Noname manuscript No.
(will be inserted by the editor)

Simultaneous and sequential approaches to joint optimization of well placement and control

T.D. Humphries · R.D. Haynes · L.A. James

the date of receipt and acceptance should be inserted later

Abstract Determining optimal well placement and control is essential to maximizing production from an oil field. Most academic literature to date has treated optimal placement and control as two separate problems; well placement problems, in particular, are often solved assuming some fixed flow rate or bottom-hole pressure at injection and production wells. Optimal placement of wells, however, does depend on the control strategy being employed. Determining a truly optimal configuration of wells thus requires that the control parameters be allowed to vary as well. This presents a challenging optimization problem, since well location and control parameters have different properties from one another.

In this paper we address the placement and control optimization problem jointly using approaches that combine a global search strategy (particle swarm optimization, or PSO) with a local generalized pattern search (GPS) strategy. Using PSO promotes a full, semi-random exploration of the search space, while GPS allows us to locally optimize parameters in a systematic way. We focus primarily on two approaches combining these two algorithms. The first is to hybridize them into a single algorithm that acts on all variables simultaneously, while the second is to apply them sequentially to decoupled well placement and well control problems. We find that although the best method for a given problem is context-specific, decoupling the problem may provide benefits over a fully simultaneous approach.

T.D. Humphries · R.D. Haynes · L.A. James
Memorial University of Newfoundland
St. John's, Newfoundland, Canada, A1C 5S7
E-mail: thumphries@mun.ca

1 Introduction

Maximizing production from an oil field is a crucial task, given the enormous financial investment at stake in any large-scale field development. Decisions regarding the placement of new wells, and control of injection and production rates at existing wells, have a significant impact on production. Poor placement and/or control of wells may result in premature water breakthrough at the production wells, or make it difficult to achieve high flow rates while maintaining reservoir pressure. The vast number of potential development scenarios drives the need for efficient, computerized optimization approaches to assist in making these decisions.

The problems of finding optimal well locations and determining optimal well control parameters are often treated separately [11]. Well placement problems involve optimizing over parameters corresponding to the positions and orientations of the injection and production wells. We limit ourselves in this paper to considering vertical wells, meaning that each well's position is parameterized simply by its (x, y) co-ordinates. A simple control scheme is typically assumed in well placement problems; for instance, injection wells are held at a fixed bottom hole pressure (BHP), while producers are held at a lower BHP in order to generate flow. Well control problems, on the other hand, focus on managing the injection and production rates at wells that are already in place. The optimization variables in this case are usually either the BHP or the flow rate for each well, which can be changed at specified time intervals. The objective function that one attempts to maximize in these problems is typically either the total amount of oil extracted, or the net present value (NPV) of the extracted oil. The NPV function is correlated with the total amount of oil extracted, but emphasizes producing more oil early in the reservoir's lifetime (due to the time value of money), and also typically includes the costs of water injection and disposal of any water produced. Heterogeneous properties of the reservoir, such as the permeability field, have a significant impact on optimal well placement; as a result, the objective function surface tends to be much rougher in well placement problems than in well control problems. This difference in the objective function surface typically leads to the use of different optimization approaches to address these two problems. Optimization studies on well placement have often focused on global algorithms with some stochastic element in order to avoid converging prematurely to local optima; well control problems have tended to make use of deterministic algorithms, based on local search techniques [11].

A unified approach to optimizing well placement and well control jointly has the potential to provide benefits over the treatment of these problems separately. In particular, the best well configuration when producers are held at some fixed BHP is not necessarily the same as the best configuration if the control can vary with time [39]. Additionally, determining the optimal placement of new wells may also require adjusting the control parameters at wells already in place. The problem of jointly optimizing well placement and control has been largely unexplored in the literature until recently. Here, we investigate approaches to addressing this problem using a two optimization algorithms: particle swarm optimization (PSO) and generalized pattern search (GPS). These approaches include hybridizing the algorithms to simultaneously optimize over both placement and control parameters, as well as applying them in sequential steps in a decoupled approach to the problem.

2 Existing research

Well placement studies have tended to use stochastic optimization approaches aimed at exploring the solution space globally. Genetic algorithms (GAs) have received the widest use [1, 5, 8, 12, 15, 25, 29, 37]. Other optimization algorithms that have been applied to the problem include simultaneous perturbation stochastic approximation (SPSA) [3], covariance matrix adaptation [6], and particle swarm optimization (PSO) [27,28]. In addition to studying suitable algorithms for the optimization problem, well placement papers have discussed other issues such as parametrization and optimal placement of nonconventional wells [8, 37], consideration of geological uncertainty when determining optimal positions [1,15], placement of well patterns rather than individual wells [28,29], and inclusion of nonlinear constraints as part of the optimization [12, 39].

A popular optimization algorithm for well control problems, on the other hand, has been the adjoint method [7, 13, 19, 31, 38]. This method, which is based on approximating the gradient of the objective function, is well-suited to the optimal control problem due to the objective function's smoothly varying nature. Implementing the adjoint method may be challenging, however, since it requires in-depth knowledge of the workings of the reservoir simulator. An alternative is to use "black box" optimization algorithms, which optimize based only on the output from simulator. Examples of black-box algorithms that have been applied to the well control problem include stochastic methods like SPSA [35] and GAs [36], as well as deterministic methods such as generalized pattern search (GPS) and Hooke-Jeeves directed search [10, 11].

