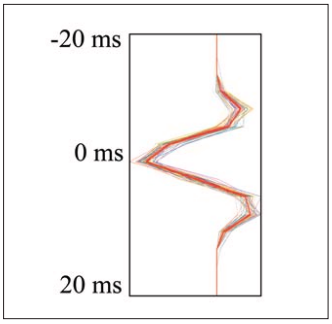

Figure 11. *Most likely wavelet (thick red) compared to the ensemble of possible wavelets. The range of the wavelets is representative of the percentage noise.*

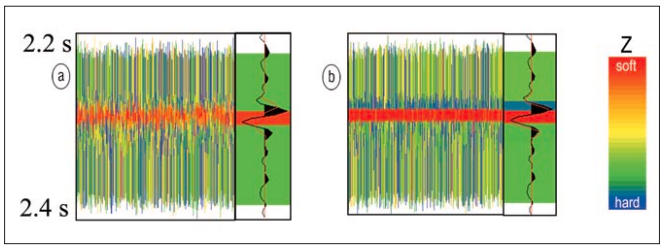

Figure 12. *Ensemble of possible models at Stybarrow-1 well compared to the average model, seismic data, and the synthetic of the average model. Shown (a) before the inversion, and (b) after the inversion. The impedance is shown as the color. The seismic data are the shaded black wiggle traces. The synthetic is the red trace.*

framework of a model that is seismically resolvable based on the sparse spike inversion. If there is not evidence of a layer on the sparse spike inversion there is little chance that it can be resolved with more sophisticated model-based inversion.

Probabilistic wavelet derivation giving critical noise level. A very important input to the model-based inversion described in the next section is the seismic noise level. The noise level is a measure of the mismatch between the synthetic seismic calculated from the well log and the seismic data. It is important that the same approximation be used in the calculation of this synthetic as will be used in the model-based inversion, because the "noise" is a process incorporating both physical noise and modeling errors caused by the approximations. We have extended the Bayesian wavelet derivation of Buland and Omre (2003) to develop a new multiple stack, multiple (possibly deviated) well tie computer program for this purpose. It optimizes over wavelets of different lengths and allows modifications to the time-to-depth mapping and lateral positioning of each well. It finds the most likely estimate of all parameters describing these variations, and their associated uncertainty (Gunning and Glinsky, 2005). The result is an esti-

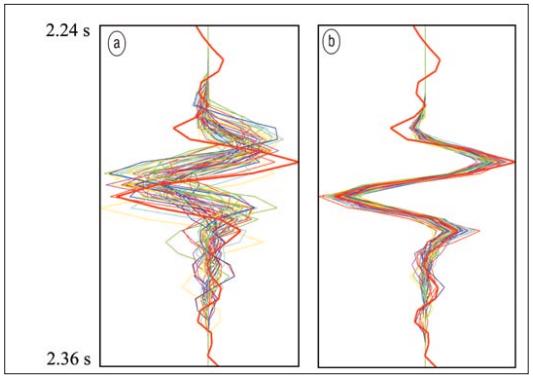

Figure 13. *The seismic data (thick red) compared to the ensemble of synthetic seismic data generated from the ensemble of models. Shown (a) before the inversion, and (b) after the inversion.*

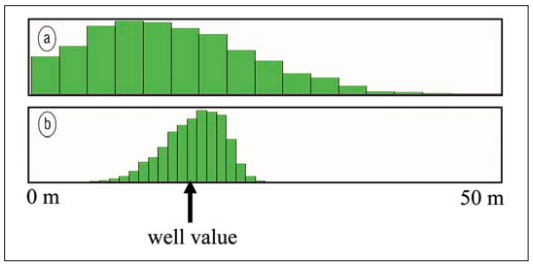

Figure 14. *Histograms of the net sand in the main sand at the Stybarrow-1 well location. Shown (a) before the inversion, and (b) after the inversion.*

