Joint Optimization of Oil Well Placement and Controls

Mathias C. Bellout*

Center for Integrated Operations in the Petroleum Industry,
Dept. of Petroleum Engineering and Applied Geophysics,
NTNU,
S.P. Andersens vei 15A, 7491 Trondheim, Norway
E-mail: mathias.bellout@ntnu.no
*Corresponding author

David Echeverría Ciaurri

T.J Watson Research Center,
IBM,
1101 Kitchawan Road, Yorktown Heights, NY 10598, USA
E-mail: decheve@us.ibm.com

Louis J. Durlofsky

Dept. of Energy Resources Engineering,
Stanford University,
Green Earth Sciences Building, Stanford, CA 94305, USA
E-mail: lou@stanford.edu

Bjarne Foss† and Jon Kleppe‡

Center for Integrated Operations in the Petroleum Industry,
†Dept. of Engineering Cybernetics,
‡Dept. of Petroleum Engineering and Applied Geophysics,
NTNU
S.P. Andersens vei 15A, 7491 Trondheim, Norway
E-mail: bjarne.foss@itk.ntnu.no
E-mail: kleppe@ipt.ntnu.no

Abstract: Well placement and control optimization in oil field development are commonly performed sequentially. In this work we propose a joint approach that embeds well control optimization within the search for optimum well placement configurations. We solve for well placement using parallelizable direct search methods based on pattern search. Control optimization is solved by sequential quadratic programming using gradients efficiently computed through adjoints. Joint optimization yields a significant increase, of up to 18% in net present value, when compared to reasonable sequential approaches. The joint approach does, however, require about an order of magnitude increase in the number of objective function evaluations compared

to sequential procedures. This increase is somewhat mitigated by the parallel implementation of several of the pattern search approaches.

Keywords: Oil production optimization, optimal oil well placement, simulation-based optimization, derivative-free optimization, adjoint methods, mixed-integer nonlinear programming.

Reference to this paper should be made as follows: Bellout, M., Echeverría Ciaurri, D., Durlofsky, L.J., Foss, B.A., and Kleppe, J. (2011) 'Joint Optimization of Oil Well Placement and Controls', *Int. J. Mathematical Modelling and Numerical Optimisation*, vol. x, no. x, pp. xxx–xxx.

Biographical notes: Mathias Bellout received his MS in Applied Physics and MS in Petroleum Engineering from the Norwegian University of Science and Technology (NTNU) in 2005 and 2008, respectively. He is currently pursuing his PhD in Petroleum Engineering and Applied Geophysics from the same institution. Since 2009 he has been affiliated with the Center for Integrated Operations in the Petroleum Industry (IO-Center) at NTNU. In 2010-2011 he was a visiting researcher at the Smart Fields Consortium in the Department of Energy Resources Engineering at Stanford University. His research interests include reservoir modeling and simulation, and derivative-free and bi-level optimization.

David Echeverría Ciaurri received the Telecommunication Engineering degree and the MS degree in Applied Mathematics from the University of Zaragoza (Spain) in 1998 and 2002, respectively. In 2007 he obtained a PhD in Numerical Analysis from the University of Amsterdam (UvA). He has worked at the Dutch National Research Institute for Mathematics and Computer Science (CWI), at the Lawrence Livermore National Laboratory (LLNL), and at Stanford University. Since 2011 he has been affiliated within the Department of Petroleum and Energy Analytics at the IBM Thomas J. Watson Research Center. His research interests are efficient optimization and modeling and signal processing.

Louis J. Durlofsky is the Otto N. Miller professor and chairman in the Department of Energy Resources Engineering at Stanford University. He was previously affiliated with Chevron Energy Technology Company. He holds a BS degree from Penn State and MS and PhD degrees from the Massachusetts Institute of Technology, all in Chemical Engineering. His research interests include optimization for subsurface flow problems, upscaling of detailed geological models for flow simulation, reduced-order modeling, geological carbon sequestration, and general reservoir simulation.

Bjarne Foss is a professor in the Department of Engineering Cybernetics at the Norwegian University of Science and Technology (NTNU). He is currently a program manager at the Center for Integrated Operations at NTNU. He received his PhD degree from the Norwegian Institute of Technology in 1987 and has been a visiting professor at the University of Texas at Austin. His main interests are in modeling, control and

optimization theory, and applications of such methods within the process and petroleum industries.

Jon Kleppe is a professor in the Department of Petroleum Engineering and Applied Geophysics at the Norwegian University of Science and Technology (NTNU). He is also Director of the Center for Integrated Operations at NTNU. He received his PhD degree in Petroleum Engineering from Texas A&M University in 1974. He has been a visiting professor at Stanford University and a visiting researcher at Elf in Pau. His research interests are reservoir simulation, improved oil recovery, and integrated operations.

1 Introduction

The development of new fields for oil and gas production is increasingly difficult and expensive. Sustaining profitable production in mature fields, where water production rates are often high, also poses a challenge. For both sets of problems, it may be difficult to achieve adequate returns on investment using traditional (heuristic) production management techniques. There is, therefore, a growing interest in the development of efficient and effective simulation-based optimization procedures for well planning and operation.

This work focuses on maximizing revenue from oil production using water-flooding by optimizing medium-to-long-term (several years) field management operations such as well placement and well control scheduling. Water-flooding, where the oil in the subsurface formation (reservoir) is driven towards production wells by a moving waterfront created by water injection wells, is a common procedure for oil production. Substantial oil volumes are often bypassed during water-flooding due to the existence of complicated geological conditions, such as high-flow regions and faults, in the reservoir. Thus, for water-flooding to be effective, the locations and control schedules of injectors and producers must be selected in an optimal manner (by control schedule we mean the well rates or pressures as a function of time). Here our objective function is net present value of the asset, though other cost functions such as total oil recovered could also be used. In either case the cost function is computed by means of the numerical solution of the system of partial differential equations that describes fluid flow in the reservoir. The required simulations are typically computationally demanding, which poses challenges for optimization.

Under current procedures, the determination of well placement and well control is generally treated in a sequential manner. In some approaches, a heuristic strategy for well control is used during the well placement optimization. One such approach, referred to as a 'reactive control' strategy, entails closing production wells according to an economic threshold that depends on the oil price and the water production cost. This approach can be reasonably effective but is clearly suboptimal as it does not adjust injection wells and handles production wells as either fully open or closed. Any approach that does not consider well location and control jointly cannot be

expected to yield optimal solutions, since it does not capture the interdependency between the well configuration and the associated controls.

In this paper we propose a joint approach for optimization of well position and control settings. In our approach the two different optimizations are considered in a nested fashion; i.e., well placement optimization alternates iteratively with control optimization. The upper-level optimization refers to well location. For a given well location, a control optimization is performed, and the resulting optimum serves as the objective function for the upper-level optimization problem. As a consequence, the solution of the upper-level problem is a local optimum not only for well placement, but also for the controls because optimality of the lower-level (well control) problem is intrinsically inherited. Hence, this joint approach can be used to compute solutions with lower objective function values (for a minimization problem) than that achieved using sequential methodologies. The computational cost associated with the joint approach is, however, much higher since every upper-level function evaluation requires the optimization of the lower-level problem.

This nested approach is consistent with the two different types of problems, well placement and control, that constitute the joint optimization. Well control optimization is commonly stated in terms of continuous variables (well flow rates or pressures) and has been observed to present smooth optimization landscapes with multiple optima but similar cost function values, see e.g., Jansen et al. (2005) and Echeverria Ciaurri et al. (2011a). On the other hand, well placement optimization can be formulated as an integer optimization problem (where integers correspond to specific grid blocks) with non-smooth objective functions (see e.g., Onwunalu and Durlofsky, 2010) containing multiple optima with significantly different cost functions. This non-smooth character is generally related to the strong variability (heterogeneity) in subsurface flow properties. Therefore, many of the existing well placement optimization procedures attempt a more global search. Consistent with these observations, well control optimization is often addressed using gradient-based techniques (where gradients are computed rapidly via adjoint procedures; see e.g., Jansen et al., 2005, or Sarma et al., 2006), while well placement optimizations usually use derivative-free algorithms or stochastic search procedures (Yeten et al., 2003; Onwunalu and Durlofsky, 2010). These latter approaches ordinarily require parallel computing implementations for efficiency. We note, however, that gradient-based techniques have been applied for well placement (e.g., Sarma and Chen, 2008; Zandvliet et al., 2008), and stochastic search has been used for well control (e.g., Echeverria Ciaurri et al., 2011a), so our observations here should not be viewed as absolute.

To our knowledge, no research has been published addressing in detail the joint optimization of oil well placement and control. There have, however, been approaches that use the reactive control strategy described above in well placement optimizations (e.g., Zandvliet et al., 2008). The work introduced in Wang et al. (2007), and later enhanced in Zhang et al. (2010) and Forouzanfar et al. (2010), aims primarily at well placement, and integrates indirect mechanisms for optimizing well controls. The method described in that work provides a comprehensive optimization framework, but it involves a number of heuristics and does not treat explicitly location and control as optimization variables. The approach presented in this paper attempts to address the complicated joint well placement and control optimization problem from a mathematically sound perspective.

Other variables besides the location and controls for each well, such as the number of wells and the length of the water-flooding process, could also be included in the optimization. These variables are more difficult to treat, however, since the dimension of the optimization problem we are considering depends on these parameters. Another important effect not included in this work is uncertainty in the reservoir model; i.e., the optimization approaches studied here do not involve stochastic programming considerations. A general method for optimizing well location under uncertainty has been developed by Wang et al. (2011), and this approach could be extended in a straightforward manner to also include well controls.

This paper is structured as follows. The governing equations for the flow of oil and water in subsurface reservoirs are given in Section 2. Next, in Section 3, we present the general problem statement and the specific formulations for the well control and the well placement parts of the optimization procedure. The joint and sequential approaches used to solve the coupled system are also described. These approaches are applied to two example cases in Section 4. The first case addresses the control optimization of one injector and four producers and the optimal positioning of the injector. For this case we are able to perform exhaustive computations, which enable clear assessments of the various optimization procedures. In the second case the well position and controls for three producers and two injectors are optimized. Section 5 provides a summary and some suggestions for future research.

2 Problem Statement

In this section we briefly describe the flow simulations used to evaluate well location and control scenarios. The optimization procedures used in this work are then presented in detail.

2.1 Governing equations for reservoir production

Hydrocarbons such as oil and gas are found within porous rock in subsurface formations. The equations that describe fluid flow in the reservoir are derived by combining expressions of mass conservation with constitutive and thermodynamic relationships. For clarity, our description here entails several simplifications (such as the assumption of incompressible flow), though in the problems considered later compressibility and other effects are included. See, e.g., Aziz and Settari (1979) or Ertekin et al. (2001) for details on the flow equations and numerical discretizations.