Some recent papers have considered optimization of well placement and well control in a more unified manner. One proposed approach [14] essentially treats the problem as one of optimizing well control rates, with a large number of injection and production wells being placed in the reservoir initially. As the optimization proceeds, flow rates at some wells are driven to zero, resulting in their removal from the simulation. Thus, the procedure determines the optimal number of injectors and producers, as well as their locations and optimal flow rates. This approach uses a modified version of the adjoint method, which incorporates a number of inequality and equality constraints to allow for the removal of wells after every iteration. A second approach uses a combination of GPS and adjoint methods in a nested optimization procedure [4]. The outer iteration of this procedure consists of using GPS to determine optimal well positions; for each configuration of wells, an inner optimization routine uses the adjoint method to determine the best control strategy. The SPSA algorithm has also been applied to problems involving both well placement and well control in [22,23], where it was found that optimizing over all variables simultaneously was preferable to applying SPSA sequentially to subproblems involving placement or control only.

The use of PSO and GPS in tandem to address the well placement and control problem has also been investigated in [18]. There it was found that hybridizing PSO and GPS provided better results than the independent application of those algorithms, and that optimizing over all variables simultaneously was preferable to sequential optimizations. The approach we use in this paper is similar to theirs, but in preliminary experiments presented in [17], we found that sequential optimization was sometimes preferable to a fully simultaneous approach. In this paper we refine the technique and perform further experiments, to determine under which circumstances one approach may be preferable to the other.

3 Optimization approach

Combining global and local optimization techniques should be advantageous when addressing well control and well placement simultaneously. We use PSO as a global optimizer in this study, and GPS for the local search. Our choice of these algorithms is motivated by the fact that both have performed well in previous production optimization studies [10, 27]; both are black-box methods that do not require in-depth knowledge of the simulator; and both are easily parallelized to help mitigate the high cost of function evaluations. We now give an overview of PSO and GPS, as well as of the specific optimization approaches used in this paper.

3.1 Particle swarm optimization

Particle swarm optimization [9, 20] is an optimization algorithm based on modeling the behaviour of a group of animals acting collectively. PSO utilizes a number of *particles* (typically 20 to 40) to explore solution space in a semi-random way. The position of particle i at iteration k, denoted $\mathbf{x}_i^{(k)}$, is a vector of size N, where N is the number of variables in the optimization problem. Every position in solution space is associated with the corresponding objective function value, and every particle remembers the best position it has found so far. Particles in the swarm also communicate with one another to share the best positions that have been found overall. Each iteration of PSO consists of determining a new position for every particle in the swarm, and then evaluating the objective function at that position. Since the objective function evaluations can be performed independently of one another, the algorithm is highly parallelizable.

Given $\mathbf{x}_i^{(k)}$, the position of the particle at iteration $k+1$ is:

$$\mathbf{x}_i^{(k+1)} = \mathbf{x}_i^{(k)} + \mathbf{v}_i^{(k+1)},$$

where the particle's velocity vector $\mathbf{v}_i^{(k+1)}$ is given by

$$\mathbf{v}_i^{(k+1)} = \iota \mathbf{v}_i^{(k)} + \mu \mathbf{r}_1^{(k)} \otimes \left(\mathbf{p}_i^{(k)} - \mathbf{x}_i^{(k)}\right) + \nu \mathbf{r}_2^{(k)} \otimes \left(\mathbf{g}_i^{(k)} - \mathbf{x}_i^{(k)}\right).$$

The velocity is a combination of three terms. The first term models the tendency of the particle to continue traveling in the direction given by its current velocity. The second term represents the tendency of the particle to move toward the best position it has found so far, denoted by $\mathbf{p}_i^{(k)}$. Finally, the third term represents the tendency of the particle to move toward the best position found by any other particle with which it communicates, denoted by $\mathbf{g}_i^{(k)}$. The constants ι, μ and ν are parameters whose values are chosen to weight these three terms appropriately. To inject randomness into the particle movement, the N-vectors $\mathbf{r}_1^{(k)}$ and $\mathbf{r}_2^{(k)}$ are generated from the uniform distribution on $(0, 1)$ at every iteration, then multiplied componentwise with the terms in brackets by the $\otimes$ operator. The algorithm iterates until some convergence criterion is met; for example, until the velocities of the particles have become sufficiently small, until the particles are sufficiently close to one another, or simply until some maximum number of iterations have been performed.

If every particle communicates with every other particle in the swarm, PSO may quickly converge to a local optimum before the solution space is fully explored. Thus, it is usually recommended that each particle communicate only with two to four other particles at any one time [9]. At every iteration of the algorithm, the neighbourhood of particles with which each particle communicates can be chosen randomly. This *random* neighbourhood topology was used for this study, as well as a swarm size of 40 particles, and parameter values of $\iota = 0.721$, and $\mu = \nu = 1.193$. These parameter values have been found to provide good convergence results in many numerical experiments [9]. It should be noted that in general, one can not guarantee that PSO converges to a global or even a local optimum; however, in practice it has proved to be effective for a wide variety of optimization problems.

