B033

Optimal Well Placement

C.L. Farmer (Oxford University), J.M. Fowkes* (Oxford University) & N.I.M. Gould (Rutherford Appleton Laboratory)

SUMMARY

One is often faced with the problem of finding the optimal location and trajectory for an oil well. Increasingly this includes the additional complication of optimising the design of a multilateral well. We present a new approach based on the theory of expensive function optimisation.

The key idea is to replace the underlying expensive function (ie. the simulator response) by a cheap approximation (ie. an emulator). This enables one to apply existing optimisation techniques to the emulator. Our approach uses a radial basis function interpolant to the simulator response as the emulator. Note that the case of a Gaussian radial basis function is equivalent to the geostatistical method of Kriging and radial basis functions can be interpreted as a single-layer neural network. We use a stochastic model of the simulator response to adaptively refine the emulator and optimise it using a branch and bound global optimisation algorithm.

To illustrate our approach we apply it numerically to finding the optimal location and trajectory of a multilateral well in a reservoir simulation model using the industry standard ECLIPSE simulator. We compare our results to existing approaches and show that our technique is comparable, if not superior, in performance to these approaches.

Introduction

The extraction of oil from the subsurface is a challenging and increasingly important problem as we begin to exhaust existing easily recoverable reserves. With the oil price recently subject to great variation there is increasing emphasis on efficient oil extraction and reservoir development. Deciding when and where to drill new injection and production wells is an important factor in an optimal reservoir development strategy. Unfortunately, such an optimal strategy may be far from obvious as often the reservoir under consideration is a highly complex system. It is natural to formulate this decision problem as a global optimisation problem on a suitable reservoir simulation model. However, in order to evaluate a single strategy one has to run the numerical simulation model which can take a considerable amount of time. We are therefore faced with a global optimisation problem where the underlying objective function is expensive to evaluate.

For the purposes of this paper we are interested in finding the optimal locations and trajectories of new injection and production wells in an existing oil reservoir. To this end, we construct a simulation model of the oil reservoir and optimise over it to determine the best well locations and trajectories. This is not a new approach and has been tried by Yeten et al. (2002), Badru and Kabir (2003), Emerick et al. (2009) and Bukhamsin et al. (2010) amongst others. A frequently encountered problem with these approaches is that typically very many simulator runs are required to achieve a good optimum and often only a local optimum is found. Our approach seeks to remedy these shortcomings by developing a method which tries to use as few simulator runs as possible while at the same time aiming to find the global optimum. It may seem that this is a very ambitious goal since any global optimisation problem is NP-hard in the general case, however we believe that our approach is the best that one can do given the limited amount of information and time (i.e. simulator runs) available.

Before proceeding further, let us define some terminology. Let the term *simulator response* denote a real-valued function with a suitable measure of simulated reservoir performance as the output and optimisation parameters governing well placement as input. For example, the output could be the total oil production of the reservoir over some prescribed time period or the net present value of the project and may include uncertainty in the reservoir geology by way of multiple realisations of the reservoir. As the simulation model is computationally expensive to evaluate one often approximates the simulator response by a surrogate model (often called an emulator or proxy). We use an extension of the geostatistical technique of Kriging (Busby et al., 2007) but other surrogate models have also been used in the literature, including neural networks and splines see Zubarev, 2009 and Yeten et al., 2005 for a comparison. Our basic approach is as follows: Construct a surrogate approximation to the simulator response using an initial set of simulator runs (often termed initial designs, see Yeten et al., 2005), optimise a suitable loss function associated with the surrogate and run the simulator at the optimal parameter configuration. Then construct a new surrogate incorporating the new simulator run and repeat the process until some prescribed number of iterations or time limit is reached. The following sections deal with constructing the surrogate, our global optimisation algorithm and the well placement problem itself. We conclude with some numerical results of our method on a simple test case.

Surrogate Approximation

The main idea behind our method is to use an inexpensive surrogate function to approximate the underlying computationally expensive simulator response. The surrogate function enables us to capture the dominant features of the simulator response surface with relatively few simulator runs. In addition, we iteratively refine our surrogate approximation as the optimisation progresses so as better to capture the behaviour of the simulator response near the optimum and in relatively unexplored regions.