We consider two-phase systems containing oil (o) and water (w). Mass conservation for each fluid i (where $i = o,\ w$) is given by:

$$\nabla \cdot \mathbf{u}_i + q_i = -\frac{\partial}{\partial t}\left(\phi S_i\right), \tag{1}$$

where $\mathbf{u}_i$ is the Darcy (or superficial) velocity, q_i represents source/sink terms, ϕ is porosity (volume fraction of the rock that can be occupied by fluids), S_i is the

volume fraction of the pore space occupied by fluid i ($S_o + S_w = 1$), and t is time. Darcy velocity is expressed as:

$$\mathbf{u}_i = -\mathbf{k}\frac{k_{r,i}}{\mu_i}\nabla p, \quad i = o,\ w,\ , \tag{2}$$

where p is pressure (here assumed the same for oil and water), μ_i is the viscosity of phase i, $\mathbf{k}$ is the absolute (rock) permeability tensor (which is essentially a flow conductivity), and $k_{r,i}(S_i)$ is the relative permeability to phase i. Relative permeability quantifies how phase i flows in the presence of another phase. It is in general a nonlinear function of saturation, which introduces an important nonlinearity into the equations.

Combining (1) and (2) yields the following flow equations:

$$\nabla\left[\mathbf{k}\frac{k_{r,i}}{\mu_i}\nabla P_i\right] - q_i = \frac{\partial}{\partial t}\left(\phi S_i\right), \quad i = o,\ w. \tag{3}$$

In (3) the pressure p and saturation S_w are the primary unknowns. These equations comprise a set of nonlinear partial differential equations, which in practice are solved numerically. Model sizes usually range from tens of thousands of grid blocks for small models, to several hundred thousand or millions of grid blocks for large models. A typical model might require several hundred time steps.

Reservoir models are coupled to well models (via the source term q_i) to enable the computation of the volumes of fluids produced and injected at each time step. See Peaceman (1978) or Ertekin et al. (2001) for a description of how this is accomplished. Current well designs may involve vertical, horizontal, deviated and multilateral wells. These wells can be controlled by specifying either rates or bottom-hole pressure (BHP). In this work, we will consider only vertical wells and use BHPs at various time intervals for the well control optimization parameters. The simulator used in this work is Stanford's General Purpose Research Simulator (GPRS; Cao, 2002).

2.2 Optimization problem

In our examples we consider oil-water systems with production driven by water-flooding. We seek to determine the optimal locations and BHP controls for a specified number of production and water injection wells using an optimization procedure based on a joint, rather than a sequential, approach.

The optimization problem studied here is defined as follows:

$$\min_{\mathbf{x}\in\mathbb{Z}^{n_1},\mathbf{u}\in\mathbb{R}^{n_2}} -\mathrm{NPV}\left(\mathbf{x},\mathbf{u}\right) \quad \text{subject to} \quad \begin{cases}\mathbf{x}_d \le \mathbf{x} \le \mathbf{x}_u \\ \mathbf{u}_d \le \mathbf{u} \le \mathbf{u}_u\end{cases}, \tag{4}$$

where $\mathbf{x}$ denotes the discrete well placement variables and $\mathbf{u}$ are the continuous well control variables. Well placement variables are intrinsically real but are often treated as integers, since reservoir simulators require wells to be assigned to discrete grid blocks in the model. Consequently, $\mathbf{x}$ is discrete-valued in many cases. All wells in this work are assumed to be vertical, hence well positions can be stated in terms of discrete areal coordinates (x, y) only. Thus $n_1 = 2(N_p + N_i)$, where N_p and N_i

are the number of production and injection wells, respectively. In more general cases, additional variables would be needed to describe well locations. For example, the optimization variables might include the perforation interval for vertical wells (if wells are not open to flow over their entire length), or the actual trajectory for deviated wells. The controls over time for each well are represented by a piecewise constant function with N_t time intervals (i.e., well controls are held constant during an interval and then jump to their value for the next interval). Hence, $n_2 = N_t(N_p + N_i)$.

In this work, we deal with bound constraints only. In order to simplify notation we introduce the well position feasible set $X = \{\mathbf{x} \in \mathbb{Z}^{n_1};\, \mathbf{x}_d \leq \mathbf{x} \leq \mathbf{x}_u\}$ and the well control feasible set $U = \{\mathbf{u} \in \mathbb{R}^{n_2};\, \mathbf{u}_d \leq \mathbf{u} \leq \mathbf{u}_u\}$. Nonlinear constraints can be handled using different techniques such as penalty functions or filter methods, as described in Echeverria Ciaurri et al. (2011a).

The objective function considered here is the undiscounted net present value (NPV) of the asset. This NPV accounts for revenue associated with the oil produced as well as for the water-handling costs incurred during production (water costs are incurred as a result of pumping and separation requirements). NPV is defined as follows:

$$\begin{aligned}\mathrm{NPV}\,(\mathbf{x},\mathbf{u}) = \sum_{k=1}^{N_s} \Bigg(& \sum_{i=1}^{Np} p_o q_o^{i,k}(\mathbf{u},\mathbf{x})\,\Delta t_k \\ & - \sum_{i=1}^{Np} c_{wp} q_{wp}^{i,k}(\mathbf{u},\mathbf{x})\,\Delta t_k - \sum_{i=1}^{Ni} c_{wi} q_{wi}^{i,k}(\mathbf{u},\mathbf{x})\,\Delta t_k \Bigg),\end{aligned} \tag{5}$$

where $q_o^{i,k}$, $q_{wp}^{i,k}$ and $q_{wi}^{i,k}$ are the flow rates of the oil, water produced and water injected for well i at the output interval k, respectively (expressed in stock tank barrels or STB per day, where 1 STB $= 0.1590\mathrm{m}^3$), and Δt_k represents the length (in days) of each of the N_s time steps in the simulation. (Note that N_s does not in general coincide with the number of controls per well, N_t.) The oil price and the cost of water produced and injected are denoted by p_o, c_{wp} and c_{wi}, respectively. Though the problem in (4) is stated jointly for $\mathbf{x}$ and $\mathbf{u}$, it has traditionally been addressed in practice in a decoupled manner (i.e., the well placement part is solved prior to, and independently of, the control optimization). In the next section we discuss some decoupled approaches and propose a methodology for addressing the problem in a joint manner.

3 Optimization Methodology

This section describes two sequential approaches and introduces a joint approach for solving the well placement and control problem given in (4). Both sequential approaches first seek optimal well placements using a predetermined control strategy, and then they optimize the controls for the wells determined in the first stage. Since the control and the well placement optimization problems possess clearly distinct characteristics, it is reasonable to address these two problems using different methodologies. Sections 3.1 and 3.2 describe the separate optimization problems and approaches corresponding to the continuous (controls) and discrete

(well placement) parts of (4). Some of the methods presented in these sections will be combined in Section 3.4, where we define our approach for the joint problem.

3.1 Well control optimization

The production optimization part of the general problem in (4) is obtained by fixing the well placement variable to $\mathbf{x}_0 \in \mathbb{Z}^{n_1}$:

$$\min_{\mathbf{u}\in U} -\mathrm{NPV}\left(\mathbf{x}_0, \mathbf{u}\right) , \tag{6}$$

and corresponds to a problem with continuous variables. The well controls $\mathbf{u} \in U \subset \mathbb{R}^{n_2}$ in this work represent BHPs. The optimization bounds define upper and lower BHP limits for both injectors and producers. Other operational constraints (e.g., minimum oil and/or maximum water production over all wells) can be addressed in an efficient manner by the filter method (Nocedal and Wright, 2006; Echeverria Ciaurri et al., 2011a). The controls for each well are defined by piecewise constant functions over N_t intervals.

Production optimization problems can be readily solved by gradient-based techniques (Nocedal and Wright, 2006). For example, the gradient-based optimization approach used in this work to solve (6) is sequential quadratic programming (SQP; Nocedal and Wright, 2006; Gill et al., 2007). Approximating gradients by, e.g., finite differences, typically requires a number of function evaluations proportional to the number of optimization variables. In addition, the quality of the approximation can depend strongly on the simulator settings. Adjoint formulations allow for efficient (though simulator-invasive) computations of gradients (Pironneau, 1974). Using an adjoint-based procedure, gradients can be computed with a total cost of roughly one solution of a linearized system of ordinary differential equations. Adjoint-based gradient estimations have recently been implemented for optimization problems in the petroleum industry (Brouwer and Jansen, 2004; Sarma et al., 2006). In this work, we use the adjoint formulation in Stanford's General Purpose Research Simulator (GPRS).

Derivative-free methods (Kolda et al., 2003; Conn et al., 2009) have also been shown to perform well for the control optimization problem (Echeverria Ciaurri et al., 2011a,b). These methods are applicable for problems with less than a few hundred optimization variables, and they can perform fairly efficiently if implemented in a distributed computing environment. We will consider derivative-free methods for the well-positioning part of the general problem in the next section.

It has been observed in previous work (Jansen et al., 2005; Echeverria Ciaurri et al., 2011a) that well control problems similar to (6) commonly display multiple local optima having similar cost function values. This suggests (though it does not prove) that local optimization approaches such as gradient-based techniques for (6) may yield solutions that are satisfactory from a global optimality point of view.

Finally, as discussed in the Introduction, reactive control can be applied, as a heuristic alternative to optimization, to address the issue of excessive water production. Under this approach, a production well is kept open (at its lower BHP limit in our implementation) until the revenue from the oil it produces no longer exceeds the cost associated with the produced water; i.e., the well is closed when

$$p_o\, q_o^{i,k} < q_{wp}^{i,k}\, c_{wp} , \tag{7}$$

where all variables are as defined previously. In practice this treatment often provides satisfactory results (and this approach is inexpensive since no optimization is required), though it is clearly suboptimal since it is based on a simple rule involving only producers. It should also be noted that even though the production strategies obtained by means of reactive control can in some cases be represented by piecewise constant functions, the lengths of the control intervals are not known a priori. Thus, reactive control strategies cannot in general be identified with elements in $\mathbb{R}^{n_2}$.

3.2 Well placement optimization

The well placement optimization part of the general problem in (4) is obtained by fixing the well control variable to $\mathbf{u}_0 \in \mathbb{R}^{n_2}$:

$$\min_{\mathbf{x}\in X} -\text{NPV}\left(\mathbf{x}, \mathbf{u}_0\right) , \tag{8}$$

and corresponds to a problem with discrete variables. In general $\mathbf{u}_0 \in \mathbb{R}^{n_2}$, but as noted above, if $\mathbf{u}_0$ corresponds to a reactive control strategy, it will not necessarily have n_2 components.