3.2 Generalized Pattern Search

Generalized Pattern Search [2,21] is an optimization algorithm that begins from a single *incumbent point* and consists of a series of *search* and *poll* steps. At every iteration k, a discrete mesh, centred at the current incumbent $\mathbf{x}^{(k)}$, is defined by:

$$M^{(k)} = \left\{\mathbf{x}^{(k)} + \Delta^{(k)} D z \; : \; \mathbf{z} \in \mathbb{N}^{n_D}\right\},$$

where $\Delta^{(k)}$ is the resolution of the mesh at iteration k, D is a matrix whose columns form the *search directions*, $\mathbb{N}$ is the set of natural numbers, and n_D is the number of search directions. The search directions must form a *positive spanning set* in solution space; i.e., one must be able to specify any point in solution space by adding together only positive scalar multiples of these directions. A common choice of search directions is:

$$\mathcal{D} = \{\mathbf{e}_1, \mathbf{e}_2, \ldots, \mathbf{e}_N, -\mathbf{e}_1, -\mathbf{e}_2, \ldots, -\mathbf{e}_N\} \tag{1}$$

where the $\mathbf{e}_n$ are the canonical basis vectors $(1, 0, 0, \ldots, 0)^T$, $(0, 1, 0, \ldots, 0)^T$, etc. Here $\mathcal{D}$ refers to the set of search directions, which form the columns of the matrix D.

The search step consists of selecting a finite number of points on $M^{(k)}$ and evaluating the objective function at each one. If any of those points improves the objective function value, the point with the best value becomes the new incumbent. The search step can employ any strategy in selecting points, and may even be omitted, if desired. If none of the points selected in the search step are better than the incumbent, then the algorithm proceeds to the poll step. The poll step consists of evaluating the objective function at all the points that are immediate neighbours of the incumbent point on the mesh $M^{(k)}$. These points are given by:

$$\left\{\mathbf{y}_j^{(k)}\right\} = \left\{\mathbf{x}^{(k)} + \Delta^{(k)}\mathbf{d}_j \,|\, \forall\, \mathbf{d}_j \in \mathcal{D}\right\}.$$

If the poll step finds one or more points with a better objective function value than the incumbent, then the point with the best value becomes the new incumbent. Optionally, $\Delta^{(k)}$ may be increased for the next iteration. If the poll step is unsuccessful, then $\Delta^{(k)}$ is reduced and another iteration begins, using the same incumbent point as before. The algorithm is considered to have converged once $\Delta^{(k)}$ is reduced beyond a specified threshold, which indicates that the current point is at least close to a local optimum. In fact, provided that the objective function is continuously differentiable, GPS is guaranteed to converge to a local optimum, at least to mesh precision [21]. Like PSO, GPS is highly parallelizable because the function evaluations required by the search and poll steps can be performed independently of one another.

3.3 Handling of bound and general constraints

Broadly speaking, there can exist two types of constraints on the optimization vector $\mathbf{x}$: *bound* and *general* constraints. Bound constraints are simple componentwise inequality constraints of the form

$$\mathbf{x}^l \le \mathbf{x} \le \mathbf{x}^u,$$

where $\mathbf{x}^l$ and $\mathbf{x}^u$ are the lower and upper bounds on $\mathbf{x}$, respectively. In the context of a reservoir optimization problem, these could be the minimum and maximum grid indices (for well placement) or upper and lower limits on the control parameters.

Both PSO and GPS can easily incorporate bound constraints. In PSO, any particles that travel outside of the bounds are projected back onto the boundary of search space. For instance, if component d of particle i's position exceeds the maximum value $\mathbf{x}_d^u$ after being updated, then the particle's position and velocity are modified as follows:

$$\mathbf{x}_{i,d} = \mathbf{x}_d^u, \qquad \mathbf{v}_{i,d} = 0$$

The velocity component is set to zero to ensure that the particle does not continue to travel in the direction that led it out of bounds. Bound constraints are treated similarly in GPS; namely, points which lie outside of search space are projected back onto the boundary during the poll step [21].

General constraints refer to any constraints on the input parameters other than simple bound constraints. Input that violates general constraints (which we refer to as *infeasible* input) can sometimes be identified prior to evaluating the objective function; for instance, if the input specifies placing two wells at the same location. Other constraints, such as an upper limit on well flow rates for wells being driven by BHP, require running the reservoir simulator to determine if they are satisfied.

A simple mechanism for PSO to handle general constraints is to allow particles to move to infeasible positions, but not store these positions in the particle's memory [16]. Thus, particles can explore search space freely, but are only attracted to positions that are feasible, in addition to providing good objective function values. This strategy requires that every particle be initialized to a feasible position, so that the particle always has at least one feasible position stored in its history. To handle general constraints in GPS, one can simply ignore any infeasible points during polling, and thus only accept feasible points which also reduce the objective function value. This approach is not ideal for general-purpose optimization, as it may prevent the algorithm from traveling through the infeasible region to find the true optimum; alternative approaches such as *filtering* are recommended instead [2,10]. We found that the first approach was sufficient for this study, however, possibly because GPS was used in conjunction with PSO, rather than as a stand-alone optimizer.

3.4 Hybrid algorithm

An optimization algorithm that hybridizes PSO and GPS has previously been proposed in [33, 34]. This algorithm, denoted PSwarm, is essentially a GPS algorithm that uses PSO as the *search* step. Thus, the algorithm behaves exactly like PSO for as long as the search step continues to find points that improve the objective function value. When this step fails to improve the solution, polling takes place around the current best position found. If the poll step finds a better solution, the current best position is updated and a new iteration of PSO begins; otherwise, the polling stencil size is reduced. The algorithm proceeds until the convergence criteria for both PSO and GPS are met; i.e., the velocity of the particles is sufficiently small, and the polling stencil size is reduced beyond a specified threshold.

