Assessing the effect of geological uncertainty on recovery estimates in shallow-marine reservoirs: the application of reservoir engineering to the SAIGUP project

John D. Matthews[1], Jonathan N. Carter[1], Karl D. Stephen[2], Robert W. Zimmerman[1], Arne Skorstad[3], Tom Manzocchi[4] and John A. Howell[5]

[1]*Department of Earth Science and Engineering, Imperial College, London SW7 2BP, UK (e-mail: j.d.matthews@imperial.ac.uk)*

[2]*Institute of Petroleum Engineering, Heriot-Watt University, Edinburgh EH9 9NX, UK*

[3]*Norwegian Computing Center, PO Box 114 Blindern, N-0314 Oslo, Norway*

[4]*Fault Analysis Group, UCD School of Geological Sciences, University College Dublin, Belfield, Dublin 4, Ireland*

[5]*Department of Earth Science/Centre for Integrated Petroleum Research, University of Bergen, Allegt. N-5007 Bergen, Norway*

ABSTRACT: Reservoir management is a balancing act between making timely operational decisions and the need to obtain data on which such decisions can be made. There is a further problem: estimates of recovery for prospective development plans are subject to uncertainty because of the uncertainty of the geological description within the simulation model.

The SAIGUP project was designed to analyse the sensitivity of estimates of recovery due to geological uncertainty in a suite of shallow-marine reservoir models. However, although it was generic, it had the hallmarks of active reservoir management, because those members of the team responsible for deriving the notional development plans for individual models via reservoir simulation, and computing the recoveries, had to work in parallel with others under time and budget constraints.

This paper describes the way the reservoir engineering was carried out to achieve these objectives, the assumptions made, the reasoning behind them, and how the principles could be used in other studies. Sample results are also presented, although the bulk of the results are presented in other papers in the project series. One surprising result was that faults that impede flow can improve recovery. The underlying physical explanation for this behaviour is provided.

KEYWORDS: *recovery estimates, uncertainty, reservoir engineering, sweep efficiency*

INTRODUCTION

Even in a mature oil province such as the North Sea, it has proved difficult to forecast recovery. Examples abound where the estimates are a factor of ten different from reality (Castle 1985; Thomas 1998). Even where large-scale uncertainty in gross-rock volume (due, perhaps, to problematic seismic data) is eliminated, the uncertainty in forecasting reserves is *c.* 30%. Much of the uncertainty is due to uncertainty in reservoir parameters rather than operational problems, although the latter cannot be ignored. With this level of uncertainty, operators may be reluctant to develop certain classes of reservoirs because of the risks of economic failure. The lead author has found, from independent reservoir studies using decline curve methods and simulation, that the fractional uncertainty in estimating remaining reserves during an active project can increase beyond 30%. The hyperbolic nature of the displacement processes in the reservoir is one reason for this. In consequence, decisions to add incremental projects to an existing development are also difficult.

The SAIGUP project (Sensitivity Analysis of the Impact of Geological Uncertainties on Production) was a generic study aimed at answering questions as to how our limited understanding of reservoir geology affects uncertainty in recovery estimates (Manzocchi *et al.* 2008*a*). Although the project was generic, it had all the essential hallmarks of active reservoir management, including tight, interrelated time-scales between members of the project team. The chosen class of reservoirs was of the shallow-marine type, but the reservoir engineering approach could be adjusted to studies on other types of reservoirs.

This paper, which is part of a thematic set, describes the reservoir engineering approach derived and adopted to meet the project requirements, the reasoning behind it, and a subset of results, with physical explanations. Fluid properties, well positions and other operational controls were chosen consistent with project objectives, and described herein.

RESERVOIR ENGINEERING AND MANAGEMENT

At any point in time, a certain amount of information is known about a reservoir. This may be used to construct a simulation model and, by exploring a range of well positions using that model, a pattern of wells (the development plan) that maximizes the economic recovery from that numerical description of the reservoir can be obtained. During the life of the reservoir, active reservoir management (Thakur 1996) can be practised by revising the simulation model to account for additional information from the reservoir, such as new well data and production history. Positions of wells still to be drilled can then be re-evaluated. The process is time consuming and, even with automated methods (such as described in Carter & Matthews 2008), it is not always possible to conclude studies before the next wells have to be drilled or equipment ordered and installed (Saleri 1993).

Within the overall project, where parts had to run in parallel over a limited span of time, active reservoir management of individual cases that would mimic this form of real-time management was impossible without sacrificing quantity. There was a further problem. In real reservoir management, only an imperfect representation of the real reservoir exists, and it is fuzzy, partially random, and incomplete (Foley *et al.* 1997). This is then converted into numerical form for simulation purposes. In contrast, in the present generic project, individual models were defined precisely through numerical data. It was therefore necessary to deliberately hide some data before attempting the reservoir engineering and management on the supplied numerical models.

GEOLOGICAL MODELS, TRUTH MODELS AND THE GENERIC STUDY

In the simulation models generated for this project (Manzocchi *et al.* 2008*a*), fault positions and properties are known exactly. In the real world, fault positions in a reservoir would have been obtained from seismic data and other means. Inferred positions would be approximate, as would their hydraulic properties. In order to create a sense of realism in the project, development plans for these truth models were derived on the false assumption that the faults in the real reservoir were sealing, and their positions were known only approximately, irrespective of data within individual truth models. Even if fault positions were known exactly in the real world at reservoir depth, there is a risk whilst drilling close to a fault of ending up on the wrong side, thus needing a side-track to return to the correct side. To cope with this risk, rather than attempt to quantify the financial aspects of a side-track, well bottom-hole locations in the simulations were chosen to be safely inside the desired fault compartment.

Using the supplied models with some of the real truth about the faults hidden in this way, a structured search was made to identify a series of pre-drilled well positions that would maximize economic recovery. However, the process was limited by heuristic rules of satisficing (Keen & Morton 1978), which require the engineer to stop the process when a satisfactory answer has been obtained rather than continue with more calculations in the hope of obtaining the optimum answer.

Four general patterns of faulting were explored within a dipping reservoir with four-way closure. These four types were labelled U for unfaulted, A for faults generally perpendicular to reservoir dip, B for faults running in several directions, and C for faults generally running updip.

The four sets of well positions obtained from this part of the study were then used for all of the geological models subsequently supplied, even though they do not contain the fault pattern appropriate for the well pattern. Whilst this quadrupled the amount of computation, insight is gained into how a poorly selected choice of wells affects both the uncertainty and absolute level of recovery estimates.