Well placement problems are in a sense more challenging than well control optimization problems because reservoir heterogeneity leads to highly non-smooth objective functions containing multiple optima (see e.g., Onwunalu and Durlofsky, 2010). Therefore, the well placement optimization problem would not appear to be as amenable to solution using gradient-based methods because these approaches can get trapped in local minima. There have, however, been procedures presented for (8) that use gradients (see e.g., Sarma and Chen, 2008 and Zandvliet et al., 2008). These methods replace the problem with a related (though not necessarily equivalent) problem that has continuous variables.

Derivative-free methods thus appear to be well suited for the solution of (8). The derivative-free methods that have been used to date for well placement are based on stochastic search procedures. Examples include genetic algorithms (Goldberg, 1989; Guyaguler et al., 2000; Yeten et al., 2003), stochastic perturbation methods (Bangerth et al., 2006), and particle swarm optimization (Clerc, 2006; Onwunalu and Durlofsky, 2010; Echeverria Ciaurri et al., 2011b). Due to their random component, these search procedures can avoid being trapped in unsatisfactory local optima. Most of these methods, however, are not supported by solid convergence theory, and consequently they contain tuning parameters that are often difficult to determine.

In this work we propose derivative-free optimization methodologies based on pattern-search (Torczon, 1997; Kolda et al., 2003; Conn et al., 2009) as a more mathematically sound alternative for well placement optimization. These methods rely on (local) convergence theory applicable to sufficiently smooth functions of continuous variables. These local convergence results can furthermore be extended to problems with discrete variables (Audet and Dennis, 2000). Examples of these techniques are Hooke-Jeeves direct search (HJDS; Hooke and Jeeves, 1961), generalized pattern search (GPS; Torczon, 1997, Audet and Dennis, 2002), mesh adaptive direct search (MADS; Audet and Dennis, 2006), and bound optimization by quadratic approximation (BOBYQA; Powell, 2009).

Pattern-search methods operate primarily through a polling procedure. Polling is accomplished by computing cost function values at solutions determined by a given stencil in the search space which is centered at the current solution. The stencil is frequently arranged along the coordinate axes, which results in a coordinate or compass search. In MADS, the stencil orientation is randomly modified after each complete polling. Pattern-search techniques are supported by local search convergence theory, but if the initial stencil size is comparable to the size of the search space (which means that, during the first iterations of the optimization, the search involves points that are distant from the initial guess), they can incorporate some global exploration features. We emphasize that, in practice, global convergence is not achieved using these procedures. However, in many well placement problems, finding a reasonable local optimum following some amount of global exploration is often sufficient.

In this work, the well placement problem is solved using HJDS, GPS and a hybrid optimization parallel search package (HOPSPACK; Plantenga, 2009). HOPSPACK is a distributed computing implementation of GPS which can be run in a so-called asynchronous mode to balance the computational load of each node in the cluster (Plantenga, 2009). HJDS is a serial computing procedure that was identified in Echeverria Ciaurri et al. (2011a) as a fairly efficient optimization procedure for oil field problems when distributed computing resources are limited or unavailable.

3.3 Sequential approaches for well placement and control optimization

As noted earlier, sequential procedures are commonly used for joint well location and well control optimization. Well placement is optimized first using some 'reasonable' control scheme. In this work we will consider two such strategies – fixed and reactive controls. The controls are then optimized for the wells positioned in the first stage. It should be noted, however, that well placement optimization results have been observed to depend to a large degree on the control scheme used (Zandvliet et al., 2008).

Fixed control strategies belong to $U \subset \mathbb{R}^{n_2}$, i.e., the same space explored in the control optimization stage. In our approach these controls correspond to the upper pressure bound $\mathbf{u}_u$ for injectors, and the lower pressure bound $\mathbf{u}_d$ for producers. This strategy provides maximum injection and total fluid production rates at all times. It is important to emphasize that, although total fluid production is maximized, oil production is in general not maximized by this strategy if water is also being produced (as it typically is). This fixed control strategy is in general suboptimal because the water front is allowed to proceed without any 'steering' (which is achieved when BHPs are varied in time). In addition, it is possible that some wells may be producing essentially all water at full capacity. The reactive control approach (described above) avoids the latter, but not the former, problem. Thus it is usually preferable to using fixed controls, but as we will see does not perform as well as the joint optimization procedure.

Algorithm (1) below shows the two basic steps in the sequential approaches. Here we use $\mathbf{x}_S^*$ and $\mathbf{u}_S^*$ to designate the optima obtained from the sequential approach. We reiterate that $\mathbf{x}_S^*$ and $\mathbf{u}_S^*$ do not coincide with the optimum of (4).

Algorithm 1 Sequential approach for well placement and control optimization

Input: initial locations $\mathbf{x}_0$ and specified control strategy (fixed or reactive) $\mathbf{u}_0$ for $N_p + N_i$ wells

Output: improved locations $\mathbf{x}_S^*$ and control strategy $\mathbf{u}_S^*$

1: Solve $\mathbf{x}_S^* = \underset{\mathbf{x}\in\mathbb{X}}{\text{argmin}} \; -\text{NPV}(\mathbf{x}, \mathbf{u}_0)$ using a pattern-search optimizer

2: Solve $\mathbf{u}_S^* = \underset{\mathbf{u}\in\mathbb{U}}{\text{argmin}} \; -\text{NPV}(\mathbf{x}_S^*, \mathbf{u})$ using a gradient-based optimizer

3.4 Joint approach for well placement and control optimization

We address the joint well placement and control problem using the following nested optimization

$$\min_{\mathbf{x}\in X} \min_{\mathbf{u}\in U} -\text{NPV}(\mathbf{x}, \mathbf{u}) \; . \tag{9}$$

In this bound-constrained optimization problem, it is relatively simple to see that the formulations in (4) and (9) are equivalent regarding the first-order optimality conditions. In accordance with the methods presented in the previous section, the outer well placement optimization in (9) is solved here by means of pattern-search optimization algorithms, while the inner control optimization is addressed through a sequential quadratic programming implementation with gradients computed efficiently using an adjoint-based scheme.

The approach in (9) may seem impractical since it requires solving a complete optimization for every cost function evaluation of the outer (upper-level) optimization problem. However, in our application a nested procedure is reasonable because of the following two observations. First, as discussed in detail earlier, the two optimizations are of different character and can be addressed using different procedures. And second, we will make use of a very efficient adjoint-based gradient computation in GPRS for the control optimization. As noted above, the bound-constrained well control optimization problem displays multiple local solutions, but with similar cost function values. Thus, there is little if any benefit from running this optimization from multiple starting points. This means it is sufficient to solve the well control optimization problem just once.

We note that the nested optimization in (9) could also be analyzed from a bilevel programming perspective (Dempe, 2002). However, bilevel optimization problems are often more complicated to study than the problem considered here since the two optimization levels are in general associated with different objective functions.

Using the formulation in (9) and the specific choice of methods for the two components of the optimization, our intent is to perform some amount of global exploration (via the use of large initial stencil size in the pattern search) in a space of dimension n_1, and not for a search in a space of dimension $n_1 + n_2$. This is an important aspect of our procedure since the computational cost associated with the global exploration of a space of dimension $n = n_1 + n_2$ grows exponentially with n (curse of dimensionality). Moreover, as we will see in Section 4.1, the function optimized in the outer approximation in (9)

$$-\text{NPV}^*(\mathbf{x}) = \min_{\mathbf{u}\in U} -\text{NPV}(\mathbf{x}, \mathbf{u}) \; , \tag{10}$$

is much smoother in $\mathbf{x}$, and as a consequence, easier to explore globally, than $\mathrm{NPV}(\mathbf{x}, \mathbf{u}_0)$, with $\mathbf{u}_0$ being a fixed control strategy. The smoothing of the optimization surface with respect to the control variable $\mathbf{u}$ occurs because the performance of wells in less promising locations can be improved, sometimes significantly, by optimizing the well controls.

The function $\mathrm{NPV}^*(\mathbf{x})$ in (10) is well defined since there exists an NPV^* for every feasible $\mathbf{x}$. We do not, however, expect there to be a unique $\mathbf{u}$ associated with NPV^*. The joint optimization approach proposed in this work can be interpreted as a well placement problem where the cost function is an optimized NPV

$$\min_{\mathbf{x}\in X} -\mathrm{NPV}^*(\mathbf{x}) \,, \tag{11}$$

with NPV^* as defined in (10).

The well control optimization required for each computation of $\mathrm{NPV}^*(\mathbf{x})$ is not solved completely in our implementation. This is motivated by the difficulty of obtaining robust stopping criteria in practical optimization problems, and the fact that an unnecessarily tight stopping criterion may result in an excessive number of cost function evaluations. In a preliminary study involving a problem of similar complexity (in terms of the well control optimization) to those studied in this work, we determined that a small number of iterations for the gradient-based optimizer yields an acceptable approximation of the optimal control strategy. Thus, during the course of the joint optimization, we use a small number of iterations for the well control problems. Then, once the optimal well locations are determined, we again run the control solution but this time with a tighter stopping criterion, which leads to a slightly improved NPV^*.

It should be stressed that the outer optimization in (11) is fully parallelizable, and indeed in this work we take full advantage of this. However, the parallel runs involve control optimizations and not simply single simulations. Therefore, the computational loads in each of the nodes can be very different, because in general, two calls to NPV^* (with different well placements) will not require the same number of simulations (even using the same number of iterations in the gradient-based optimizer). This issue can be alleviated to some extent by means of asynchronous distributed computing approaches (see Griffin and Kolda, 2007 or Griffin et al., 2008, for an example within the context of pattern search).

In the remainder of the paper, the sequential optimization methodologies with fixed and reactive control strategies, and the joint technique, are denoted as *sequential fixed*, *sequential reactive* and *joint* approaches, respectively.

4 Example Cases

In this section we apply the methodologies described in Section 3 to two examples. As indicated above, each control optimization problem is solved by means of a gradient-based optimizer, and the well placement portion of the optimizations is handled using three different pattern-search based algorithms, namely, Hooke-Jeeves direct search (HJDS), generalized pattern search (GPS), and the hybrid optimization parallel search package (HOPSPACK). GPS and HOPSPACK were

Figure 1 Permeability field (mD) used for the two cases in Section 4 ($\log k$ is displayed). Geological heterogeneity is clearly evident.

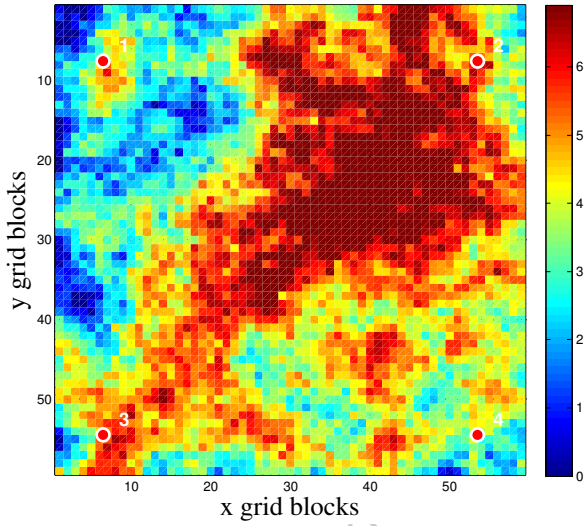

Table 1 Model and optimization parameters for the two examples.