The surrogate is constructed as follows. Let $\{\pi_k\}_{k=1}^{M}$ be a polynomial basis of degree d on $\mathbb{R}^n$, so that $M = \frac{(d+n)!}{n!d!}$ and let $\mathcal{D} \subset \mathbb{R}^n$ be a compact subregion of interest. Suppose that the simulator response

$f\colon \mathcal{D} \to \mathbb{R}$ is a realisation of a stochastic process $F\colon \mathcal{D} \times \Omega \to \mathbb{R}$ of the form

$$F(x) = \sum_{k=1}^{M} \mu_k \pi_k(x) + Z(x)$$

where the first term specifies the mean structure and Z denotes a Gaussian stochastic process with zero mean and covariance

$$\operatorname{cov}(Z(x), Z(y)) = \sigma^2 R(x, y)$$

between $Z(x)$ and $Z(y)$. Here σ^2 denotes the process variance and

$$R(x, y) = \varphi(\|x - y\|)$$

is the correlation between $Z(x)$ and $Z(y)$. We take the correlation function $\varphi(\cdot)$ to be a radial basis function where $\|\cdot\| = \|W\cdot\|_2$ is the weighted ℓ_2-norm with weight matrix W (which is usually taken to be diagonal). Examples of radial basis functions include

$$\begin{aligned}
\text{Surface Spline:}\quad & \varphi(r) = \begin{cases} r^k & \text{if } k \text{ is odd} \\ r^k \log r & \text{if } k \text{ is even,} \end{cases} \\
\text{Multiquadric:}\quad & \varphi(r) = (r^2 + \gamma^2)^\beta \qquad \beta > 0,\ \beta \notin \mathbb{N}, \text{and} \\
\text{Gaussian:}\quad & \varphi(r) = \exp(-\gamma^2 r^2),
\end{aligned}$$

where γ is a nonzero constant referred to as the shape parameter (see Wendland, 2005). Assume there are N samples $y = (y_1, \ldots, y_N)^T$ of f at the corresponding sample points $x_1, \ldots, x_N$. The surrogate is taken to be the posterior mean $s(x)$ of the stochastic process F given the response function samples y. It can be shown (see Busby et al., 2007) that the posterior mean $s(x)$ of F given y is

$$s(x) := E[F(x)|y] = \sum_{k=1}^{M} \mu_k \pi_k(x) + \sum_{j=1}^{N} \lambda_j \varphi(\|x - x_j\|)$$

where μ_k, λ_j are coefficients. The coefficients are determined by solving the linear interpolation system

$$y_i = \sum_{k=1}^{M} \mu_k \pi_k(x_i) + \sum_{j=1}^{N} \lambda_j \varphi(\|x_i - x_j\|), \qquad i = 1, \ldots, N$$

along with the additional conditions

$$\sum_{j=1}^{N} \lambda_j \pi_k(x_j) = 0, \qquad k = 1, \ldots, M$$

which complete the system and ensure that polynomials of degree less than or equal to d are interpolated exactly. In matrix form this gives the non-singular (provided d is sufficiently large and $\{x_i\}_{i=1}^N$ is a unisolvent set, see Wendland, 2005) saddle-point system

$$\begin{pmatrix} R & P \\ P^T & 0 \end{pmatrix} \begin{pmatrix} \lambda \\ \mu \end{pmatrix} = \begin{pmatrix} y \\ 0 \end{pmatrix}$$

where $R = (R(x_i, x_j))$ and $P = (\pi_j(x_i))$. The posterior variance $e^2(x)$ of the stochastic process F given the response function samples y can be used as a measure of error in the surrogate approximation. One can show (see Sacks et al., 1989; Schonlau, 1997) that the posterior variance $e^2(x)$ of F given y is

$$e^2(x) := V[F(x)|y] = \sigma^2 \left[\varphi(0) - \begin{pmatrix} r(x) \\ p(x) \end{pmatrix}^T \begin{pmatrix} R & P \\ P^T & 0 \end{pmatrix}^{-1} \begin{pmatrix} r(x) \\ p(x) \end{pmatrix} \right]$$

where $r(x) = (\varphi(\|x - x_j\|))$ and $p(x) = (\pi_j(x))$. The process variance σ^2 is determined using maximum likelihood estimation (see Busby et al., 2007) and given by

$$\sigma^2 = \frac{1}{N}(y - P\mu)^T R^{-1}(y - P\mu).$$

Note that the geostatistical method of Kriging is simply the above approach with a Gaussian radial basis function and diagonally weighted ℓ_2-norm. There are two main approaches one can use to find the weights in the weight matrix W, which we will now assume to be diagonal. The first approach consists of choosing W to maximise the likelihood of the observed data $x_1, \dots, x_N$ and leads one to choose W to be the maximiser of the log-likelihood function (see Busby et al., 2007)

$$\ell(W) = -\frac{1}{2}(N \log \sigma^2 + \log \det(R))$$

which we optimise using a general purpose global optimisation algorithm (DIRECT, Jones et al., 1993). The second approach is to use leave-one-out cross-validation (see Rippa, 1999) and leads one to choose W to minimise the ℓ_2-norm of the cross-validation error $\epsilon(W) \in \mathbb{R}^N$. The k-th element of the cross-validation error $\epsilon(W)$ is the error at the validation point x_k, given by

$$\epsilon_k(W) = \frac{\lambda_k}{A^{-1}_{k,k}}$$

where $A^{-1}_{k,k}$ is the k-th diagonal element of the inverse of the interpolation matrix $A = \left(\begin{smallmatrix} R & P \\ P^T & 0 \end{smallmatrix}\right)$.