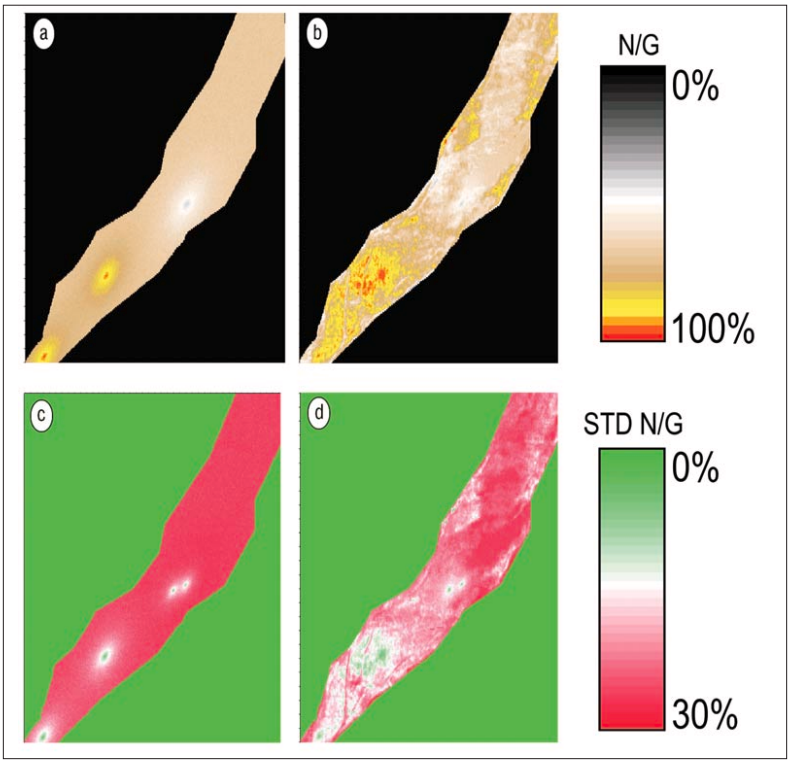

Figure 15. *N/G maps (a) before and (b) after the inversion for the main sand. The map before the inversion is kriged to be consistent with the well information. Standard deviation maps for the N/G (c) before and (d) after inversion.*

mate of the most likely wavelet along with the match of the synthetic to the seismic data (Figures 9 and 10). For Stybarrow, the wavelet was derived simultaneously for four different wells. In Figure 10 the most likely probabilistic wavelet is compared to the standard correlation wavelet.

Note that the length of the probabilistic wavelet is much shorter than the correlation wavelet. This is a very common occurrence. We have consistently found that standard practice is to use wavelets that are too long. The extra side lobes are equivalent to extra fitting parameters that can be shown to be statistically insignificant, and are best suppressed.

The postwell-tie wavelet uncertainty is displayed in Figure 11 which shows the range of possible wavelets which all generate acceptable matches to the seismic. The noise level has a comparably narrow distribution, whose mean gives an estimate of the noise level required for the inversion. For these data, the noise level is 17% of the amplitude of the peak reservoir reflector, which is a 15 dB SNR.

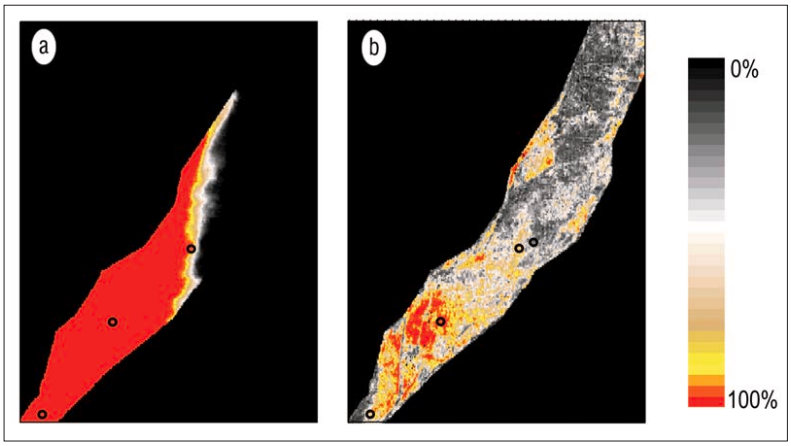

Figure 16. *Fluid probability maps (a) before and (b) after inversion for the main sand. The map before inversion was constructed using known fluids and pressures from the wells combined with depth uncertainty of 3 m. The map after inversion did not use the map shown in (a). It used a constant map of 50%, so that one can see what the seismic data alone predict for the fluid contact shown in (a).*

Probabilistic model-based inversion with uncertainty. The starting point for the model-based inversion is a layer-based model built with seismically resolvable layers. This model is based on interpreted horizons from the acoustic impedance volume as previously discussed. Typically this model is built at a resolution of the stratigraphic seismic loop scale. The geologic time span of this scale is dependent on the sedimentary deposition rates and the bandwidth of the seismic data. It is only built over the region of interest but includes nearby reflectors that may produce seismic amplitudes in the region of interest. For Stybarrow this model consisted of six layers: the upper bounding shale, a small sand above the main sand, a thin intervening shale, the main sand, a seismically hard bioturbated zone below the main sand, and a lower bounding shale. These layers were identified on the acoustic impedance volume shown in Figure 4.