In this paper we have made the following modifications to PSwarm in order to adapt it to the simultaneous well placement and control problem:

1. We have extended the PSO and GPS components of the algorithm to handle general constraints, as described in the previous section. The PSwarm algorithm as described in [33] handles linear constraints,

but not general constraints of the type seen in this problem.
2. We have replaced the global communication topology used by the search step of PSwarm with the random variable neighbourhood topology. Each particle's communication neighbourhood consists of itself and two other particles, which are selected randomly at each iteration.
3. We have chosen to skew the sampling of control parameters when initializing the swarm. PSO is typically initialized by assigning a position sampled uniformly within the search space to each particle in the swarm. In the context of this problem, however, we can be reasonably certain that for production wells, BHP values towards the low end of the range will tend to increase oil production, while BHP values towards the high end of the range will do the same for injection wells. We take advantage of this *a priori* knowledge in order to accelerate the convergence of the algorithm.
4. We have investigated the effect of allowing the search step (PSO) to fail several times consecutively before a poll step is performed. As a result of the high cost of function evaluations in our problem, the poll step is more computationally expensive than in many other optimization problems. Performing the poll step less frequently may allow us to reduce the computational cost of the algorithm.
5. We have investigated the use of specifically selected direction vectors $\mathcal{D}$ to use during the poll step. The standard GPS search directions (Eq. (1)) are fairly incremental, particularly with respect to control parameters, since each variable corresponds only to the BHP value at a single well for a single time interval. We may be able to achieve a larger improvement in a single step by choosing search directions that raise or lower BHP in multiple years. The specific choice of search directions is discussed in the next section.

4 Experiments

We now describe several experiments that were used to test the performance of the different optimization approaches. All experiments were performed using the Matlab Reservoir Simulation Toolbox (MRST) [24, 32] as the reservoir simulator. MRST is an open-source simulator implemented in Matlab, which includes routines for processing and visualizing unstructured grids, as well as several solvers for single and two-phase flow. The flow and transport equations are solved in alternating steps in order to determine the phase pressures, flow rates and saturations at every time point. Modeling of simple vertical and horizontal wells is provided using the Peaceman model [30].

Table 1 Economic parameters used in all experiments.

Parameter	Value
c_o	\$80/bbl
$c_{w,disp}$	\$12/bbl
$c_{w,inj}$	\$8/bbl
r	10% or 0%
Max water cut	78%

The objective function we used in these experiments was the net present value (NPV) over the entire production period $[0, T]$. The NPV was computed as in [3], with

$$NPV(\mathbf{x}) = \int_0^T \left\{ \sum_{n \in prod} \left[c_o q_{n,o}^-(t) - c_{w,disp} q_{n,w}^-(t) \right] - \sum_{n \in inj} c_{w,inj} q_{n,w}^+(t) \right\} (1+r)^{-t} \, dt. \tag{2}$$

The parameters c_o, $c_{w,disp}$ and $c_{w,inj}$ represent the price per barrel of produced oil, disposal cost per barrel of produced water, and cost per barrel of injected water, respectively. The functions $q_{n,o}^-(t)$ and $q_{n,w}^-(t)$ are the production rates (barrels/day) of oil and water, respectively, at well n, while $q_{n,w}^+(t)$ is the water injection rate at well n. These rates are implicitly functions of the optimization vector $\mathbf{x}$, since they depend on the well positions and prescribed BHPs. The yearly interest rate is specified by r. We used the parameter values provided in Table 1 for all experiments. This choice of values meant that production became unprofitable once the water cut at a well reached roughly 78%. This threshold value is often as high as 90 or 95% in practice; we chose a lower value to ensure that shutting in a well was the optimal choice in some experiments.

Experiment 1

The first experiment used a simple 2D reservoir model, consisting of 50×60 grid cells measuring 32×32×10 m (total field size: 1600×1920×10m). The permeability and porosity fields (Fig. 1) were taken from the third layer of the SPE10 Benchmark model. The reservoir was initially saturated uniformly with an 80/20 mix of oil to water. The optimization problem was to place two injection and two production wells in the reservoir, all of which were subject to control via BHP. The production period was 10 years, and the BHP at each well could be altered every 2 years. Thus, there were 28 variables

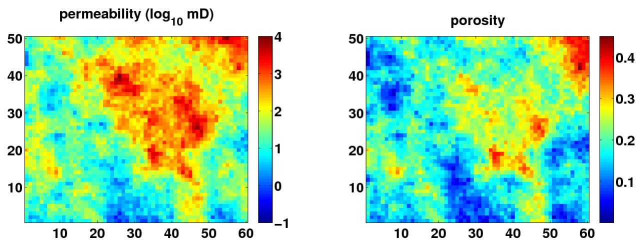

Fig. 1 Permeability and porosity fields used in first experiment. Permeability values are shown on a logarithmic scale.

Table 2 Parameters (top) and constraints (bottom) used in Experiment 1.