We will use the decision theoretic expected improvement approach for updating the surrogate approximation as outlined in Jones et al. (1998). Define the improvement $I(x)$ a new response function evaluation at x achieves over the current smallest response function value $y_{min} := \min_{1 \le i \le N} y_i$ by

$$I(x) = \max\{0, y_{min} - F(x)\}.$$

The next sample point is then chosen to maximise the *expected* improvement with respect to the posterior distribution $F|y$. It can be shown (see Jones et al., 1998) that the expected improvement is given by

$$E[I(x)] = e(x)\left[z(x)\Phi(z(x)) + \phi(z(x))\right]$$

where Φ and ϕ are the cumulative distribution function and probability density function respectively of the standard normal distribution and

$$z(x) = \frac{y_{min} - s(x)}{e(x)}.$$

In order to maximise the expected improvement using our preferred optimisation algorithm we require bounds on the posterior mean $s(x)$ and posterior variance $e^2(x)$. We will discuss how these bounds are obtained in the next section along with our choice of global optimisation algorithm.

Global Optimisation

As part of the solution to the general problem it is necessary to solve the ancillary global optimisation problem

$$\min_{x \in \mathcal{D}} l(x)$$

where $l \colon \mathcal{D} \subset \mathbb{R}^n \to \mathbb{R}$ is a specially constructed objective function (referred to in the decision theory literature as a loss function, see Berger, 1985) and $\mathcal{D}$ is a n-dimensional rectangle (or box). We present the canonical branch and bound algorithm (Horst, 1986) with bounds from Jones et al. (1998) and Gutmann (2001). Before describing the algorithm in detail we first need some notation. Let $\mathcal{B} \subset \mathcal{D}$ denote an n-dimensional box and let $x_{\mathcal{B}}$ be the centre of $\mathcal{B}$ unless $\mathcal{B}$ has been locally searched, in which case it is the feasible point found by a constrained local search procedure (see step 1e in the algorithm below). Define $\alpha(\mathcal{B}), \beta(\mathcal{B})$ to be lower and upper bounds respectively on the global minimum of $l(x)$, i.e. $\alpha(\mathcal{B}), \beta(\mathcal{B})$ satisfy

$$\alpha(\mathcal{B}) \leq \min_{x \in \mathcal{B}} l(x) \leq \beta(\mathcal{B}).$$

Suitable choices for these bounds will be discussed below. We then follow the branch and bound algorithmic framework set out in Balakrishnan et al. (1991):

Branch and Bound Algorithm

```
0. Initialisation:
    (a) Set $k = 0$ and $s = 0$.
    (b) Let $\mathcal{L}_0$ be the initial list of boxes and set $\mathcal{L}_0 = \{\mathcal{D}\}$.
    (c) Let $L_0 = \alpha(\mathcal{D})$ be the initial lower bound for $\min_{x\in\mathcal{D}} l(x)$.
    (d) Let $U_0 = \beta(\mathcal{D})$ be the initial upper bound for $\min_{x\in\mathcal{D}} l(x)$.
1. While $U_k - L_k > \varepsilon$, repeat the following procedure:
    (a) Remove from $\mathcal{L}_k$ boxes $\mathcal{B} \in \mathcal{L}_k$ such that $\alpha(\mathcal{B}) > U_k$.
    (b) Choose $\mathcal{B} \in \mathcal{L}_k$ such that $\alpha(\mathcal{B}) = L_k$.
    (c) Bisect $\mathcal{B}$ along its longest edge into $\mathcal{B}_I$ and $\mathcal{B}_{II}$. Set $\mathcal{L}_{k+1} := \mathcal{L}_k \cup \{\mathcal{B}_I, \mathcal{B}_{II}\}$ and remove $\mathcal{B}$
        from $\mathcal{L}_{k+1}$.
    (d) If any boxes have been discarded set $s = 0$, otherwise set $s = s + 1$.
    (e) If $s > 2$ run an approximate constrained local search algorithm on all previously unsearched
        boxes $\mathcal{B}$ in $\mathcal{L}_{k+1}$, update $x_{\mathcal{B}}$ to be the minimiser found by the local search and set $s = 0$.
    (f) Set $L_{k+1} := \min_{\mathcal{B}\in\mathcal{L}_{k+1}} \alpha(\mathcal{B})$.
    (g) Set $U_{k+1} := \min_{\mathcal{B}\in\mathcal{L}_{k+1}} \beta(\mathcal{B})$.
    (h) Set $k = k + 1$.
```