CHOICE OF PRODUCTION MECHANISM, RESERVOIR FLUIDS AND FLUID VOLUMES

The working hypothesis for the generic study was that each target reservoir formed part of a larger development that provided shared drilling and production facilities. Otherwise individual wells may have been sub-economic because of small volumes of oil within individual drainage radii and/or low flow rates. This was a collective decision to simplify the project, but was not deemed to be a fundamental hindrance to understanding the effect of the uncertainty of geology on recovery estimates.

Even in a generic project, the reservoir engineers were not only faced with the choice of well positions, but also the depletion mechanism. For the type of reservoir on which the project focused, natural depletion would produce only 2–6% of the oil, because of the lack of a large aquifer. Gas injection, apart from being costly in a European environment, would result in long computing times and possibly be inappropriate in such geologically complex reservoirs. Waterflooding, therefore, became the chosen production method, since this seemed likely to highlight the effects of geological uncertainty on recovery estimates, because a wide range of recoveries might be expected. If a particular development plan derived on the basis of an incomplete description of the truth model proves, by happy circumstance, to be almost optimum for a particular truth model, the recovery might be as high as *c.* 55%. This figure represents almost 90% of the mobile oil in an individual reservoir model. Lower recoveries would be expected if inter-well connectivity is poor, and this depends on the geological characteristics of the models.

In a heterogeneous reservoir, the recovery will generally improve if the mobility ratio is favourable (Craig 1980). In this way the recovery will be less sensitive to the detailed geology. If the mobility ratio is unfavourable, then the water is likely to channel through a few high permeability zones. A strong relationship between some types of geology and recovery would then be observed. However, this would leave a large fraction of mobile oil unproduced, thus limiting our understanding of the influence of the full range of geological variation. In addition, coning at wells may occur, requiring effort to be invested in the numerical aspects of simulating what is a local well phenomenon. Therefore, a compromise was agreed within the project to set the fluid properties such that the mobility ratio was approximately unity.

The reservoirs of the Brent province, Northern North Sea (Morton *et al.* 1992) were considered to be prototypical of the reservoirs being considered, but not in their entirety. The reservoir dip, their fault patterns and certain parts of the stratigraphy, such as Ness and Etive units (Flint *et al.* 1997), are common, though not their detailed layering. In these reservoirs the oil varies from a light oil to a condensate, with occasional areal- and depth-dependent variations. The properties of the oil in the reservoir models were based on those of a uniform light oil. It was necessary to make a number of compromises to the fluid properties, so as to meet the project's overall aims. The main adjustment was to the compressibilities of the fluids and rocks within the reservoir, which are only about 20–30% of typical values. This reduces the importance of expansion drive and increases the understanding of the effect of inter-well connectivity and tortuosity of flow paths. Any further reduction

Table 1. *The fluid and rock properties used in the SAIGUP simulations*

	Formation vol. factor (vol/vol)	Viscosity (cP)	Density (in reservoir) (kg m^{-3})	Compressibility (bar^{-1})
Oil	1.5	1.0	721	10^{-4}
Water	1.0	0.4	1000	10^{-5}
Rock	—	—	—	10^{-5}

in compressibilities could have compromised the operation of the simulator when solving for pressures by mathematically implicit methods. The values used in the simulations are given in Table 1.

Most reservoirs show a variation of oil saturation with height for a common facies (the transition zone). This has generally been referred to as the capillary transition zone, although the bulk of it is not a capillary phenomenon, it is diagenetic (Matthews 2004). Traditionally, reservoir models were generally initialized with a capillary drainage curve to represent this variation with height. For completeness in simulation models where waterflooding takes places, scanning curves between drainage and imbibition relative permeabilities are then required (Eigestad & Larsen 2000). For simplicity, the bulk of the SAIGUP simulations were initialized with connate water saturations that were independent of height. This ensured that all simulation models mimic the initial dry oil production observed from real reservoir transition zones, without the need to invoke arbitrary sets of mathematically scaled curves, with their attendant increase in computing time.

The depth of the oil–water contact (OWC) was chosen so that at the crestal region of the reservoir the whole of the reservoir thickness was above the OWC, whilst the contour of the OWC on the top surface remained within the reservoir model. The OWC was finally set at the mid-point between the crest and the deepest part of the aquifer.

In regard to properties below the OWC, no oil was introduced below that main contact, either in the form of palaeo-oil or pockets of mobile oil trapped below low permeability and/or high capillary pressure regions during oil emplacement. With these assumptions, the volume of stock tank oil-in-place (STOIIP) was around 60×10^6 m^3. Furthermore, even though many reservoirs show deterioration in porosity and permeability below the OWC (to the point at which some cannot sustain water injection there), no such diagenetic effects were included in the simulation models.

In a real reservoir development, there is also uncertainty in the assessed values of porosity and water saturation. An imprecise estimate of STOIIP will result, and such uncertainty will also have an effect on the uncertainty in recovery estimates. These effects were not included in the project. Methods are now becoming available that will allow such estimates to be made (Matthews 2004) and these need to be followed up in simulation models.

WELL LOCATIONS, TYPES AND WELL CONTROLS

Many of the principles of well placement for reservoir development have their origins in the days preceding digital reservoir simulation (Craig 1980). Strenuous attempts have since been applied to those basic ideas that microscopic sweep effects, practical water-cut issues, areal sweep, vertical sweep, gravity effects and rate dependency can be uncoupled in the hope of obtaining a basic understanding of how to position and operate wells in a particular reservoir context. The problem has proved to be intractable by this route.

Simulation and operational practice, based on a wide range of reservoirs that have a definite crestal position, have shown that a rolling line drive towards the crest is an adequate way of recovering oil. However, aggressive faulting disrupts this ideal. Where faulting is dominant, identification of reservoir compartments and providing producer–injector pairs in each compartment is essential (Nadir 1981). This means that areal sweep efficiency is controlled by well positions. Vertical sweep efficiency is controlled by perforation and workover policies in individual wells. In this way, higher recoveries can be achieved from reservoirs having layers even with sharp permeability contrasts (Thomas & Bibby 1993).

Industry practice is moving towards the use of dedicated wells that target specific areas of the reservoir that may not have been drained (Sackmaier 2002). This advanced level of reservoir management could not be handled within the SAIGUP project. So, vertical wells were chosen rather than horizontal, branched or designer ones. In fact, there are occasions when vertical wells may be optimum (Matthews *et al.* 1992).

Simple well and production controls were used in the development plans.