The wells in this field and other relevant wells in the Exmouth sub-basin were analyzed to produce standard rock physics correlations of the acoustic properties (compressional velocity, shear velocity, and density) for the end member sands and shales. These correlations are expressed as a set of linear relationships with uncertainty. The fundamental properties of each layer are the net-to-gross ratio (N/G), layer top and base, and the type fluid in the pore volume of the sand. The mathematical machinery is embedded in the inversion to calculate the synthetic seismic given the basic properties for each layer (Gunning and Glinsky, 2004) as follows: (1) the acoustic properties of the end member sands and shales are sampled accord-

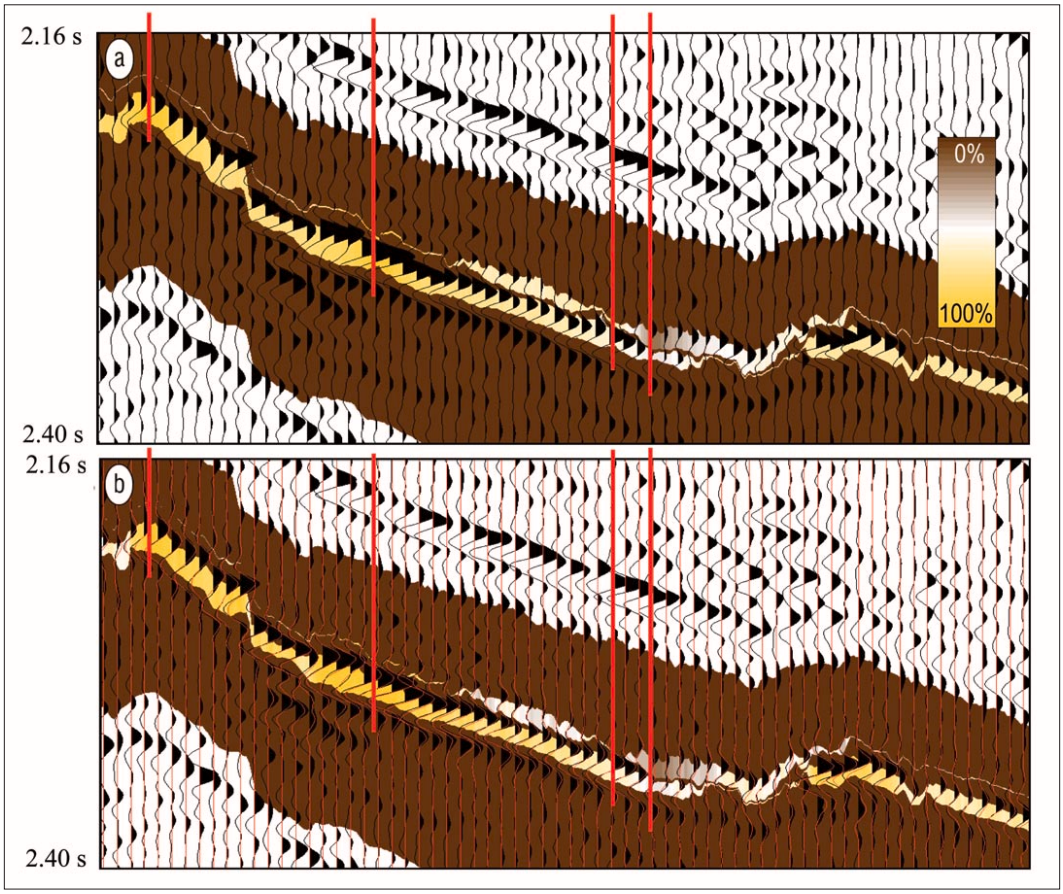

Figure 17. Cross-sections through the well locations showing the N/G of the model and the seismic data (black shaded wiggles) and the well locations (thick red lines). Shown (a) before the inversion, and (b) after the inversion. The synthetic of the model is shown as the thin red wiggles.