The idea behind the algorithm is to recursively partition the domain $\mathcal{D}$ into sub-boxes until a box of sufficiently small size containing the global minimum of $l(x)$ over $\mathcal{D}$ is found. Since it is possible to obtain bounds on the minimum of $l(x)$ over any box in $\mathcal{D}$, they can be used to discard boxes which cannot contain the global minimum, i.e. boxes whose lower bound is greater than the smallest upper bound. The algorithm is accelerated by running constrained local searches on suitable boxes to obtain more accurate upper bounds. This is achieved through the use of a heuristic from Pedamallu et al. (2008) which suggests running local searches on all previously unsearched boxes if no boxes are discarded after two successive iterations of the algorithm. We use a conjugate gradient based active set method by Hager and Zhang (2006) for the local searches but in principle one can use any constrained local search algorithm.

For our particular example we let $l(x)$ be the negative expected improvement

$$l(x) = -e(x)\left[z(x)\Phi(z(x)) + \phi(z(x))\right]$$

but note that many other choices of loss function are possible (see Jones, 2001; Sasena, 2002). Define the upper bound $\beta(\mathcal{B})$ to be l evaluated at $x_{\mathcal{B}}$ i.e.

$$\beta(\mathcal{B}) = l(x_{\mathcal{B}}).$$

To obtain a lower bound on $l(x)$ it suffices to obtain a lower bound on $s(x)$ and an upper bound on $e(x)$. For the surrogate

$$s(x) = \sum_{k=1}^{M} \mu_k \pi_k(x) + \sum_{j=1}^{N} \lambda_j \varphi(\|x - x_j\|)$$

we find a lower bound $p(\mathcal{B})$ for the minimum of the polynomial term over a box $\mathcal{B}$ using interval arithmetic (Neumaier, 2004) or polynomial optimisation (Lasserre, 2001) and bound each radial basis function term over $\mathcal{B}$ using quadratic functions (as in Jones et al., 1998)

$$a_j + b_j\|x - x_j\|^2 \le \varphi(\|x - x_j\|) \le A_j + B_j\|x - x_j\|^2$$

for which details are given in Gutmann (2001). Now we can define the lower bound $\alpha_s(\mathcal{B})$ for the global minimum of $s(x)$ over a box $\mathcal{B}$ as

$$\alpha_s(\mathcal{B}) = p(\mathcal{B}) + \min_{x \in \mathcal{B}} \left\{ \sum_{\substack{j=1 \\ \lambda_j > 0}}^{N} \lambda_j (a_j + b_j\|x - x_j\|^2) + \sum_{\substack{j=1 \\ \lambda_j < 0}}^{N} \lambda_j (A_j + B_j\|x - x_j\|^2) \right\}.$$

Note that the minimum above is easy to obtain as the sum can be rewritten in terms of one dimensional quadratic functions and thus minimised componentwise. We now turn our attention to $e(x)$ which is defined as

$$e^2(x) = \sigma^2 \left[\varphi(0) - L(\xi(x))\right]$$

where

$$L(\xi) = \xi^T A^{-1} \xi, \quad A = \begin{pmatrix} R & P \\ P^T & 0 \end{pmatrix}, \quad \xi(x) = (\, r(x) \; p(x)\,)^T.$$

The key idea here is to underestimate $L(\xi)$ by a convex relaxation $C(\xi)$. Following Jones et al. (1998) and Gutmann (2001), define an upper bound $\beta_e(\mathcal{B})$ for $e(x)$ over $\mathcal{B}$ as

$$\beta_e(\mathcal{B}) = \max_{\xi \in [l,u]} \sigma \left[\varphi(0) - C(\xi)\right]^{1/2}$$

where

$$C(\xi) := L(\xi) - \lambda_{min}(\xi - l)^T(\xi - u)$$

is the convex relaxation of $L(\xi)$. Here λ_{min} denotes the smallest eigenvalue of A^{-1} and

$$l_j = \begin{cases} \varphi\left(\sqrt{r_j^l(\mathcal{B})}\right) & \text{for } j = 1, \dots, N \\ p_j^l(\mathcal{B}) & \text{for } j = N+1, \dots, N+M \end{cases}$$

$$u_j = \begin{cases} \varphi\left(\sqrt{r_j^u(\mathcal{B})}\right) & \text{for } j = 1, \dots, N \\ p_j^u(\mathcal{B}) & \text{for } j = N+1, \dots, N+M \end{cases}$$