Parameter	Value
Grid cell dimensions	32 × 32 × 10 m
Fluid viscosities μ_o and μ_w	2.4 and 1.0 cp
Fluid densities ρ_o and ρ_w	835 and 1000 kg/m^3
Initial reservoir pressure	260 bars
Injector BHP range	275–450 bars
Producer BHP range	100–250 bars
Production period	10 years
Control interval	2 years
Maximum water injection rate	1000 m^3/day
Maximum fluid production rate	1000 m^3/day
Minimum distance between wells	250 m

being considered in total (2 positional variables and 5 control variables per well). The experimental parameters are summarized in Table 2.

We considered three variations of the optimizationl problem:

– Case 1A: no constraints on injection and production rates; discounting rate r of 10%.
– Case 1B: no constraints on injection and production rates; no discounting.
– Case 1C: maximum flow constraints on the injection and liquid production rates as described in Table 2; discounting rate r of 10%.

The goal of considering these three subproblems was to see how the different conditions affected the optimal solutions to the problem, as well as the effectiveness of different optimization approaches. Three main optimization approaches were applied to each problem. The fact that every optimization approach that we considered included a stochastic component necessitated performing multiple runs of each approach, in order to assess the average performance. Each approach was, therefore, applied 20 times to the appropriate problems.

The first approach was to apply PSO to the vector of all decision variables (well locations and control parameters) simultaneously. PSO was run for up to 250 iterations, or until the average velocity of the particle swarm decreased below a certain threshold. We subsequently used GPS to poll repeatedly around the best solution found by each run of PSO, in order to see if it could be improved further. This second step was not considered to be part of the optimization approach, but rather as a test to see how close the solutions found by PSO were to being locally optimal.

The second approach was to apply 200 iterations of the hybrid algorithm described in the previous section. The following three variants of the hybrid algorithm were applied:

– `hybrid-1`: Polling was performed every time the PSO step failed to improve the solution. We used the standard search directions (Eq. (1)).
– `hybrid-5`: Polling was performed only after PSO failed five times consecutively. We used the standard search directions.
– `hybrid-5S`: Polling was performed only after PSO failed five times consecutively. We used special search directions.

The special search directions used by the `hybrid-5S` approach for this problem are illustrated in Fig. 2. Every row of the matrix shown corresponds to one of the 28 variables, and every column represents one search direction. Consider the first 14 directions shown, which act on the first seven variables only. These variables correspond to the two positional parameters (x and y co-ordinates) and five control parameters for the first injector. The key difference from the standard search directions is that we allow the BHP of the injector to be raised for more than one time period simultaneously. The BHP of the injector can only be lowered, however, for one time period at a time. The rationale is that raising the BHP in an injector increases flow, and is thus more likely to raise oil production than lowering the BHP. For producers, the opposite is true, and so in that case we allow the BHP to be lowered for multiple time periods simultaneously.

In all the hybrid approaches, we scaled the directions corresponding to positional parameters independently of the control variables, so that the x or y co-ordinates of a well were only ever perturbed by one grid space during polling. The idea is that the optimization of the well positions is primarily achieved by the PSO step. Well positions should only need to be perturbed slightly during the poll step, which is aimed mainly at optimizing the controls.

The third approach we considered was to decouple the placement and control components of the problem. The first step of this approach consisted of treating the problem strictly as a well placement problem, by assuming that the producers were held at some fixed BHP throughout the production period. We used up to 200 iterations of PSO to determine the optimal well positions under this assumption. Once optimal positions

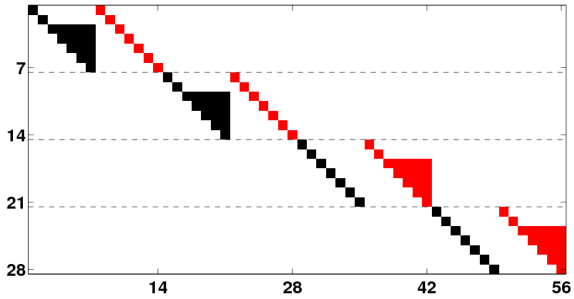

Fig. 2 Specialized search directions used by the `hybrid-5S` approach for Experiment 1. A red entry corresponds to a value of -1, and a black entry to +1. Dashed horizontal lines separate variables corresponding to each of the the four wells.

were found, we allowed the controls to vary, and optimized the control using GPS with standard search directions. The positions could also be incrementally adjusted in this second step. This second step ensured that the solutions found by the decoupled approach were locally optimal. The advantage of the decoupled approach is that it splits the problem into two smaller problems which are easier to solve than the full problem. A potential disadvantage is that we may find suboptimal solutions by not optimizing over all variables at the same time.

The decoupled approach requires defining fixed BHP values to assign to the injection and production wells during the first (well placement) phase of the approach. Our default choice was to use the maximum BHP value (450 bars) for injectors, and the minimum BHP value (100 bars) for producers. For Case 1C, where there were maximum constraints on the injection and production rates, we also tried a second variant where a BHP value of 425 bars was used for injectors, and 125 bars for producers. The rationale for this modification, which we denote by `decoupled-M`, is that using the maximum and minimum BHP values when placing wells will produce the highest flow rates possible for a given configuration as wells. If these flow rates exceed the maximum flow constraint, then these configurations will be considered as infeasible during the first stage of optimization, even if they could be made feasible by adjusting the BHP values. The net effect is that the well positions found in the first phase may tend to place wells farther apart than necessary, or in regions of lower permeability, in order to satisfy the flow constraints. Thus, by choosing BHP values that are slightly below or above the maximum and minimum values, respectively, we may be able to find better well positions.