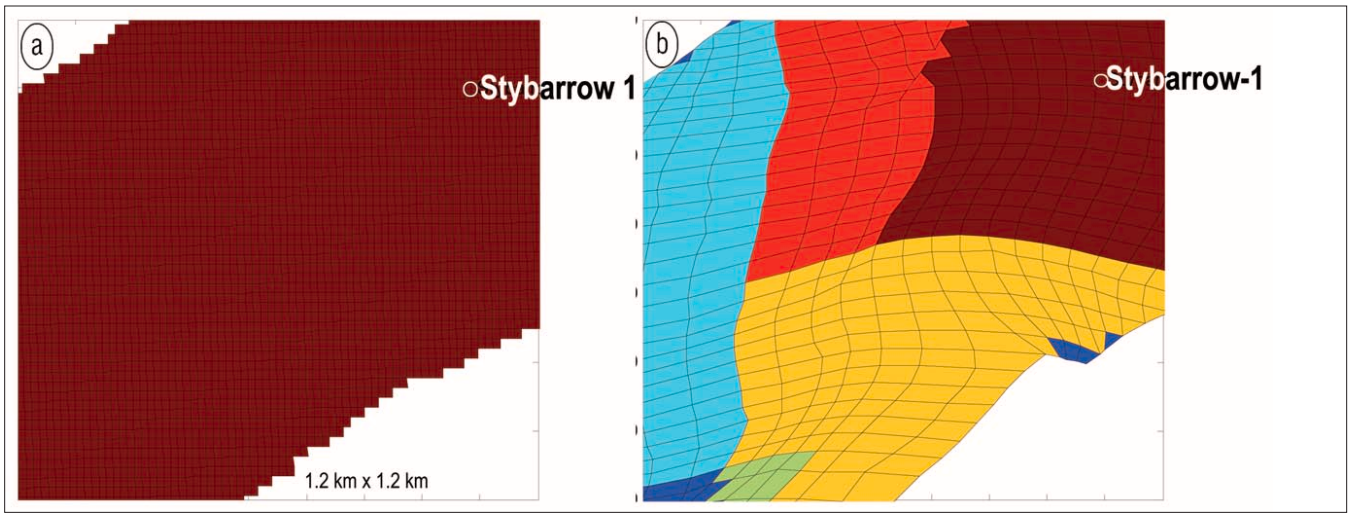

Figure 18. (a) A section of the seismic grid near Stybarrow-1. (b) The same section of reservoir simulation grid. Each fault block is a different color.

ing to the regional rock physics correlations, (2) the fluid is then substituted into the sand using the Gassmann relationship, (3) the sand and shale are mixed together with the specified N/G using the laminated mixing model (Backus average), and (4) the synthetic is then calculated for the appropriate offsets using a convolution model with reflection coefficients based on the small angle and small contrast approximations.

We want to emphasize that uncertainties are estimated for all properties that have a significant contribution to final uncertainty. This includes things like N/G, fluid type (i.e., oil, gas, brine, and low-saturation gas) acoustic fluid properties, fluid saturations, and the interpreted horizons. The inversion starts by combining all these uncertainties into an ensemble of models that sample the range of possibilities, but ignores the match of the synthetic seismic of the model to the seismic data. A set of these models before inversion can be seen in Figure 12a. The synthetics of this ensemble are shown in Figure 13a compared to seismic data. Note the lack of agreement. The seismic inversion machinery produces an ensemble of models whose synthetic seismic match the seismic data to within the estimated noise level. The result of the inversion can be seen in Figures 12b and 13b. Note the much smaller variation in the models and the much better agreement of the synthetic seismic to the seismic data after inversion.

Very important quality checks of the inversion must be performed. The first is a set of "hold one out" tests at the known well locations. The inversion is performed at the well locations without constraining the inversion to the correct values of fluid type, N/G, and gross thickness. The results of the inversion are then compared to the well data. For Stybarrow-1, some