Parameter	Case 1 (Section 4.1)	Case 2 (Section 4.2)
Cell size	130 ft × 130 ft × 20 ft	50 ft × 50 ft × 50 ft
Production time frame	2190 days	2920 days
Oil price (p_o)	80 \$/bbl	80 \$/bbl
Water production cost (c_{wp})	10 \$/bbl	20 \$/bbl
Water injection cost (c_{wi})	10 \$/bbl	20 \$/bbl
Injector BHP upper & lower bounds	5200 and 4100 psia	6000 and 4100 psia
Producer BHP upper & lower bounds	3500 and 1000 psia	3500 and 1000 psia

implemented within a distributed computing framework consisting of 8 and 20 + 1 computing cores, respectively (in HOPSPACK one of the cores is used for coordination tasks). We reiterate that, in the parallel implementations, each processor handles the full well control optimization, not just a single simulation run.

The two cases considered are based on a reservoir discretized on a 60 × 60 two-dimensional grid. The permeability and porosity fields are portions of layer 21 of the SPE 10 model (Christie and Blunt, 2001). These fields display strong variability in properties, as can be seen for permeability in Figure 1. In both examples there are five wells (one injector and four producers in the first case and two injectors and three producers in the second case). The controls for all five wells are optimized in both examples. In the first example, only one well (the injector) location is

optimized, while in the second case all well locations are optimized. The key model and optimization parameters for both cases are shown in Table 1.

The gradient-based optimizer used for well control optimization is SNOPT (Gill et al., 2007), which is based on sequential quadratic programming (Nocedal and Wright, 2006). The initial guess in all situations is obtained by setting the injector and producer BHPs at their upper and lower bounds, respectively. This configuration provides maximum flow rates. Since the bound-constrained control optimization problem displays multiple optima but with similar cost function values, the selection of the starting point is not expected to impact the quality of the optimized solution. The stopping criteria selected for the control optimization are based on the major optimality tolerance (a value of 10^{-6} in all situations) and on the maximum number of major iterations allowed. For more details on these stopping criteria, see Gill et al. (2007). In most cases it is the maximum number of major iterations that terminates the optimization. During a major iteration, a quadratic programming problem is solved, and this usually requires several reservoir simulations.

As explained in Section 3.4, the control optimization required for finding NPV*, when called from the outer well placement loop, is not solved to full accuracy. The maximum number of major iterations is equal to eight (a relatively small number) in most cases, though in Section 4.2 results are also presented using a value of four. The control optimizations performed at the last iteration of both the sequential and joint approaches aim at a more precise solution. For these optimizations the maximum number of major iterations is set to 32.

The optimizations for the well placement problem are expected to depend on the initial guess, since in general these problems are markedly nonconvex (unlike the control optimization problem). For this reason, we perform optimization runs starting from different points. The initial stencil size in all cases is 16, which is slightly less than 1/3 of the feasible search space. A stencil of this size leads to some amount of global exploration since regions far from the stencil center are evaluated. All pattern-search algorithms terminate the optimization when the stencil size is equal to one, and the cost function value corresponding to the stencil center is lower than the cost function value associated with any other stencil point. This termination condition defines the notion of local optimality that will be considered for the discrete variables.

4.1 Optimization of injector location and control of all wells

4.1.1 Case description

In this case we consider four producers, fixed at the corners of a square, along with one injector. The four producers (designated by red circles) are positioned as shown in Figure 1. The production wells are located somewhat away from the reservoir boundaries, which are prescribed to honor no-flow conditions. The injector can be positioned anywhere inside of the square (40×40 grid blocks) defined by the producers. The control strategies for all five wells and the location of the injector will be optimized. These strategies refer to a production time frame of six years, and except for reactive control, the strategies are divided into ten intervals of 219 days each (during each time interval the BHPs are held constant). Hence, for this problem, $n_1 = 2$ and $n_2 = 50$.

Table 2 Injector well location and NPV for the best solution obtained for the exhaustive explorations. The * indicates that an additional gradient-based control optimization is performed.

Approach	Location [x,y]	NPV [$MM]
fixed	[18,26]	976
sequential fixed*		1091
reactive	[17,42]	1061
sequential reactive*		1074
joint	[12,36]	1135
joint*		1137

In the next section we will perform an approximation of the exhaustive search of the optimization spaces corresponding to the sequential and joint approaches described in Sections 3.3 and 3.4, respectively. Thereafter, we will use this example to compare some of the optimization techniques discussed above.

4.1.2 Exhaustive search results

Due to the low value of n_1 in this example, it is feasible to exhaustively explore the discrete space X both for $\text{NPV}(\mathbf{x}, \mathbf{u}_0)$, with $\mathbf{u}_0$ being a fixed or a reactive control strategy, and for $\text{NPV}^*(\mathbf{x})$ (where $\mathbf{u}$ is determined from optimization). This discrete space exhaustive search requires $40 \times 40 = 1600$ simulations for the sequential cases, and 1600 control optimizations for the joint approach. We note that, since the cost function appears to be close to convex in $\mathbf{u}$, we expect the exhaustive exploration of $\text{NPV}^*(\mathbf{x})$ to be a reasonable approximation of a global exhaustive search for the complete optimization space in (4). This type of exhaustive search is already impractical for the example in Section 4.2, where n_1=10.

In Figure 2 we present results for the three exhaustive explorations corresponding to a fixed control strategy (injectors at maximum BHP, producers at minimum BHP), the reactive control strategy, and the optimized control strategy. It is clear that the surface associated with the fixed control strategy is much rougher than the surfaces obtained with the other strategies. This demonstrates that it is possible to somewhat compensate for less promising well locations with a proper control strategy (in terms of net present value). As a consequence, the associated optimization landscape $\text{NPV}^*(\mathbf{x})$ can be expected to be smoother than the landscape corresponding to $\text{NPV}(\mathbf{x}, \mathbf{u}_0)$, for $\mathbf{u}_0$ a fixed strategy. This suggests, consistent with our earlier discussion, that the joint optimization landscape may be somewhat easier to explore globally.

The well locations with the highest net present value resulting from the three exhaustive explorations are given in Table 2. The 'fixed' and 'reactive' results are for the best wells in Figure 2a and 2b. The 'sequential fixed*' and 'sequential reactive*' results additionally apply gradient-based optimization for the well controls using the positions found in the exhaustive search. This optimization is performed with a tight tolerance (a maximum number of major iterations of 32), which is why we include the * designation. For the joint optimization, in the exhaustive search we use a maximum number of major iterations of 8 (these results are designated 'joint'

Figure 2 Exhaustive search results for (a) NPV $(\mathbf{x}, \mathbf{u}_0)$ with $\mathbf{u}_0$ a fixed control strategy (with BHPs set to provide maximum flow rates), (b) NPV $(\mathbf{x}, \mathbf{u}_0)$ with $\mathbf{u}_0$ a reactive control strategy, and (c) NPV* $(\mathbf{x})$. Production and injection wells are represented as red and blue circles, respectively. The dark blue region near the boundaries is infeasible. The scale indicates 10^6 \$.

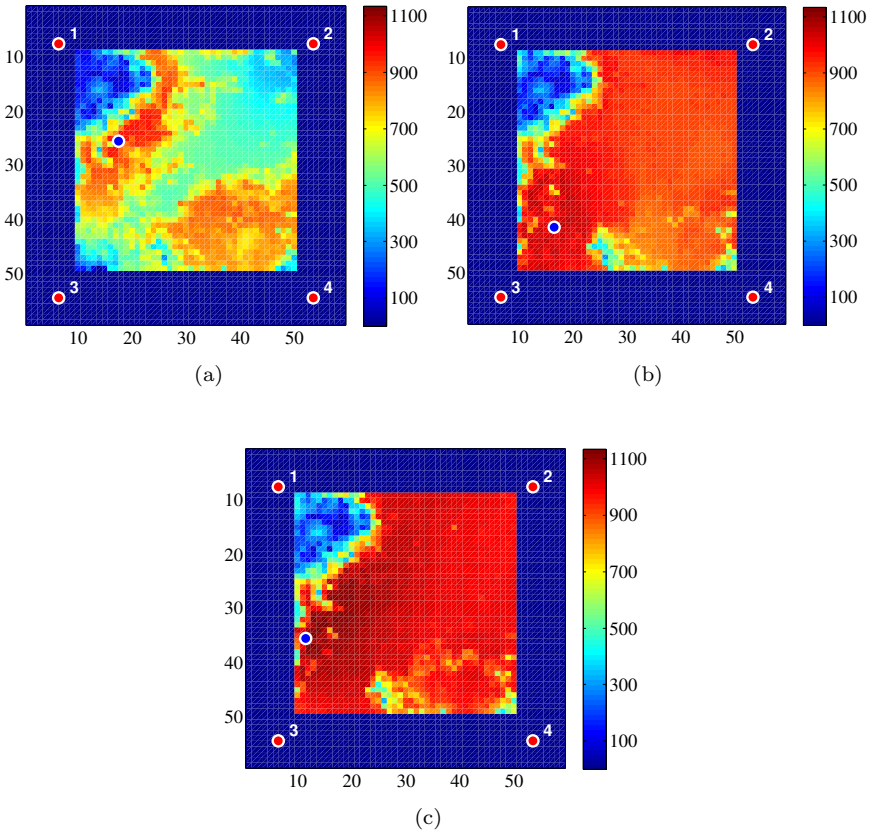

in the table). Using the best well found during the exhaustive search, we again run the control optimization, this time using 32 major iterations. These results are designated 'joint*'.

As expected, the joint scheme outperforms the sequential methodologies, even after the additional control optimization step. The joint approach yields an increase of 4.2% and 5.9% in NPV with respect to the sequential fixed and reactive approaches. In this simple example, these improvements correspond to \$45 million and \$62 million. These amounts, as will be seen in the next example, can be even greater in larger and more realistic problems. It is interesting to note that while the reactive approach obtains a better solution than the fixed scheme before the final control optimization, the situation changes after the control optimization. This reiterates that the control optimization can somewhat compensate for well locations that are suboptimal in terms of NPV.

Figure 3 Cumulative production and injection profiles for the well location and controls (after the additional control optimization) corresponding to the highest NPV solution: (a) sequential fixed, (b) sequential reactive, and (c) joint approach.