where $p_j^l(\mathcal{B}), p_j^u(\mathcal{B})$ are lower and upper bounds on the minimum of the polynomial basis terms π_{j-n} over $\mathcal{B}$ found using interval arithmetic or polynomial optimisation and $r_j^l(\mathcal{B}), r_j^u(\mathcal{B})$ are exact lower and upper bounds for the minimum of $\|x - x_j\|^2$ over $\mathcal{B}$ as in Gutmann (2001). Note that $C(\xi)$ is convex as its Hessian, given by $2(A^{-1} - \lambda_{min} I)$, is positive semidefinite and it clearly underestimates $L(\xi)$. Thus we have relaxed the problem to a convex quadratic programming problem which can be solved efficiently; we use an implementation of the algorithm by Goldfarb and Idnani (1983). Combining the lower bound $\alpha_s(\mathcal{B})$ and upper bound $\beta_e(\mathcal{B})$ gives the lower bound $\alpha(\mathcal{B})$ as

$$\alpha(\mathcal{B}) = \beta_e(\mathcal{B}) \left[\zeta(\mathcal{B})\Phi(\zeta(\mathcal{B})) + \phi(\zeta(\mathcal{B}))\right]$$

where

$$\zeta(\mathcal{B}) = \frac{y_{min} - \alpha_s(\mathcal{B})}{\beta_e(\mathcal{B})}.$$

Well Placement

For the purposes of this paper we consider arbitrary multilateral wells with up to n laterals on a Cartesian simulation grid. The aim is to find optimal completion locations and trajectories for these wells. For simplicity, we assume the main wellbore is completed along a single continuous section of the well and any laterals are completed along their entire length. We parameterise the wells similarly to Yeten et al. (2002) in continuous grid coordinates, i.e. if a point has coordinates $(1.5, 2.5, 3.5)$ it is in the centre of the $(1, 2, 3)$ grid block. Let h_0 denote the grid coordinates of the mainbore completion heel and t_0 of the mainbore completion toe. Also, let $l_i \in [0, 1]$ denote the relative position of the i-th lateral heel on the mainbore and t_i denote the grid coordinates of the i-th lateral toe. The coordinates of the i-th lateral heel can then be calculated as $h_i = h_0 + l_i(t_0 - h_0)$. This parameterisation is illustrated in Figure 1 below. Denote the input parameters by $x = (h_0, t_0, l_1, t_1, \ldots, l_n, t_n)$ and let f be the simulator response to these parameters. The optimisation problem then becomes

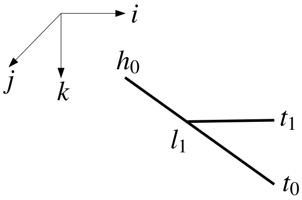

Figure 1: *Parameterisation of a multilateral well with a single lateral off the mainbore in continuous grid coordinates i, j, k.*

$$\min_{x \in \mathcal{D}} f(x)$$

i.e. minimise the simulator response over a suitably defined rectangular coordinate range $\mathcal{D}$. As the simulator response is an expensive function we approximate it using a surrogate function $s(x)$ and associated error $e(x)$ as described in the previous section on Surrogate Approximation. We iteratively improve the surrogate approximation by adding the optimiser of the associated loss function as described in the previous section on Global Optimisation. Our complete framework is as follows:

Well Placement Optimisation Framework

0. Initialisation:

- *(a) Set $t = 0$ and let T be a prescribed number of total iterations.*
- *(b) Let $f : \mathcal{D} \subset \mathbb{R}^n \to \mathbb{R}$ be the simulator response on a rectangular domain $\mathcal{D}$.*
- *(c) Let $\mathcal{X}_0 \subset \mathcal{D}$ denote an initial set of points where f has been sampled.*

1. While $t < T$, repeat the following procedure:

- *(a) Construct a surrogate approximation s_t to the simulator response f at $\mathcal{X}_t$ along with the corresponding error e_t.*
- *(b) Construct the corresponding loss function $l_t(x) = -e_t(x)\,[z_t(x)\Phi(z_t(x)) + \phi(z_t(x))]$, where $z_t(x) = [y_{min} - s_t(x)]/e_t(x)$.*
- *(c) Optimise $l_t(x)$ using the prescribed branch and bound algorithm to obtain a minimiser x_t and evaluate the simulator response at x_t.*
- *(d) Set $\mathcal{X}_{t+1} := \mathcal{X}_t \cup \{x_t\}$.*
- *(e) Set $t = t + 1$.*

The algorithm stops after a prescribed number of iterations T, but it is of course possible to use other stopping criteria (see Sasena, 2002) with perhaps the most obvious being to stop after a certain amount of elapsed time. It is assumed that the domain $\mathcal{D}$ of the simulator response is rectangular so that we can dispense with the complexity that arises when one has to incorporate constraints into a branch and bound algorithm (see Pedamallu et al., 2008). The initial set of points $\mathcal{X}_0$ is typically chosen to be space filling, see Yeten et al. (2005) and Koehler and Owen (1996) for further details on this. As mentioned previously, it is possible to use other loss functions (such as $l_t \equiv s_t$) and we refer the interested reader to Jones (2001) and Sasena (2002) for examples.