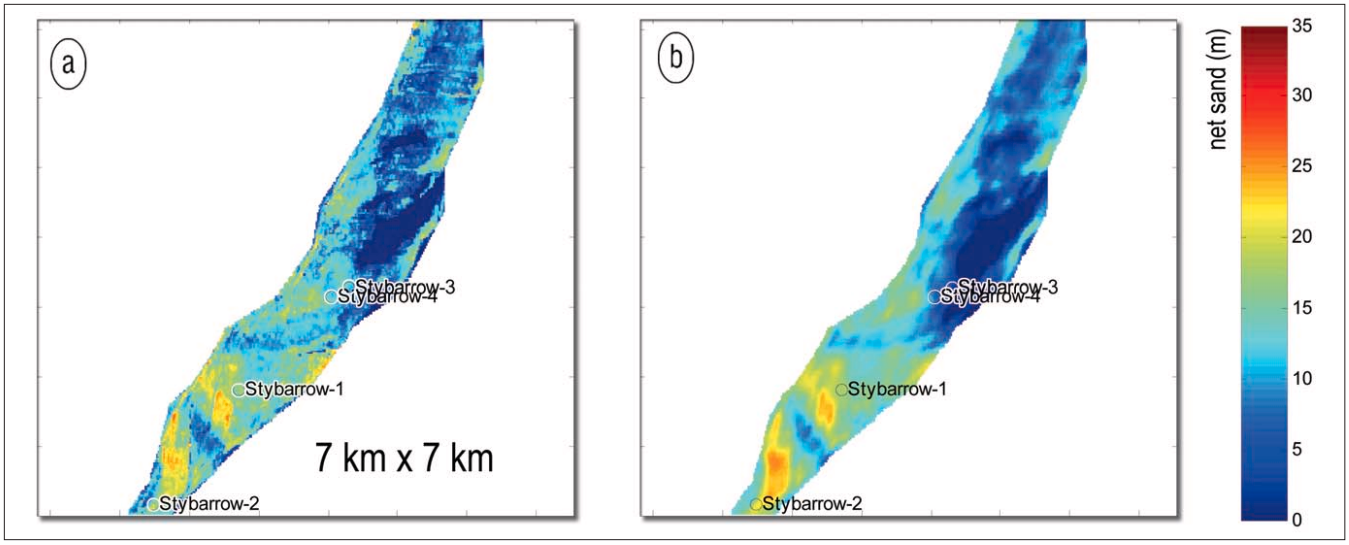

Figure 19. *Average net sand map (a) after the inversion and (b) after it has been massaged.*

Figure 20. *Realizations of the net sand maps generated by the massaging to the reservoir simulation grid. They are ordered from the least to the most net sand above the fluid contact.*