- Production wells were controlled on a fixed bottom-hole pressure, 50 bars below the initial reservoir pressure, at the OWC depth. No tubing head pressure controls were included.
- Production wells had a minimum allowed flow rate of 50–100 m^3 per day, which varied depending on the location of the well within the oil column.
- The whole of the oil column was used for the initial perforations.
- Production wells were worked over whenever the water-cut reached 30%, since at this point well-lift requirements would typically halt the natural flow.
- Injection wells were on voidage replacement, subject to a maximum pressure of 50 bars above the initial reservoir pressure, mimicking typical fracture and equipment limitations.
- In injection wells, the whole of the reservoir interval was used for perforations.
- A skin of zero was consistently assumed for all wells.
- Any recovery after 30 years or below 200 m^3 per day was judged to be uneconomic.
- No water rate constraints from the whole field.

So as to be consistent with the declared objective of hiding data in the truth models from the engineer, wells were not placed closer than one complete grid-block from any of the mapped faults, even though that could have increased recovery in the simulation (the grid-block size is 75 × 75 m). This meets the criteria discussed above on fault positions and drilling targets.

The unfaulted models (U)

The overall structure of the reservoir is a tilted fault block with four-way closure (Manzocchi *et al.* 2008*a*), with the oil being

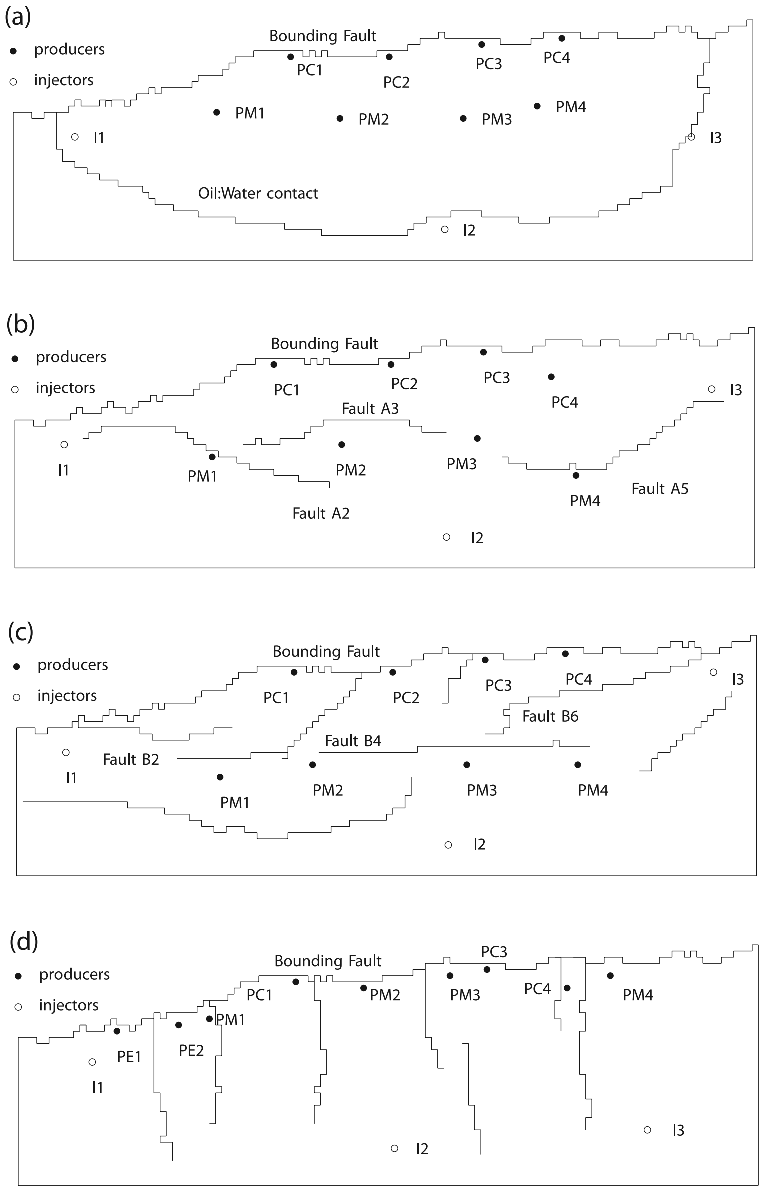

Fig. 1. Development well positions for (**a**) the unfaulted model; (**b**) fault pattern A; (**c**) fault pattern B; (**d**) fault pattern C.

trapped in the uplifted section. Figure 1a shows the major fault, the model boundary and the OWC. Three injectors were located close to this contact and four producers were set against the crest. These positions are based on conventional reservoir engineering rules-of-thumb. Four mid-row producers were set at intermediate positions between injectors and crestal producers. The aim was to provide early higher production rates, thus adding to value, though not necessarily increasing final recovery. The recovery was found not to be overly sensitive to individual well positions, provided that the general pattern of rolling line drive was retained. The pattern of wells was called PDO U during the project.

Fault model A

The faults in model A (Fig. 1b) are dominantly along contour lines. Thus, if they are sealing they will prevent direct flow of water from the injectors to the crestal producers, if the pattern of wells derived for the unfaulted reservoir was used. For this reason, individual well positions were chosen to be in the middle of fairways, or between faults to give maximum connectivity to the surrounding oil. The three longest faults (Fig. 1b) have throws exceeding 10 m.

Faults may or may not be sealing. The throw of a fault, on its own, is insufficient to predict whether there is any hydraulic impedance across the fault. Probability statements can be made about individual faults within a particular reservoir setting, but until additional reservoir information is available, such as perched contacts or different fluid compartments being detected, only active reservoir development will establish the true nature of individual faults.

Geometrically, the pattern of wells looks acceptable in terms of obtaining reasonably high recoveries from such a reservoir even if the faults are fully sealing, but this can be confirmed only by running the simulator with individual models. However, the crestal area is potentially almost isolated from the deeper parts of the reservoir. In particular, the area around well PC1 is at a disadvantage. There is only a small gap between fault A2 and the end of fault A3, and this gap may not actually exist, as it could be an artefact of the seismic mapping. The area around well PC4 may also not be swept because of a large number of small faults there, not shown on the figure (see Manzocchi *et al.* 2008*b* for more detailed maps). The OWC is not shown on the figure; it varies slightly from that shown in Figure 1a because the faults cause the grid-blocks to be at different depths. Perched OWCs were not assumed. The pattern of wells was called PDO A during the project.