Numerical Example

We will now demonstrate the effectiveness of our well placement optimisation framework by means of a simple numerical example. Consider the case of finding the optimal location and trajectory of a single multilateral oil producing well in a synthetic reservoir simulation model. We will use the Snark simulation model, pictured in Figure 2, which consists of a 24 × 25 × 12 corner point simulation grid representing 12 geological layers with three faults and an analytical aquifer at the southern end of the model. The simulation will be performed using the industry standard ECLIPSE simulator from Schlumberger. The aim is to find the optimal mainbore completion along with up to three laterals. We will compare two different methods for achieving this aim, our surrogate based optimisation framework and a direct approach using genetic algorithms (c.f. Yeten et al., 2002). Let the simulator response f be the negative total oil production from the simulation model over a four year period. Note that we do not consider other reservoir performance factors here as our primary focus is on the optimisation problem. We will start with a maximin latin hypercube sample (see Koehler and Owen, 1996) of 20 initial well configurations and we limit ourselves to 200 simulator runs. For our optimisation framework we will use the cubic spline radial basis function $\varphi(r) = r^3$ with a diagonally weighted ℓ_2-norm and linear polynomial term, as it is numerically more stable than Kriging using the Gaussian radial basis function (see Section 12.2 of Wendland, 2005). Note that the linear polynomial term in the surrogate can be trivially optimised componentwise over a box, so in this case there is no need to use interval arithmetic or polynomial optimisation. We use our own MATLAB based implementation of the optimisation framework and the genetic algorithm provided by MATLAB with default settings. The simulated oil production of the best well found after a given number of simulator runs by our framework and the genetic algorithm is shown in Figure 3. One can see that even for this simple example our framework finds a multilateral well with consistently better oil production after a given number of simulator runs than the genetic algorithm. The trajectory of the optimal well as found by our optimisation framework is shown in Figure 4. This is

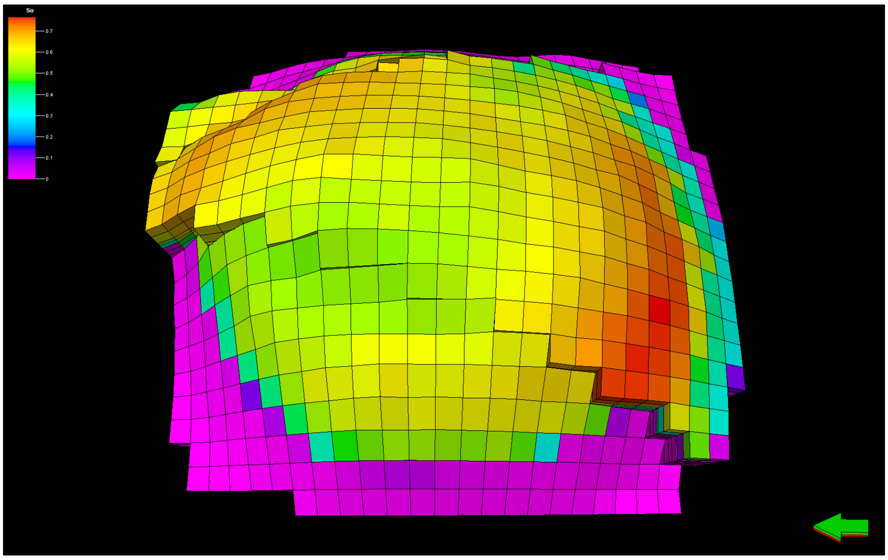

Figure 2: *A top-down view of the Snark simulation model showing the initial oil saturation ranging from* 0% *(purple) to* 80% *(red) as indicated by the colourmap in the top left-hand corner.*

essentially a horizontal well which maximises contact with areas of high oil saturation and is the type of optimal well one would expect given our choice of simulation model.

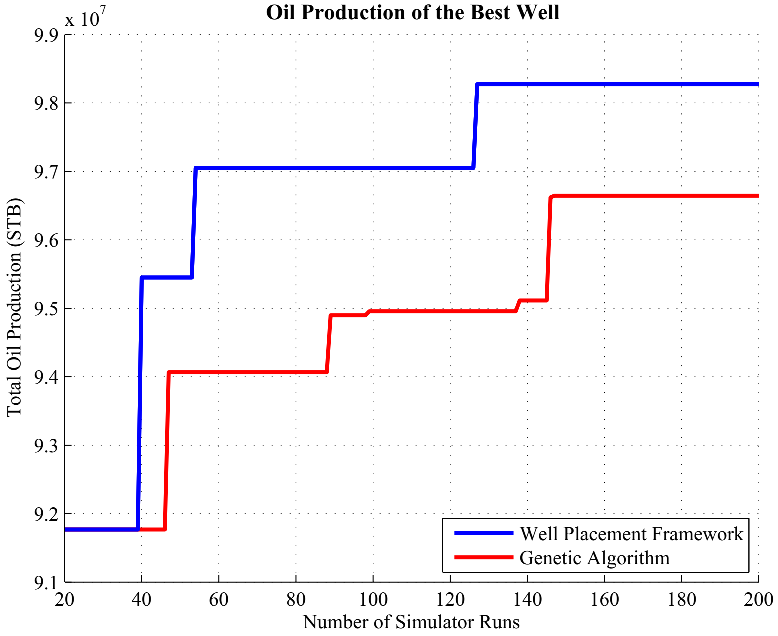

Figure 3: *Simulated oil production of the best multilateral well found after the given number of simulator runs by our framework and the genetic algorithm.*