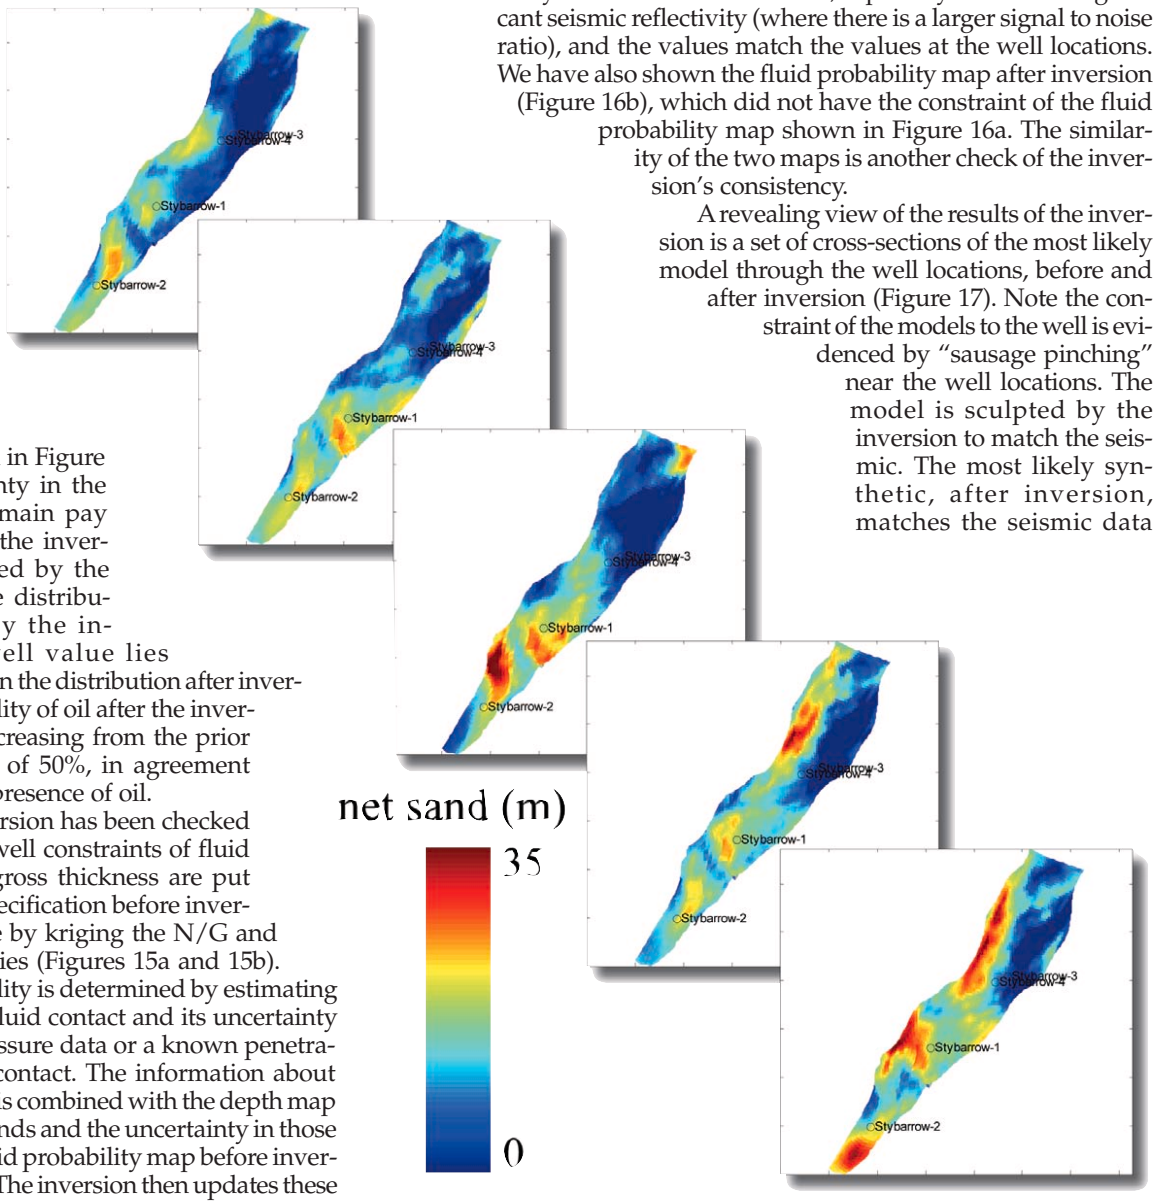

results are shown in Figure 14. The uncertainty in the net sand in the main pay was reduced by the inversion, as evidenced by the narrowing of the distribution function by the inversion. The well value lies comfortably within the distribution after inversion. The probability of oil after the inversion was 97%, increasing from the prior "agnostic" value of 50%, in agreement with the known presence of oil.

Once the inversion has been checked at the wells, the well constraints of fluid type, N/G and gross thickness are put into the model specification before inversion. This is done by kriging the N/G and thickness properties (Figures 15a and 15b). The fluid probability is determined by estimating the depth of the fluid contact and its uncertainty from the well pressure data or a known penetration of the fluid contact. The information about the contact depth is combined with the depth map of the reservoir sands and the uncertainty in those maps to give a fluid probability map before inversion (Figure 16a). The inversion then updates these properties (Figures 15c and 15d). Note that the uncertainty away from the well is reduced, especially in areas of significant seismic reflectivity (where there is a larger signal to noise ratio), and the values match the values at the well locations. We have also shown the fluid probability map after inversion (Figure 16b), which did not have the constraint of the fluid probability map shown in Figure 16a. The similarity of the two maps is another check of the inversion's consistency.

A revealing view of the results of the inversion is a set of cross-sections of the most likely model through the well locations, before and after inversion (Figure 17). Note the constraint of the models to the well is evidenced by "sausage pinching" near the well locations. The model is sculpted by the inversion to match the seismic. The most likely synthetic, after inversion, matches the seismic data

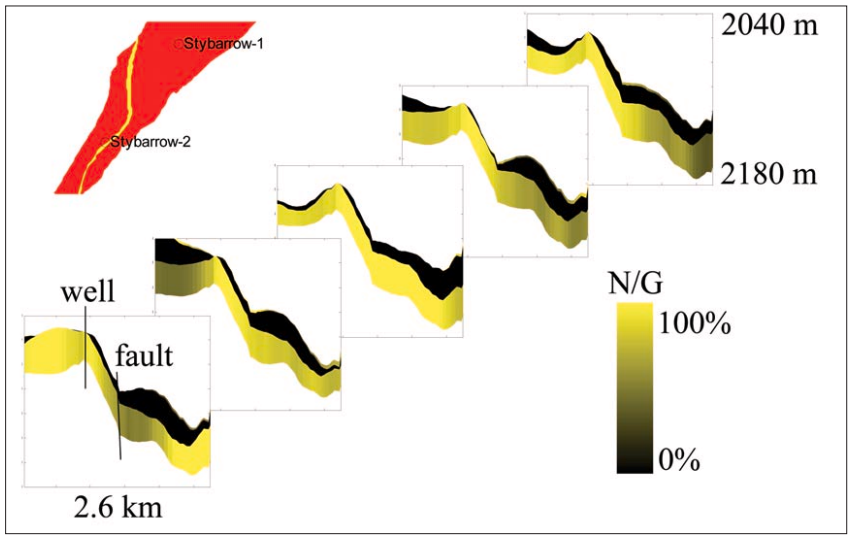

Figure 21. *Depth cross-sections through the reservoir model corresponding to the same realization shown in Figure 20. The location of Stybarrow-2 and a fault are indicated in the last cross-section. The location of the cross section is shown in the upper left hand corner as the yellow line.*

quite well over the modeled area throughout the cross section.