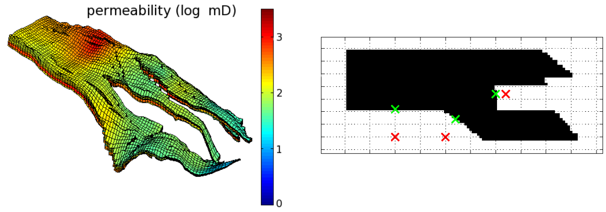

Fig. 3 Left side: reservoir geometry of the *Ile* formation from the Norne field. Log of permeability (mD) in the $x-y$ directions is shown. Right side: Projection of invalid vertical well locations in the (x, y) plane onto the nearest valid coordinates.

Experiment 2

The second experiment used a reservoir model provided by the Norwegian University of Science and Technology (NTNU) as part of the Norne benchmark case [26]. The full model of the Norne field is a 46×112×22 grid consisting of 44,927 active cells. The reservoir model is subdivided into four different formations from top to base, denoted *Garn*, *Ile*, *Tofte* and *Tilje*. In order to reduce simulation time, we extracted the seven layers corresponding to the *Ile* formation to provide a smaller reservoir model, consisting of 15,004 active cells. The porosity of the reservoir ranged between 25–30% and the permeability from 20 to 2500 mD. The reservoir geometry is shown in Fig. 3 (left image). The initial saturation was assumed to be 80% oil and 20% water, as in Experiment 1.

The reservoir's irregular shape meant that wells whose positional co-ordinates fell within the bounds prescribed by the grid might not correspond to valid locations in the reservoir. Thus, any positional co-ordinates in the (x, y) plane which did not correspond to a valid reservoir location were projected onto the nearest active cell during the optimization. This process is illustrated in Fig. 3 (right image). Black cells indicate grid locations which pass through at least one active cell in the z-direction. The three red $\times$ symbols indicate positions that are invalid, which were projected onto the nearest valid location (indicated by the green $\times$ symbols).

The goal of Experiment 2 was to test the optimization approaches on a more complex simulation with a larger number of variables. The experiment consisted of placing eleven wells (four injectors, seven producers) in this field, and optimizing production over a 15-year time period. The experimental parameters are summarized in Table 3. As in Experiment 1, there were 2 positional and 5 control parameters associated with each well, meaning that there were 77 variables in total. The same economic parameters were used as in Experiment 1 (see Table 1). For this experiment, only two

Table 3 Parameters (top) and constraints (bottom) used in Experiment 2.

Parameter	Value
Fluid viscosities μ_o and μ_w	2.4 and 1.0 cp
Fluid densities ρ_o and ρ_w	835 and 1000 kg/m^3
Initial reservoir pressure	340 bars
Injector BHP range	350–500 bars
Producer BHP range	150–325 bars
Production period	15 years
Control interval	3 years
Maximum water injection rate	5000 m^3/day
Maximum fluid production rate	3000 m^3/day

cases were considered; Case 2A, where there were no constraints on production, and Case 2B, which used the constraints given at the bottom of Table 3. A discounting rate of 10% was used for both cases.

Based on the results of Experiment 1, we limited the number of optimization approaches that we investigated to the two that were found to be most effective: `hybrid-5S`, and one of the two decoupled approaches. For Case 2A, we used the `decoupled` approach, where the well BHPs were held at their maximum and minimum values; for Case 2B, we used the `decoupled-M` modification with injectors held at a BHP of 450 bars and producers held at 200 bars. We performed only five runs of each optimization approach, due to the high computational cost of this experiment. The hybrid algorithm was run for up to 300 iterations per run, while the decoupled approach involved up to 250 iterations of PSO during the well placement phase, followed by running GPS until convergence.

5 Results

The results of both experiments are shown in Table 4. The section entitled "NPV" shows the average, best and worst NPV values over the multiple runs of each approach that were performed for each test case (twenty runs for Experiment 1, and five runs for Experiment 2). The section entitled "reliability" quantifies the reliability of each approach, by indicating how often the solutions found were within 10% and 5% of the best solution found overall. A reliability value of 0.55 in the 10% column, for instance, indicates that 11 out of the 20 solutions found by that method had an NPV within 10% of the best overall. The section entitled "after GPS" shows the average NPV after the GPS algorithm was applied to each solution found by PSO and the hybrid algorithm, as well as the percentage improvement ($\Delta\%$) compared to the original average. These values indicate how close, on average, the solutions found by each algorithm were to being locally optimal. These values were not calculated for any of the decoupled approaches, since the solutions found by those approaches were guaranteed to be locally optimal. Plots of the mean convergence of the respective algorithms versus the number of function evaluations (fevals) for both experiments are shown in Fig. 4.

6 Discussion

Experiment 1

Some general observations can be drawn from the results shown in Table 4. First, the decoupled approach was more reliable than the simultaneous approaches (i.e. the PSO and hybrid algorithms). In all three test cases, every solution found by the decoupled method was within 10% of the best overall; the hybrid approaches typically scored between 0.8 and 0.95, and PSO's reliability was as low as 0.55 in two out of the three test cases. The decoupled approach also gave the most results within 5% of the best overall for test cases 1B and 1C. On average, the NPV of solutions found by the decoupled approach was better than the other methods for Cases 1B and 1C.