Fault model B

Fault pattern B has faults that are generally perpendicular to each other (Fig. 1c). Notionally, this could provide problems for depleting the reservoir, because large numbers of isolated fault compartments could exist. The situation is probably less severe than this, as the faults may not be sealing and thus the compartments are not so severely isolated. Injector and producer pairs are not, therefore, necessary in each fault block. It is possible to pick a series of well positions that are not isolated and are not far removed from the positions adopted in model A. The pattern of wells was called PDO B during the project.

An attempt was made to use automatic computing techniques to identify optimum well positions for a reservoir model having fault pattern B which, because of the compartmentalized nature of the faulting, often gave low recoveries (Carter & Matthews 2008). Rather surprisingly, a crestal position for an injector was identified as improving recovery. This deserves further exploration because whilst it looks attractive, since the result was derived from a truth model, the reality may be different.

Fault model C

In model C, the bulk of the faults are perpendicular to the main fault, namely in the dip-direction (Fig. 1d). This makes it more difficult to develop a pattern of wells. The concept of mid-row and crestal producers, used in the three previous models, had to be abandoned. There are twelve faults running against the crest of the reservoir. Without increasing the number of producers to twelve, calculated recoveries in test cases were poor. To provide a reasonable economic comparison with the other three models, the number of producers was restricted to nine. The wells were set in the nine compartments that held the largest volumes of oil. The pattern of nine wells was called PDO C during the project.

The production line

After about one year into the project, no further work was carried out on optimizing the well positions. The purpose was to allow the production line of calculations from geological models to results to start in earnest.

The simulations were run on a 24-node computing cluster constructed from Sun Ultra 5 workstations, and using the MORE simulator (Young & Hemanth-Kumar 1991), completing almost 4100 simulations per month. This number would have been much smaller had opportunities to speed calculations without loss of essential accuracy not been explored (see below). By the end of the project, around 35 000 simulations had been completed. The reservoir engineering then concentrated on trying to explain typical results, and examine side issues.

PRELIMINARY RESULTS

The paper digresses before considering the final choice of data for the simulations. Although formal analysis of the entire suite of simulations is recorded (Skorstad *et al.* 2008), four aspects of the results are discussed now because they provide either physical insight into the controls on recovery predictions and/or an anchor in the world of reality in what could otherwise become merely a statistical analysis.

The effect of the time-scales

In some simulations, oil production had not ceased at thirty years. Had the field been allowed to continue in production, much of the mobile oil would have been produced. The decision to ignore production possibilities after thirty years was taken on the basis that such numerical data add nothing to understanding the practical issues with production from shallow-marine reservoirs. Thirty years after a project had begun, investment in rejuvenating the production facilities would almost certainly be required, apart from the discounted value of the production on which the initial investment decisions were made being virtually nil by this time.

For comparison purposes, some of the production profiles were discounted to provide a better understanding of the effects of the geological characteristics of the models.

The effect of workovers

Well workovers at water-cuts of 30% added very little to recovery – typically a fraction of one percent of the STOIIP, compared with shutting in wells once the water-cut limit was reached. Examining the calculated flow paths within the simulations showed that vertical communication was good over most of the reservoir, and that there were no major contrasts in horizontal permeability at different vertical positions. The simulations, therefore, showed that recovery was controlled by a series of tortuous paths for the flood fronts in the reservoir, which are not influenced significantly by well perforations or workovers. This is a reflection of the geological models.

The effect of faults

The effect of faults was surprising. The results showed that as the hydraulic impedance of faults – introduced by transmissibility multipliers of less than unity – lying between injectors and producers was increased, recovery increased. This seemed counter-intuitive. In other calculations, increases in the hydraulic impedance between geological sequences also increased recovery.

The reason for the increases is subtle and may be explained by considering the simpler case of areal sweep efficiency in a homogeneous flat reservoir. At unit mobility ratio, streamline theory may be used to track the flow of water from injector to producer. If the permeability is isotropic, the streamlines follow a particular pattern and thus provide a particular recovery at a chosen water-cut. If the permeability is anisotropic (with permeability decreased in the direction of the interwell vector) then the streamlines are forced to spread out more, the swept area is increased and recovery is thereby improved for that chosen water-cut. If, as an alternative to anisotropic permeability, a fault lying between an injector and producer impedes flow between the wells, then the streamlines also spread out and improve recovery. However, the flow rate is reduced for a fixed pressure drop between injector and producer. In consequence, the recovery, discounted by 10–15% per annum, typically decreases.

Fuller details, including the mathematical analysis, is given in Manzocchi *et al.* (2008*b*).

Variety in recovery

Figure 2 shows a ranked plot of cumulative recovery obtained from *c.* 950 versions of a representative sedimentological

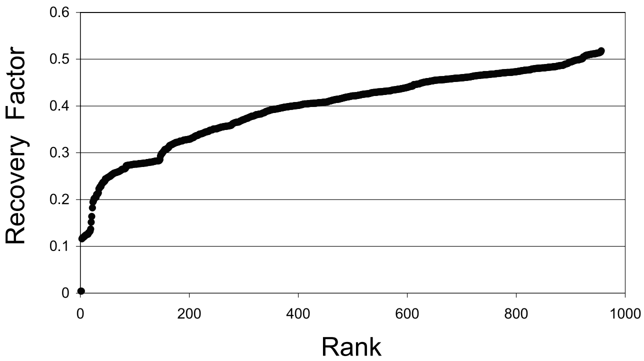

Fig. 2. Recovery factors of 956 distinct faulted or unfaulted versions of the median sedimentological model, simulated using the well configuration designed around the unfaulted structure.

model. These versions contain different relative permeability cases and different fault patterns and properties and have been simulated with different development plans. The recovery range is almost zero to 52% of STOIIP, with 95% of the results in the working range of 25% upwards. The fact that the recovery is almost linear over this range with rank indicates that there are no significant peaks and troughs in recovery, which would have suggested that the project has over-emphasized or under-emphasized a particular range of recoveries by accident.

Note that the amount of mobile oil in the models controls the upper limit of recovery to slightly in excess of 60% STOIIP.

RELATIVE PERMEABILITY AND CAPILLARY PRESSURE

An important issue was the choice of relative permeability and capillary pressure curves to use in the modelling. In practice it is only possible to determine experimentally the relative permeability and capillary pressure curves at a very limited number of positions in a reservoir. The engineer then needs to upscale the experimental results, made at the core scale, to the grid-block scale. An appropriate set of curves then needs to be allocated to every grid-block in the simulation model. There is plenty of opportunity for uncertainty to be introduced in this process, and the question is how important might the choices made be on simulated production. The extreme situation is when no experimental results are available, and the engineer is forced to use smooth ‘engineering’ curves. This section describes the derivation of the final set of curves used in the modelling before considering the impact that the extreme situation would have on the production profiles from the SAIGUP reservoir models.