The final test of the inversion is a well prediction ahead of the drill bit. The inversion was done before the Stybarrow-3 and 4 wells were drilled. The net sand predictions at these two locations were 9.1±6.4 m and 12.3±4.3 m, respectively, but the predictive distributions contain significant peaks at zero net sand. The results were 2.0 m and 7.6 m, respectively. Without the uncertainty, one may be tempted to say that the inversion does not match the results, but since the results were within about a deviation of the estimates one must conclude that the results are reasonably likely outcome given the predictions (statistically there is a 40% probability that the well values could have had a larger deviation from the prediction than observed). Without the uncertainty of the prediction, one can not tell whether a prediction was good or not. As an example, if one predicts 100 m of sand and drills a well that encounters 105 m of sand, this would not be a good prediction if the uncertainty were 1 m or 30 m. In the first case, there is an obvious bias in prediction. In the second case, there is probably an overestimate of the uncertainty. If the uncertainty were 5 m, this would have been a very good prediction.

We want to emphasize a very important feature of this type of model-based inversion—the fundamental properties that are constrained and estimated are quantities that one understands and wants to know. The questions that one answers to constrain the model are things like "Where is the fluid contact?", "How accurate is the interpretation?", "What is the geologic model of N/G distribution?"—not esoteric questions like "What are the range of expected Poisson ratios and bulk moduli of the rocks?" The results given are things that one wants to know such as the fluid in the rock, the porosity of the sand, and the amount of net sand—not merely the acoustic and shear impedance of the rock. The transformation of the latter, especially its uncertainty, to the parameters that one needs to know is messy and better off incorporated directly in the forward inversion model. By imbedding these transformations in the inversion, the user of this technology can focus on the important questions, not geophysical minutia.

"Massaging" the results into the reservoir simulation model. The output of the model-based inversion is an ensemble of models for each trace on the seismic grid. This grid is regularly spaced and orthogonal in the x and y directions. Although significant correlation is built into the models between layers and properties at a specific location there is no correlation built into the model between traces. The lack of lateral correlation in the model before inversion allows a short scale lateral noisiness to appear in the model after inversion, which is geologically unrealistic. Not only is this noise not physical, it also will create problems in the numerical solutions that are part of the reservoir simulation. The lack of lateral correlation also has a significant effect on the uncertainty in the volumetric estimation. When stochastic maps of properties such as net sand are integrated, the standard deviation of the integrand will be proportional to the main lateral correlation length. An upper bound on the volumetric uncertainty will be the integral of the standard deviation, which assumes an infinite correlation length. The uncertainty will decrease from this upper bound by the square root of the area that is integrated, divided by the square of the correlation length. In essence, the smaller the main lateral correlation length, the smaller the variation in the volumetrics.

The reservoir simulation is done on an often highly irregular grid in space. This grid conforms to the geologic faulting, has less lateral resolution than the seismic grid, and more vertical resolution (Figure 18). The "massaging" step addresses the lateral correlation issue and mapping onto the irregular simulation grid, while the vertical downscaling step is addressed in the next section.

The process of massaging the results couples the interlayer and inter-property correlations from the inversion with the lateral correlation via a geostatistical kriging of the results from the seismic grid onto the reservoir simulation grid (Gunning et al., 2006). Well constraints on properties such as N/G and gross thickness are honored. Selected properties such as gross thickness are not kriged across faults at the user's discretion. The reservoir simulation grid is built by standard means and stored in a standard industry format. The result is an ensemble of reservoir simulation models.

When the output of the Stybarrow model-based inversion is massaged, the result is shown in Figures 19 through 21. Note how the average net sand map has been smoothed in Figure 19. A correlation length of 1500 m was used in the SW-to-NE direction and 750 m in the SE-to-NW direction. This correlation length was based on variation seen in the wells and geologic analogues. Realizations of net sand maps for the main sand are shown in Figure 20. They are ordered according to the amount of net sand above the fluid contact. Note the difference between adjacent realizations. The same set of realizations is shown in cross section. Note how the fault discontinuity and well constraints are honored. These cross sections are in depth, in contrast to the time cross sections shown in Figure 17. Although the inversion is done in time (to minimize the effects of wavelet stretch), each model in the ensemble can be displayed in either time or depth. The transformation between the two is determined by the velocities of the layers. It is not necessarily true that the ensemble average time multiplied by the average veloc-