This result indicates that reducing the size of the problem by focusing on well placement first, while assuming that wells are held at or near the extreme BHP values, may help to ensure that one obtains a "good" solution. The reason may be that this approach allows us to explore a number of well placement possibilities while holding the controls at a configuration that is generally likely to result in higher production. In the simultaneous approaches, good well positions may be missed by the algorithm if the well controls are poorly configured. It is worth noting, however, that for Case 1A, the decoupled approach had the poorest performance of any of the methods, in terms of the best solution found.

Applying standard PSO to the problem gave the worst results of any of the methods; it had by far the lowest reliability ratings and average NPV for Cases 1A and 1C. For Case 1B, its performance was comparable to that of the hybrid approaches, but worse than the decoupled approach. This would seem to indicate that a purely stochastic approach is insufficient in addressing the combined well placement and well control problem. The "After GPS" column of Table 4 indicates that the solutions found by PSO were usually not even locally optimal and could be improved significantly (from between 2 to 5%, on average) by applying GPS subsequently. The solutions found by the hybrid algorithms, on the other hand, tended to be close to locally op-

Table 4 Results of first and second experiments. The best results for each experiment are highlighted in bold font. Columns headed "Field", "Constr." and "Discount." are included to differentiate the experiments with respect to the field used, presence or absence of maximum constraints on flow, and discounting rate used to calculate NPV. The constraints for Cases 1C and 2B are described in Tables 2 and 3, respectively.

Case	Field	Constr.	Discount.	Algorithm	NPV Avg. ($\$\times10^8$)	NPV Best ($\$\times10^8$)	NPV Worst ($\$\times10^8$)	Reliability 10%	Reliability 5%	After GPS Avg. ($\$\times10^8$)	After GPS $\Delta\%$
1A	SPE10	no	10%	PSO	5.93	6.38	5.55	0.55	0.20	6.19	4.45
				`hybrid-1`	**6.17**	**6.51**	5.77	0.90	**0.50**	6.18	0.07
				`hybrid-5`	6.15	6.44	5.76	0.85	**0.50**	6.18	0.56
				`hybrid-5S`	6.12	6.46	5.79	0.85	0.30	6.15	0.40
				`decoupled`	6.15	6.30	**5.97**	**1.00**	0.20	—	—
1B	SPE10	no	0%	PSO	8.18	8.57	7.50	0.90	0.50	8.35	2.22
				`hybrid-1`	8.23	8.63	7.49	0.80	0.60	8.25	0.21
				`hybrid-5`	8.15	8.59	7.69	0.90	0.45	8.20	0.56
				`hybrid-5S`	8.23	8.59	7.66	0.95	0.60	8.25	0.27
				`decoupled`	**8.35**	**8.64**	**8.04**	**1.00**	**0.65**	—	—
1C	SPE10	yes	10%	PSO	5.58	6.06	5.05	0.50	0.15	5.87	5.41
				`hybrid-1`	5.74	6.02	5.35	0.85	0.25	5.75	0.11
				`hybrid-5`	5.73	6.00	5.51	0.95	0.25	5.76	0.45
				`hybrid-5S`	5.78	**6.15**	5.44	0.80	0.35	5.88	1.70
				`decoupled`	5.89	6.05	5.63	**1.00**	0.50	—	—
				`decoupled-M`	**5.99**	6.14	**5.80**	**1.00**	**0.80**	—	—
2A	Norne	no	10%	`hybrid-5S`	112	**117**	109	**1.00**	0.40	113	0.22
				`decoupled`	**113**	**117**	**110**	**1.00**	**0.80**	—	—
2B	Norne	yes	10%	`hybrid-5S`	97.2	98.4	95.9	0.00	0.00	97.6	0.35
				`decoupled-M`	**106**	**112**	**102**	**1.00**	**0.60**	—	—

tima,l and were typically only improved slightly by this subsequent application of GPS.

Case 1C was the only test case of Experiment 1 to feature constraints on the injection and production rates. In this case we found that the `decoupled-M` approach outperformed the `decoupled` approach in every measure of performance. This would seem to validate our hypothesis that for constrained problems, choosing BHP values that are slightly below or above the maximum and minimum values when placing wells is preferable to using the maximum and minimum values during the well placement phase. The convergence plot (Fig. 4, rightmost plot) shows that the NPVs of the solutions found by `decoupled-M` are much lower than those found by `decoupled` during the well placement phase. During the control optimization phase, however, the GPS algorithm was able to significantly improve the solutions, to the point that they surpass those found by `decoupled`. Examining the best solutions found by either approach for this test case indicated that the placement of wells was similar, but that the `decoupled-M` approach was able to place wells in regions of higher permeability, which improved production by reaching the maximum flow rate more quickly. This placement of wells was not found by the `decoupled` approach since it caused a constraint violation when BHPs were held fixed at the extreme values.

In terms of the solutions found by each method, the three variants of the hybrid approach (`hybrid-1`, `hybrid-5` and `hybrid-5S`) were fairly comparable. Across the three test cases, none of the three were markedly more reliable or provided better NPVs, on average. The convergence plots shown in Fig. 4 do indicate that generally speaking, `hybrid-1` required more function evaluations to arrive at a comparable solution to the other two approaches. This indicates that the increased polling frequency of this method increased the computational cost of the method without significantly improving its performance. Between `hybrid-5` and `hybrid-5S`, we observed the greatest difference in Case 1C, where the `hybrid-5S` method tended to converge more rapidly. This method also had somewhat better performance for Case 1B, while the performance of these two approaches for Case 1A was essentially the same. We conclude that allowing the pattern search to raise or lower BHPs over multiple years at a time did accelerate the convergence of the algorithm to some degree.