Initial modifications to the relative permeability and capillary pressure curves

All sedimentological models used in the SAIGUP study were built with six facies associations (Howell *et al.* 2008), and were simulated using six relative permeability curves selected on the basis of the dominant facies in each upscaled grid block. With the exception of the models discussed in this paper and by Stephen *et al.* (2008), all were simulated using the same set of curves. Stephen *et al.* (2008) were concerned specifically with uncertainties associated with different assumptions made in the construction of the relative permeability and capillary pressure curves, while the work described here is concerned with preparing a base-case set of curves for general use in the project.

Table 2 summarizes the sets of curves considered in this paper. The curves called ‘initial’ (Table 2) are those provided to

Table 2. *Summary of the six sets of curves considered in this paper*

Simulation case	Relative permeability	Capillary pressure
Initial curves (001)	High resolution (red)	High resolution (red)
Hand-smoothed (010)	Smoothed (black)	Smoothed (black)
Base-case (000)	Smoothed (black)	None
Engineering 1 (011)	Quadratic (blue)	None
Engineering 2 (012)	Quadratic (blue)	Quadratic (blue)
Engineering 3 (013)	Quadratic (blue)	Linear (green)

The three-digit number in the first column is the code name used in the project, and the colours listed in the other two columns refer to those used in Figure 3. The bulk of the SAIGUP modelling used the base-case curves.

the reservoir engineering team as a starting point. These were preliminary versions of the curves described as ‘base-case’ by Stephen *et al.* (2008) and were generated through a geologically based upscaling approach. Preliminary curves were used because of time constraints, a further example of the active nature of the reservoir management portion of the project work. The red curves on Figure 3 show examples of the initial water and oil relative permeability and capillary pressure curves for one of the facies (the offshore transition zone). Examples of

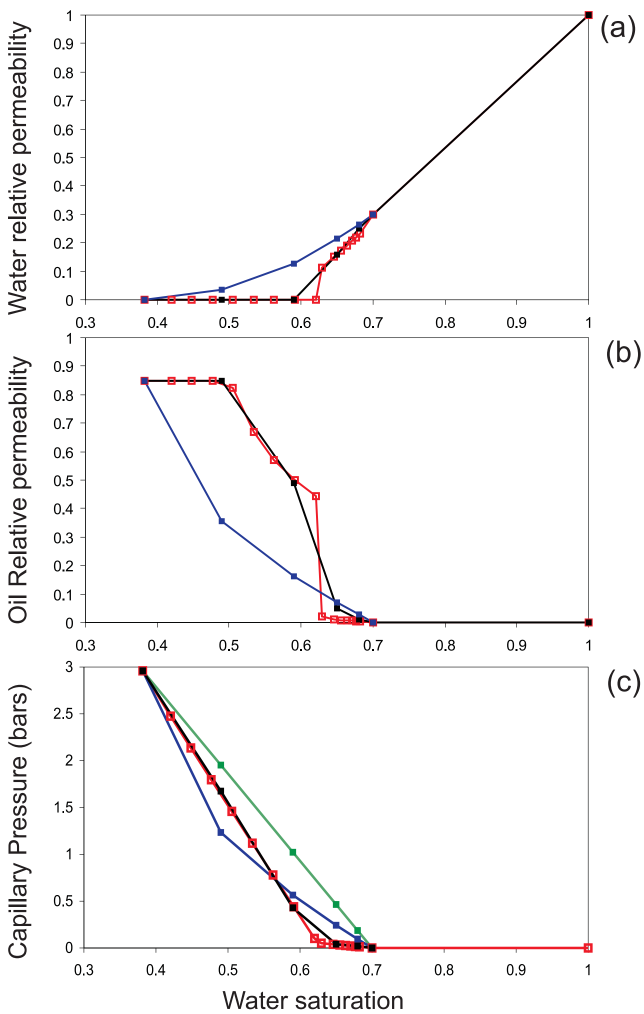

Fig. 3. Examples of (**a**) water, (**b**) oil relative permeability and (**c**) capillary pressure curves for one of the six facies. Red: initial high resolution curves provided to the engineering team. Black: hand-smoothed version of the initial curves. Blue: Quadratic engineering curves. Green: Linear engineering curve (capillary pressure only). See Table 2 for how these curve types are combined.

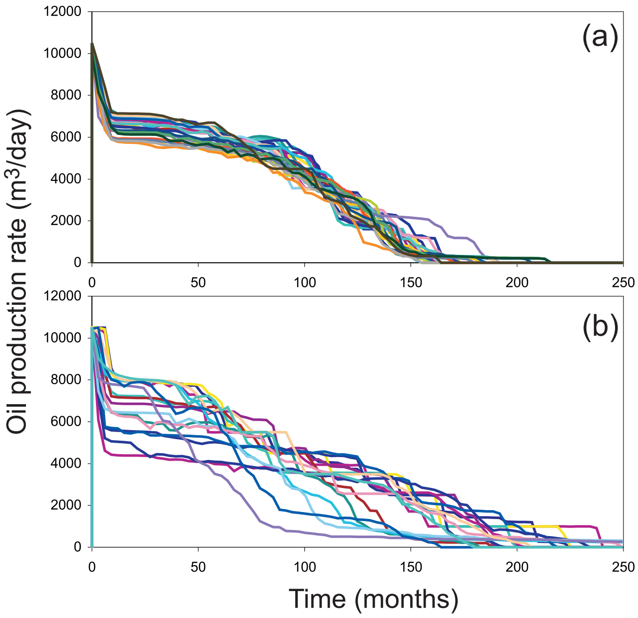

Fig. 4. Oil production rates for various models containing the initial relative permeability curves, simulated using the well configuration designed around the unfaulted structure: (**a**) for 28 faulted versions of the median sedimentological model; (**b**) for 18 distinct unfaulted sedimentological models.

production profiles generated using these curves are shown in Figure 4.

The nature of relative permeability and capillary pressure curves has a major effect on the time taken to complete a simulation. For example, as noted by Mattax & Dalton (1990), abrupt changes in relative permeability can cause the simulator to take two to three times as long compared with smoother curves. Clearly, therefore, smoothing the curves might allow a larger number of simulations to be completed within the project time, thus gaining additional information. However, there would be an associated loss of accuracy in individual results, which may or may not be acceptable. The relative permeability and capillary pressure curves were hand-modified with the aim of honouring the shape as much as possible, but to smooth out some of the more rapid changes in slope. These curves are shown in black on Figure 3 (note that the MORE simulator uses a linear scheme to interpolate between points).