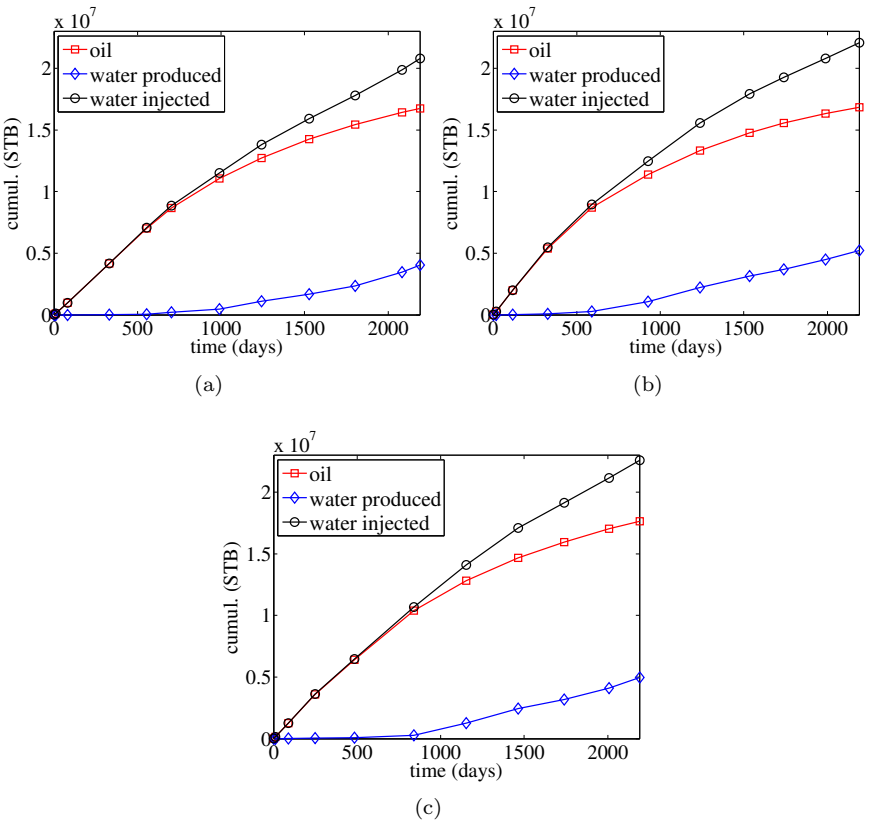

In Figure 3 we show, for the three exhaustive explorations performed (plus the additional well control optimization), the cumulative injection and production profiles for the configurations with the highest NPV. From these plots, it is evident that the joint optimization provides more cumulative oil than the other two procedures. The joint optimization scenario also involves more water injection than the other scenarios, but this is more than compensated for by the increase (of about 5.1%) in cumulative oil. Figure 4 displays the oil saturation distributions at the end of the production time frame for the three optimizations. These plots illustrate how the different approaches perform in terms of reservoir 'sweep' efficiency. It is evident that there is less bypassed oil in the joint approach than in the sequential approaches.

The well controls (BHPs) corresponding to the highest NPV solutions are shown in Figure 5. The BHPs for the injectors (blue curves) for the various optimizations are in the top row and the next four rows (red curves) represent the producers.

Figure 4 Oil saturation distribution at the end of the production time frame for the well location and controls (after the additional control optimization) corresponding to the highest NPV solution: (a) sequential fixed, (b) sequential reactive, and (c) joint approach. Injection and production wells are represented as blue and red circles, respectively. Oil saturation refers to the oil fraction in each simulation grid block (blue indicates water and red indicates oil).

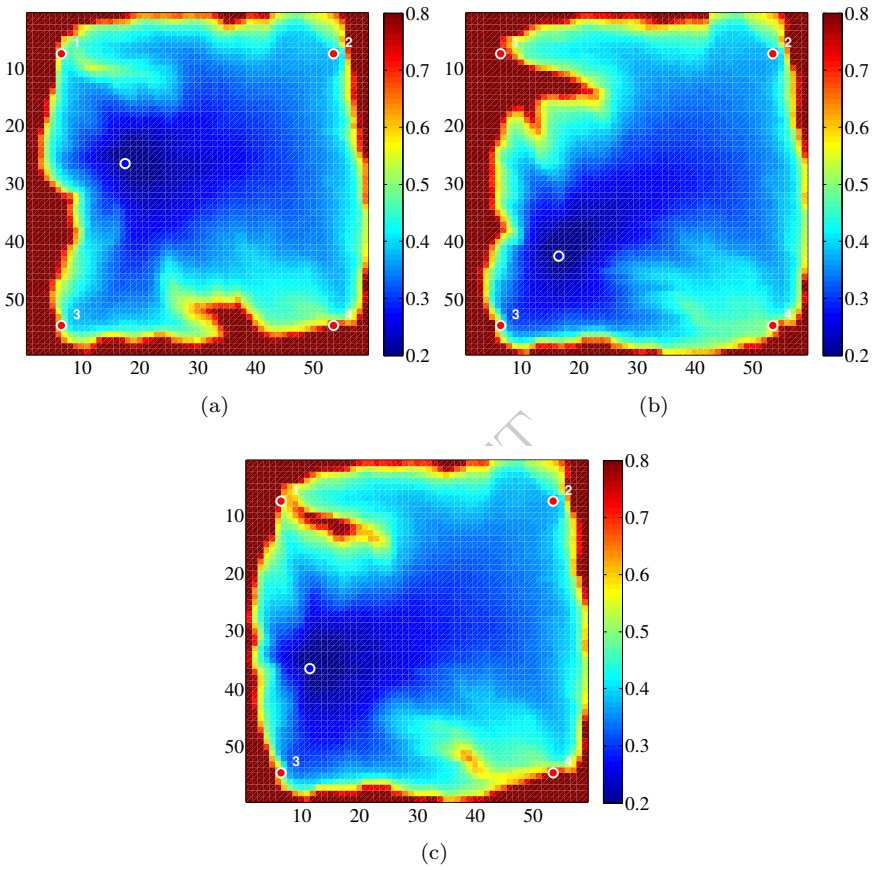

Upper and lower BHP bounds are indicated by dashed lines. The time axes span the entire production period (2190 days). Note that the BHPs for Producers 1 and 4 stay at the minimum BHP limit in all cases, presumably because these wells are outside of the large (diagonally-oriented) high-permeability region evident in Figure 1. The BHPs for Producers 2 and 3 are, by contrast, away from the lower BHP limit, for at least some of the simulation, for all three optimization schemes. This is likely due to the fact that these wells, along with the injector, fall within the high-permeability region. If these two wells produced at their lower limits for the full simulation, significant water production would result. In order to avoid this, the optimizations reduce the flow rates (which leads to water breakthrough at later times) for these two wells.

Figure 5 Injection and production well controls (BHPs) corresponding to the highest NPV solution: (left) sequential fixed, (center) sequential reactive, and (right) joint approach. Top figure corresponds to injector and next four figures to the producers.

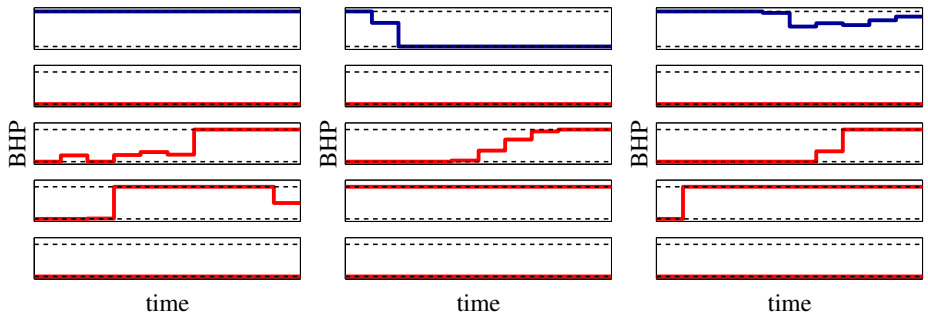

4.1.3 Optimization solutions

The results in Table 2 required an exhaustive search, which is not feasible in practical situations. In this section, rather than search exhaustively, we apply pattern-search optimization for the well location part of the problem. The control optimization is again handled via gradient-based optimization, with all derivatives computed efficiently using adjoint-based procedures.

We reiterate that most derivative-free optimization techniques (such as pattern-search algorithms) can be readily applied to problems with discrete optimization variables, and that these methodologies have been observed to perform satisfactorily on relatively non-smooth cost functions such as that in Figure 2(a). Although the sequential reactive and joint strategies displayed relatively smooth cost functions (Figures 2(b) and (c)), the degree of smoothness observed for high-dimensional searches may differ from that for these low-dimensional ($n_1 = 2$) cases. In any event, as we will see below, all of the derivative-free algorithms considered yield solutions that are on average relatively close, in terms of NPV, to the results from the exhaustive explorations.

As mentioned earlier, the pattern-search algorithms considered here are Hooke-Jeeves direct search (HJDS), generalized pattern search (GPS) and the hybrid optimization parallel search package (HOPSPACK). These algorithms (all of which are supported by local convergence; see e.g., Torczon, 1997) rely on the same principles, and this facilitates meaningful comparisons. Pattern-search optimization is based on evaluating a stencil whose size decreases along iterations (the reduction in the stencil size is performed when all the stencil points have a higher cost function than the stencil center). The stencil used in all cases here has $2\,n_1$ points distributed along the coordinate axes from the stencil center (as in a compass). The initial stencil size is always equal to 16, and this value allows a rough exploration of the search space (for any initial guess) since the lower and upper bounds for $\mathbf{x}$ are 10 and 50 for this case, and 6 and 55 for the second case. The sequence of stencil sizes $\{16, 8, 4, 2, 1\}$ is consistent with the optimization variables being

discrete. Hence, all the algorithms stop when the stencil size is equal to 1, and the stencil center cost function value improves on every stencil point. Upon termination, the solution obtained is a (discrete) local optimizer for the $2\,n_1$-point (compass) neighborhood.

Hooke-Jeeves direct search does not compute the cost function for all $2\,n_1$ stencil points. As soon as a point in the current stencil improves on the cost function value for the stencil center point, the stencil is moved to a new center (this strategy is known as opportunistic polling). This makes HJDS a serial strategy that can be attractive when distributed computational resources are scarce, or when commercial software licensing issues limit massive parallelization. Since both GPS and HOPSPACK evaluate the $2\,n_1$ points for every stencil, the use of distributing computing is very beneficial for these algorithms.