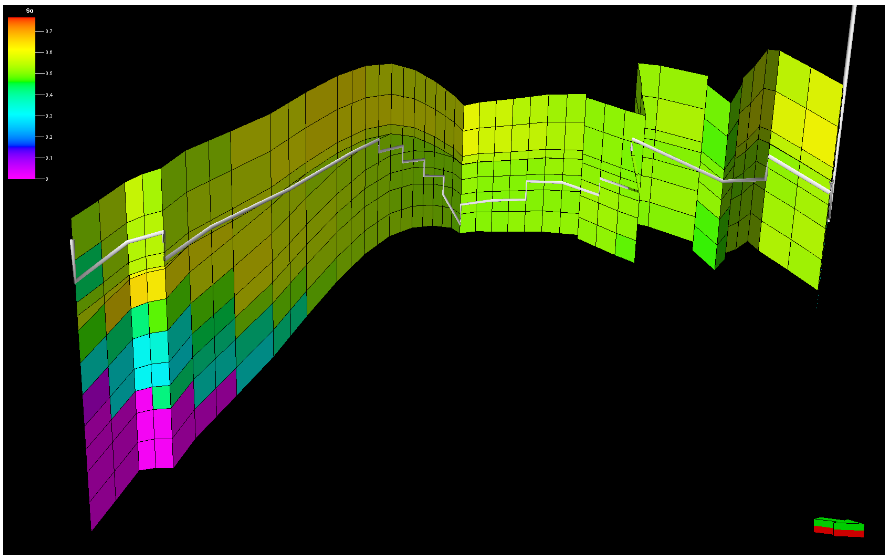

Figure 4: *The optimal well trajectory as found by our optimisation framework pictured on a vertical slice through the model. The colours depict the initial oil saturation ranging from* 0% *(purple) to* 80% *(red) as indicated by the colourmap in the top left-hand corner.*

Conclusions

We have presented a new approach to the multilateral well placement problem based on the theory of expensive function optimisation. Our approach consists of a general surrogate-based framework for the optimisation of a number of injection and production wells with multiple laterals along with a branch and bound global optimisation algorithm. We have tested our method on a simple reservoir simulation model and results indicate that our approach is comparable, if not superior, to that of genetic algorithms used in previous studies.

Acknowledgements

We wish to thank Schlumberger for their helpful assistance and for providing us with access to their reservoir simulation software for the numerical examples. We would also like to thank Mohammad Farshi for his advice on interfacing our code with the reservoir simulator. This research was supported through an EPSRC Industrial CASE studentship in conjunction with Schlumberger.

References

Badru, O. and Kabir, C.S. [2003] Well Placement Optimization in Field Development. *SPE 84191 presented at the SPE Annual Technical Conference and Exhibition, Denver, Colorado, 5–8 October*, http://dx.doi.org/10.2118/84191-MS.

Balakrishnan, V., Boyd, S. and Balemi, S. [1991] Branch and Bound Algorithm for Computing the Minimum Stability Degree of Parameter-Dependent Linear Systems. *International Journal of Robust and Nonlinear Control*, **1**(4), 295–317, http://dx.doi.org/10.1002/rnc.4590010404.

Berger, J.O. [1985] *Statistical Decision Theory and Bayesian Analysis*. Springer, 2nd edn., ISBN 978-0-387-96098-2.

Bukhamsin, A.Y., Aziz, K. and Farshi, M.M. [2010] Optimization of Multilateral Well Design and Location in a Real Field using a Continuous Genetic Algorithm. *SPE 136944 presented at the SPE/DGS Saudi Arabia Section Technical Symposium and Exhibition, Al-Khobar, Saudi Arabia, 4–7 April*, http://dx.doi.org/10.2118/136944-MS.

Busby, D., Farmer, C.L. and Iske, A. [2007] Hierarchical Nonlinear Approximation for Experimental Design and Statistical Data Fitting. *SIAM Journal on Scientific Computing*, **29**(1), 49–69, http://dx.doi.org/10.1137/050639983.