Table 3 compares five measures of reservoir performance for 1260 versions of models using the initial and the hand-smoothed curves. The performance of the models are virtually identical for all measures (models containing the hand-smoothed curves produce approximately 0.7% more oil), but the simulations using the hand-smoothed curves take, on average, about half as long to run (Fig. 5). Figure 6 compares production profiles for the two compilations of models shown in Figure 4, and shows extremely close matches between models containing the hand-smoothed and initial curves. Examination on a model-by-model and well-by-well basis shows slightly larger changes, but nothing that would give cause for concern. This slight loss in accuracy would therefore allow twice as many models addressing other issues – a significant gain.

Table 3. *Comparisons between the performance of 1260 models simulated using the initial, hand-smoothed and base-case sets of curves*

	Hand-smoothed curves/initial curves	Base-case curves/ hand-smoothed curves
STOIIP	0.9988 ± 0.0002	1.124 ± 0.012
Total oil production	1.006 ± 0.008	1.205 ± 0.035
Discounted oil production	1.007 ± 0.007	1.131 ± 0.033
Recovery factor	1.007 ± 0.008	1.072 ± 0.024
Discounted recovery factor	1.008 ± 0.007	1.006 ± 0.023

Average performance ratios are quoted with the error as ± 1 standard deviation. Discounted oil production and discounted recovery factor are calculated using a discount factor of 10% per year.

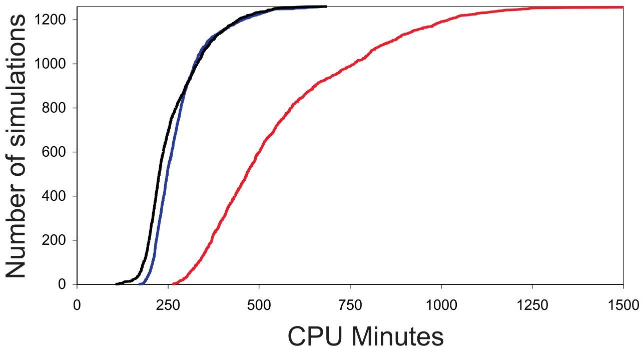

Fig. 5. Simulation run times for 1260 models simulated using the initial (red), hand-smoothed (blue) and base-case (black) cases.

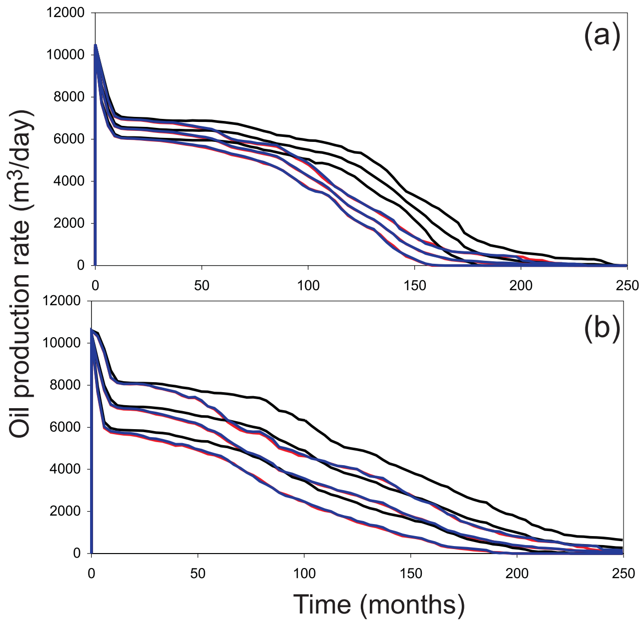

Fig. 6. Average, and average ± 1 standard deviation, oil production rates for the same two compilations of models as shown in Figure 4, simulated using the initial (red), hand-smoothed (blue) and base-case (black) cases. See text and Table 2 for details of these cases.

Impact of omitting capillary pressure

Typical simulators use a dataset nominally called 'capillary pressure' to cover two separate aspects of the simulation process. First, the data modify the viscous pressure gradient experienced when water is displacing oil. However, zero capillary pressure has very little effect on the displacement efficiency compared with using known imbibition data. The 'capillary pressure' data are also used to initialize the oil saturation through the transition zone based on the paradigm of Muskat that the transition zone is solely a capillary phenomenon (Muskat 1949). However, this simplification does not result in the simulator mirroring the saturation profile obtained from electric logs or the general lack of mobility of the water phase of many reservoirs (Matthews 2004). To cope with this discrepancy, an engineer may use either end-point scaling, non-equilibrium initialization or input a wide range of differing relative permeability data as a function of height above the OWC.

As discussed above, it was decided not to include capillary pressure derived from the mathematical core curves in the simulation models. The final set of relative permeability curves used in the bulk of the SAIGUP modelling were therefore the hand-smoothed curves, but omitting capillary pressure (this combination is referred to as 'base-case' in Table 2 and elsewhere in this paper). Figure 6 and Table 3 show clearly that the omission of capillary pressure has quite a significant effect on production from the reservoir models.

One effect on the reservoir of replacing the zero capillary pressure data input to the simulator with non-zero data is to reduce the volume of oil present. Table 3 shows a decrease of *c.* 12% in the suite of cases simulated. However, the recovery declined by *c.* 20%. The reason for this almost doubling of the expected decline in recovery is that the oil saturations were reduced in the transition zone. Therefore, without end-point scaling or inputting multiple sets of relative permeabilites, the fraction of the oil volume in the reservoir model that is mobile is reduced, having a direct non-geological impact on recovery. Whether this second effect is an accurate representation of the real physics of the transition zone is now an area of active research (Matthews *et al.* 2007).

In terms of the fraction of oil recovered, there is a decline of 7%. (i.e. from 47% to 43%). The extra recovery in the case with the zero capillary pressure data is produced by a later decline in production rate rather than by higher early production rates (Fig. 6). The increases in discounted recovery factor are therefore lower than increases in non-discounted recovery factor when capillary pressure is omitted (Table 3), since these more economically relevant measures weight earlier production preferentially. In terms of discounted recovery factor, the mean effect of omitting capillary pressure is of a similar magnitude as the mean effect of hand-smoothing the curves (i.e. about 0.6%; Table 3).