In this example the three cost functions ($\mathrm{NPV}\,(\mathbf{x}, \mathbf{u}_0)$, with $\mathbf{u}_0$ corresponding to all wells at their BHP limits, and to a reactive strategy, and $\mathrm{NPV}^*\,(\mathbf{x})$) are based on a lookup table constructed with the results from the exhaustive explorations. Therefore, for this case, GPS and HOPSPACK do not take real advantage of being implemented in parallel. In the example in Section 4.2, this feature will be effectively exploited.

Separately, and this is applicable to pattern-search methods in general where the stencil only changes its size along iterations, some points in the optimization are revisited at different times. The cost function computation in these cases can be avoided if all (or just a number of) evaluations are stored in a cache. In this work caches are implemented for the three pattern-search algorithms considered.

The results from the three approaches, together with the NPVs obtained in the exhaustive explorations, are summarized in Table 3. The NPVs for the exhaustive explorations are taken from Table 2. Because different initial guesses result in different locally optimal solutions, we run each optimization 12 times, starting at different initial points. Each pattern-search run is followed by a gradient-based control optimization with tight tolerances (as above, * denotes the use of a maximum of 32 major iterations). The NPVs, expressed in 10^6\$, are averaged over the 12 runs. The average number of iterations $\mathrm{n_{ps}}$ for each pattern-search procedure is also reported. It is important to note that for the sequential fixed and reactive approaches this number is equivalent to the average number of reservoir simulations needed for the entire optimization process. However, for the joint approach it indicates the (average) number of control optimizations required in the complete search. In this example, each control optimization requires on average 14 reservoir simulations.

The differences (in terms of NPV) between the results obtained by the sequential and joint approaches before performing the additional control optimization are somewhat larger than the corresponding results for the exhaustive explorations (shown in Table 2). This may be because the cost function for the joint approach is globally smoother, which makes it easier to optimize. The additional control optimization to some extent reduces the discrepancies in the results. Before the control optimization step, the average optimized NPV by the joint approach is 24.1% and 9.8% larger than the average optimized NPVs from the fixed and reactive approaches, respectively. After the additional optimization, these percentages decrease to 10.3% and 6.1%.

Table 3 Average NPV (over 12 runs, expressed in $MM) for the optimal location of one injector and control of five wells. For the sequential approaches, n_{ps} is equivalent to the average number of reservoir simulations needed in the entire optimization process. For the joint approach, n_{ps} indicates the average number of control optimizations required in the complete search.

Approach	HJDS		GPS		HOPSPACK		Exhaustive
	NPV	n_{ps}	NPV	n_{ps}	NPV	n_{ps}	NPV
fixed	901	42	883	29	891	23	976
sequential fixed*	1015		994		1002		1091
reactive	1003	33	1015	25	1004	21	1061
sequential reactive*	1034		1053		1044		1074
joint	1117	47	1109	32	1093	25	1135
joint*	1118		1110		1094		1137

It is not clear from the results in Table 3 if one pattern-search algorithm in preferable over the other two. GPS and HOPSPACK are slightly faster than HJDS, but they yield lower average cost function values. As noted earlier, however, GPS and HOPSPACK can be accelerated, in terms of clock time, if a cluster is available (and in this situation, they will outperform HJDS). In this relatively simple case ($n_1 = 2$, and cost function computed via a lookup table) the performance of GPS and HOPSPACK seems to be comparable. Differences between the various approaches and algorithms will be more evident in the next example, which is more realistic and more complex.

4.2 Optimal location and control of five wells

4.2.1 Case description

In this example we optimize both the location and control of two injectors and three producers. Some of the reservoir parameters are different than those used in Section 4.1. Specifically, the reservoir area is reduced, the production time frame is longer, and the costs for injected and produced water are doubled (the corresponding model and optimization parameters are given in Table 1). A water-flooding configuration with two injection wells is richer in terms of variety of sweeping strategies than an arrangement with only one injector (as considered previously). Thus we aim at increasing the diversity of production scenarios, which renders the search more challenging. In addition, our reactive control strategy (which addresses only producers) is now less appealing because water injection is costly.

The number of grid blocks in the reservoir model, and the permeability and porosity values for each grid block, are the same as before (see Figure 1). Because all five well locations are optimized we now have $n_1 = 10$. The production time frame is again divided into ten intervals (hence, $n_2 = 50$). As in the previous example, the gradient-based optimization algorithm embedded in the joint approach is SNOPT, and the pattern-search methods considered for the well placement search are HJDS, GPS, and HOPSPACK. Both GPS and HOPSPACK are implemented

Table 4 Results for NPV (expressed in $MM) and total number of simulations n_{sim} for the second example (optimization of the location and control of five wells). GPS is used for the well location optimization. The highest NPV for the nine runs for each approach is underlined.

Run	seq. fixed$^*_{GPS}$		seq. reactive$^*_{GPS}$		joint$^*_{4,GPS}$		joint$^*_{8,GPS}$	
#	NPV	n_{sim}	NPV	n_{sim}	NPV	n_{sim}	NPV	n_{sim}
1	<u>336.6</u>	321	334.5	295	347.4	1075	<u>385.1</u>	2683
2	300.5	505	354.8	422	353.0	2021	355.2	3276
3	328.9	426	314.7	310	329.1	1770	346.7	3327
4	328.0	511	192.7	240	325.8	1922	372.3	2481
5	326.7	477	240.9	377	<u>355.2</u>	1936	354.8	5003
6	294.9	468	253.3	361	336.6	2741	336.0	3278
7	263.4	423	345.4	329	344.7	2031	360.0	3941
8	256.8	644	279.6	420	339.5	2602	357.2	4187
9	293.7	587	<u>358.8</u>	447	330.6	1938	358.0	4855
Mean	303.3	485	297.2	356	340.2	2004	358.4	3670
σ	29.2	94	58.6	68	10.6	479	14.0	890

within a distributed computing framework. While HOPSPACK uses 21 cores (one core is dedicated to the coordination of the concurrent jobs, leaving effectively 20 computing cores), GPS, due to a limited number of licenses, is applied only on eight cores. The markedly non-convex character of the well placement optimization is dealt with by running the optimizations nine times with different initial guesses. These initial guesses were not randomly selected – rather, they correspond to well placements that are reasonable from a reservoir engineering perspective.

4.2.2 Optimization solutions

In these optimization runs, the parameters for GPS are the same as in Section 4.1, i.e., the sequence of stencil sizes is $\{16, 8, 4, 2, 1\}$. The control optimization in the joint approach is solved with two different values (4 and 8) for the maximum number of major iterations. In all cases, one supplementary control optimization is performed with a maximum number of major iterations of 32 (one optimization for the sequential fixed approach needed 64 iterations because convergence was not obtained after 32 iterations).

The results using GPS in the well placement optimization part for all of the approaches and each of the nine different initial guesses for well location are presented in Table 4 (joint$^*_{4,GPS}$ and joint$^*_{8,GPS}$ denote the joint approach with the maximum number of major iterations in the gradient-based control optimization equal to 4 and 8, respectively). The final control optimization is performed in all cases, and the simulations required for this step are included in n_{sim}.

We observe that the average (maximum) NPV for joint$^*_{8,GPS}$ over the nine runs is 5.3% (8.4%) higher than for joint$^*_{4,GPS}$. This observation is consistent with the much larger number of simulations performed in joint$^*_{8,GPS}$, and indicates that a maximum number of major iterations of 4 tends to terminate the optimization prematurely. An insufficient maximum number of major iterations may yield a

clearly suboptimal solution, and an excessively large value could lead to excessive computational requirements. Therefore, a tuning process for this parameter might be beneficial when applying the joint approach. For the remainder of this section, our remarks on the joint approach will refer to the case with maximum number of major iterations of 8.

In terms of NPV, the sequential fixed and reactive strategies clearly underperform the joint approach. The average (maximum) NPV over all of the runs obtained with the joint approach is 18.2% (14.4%) and 20.6% (7.3%) higher than with the sequential fixed and reactive schemes, respectively. The average number of simulations required by $\text{joint}^*_{8,\text{GPS}}$ is, however, about one order of magnitude higher than that needed by the sequential methodologies. Along these lines, it is important to realize that the maximum NPVs reported in Table 4 for the sequential strategies are based on a fraction of the computational effort dedicated to the joint approach. Thus, in order to complement the results in the table, we tested the sequential reactive scheme for 100 new random initial well locations (in that manner, the associated total computational cost is comparable to that for $\text{joint}^*_{8,\text{GPS}}$). The average and maximum NPV over these 100 runs are \$288.0 million and \$353.8 million, respectively. These values are lower than the corresponding values in Table 4 (\$297.2 million and \$358.8 million), which is consistent with the fact that the nine initial well placements were not selected randomly but rather based on engineering judgement. The key observation, however, is that, even when we compare based on the same number of total simulation runs, $\text{joint}^*_{8,\text{GPS}}$ still outperforms the sequential reactive scheme.

As can be seen in Table 4, an advantage of the joint approach is that it results in smaller standard deviation σ of the NPV than the sequential methodologies. This fact is consistent with the smoothing of the well placement optimization landscape observed for the joint strategy (which was illustrated earlier in Figure 2). Note further that the results for the sequential reactive approach are not in this case as close to those for the joint strategy as in the previous example (indeed, here they are more comparable to those for the sequential fixed approach). This may be explained by the increased complexity of this problem and by the elevated cost of injected water.

The optimization results for the sequential and joint approaches are further illustrated in Figures 6(a) and (b), where the evolution of the objective function (NPV) versus the number of forward simulations is represented for each of the runs. The corresponding averages over the nine runs are plotted in Figures 6(c) and (d). In order to enable clear comparisons, all figures use the same vertical scale. We note that, prior to the supplementary (final) control optimization, all solutions for the sequential fixed scheme have NPVs lower than \$200 million. It is evident in Figure 6(a) that the additional control optimization is crucial in the sequential approaches. We reiterate that in both the fixed and reactive strategies the water injectors operate at maximum BHP, and this may negatively impact the objective function. Hence, the supplementary control optimization can again be seen as a means to compensate for suboptimal well locations. The lower standard deviation in the joint approach compared to the sequential strategies is also evident in Figure 6.

The oil saturation distributions corresponding to the solutions with maximum NPV, at the end of the simulation time frame, are presented in Figure 7 (injection and production wells appear as blue and red circles, respectively). The amount

Figure 6 Evolution of the objective function (NPV) versus number of simulations: (a) sequential fixed and reactive approaches for all nine runs, (b) joint approach for all nine runs. The corresponding averages over the nine runs are represented in (c) and (d). All runs include a final control optimization.