Emerick, A.A. et al. [2009] Well Placement Optimization using a Genetic Algorithm with Nonlinear Constraints. *SPE 118808 presented at the SPE Reservoir Simulation Symposium, The Woodlands, Texas, 2–4 Feburary*, http://dx.doi.org/10.2118/118808-MS.

Goldfarb, D. and Idnani, A. [1983] A Numerically Stable Dual Method for Solving Strictly Convex Quadratic Programs. *Mathematical Programming*, **27**(1), 1–33, http://dx.doi.org/10.1007/BF02591962.

Gutmann, H.M. [2001] *Radial Basis Function Methods for Global Optimization*. Ph.D. thesis, DAMTP, University of Cambridge.

Hager, W.W. and Zhang, H. [2006] A New Active Set Algorithm for Box Constrained Optimization. *SIAM Journal on Optimization*, **17**(2), 526–557, http://dx.doi.org/10.1137/050635225.

Horst, R. [1986] A General Class of Branch-and-Bound Methods in Global Optimization with Some New Approaches for Concave Minimization. *Journal of Optimization Theory and Applications*, **51**(2), 271–291, http://dx.doi.org/10.1007/BF00939825.

Jones, D.R. [2001] A Taxonomy of Global Optimization Methods Based on Response Surfaces. *Journal of Global Optimization*, **21**(4), 345–383, http://dx.doi.org/10.1023/A:1012771025575.

Jones, D.R., Perttunen, C.D. and Stuckman, B.E. [1993] Lipschitzian Optimization Without the Lipschitz Constant. *Journal of Optimization Theory and Applications*, **79**(1), 157–181, http://dx.doi.org/10.1007/BF00941892.

Jones, D.R., Schonlau, M. and Welch, W.J. [1998] Efficient Global Optimization of Expensive Black-Box Functions. *Journal of Global Optimization*, **13**(4), 455–492, http://dx.doi.org/10.1023/A:1008306431147.

Koehler, J.R. and Owen, A.B. [1996] Computer Experiments. In: Ghosh, S. and Rao, C.R. (Eds.) *Handbook of Statistics*. Elsevier, vol. 13, 261–308, http://dx.doi.org/10.1016/S0169-7161(96)13011-X.

Lasserre, J.B. [2001] Global Optimization with Polynomials and the Problem of Moments. *SIAM Journal on Optimization*, **11**(3), 796–817, http://dx.doi.org/10.1137/S1052623400366802.

Neumaier, A. [2004] Complete Search in Continuous Global Optimization and Constraint Satisfaction. *Acta Numerica*, **13**, 271–369, http://dx.doi.org/10.1017/S0962492904000194.

Pedamallu, C.S., Özdamar, L., Csendes, T. and Vinkó, T. [2008] Efficient Interval Partitioning for Constrained Global Optimization. *Journal of Global Optimization*, **42**(3), 369–384, http://dx.doi.org/10.1007/s10898-008-9297-7.

Rippa, S. [1999] An Algorithm for Selecting a Good Value for the Parameter c in Radial Basis Function Interpolation. *Advances in Computational Mathematics*, **11**(2), 193–210, http://dx.doi.org/10.1023/A:1018975909870.

Sacks, J., Welch, W.J., Mitchell, T.J. and Wynn, H.P. [1989] Design and Analysis of Computer Experiments. *Statistical Science*, **4**(4), 409–435, http://dx.doi.org/10.1214/ss/1177012413.

Sasena, M.J. [2002] *Flexibility and Efficiency Enhancements for Constrained Global Design Optimization with Kriging Approximations*. Ph.D. thesis, University of Michigan.

Schonlau, M. [1997] *Computer Experiments and Global Optimization*. Ph.D. thesis, University of Waterloo.

Wendland, H. [2005] *Scattered Data Approximation*. Cambridge Monographs on Applied and Computational Mathematics, Cambridge University Press, ISBN 978-0-521-84335-5.

Yeten, B., Castellini, A., Guyaguler, B. and Chen, W. [2005] A Comparison Study on Experimental Design and Response Surface Methodologies. *SPE 93347 presented at the SPE Reservoir Simulation Symposium, The Woodlands, Texas, 31 January – 2 Feburary*, http://dx.doi.org/10.2118/93347-MS.

Yeten, B., Durlofsky, L.J. and Aziz, K. [2002] Optimization of Nonconventional Well Type, Location and Trajectory. *SPE 77565 presented at the SPE Annual Technical Conference and Exhibition, San Antonio, Texas, 29 September – 2 October*, http://dx.doi.org/10.2118/77565-MS.

Zubarev, D.I. [2009] Pros and Cons of Applying Proxy-models as a Substitute for Full Reservoir Simulations. *SPE 124815 presented at the SPE Annual Technical Conference and Exhibition, New Orleans, Louisiana, 4–7 October*, http://dx.doi.org/10.2118/124815-MS.