The production profiles for the two compilations indicate more variability among the 18 different unfaulted sedimentological models than among the 28 structural versions of a single sedimentological model (Figs 4, 6). Qualitatively, it seems that omitting capillary pressure from the models has a larger effect on the production profiles than the choice of fault model (Fig. 6a) and, perhaps, a similar magnitude of effect as choosing an entirely different sedimentological model (Fig. 6b). This is consistent with the results of Stephen *et al.* (2008), which indicate that the shape of the capillary pressure curve is the single most significant control on recovery factor examined throughout the modelling.

Impact of simpler engineering curves

The blue and green curves in Figure 3 have the same end-point properties as the curves derived from geologically based upscaling discussed above, but have relative permeability and capillary pressure expressed using quadratic or linear relationships at intermediate saturations. Such 'engineering' curves might be used in the absence of lamina-scale geological and petrophysical information and fine-scale upscaling, and it is interesting to consider what the effects of these simplifications might be.

Figure 7 shows four sets of curves for the two compilations of production responses used in Figure 4. The black curve uses the base-case function discussed above (i.e. the hand-smoothed curves with no capillary pressure), while the pink curves also

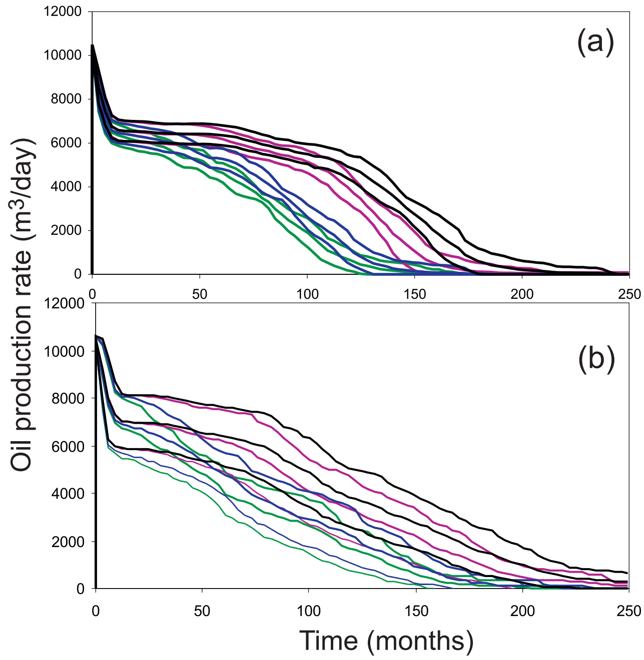

Fig. 7. Average, and average ± 1 standard deviation, oil production rates for the same two compilations of models as shown in Figure 4, simulated using the base-case (black), engineering 1 (pink), engineering 2 (blue) and engineering 3 (green) cases. See text and Table 2 for details of these cases.

omit capillary pressure but use the quadratic relative permeability functions. These results show that the choice of curves has little effect on oil production rate during the early production period. This is surprising since the quadratic curves have lower oil relative permeabilities, and higher water permeabilities, than the hand-smoothed curves. What appears to be happening is that the geometry, structure and sedimentology of the reservoirs have a dominating influence on what amounts to an upscaled relative permeability appropriate to the whole of the reservoir and its development. This overall way of understanding reservoir production has been used to effect before in understanding reservoir performance and optimizing it (Matthews *et al.* 1992). Within the production profile, there is a close analogy to a Buckley-Leverett displacement in a one-dimensional problem. There is a period of dry oil, the arrival of a shock front, and then a period of increasing water-cut. The profiles separate distinctly after about four years, and there is a moderate reduction in production rate in the mid years, which typically causes a reduction of about 0.06 in the recovery factor.

Inclusion of capillary pressure (Fig. 7) has the principal effect of reducing the oil volume present, but secondary effects of capillary pressure also result in lower recovery factors, as discussed above. Figure 8 compares recovery factors measured in 504 distinct models using the simplified engineering curves with recovery factors obtained in the same models but using the base-case curves. The graph indicates that although there are significant differences in recovery factor, these differences are relatively consistent for each model. Decision making in these reservoir models could therefore be conducted using the simplified data, because the optimum choice of development and operational decisions only depend on the relative performance of the options, and not the absolute.

Discussion

This section has described the derivation of the base-case relative permeability curves used in the modelling, and has

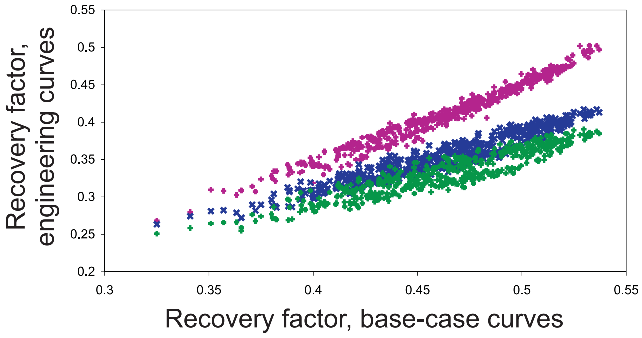

Fig. 8. Recovery factors obtained for 504 distinct models when simulated using the engineering curves compared to those obtained using the base-case curves. Pink: engineering 1 curves; blue: engineering 2 curves; green: engineering 3 curves. See text and Table 2 for details of these cases.

demonstrated that in this case there is significant benefit to be gained from slight modifications to the curves. However, no systematic study has been completed of how one should modify the curves to achieve the required result. Nor has an attempt been made to verify that the revised curves would perform equally well on another simulator.

The final curves used in the main SAIGUP study are based on a preliminary version of the detailed curves prepared by Stephen *et al.* (2008) and omit capillary pressure. As discussed by Manzocchi *et al.* (2008*a*), use of a single set of curves appropriate for a particular set of assumptions in the bulk of the modelling was a deliberate decision, allowing investigation of assumptions made in the generation of these curves to be addressed in parallel to the main modelling programme.

Skorstad *et al.* (2008) and Stephen *et al.* (2008) assessed the relative effects of a number of geological parameters on the variability in oil production from the SAIGUP models. The first analysis used the simplified base-case curves discussed here, while the second analysis focused on relative permeability and capillary pressure issues and therefore used curves derived from a systematic upscaling programme. The models analysed by Skorstad *et al.* (2008) took 290 ± 65 CPU minutes to simulate, while those of Stephen *et al.* (2008) required several additional months of high-resolution modelling to prepare, and then took 408 ± 165 CPU minutes to complete. The fact that both analyses obtained exactly the same uncertainty rankings for factors in common (see Manzocchi *et al.* 2008*a* for details) vindicates the decision to proceed with simplified versions of preliminary relative permeability curves in the main suite of simulation models.