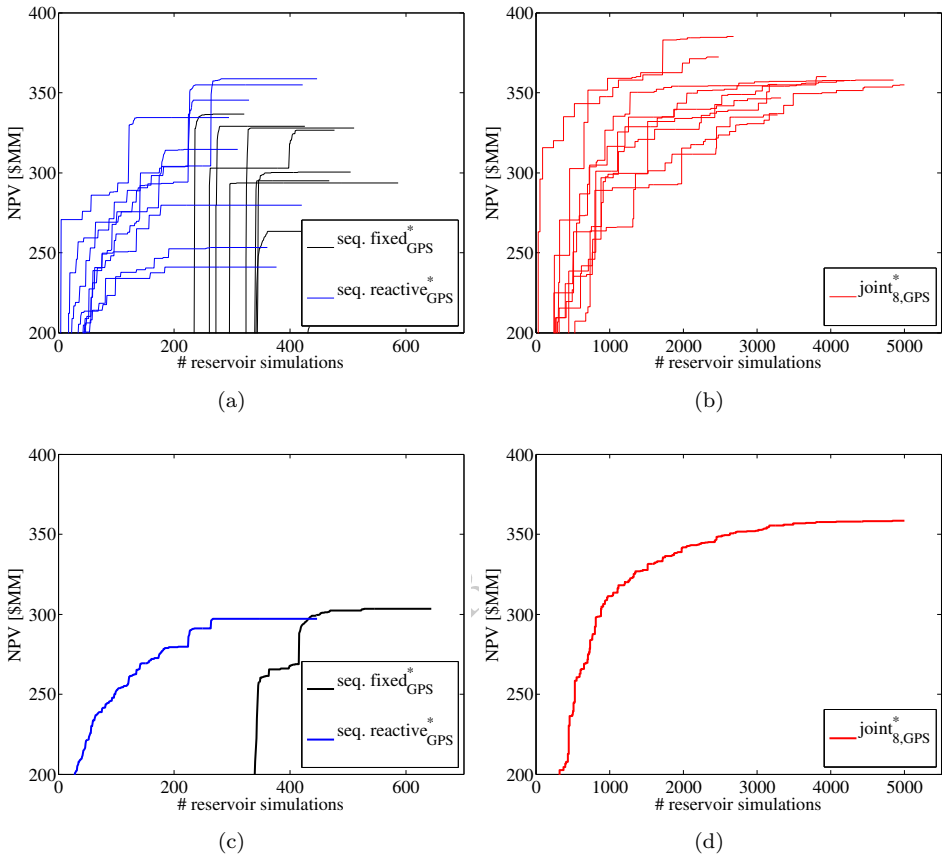

of bypassed oil is noticeably less for the joint approach than for the sequential strategies. The well locations obtained generally tend to be toward the boundaries of the domain. In a few cases, some wells are placed very close to each other, as can be seen e.g., in Figure 7(b) for the sequential reactive approach. This type of solution might not be acceptable in practice, and can be prevented in the optimization by including (nonlinear) constraints that ensure a minimum distance between wells. The computation of these constraints do not involve time-consuming function evaluations, and for that reason, they are not as complicated to handle as other simulation-based constraints that may be present.

The results obtained for HJDS, GPS and HOPSPACK are shown in Table 5 for the same nine initial well locations considered in Table 4. The settings and stopping criteria for these derivative-free optimizers are the same as were used for GPS. In all cases the maximum number of major iterations in the gradient-based

Figure 7 Oil saturation distribution at the end of the production time frame for the well controls and locations corresponding to the run from Table 4 with maximum NPV: (a) sequential fixed, (b) sequential reactive, and (c) joint approach. Injection and production wells are represented as blue and red circles, respectively.

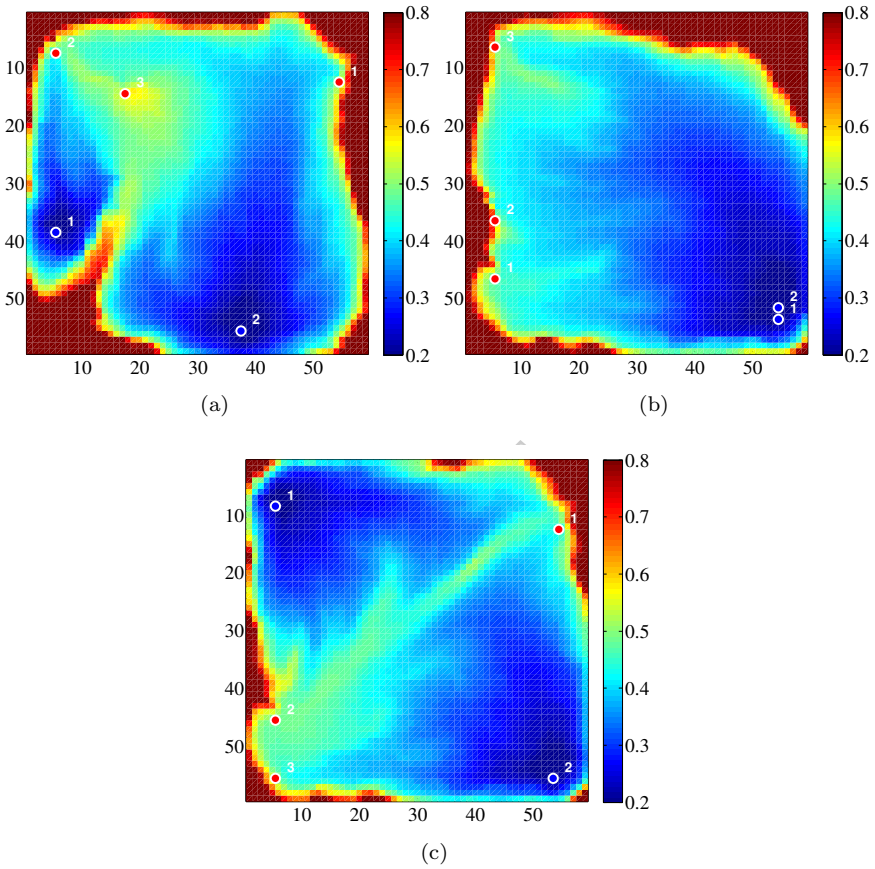

control optimization is equal to 8 (and again all runs include an additional control optimization with a maximum number of major iterations of 32). The total number of control optimizations solved (n_{ps}) coincides with the number of times the function NPV* is called within each pattern-search algorithm, and can also be assumed to be roughly proportional to the total computing cost. Each call to NPV* involves approximately 12-15 forward simulations, so the total number of simulations is around 4000 for the different approaches (consistent with the 3670 value given in Table 4). It is useful to express the results in terms of n_{ps} (rather than the total number of simulations n_{sim}) when the optimizations are parallelized.

In the absence of distributed computing resources, HJDS performs marginally better than GPS and HOPSPACK. However, these two algorithms are preferable to HJDS once they are implemented in parallel. The effect of distributed computing

Table 5 Results for NPV (expressed in $MM) and total number of control optimizations solved (n_{ps}) for the second example using Hooke-Jeeves direct search (HJDS), generalized pattern search (GPS) and a hybrid optimization parallel search package (HOPSPACK). The highest NPV for the nine runs for each approach is underlined.

Run	$joint^*_{HJDS}$		$joint^*_{GPS}$		$joint^*_{HOPS}$	
#	NPV	n_{ps}	NPV	n_{ps}	NPV	n_{ps}
1	<u>386.6</u>	331	<u>385.1</u>	188	<u>386.6</u>	222
2	380.8	391	355.2	225	388.0	333
3	327.6	216	346.7	225	358.5	375
4	386.4	316	372.3	175	380.0	358
5	377.5	321	354.8	331	343.1	361
6	344.9	456	336.0	244	333.9	286
7	377.9	556	360.0	285	353.9	331
8	313.8	441	357.2	311	350.6	344
9	371.9	306	358.0	332	358.0	468
Mean	363.0	370	358.4	257	361.4	342
σ	27.3	101	14.0	60	19.3	66

Figure 8 Evolution of NPV averaged over the nine runs versus the equivalent number of control optimizations for the three pattern-search optimization algorithms. The number of equivalent control optimizations solved is the total number of optimizations divided by an estimate of the speedup. The speedup factors estimated for GPS and HOPSPACK are 4.1 and 6.4, respectively.

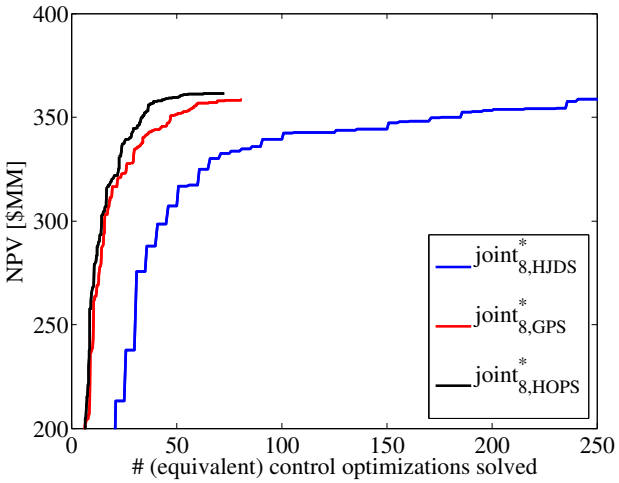

on GPS and HOPSPACK is shown in Figure 8. In that figure the vertical axis represents the evolution of NPV averaged over all nine runs, and the horizontal axis corresponds to the equivalent number of control optimizations solved for

each pattern-search optimization algorithm. The number of equivalent control optimizations solved is defined as the total number of optimizations divided by an estimate of the speedup obtained through parallelization. We note that HJDS is inherently serial, and for that reason the number of equivalent optimizations coincides with the total number of optimizations solved. For all algorithms, the horizontal axis in Figure 8 is roughly proportional to total clock time. Though GPS and HOPSPACK are parallelized on 8 and 20 computing cores, respectively, the speedup factors estimated for these procedures are 4.1 and 6.4, respectively. As a consequence, as can be seen in Figure 8, HOPSPACK outperforms GPS in terms of total elapsed time.

It is worth noting that the ratio of the two speedup factors is different than the ratio of the numbers of computing cores used for the two algorithms. This discrepancy is related to an observed increase in the reservoir flow simulation clock time with the number of nodes used. This may be explained by the concurrent sharing of common libraries by the parallelized simulations or by excessive input/output data traffic within the cluster. We can thus expect that in some practical applications there might be an optimal number of nodes to use in a distributed computing framework (in other words, a larger number of nodes does not always provide a higher speedup factor). From Figure 8 it can also be concluded that HJDS could be an alternative to the other two derivative-free methods if distributed computing resources are limited or unavailable, particularly if the algorithm need not be run to full convergence.