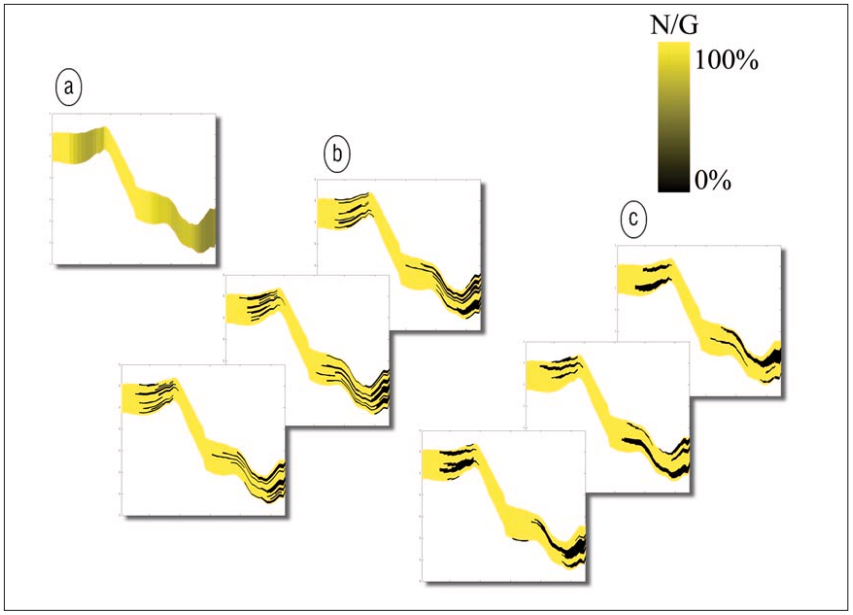

Figure 22. *Depth cross sections through the "decorated" and "enforced" reservoir model (a) massaged realization showing N/G before decoration, (b) three versions of this massaged realization decorated with proportional beds, (c) three versions decorated with offlapping beds.*

ity of a layer will equal the average thickness of that layer. This is the origin of the irreversibility of average time to depth mappings. This dilemma is resolved by looking at the individual model realizations.

"Decoration" and "enforcement"—adding the subseismic structure. The remaining task is to add the stratigraphic first- and second-order subseismic structure needed by the reservoir simulation (Willis and White, 2000). The approach is to "decorate" the model with this structure according to several geologic templates of stratigraphic style (each with distinctly different topology) specified by template models built with industry standard reservoir modeling software. The decoration is done in a way that both the N/G and gross thickness will closely match, but not exactly, the values demanded by the stratigraphic seismic loop scale average given by a specific realization given by the massager. A subsequent "enforcement" step will be done to ensure that the N/G and gross thickness exactly match the values given by the realization generated by the massager, as well as match the stratigraphic first- and second-order subseismic structure seen at the wells. Preliminary results using a truncated Gaussian simulation are shown in Figure 22. One random realization of massaged model-based inversion is decorated with two different subseismic stratigraphic styles. The first is proportional bedding. The second is offlap bedding. Future work will consider alternatives to the Gaussian simulation such as Markov random fields, object models, and other indicator methods.

Conclusions. We now have an ensemble of reservoir simulation models that are consistent with the seismic data, well information, and geologic concepts. These models are analyzed to give volumetric distributions and maps of the minimum amount of expected sand (P10 net sand map, that is, 10% chance of having less net sand, on a point to point basis). The latter is being used to determine well locations, because a minimum amount of sand is needed to complete the well.

The most important output of reservoir modeling and simulation are the possible production profiles from a field, given the uncertainty inherent in the data. This is the information that is needed to evaluate the economics and determine well locations. After the application of the methods described in this paper, the range of production profiles should better reflect the uncertainty inherent in the subsurface data and thus lead to a better appreciation of the field's potential production and risks for development.

Suggested reading. "Hierarchy of deep-water architectural elements with reference to seismic resolution" by Prather et al. (Proceedings of GCSSEPM Foundation Annual Research Conference, 2000). "Stybarrow oil field—from seismic to production, the integrated story so far" by Ementon et al. (SPE Asia Pacific Oil and Gas Conference, *Expanded Abstracts*, 2004). "Application of integrated risking on a South African prospect" by Glinsky et al. (EAGE *Expanded Abstracts*, 2004). "Bayesian wavelet estimation from seismic and well data" by Buland and Omre (GEOPHYSICS, 2003). "Delivery: An open source model-based Bayesian seismic inversion program" by Gunning and Glinsky (*Computers and Geosciences*, 2004), "Delivery-extractor: a new open source wavelet extraction and well tie program" by Gunning and Glinsky (*Computers and Geosciences*, in press), and "Delivery massager: an open source preparation of model-based seismic inversion for reservoir simulation" by Gunning et al. (in preparation for *Computers and Geosciences*, 2006), "Quantitative outcrop data for flow simulation" by Willis and White (*Journal of Sedimentary Research*, 2000). TLE

Acknowledgments: The authors acknowledge the financial support of the BHP Billiton technology program, and thank BHP Billiton and Woodside Petroleum for permission to publish the results. They also thank Alan Curtis, Chris Lerch, Neil Ementon, and Guy Duncan for many useful discussions.

Corresponding author: mglinsky@bhpbilliton.com