SUMMARY AND CONCLUSIONS

This paper described the reservoir engineering procedures and decisions used in the SAIGUP project. Although the project was generic, it had all the essential hallmarks of active reservoir management, including tight, interrelated time-scales between members of the project team. The chosen class of reservoirs was the shallow-marine type. Numerous reservoirs were simulated, having different degrees of faulting and heterogeneity.

It has been shown how reservoir engineering, conducted in an environment almost typical of active management, may be applied to a specific generic project to allow an assessment of the effects of geological uncertainty on estimates of recovery. To summarize some of the main features of this work:

- fluid properties were chosen similar to those of a light oil, except for compressibilities;
- no transition zone was included;
- four patterns of wells were developed that provide near-optimum recovery from four different realizations of fault patterns;
- simplifications in relative permeabilities were achieved without compromising results, but achieving a greater turnover of calculations;
- around 35 000 calculations were performed, and the results forwarded for statistical analysis;
- a surprising effect of faults was observed and explained, whereby recovery increases when impedance is increased for faults lying between injectors and producers;
- further work was identified relating to the volume of oil-in-place, and how it affects estimates of recovery;
- other types of reservoirs also need to be considered.

The European Commission is thanked for partly funding SAIGUP under the EU fifth Framework Hydrocarbons Reservoir Programme. A number of members of the earlier PUNQ project provided useful starter ideas on reservoir engineering for SAIGUP. We are very grateful to Badley Geoscience, BG International, Roxar and Shell for their support of the SAIGUP project and for their sponsorship of the production of this thematic set.

REFERENCES

Carter, J.N. & Matthews, J.D. 2008. Optimization of a reservoir development plan using a parallel genetic algorithm. *Petroleum Geoscience,* **14**, 85–90.

Castle, G.R. 1985. Assessment of North Sea Field Performance. *Petroleum Economist,* **October**, 358–359.

Craig, F.F. 1980. *The Reservoir Engineering Aspects of Waterflooding.* SPE Monograph, **30**. Society of Petroleum Engineers, Dallas.

Eigestad, G. T. & Larsen, J. A. 2000. Numerical Modelling of Capillary Transition Zone. Paper SPE 64374, presented at the Asia Pacific Oil and Gas Conference and Exhibition, Brisbane, 15–18 October.

Flint, S., Howell, J.A., Knight, S. & Tilbrook, A. 1997. Brent Group of Northwest Hutton Field, UK North Sea: Fill Characteristics, Geometries and Comparison with Outcrop Analogues, Shallow Marine and Nonmarine Reservoirs. *Proceedings of the Gulf Coast Section Society of Economic Paleontologists Research Conference*, 107–116.

Foley, L., Ball, L., Hurst, A., Davis, L. & Blockley, D. 1997. Fuzziness, incompleteness and randomness: classification of uncertainty in reservoir appraisal. *Petroleum Geoscience,* **3**, 203–209.

Howell, J.A., Skorstad, A., MacDonald, A., Fordham, A., Flint, S., Fjellvoll, B. & Manzocchi, T. 2008. Sedimentological parameterization of shallow-marine reservoirs. *Petroleum Geoscience,* **14**, 17–34.

Keen, P. & Morton, M. 1978. *Decision Support Systems – An Organisational Perspective.* Addison-Wesley, Boston.

Manzocchi, T., Carter, J.N., Skorstad, A. *et al.* 2008*a*. Sensitivity of the impact of geological uncertainty on production from faulted and unfaulted shallow-marine reservoirs: objectives and methods. *Petroleum Geoscience,* **14**, 3–15.

Manzocchi, T., Matthews, J.D., Strand, J.A. *et al.* 2008*b*. A study of the structural controls on oil recovery from shallow-marine reservoirs. *Petroleum Geoscience,* **14**, 55–70.

Mattax, C.C. & Dalton, R.L. (eds) 1990. *Reservoir Simulation.* SPE Monograph, **13**. Society of Petroleum Engineers, Dallas.

Matthews, J.D., Carter, J. N. & Dake, L. P. 1992. Investigation of Optimum Well Spacing for North Sea Eocene Reservoirs. Paper SPE 25030, presented at the SPE European Petroleum Conference, Cannes, 16–18 November.

Matthews, J.D. 2004. *Geological and Physical Assessment of the Oil Reservoir Transition Zone.* PhD thesis. Imperial College, London.

Matthews, J. D., Carter, J. N., & Zimmerman, R. W. 2007. Using Drill-Stem and Production Tests to Model Relative Permeabilites. Paper SPE 107209, presented at the Offshore Europe conference, Aberdeen, 4–7 September.

Morton, A.C., Haszeldine, R.S., Giles, M.R. & Brown, S. 1992. (eds) Geology of the Brent Group. *Geological Society, London, Special Publications,* 61.

Muskat, M. 1949. *Physical Principles of Oil Production.* McGraw Hill, New York.

Nadir, F. 1981. Thistle Field Development. *Journal of Petroleum Technology,* **33**, 1828–1834.

Sackmaier, W.E. 2002. Multi-target Wells: A New Concept to Improve Well Economics. *Petroleum Geoscience,* **8**, 3135.

Saleri, N.G. 1993. Reservoir Performance Forecasting: Acceleration by Parallel Planning. *Journal of Petroleum Technology,* **July**, 652–657.

Skorstad, A., Kolbjørnsen, O., Manzocchi, T., Carter, J.N. & Howell, J.A. 2008. Combined effects of structural, stratigraphic and well controls on production variability in faulted shallow-marine reservoirs. *Petroleum Geoscience,* **14**, 45–54.
Stephen, K.D., Yang, C., Carter, J.N., Howell, J.A., Manzocchi, T. & Skorstad, A. 2008. Upscaling uncertainty analysis in a shallow-marine environment. *Petroleum Geoscience,* **14**, 71–84.
Thakur, G.C. 1996. What is Reservoir Management? *Journal of Petroleum Technology,* **June**, 520–525.
Thomas, J.M. & Bibby, R. 1993. The Depletion of the Rannoch/Etive Sand Unit in Brent Sands Reservoirs in the North Sea. *In:* Linville, B. (ed.) *Reservoir Characterisation III.* Penwell Books, Tulsa, 675–713.
Thomas, J.M. 1998. Estimation of Ultimate Recovery for UK Oil Fields: The Results of the DTI Questionnaire and a Historical Analysis. *Petroleum Geoscience,* **4**, 157–163.
Young, L.C. & Hemanth-Kumar, K. 1991. High Performance Black Oil Calculations. *SPE Computer Applications,* **4** (2), 10–15.

Received 29 January 2007; revised typescript accepted 26 October 2007.