5 Concluding Remarks

In this work we considered the joint optimization of oil well placement and well controls. These two problems, though clearly coupled, have been treated as separate optimizations in most previous studies. We devised a nested optimization approach where the outer (high-level) optimization addresses the well placement problem. For each well configuration, the optimization cost function is defined as the optimal objective function value after performing a well control optimization for the particular well arrangement. Since well control optimization often displays a more convex character than well placement optimization, the former optimization can be approached from a more local, and thus efficient, perspective than the latter optimization. Therefore, in the well control optimization we apply a gradient-based procedure, with gradients provided by an adjoint solution. For the well placement part of the optimization, several pattern search algorithms were considered. Although these are local optimizers, by using a large initial stencil size we achieve some amount of global search.

We considered two optimization problems involving different numbers of optimization variables. Three basic optimization strategies were considered – two of these were sequential schemes that involved particular assumptions regarding the well controls (specifically, fixed and reactive controls) used during the well location optimizations, and the third was the joint optimization procedure. In all cases, after the basic optimization had converged, we performed an additional well control optimization (for the optimized well locations) using the gradient-based procedure with a large number of iterations. In the first example, the location of only one

well was optimized, so we were able to perform an exhaustive search. This enabled a clear assessment of the performance of the different optimization methods. The exhaustive search results showed that the optimization landscape corresponding to the well location in the joint approach was smoother (suggesting that global exploration can be more readily accomplished in this case) than the optimization landscape for the sequential fixed strategy.

The joint procedure was shown to consistently outperform the sequential schemes in terms of the optimized cost function (net present value in our examples). For the second (more challenging and more realistic) example, the increase in net present value obtained by the joint approach exceeded that achieved by the sequential methodologies by around 20% on average. The joint approach does, however, require around an order of magnitude more reservoir simulations than are required for the sequential approaches. This high computational demand can be mitigated through use of parallel implementations of the pattern search algorithms. Two of the algorithms considered, generalized pattern search (GPS) and hybrid optimization parallel search package (HOPSPACK), parallelize readily and such implementations were in fact applied.

The joint optimization procedure presented here can be extended in several interesting directions. For example, surrogate models could be used to accelerate the optimizations. Specifically, in some of the computations, the cost function value could be estimated using a sequential reactive strategy. This approach would be most effective if the particular reactive strategy is 'tuned' (including some treatment for injection wells) based on the joint optimization results. Further effort should also be expended toward including inter-well distance constraints and nonlinear simulation-based production constraints (such as maximum fraction of water in the produced fluid), possibly through use of a filter method (see e.g., Echeverria Ciaurri et al., 2011a). It will also be useful to consider global exploration techniques such as particle swarm optimization (Eberhart et al., 2001) or genetic algorithms (Goldberg, 1989) for the well placement part of the optimization. Uncertainty in the reservoir model should also be included in the optimization using, for example, the stochastic procedure recently presented by Wang et al. (2011). Developments along some of these lines are currently underway and will be the subject of future publications.

Acknowledgements

We are grateful to the Center for Integrated Operations in the Petroleum Industry at NTNU for partial funding of this work. We also thank the Stanford Smart Fields Consortium and the Department of Energy Resources Engineering for hosting the first author as a visiting researcher.

References

Audet, C. and Dennis Jr., J.E. (2000) 'Pattern search algorithms for mixed variable programming', *SIAM J. on Optimization*, vol. 11, pp. 573–594.

Audet, C. and Dennis Jr., J.E. (2002) 'Analysis of generalized pattern searches', *SIAM Journal on Optimization*, vol. 13, no. 3, pp. 889–903.

Audet, C. and Dennis Jr., J.E. (2006) 'Mesh adaptive direct search algorithms for constrained optimization', *SIAM Journal on Optimization*, vol. 17, no. 1, pp. 188–217.

Aziz, K. and Settari, A. (1979) *Petroleum Reservoir Simulation*, Applied Science Publishers, New York, NY.

Bangerth, W., Klie, H., Wheeler, M., Stoffa, P., and Sen, M. (2006) 'On optimization algorithms for the reservoir oil well placement problem', *Computational Geosciences*, vol. 10, no. 3, pp. 303–319.

Brouwer, D.R. and Jansen, J.D. (2004) 'Dynamic optimization of waterflooding with smart wells using optimal control theory', *SPE Journal*, vol. 9, no. 4, pp. 391–402, SPE 78278.

Cao, H. (2002) *Development of Techniques for General Purpose Simulators*, Ph.D. thesis, Dept. of Petroleum Engineering, Stanford University, Stanford, CA.

Christie, M.A. and Blunt, M.J. (2001) 'Tenth SPE comparative solution project: A comparison of upscaling techniques', *SPE Reservoir Evaluation & Engineering*, vol. 4, pp. 308–317, SPE 66599.

Clerc, M. (2006) *Particle Swarm Optimization*, Wiley-ISTE, London.

Conn, A.R., Scheinberg, K., and Vicente, L.N. (2009) *Introduction to Derivative-Free Optimization*, MPS-SIAM Series on Optimization, Society for Industrial and Applied Mathematics and the Mathematical Programming Society, Philadelphia, PA.

Dempe, S. (2002) *Foundations of Bilevel Programming*, Kluwer Academic Publishers, New York, NY.

Eberhart, R.C., Kennedy, J., and Shi, Y. (2001) *Swarm Intelligence*, Morgan Kaufmann, Waltham, MA.

Echeverria Ciaurri, D., Isebor, O.J., and Durlofsky, L.J. (2011a) 'Application of derivative-free methodologies for generally constrained oil production optimization problems', *International Journal of Mathematical Modelling and Numerical Optimisation*, vol. 2, no. 2, pp. 134–161.

Echeverria Ciaurri, D., Mukerji, T., and Durlofsky, L.J. (2011b) 'Derivative-free optimization for oil field operations', in X.S. Yang and S. Koziel, editors, 'Computational Optimization and Applications in Engineering and Industry', vol. 359 of *Studies in Computational Intelligence*, pp. 19–55, Springer, Berlin / Heidelberg.

Ertekin, T., Abou-Kassem, J.H., and King, G.R. (2001) *Basic Applied Reservoir Simulation*, vol. 7 of *SPE Textbook Series*, SPE, Richardson, TX.

Forouzanfar, F., Li, G., and Reynolds, A.C. (2010) 'A two-stage well placement optimization method based on adjoint gradient', in 'Proceedings of the 2010 SPE Annual Technical Conference and Exhibition', Florence, Italy, SPE 135304.

Gill, P.E., Murray, W., and Saunders, M.A. (2007) *User's Guide for SNOPT Version 7: Software for Large Scale Nonlinear Programming*, Stanford University, Stanford, CA, and University of California, San Diego, CA.

Goldberg, D.E. (1989) *Genetic Algorithms in Search, Optimization and Machine Learning*, Addison-Wesley Professional, Boston, MA.

Griffin, J.D. and Kolda, T.G. (2007) 'Nonlinearly-constrained optimization using asynchronous parallel generating set search', Tech. Rep. SAND2007-3257, Sandia National Laboratories, Livermore, CA.

Griffin, J.D., Kolda, T.G., and Lewis, R.M. (2008) 'Asynchronous parallel generating set search for linearly-constrained optimization', *SIAM Journal on Scientific Computing*, vol. 30, no. 4, pp. 1892–1924.

Guyaguler, B., Horne, R.N., Rogers, L., and Rosenzweig, J.J. (2000) 'Optimization of well placement in a Gulf of Mexico waterflooding project', in 'Proceedings of the 2000 SPE Annual Technical Conference and Exhibition', Dallas, TX, SPE 63221.

Hooke, R. and Jeeves, T.A. (1961) 'Direct search solution of numerical and statistical problems', *Journal of the ACM*, vol. 8, no. 2, pp. 212–229.

Jansen, J.D., Brouwer, D.R., Naevdal, G., and van Kruijsdiik, C.P.J.W. (2005) 'Closed-loop reservoir management', *First Break*, vol. 23, pp. 43–48.

Kolda, T.G., Lewis, R.M., and Torczon, V. (2003) 'Optimization by direct search: New perspectives on some classical and modern methods', *SIAM Review*, vol. 45, no. 3, pp. 385–482.

Nocedal, J. and Wright, S.J. (2006) *Numerical Optimization*, Springer, 2nd Edition, New York, NY.

Onwunalu, J.E. and Durlofsky, L.J. (2010) 'Application of a particle swarm optimization algorithm for determining optimum well location and type', *Computational Geosciences*, vol. 14, no. 1, pp. 183–198.

Peaceman, D.W. (1978) 'Interpretation of well-block pressures in numerical reservoir simulation', *SPE Journal*, vol. 18, no. 3, pp. 183–194, SPE 6893.

Pironneau, O. (1974) 'On optimum design in fluid mechanics', *J. Fluid Mech.*, vol. 64, pp. 97–110.

Plantenga, T.D. (2009) 'HOPSPACK 2.0 user manual', Tech. Rep. SAND2009-6265, Sandia National Laboratories, Livermore, CA.

Powell, M.J.D. (2009) 'The BOBYQA algorithm for bound constrained optimization without derivatives', Tech. Rep. NA2009/06, Department of Applied Mathematics and Theoretical Physics, Cambridge, UK.

Sarma, P. and Chen, W.H. (2008) 'Efficient well placement optimization with gradient-based algorithm and adjoint models', in 'Proceedings of the 2008 SPE Intelligent Energy Conference and Exhibition', Amsterdam, The Netherlands, SPE 112257.

Sarma, P., Durlofsky, L.J., Aziz, K., and Chen, W.H. (2006) 'Efficient real-time reservoir management using adjoint-based optimal control and model updating', *Computational Geosciences*, vol. 10, pp. 3–36.

Torczon, V. (1997) 'On the convergence of pattern search algorithms', *SIAM Journal on Optimization*, vol. 7, no. 1, pp. 1–25.

Wang, C., Li, G., and Reynolds, A.C. (2007) 'Optimal well placement for production optimization', in 'Proceedings of the 2007 SPE Eastern Regional Meeting', Lexington, KY, SPE 111154.

Wang, H., Echeverria Ciaurri, D., Durlofsky, L.J., and Cominelli, A. (2011) 'Optimal well placement under uncertainty using a retrospective optimization framework', in 'Proceedings of the 2011 SPE Reservoir Simulation Symposium', The Woodlands, TX, SPE 141950.

Yeten, B., Durlofsky, L.J., and Aziz, K. (2003) 'Optimization of nonconventional well type, location, and trajectory', *SPE Journal*, vol. 8, no. 9, pp. 200–210, SPE 86880.

Zandvliet, M., Handels, M., van Essen, G., Brouwer, R., and Jansen, J.D. (2008) 'Adjoint-based well-placement optimization under production constraints', *SPE Journal*, vol. 13, no. 4, pp. 392–399, SPE 105797.

Zhang, K., Li, G., Reynolds, A.C., Yao, J., and Zhang, L. (2010) 'Optimal well placement using an adjoint gradient', *Journal of Petroleum Science and Engineering*, vol. 73, no. 3-4, pp. 220–226.