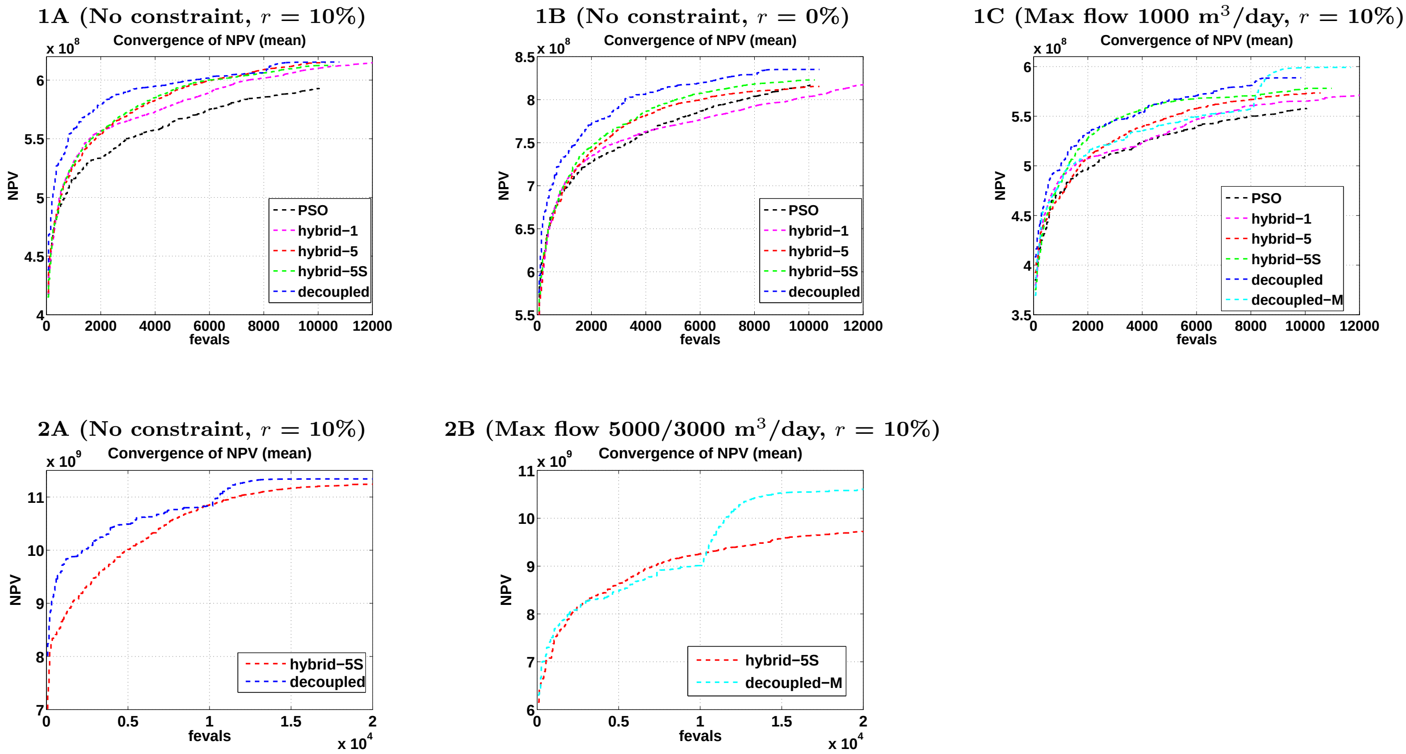

Fig. 4 Plots showing the convergence of different optimization approaches for the three problems of Experiment 1 and the two problems considered in Experiment 2. The best NPV found is shown as a function of the number of reservoir simulations (fevals), averaged over all runs for each approach. Note that the scale of the y-axis differs between plots.

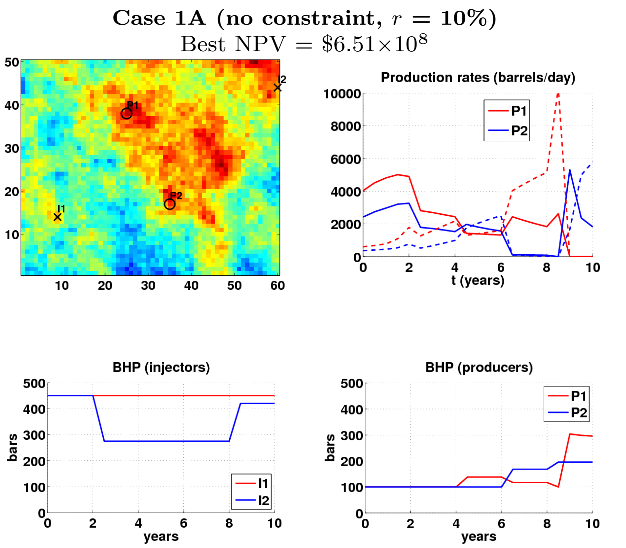

Fig. 5 Best solution found for Case 1A. The top left plot shows positions of the four wells overlaid on the log-perm field, with ∘ denoting a producer, and × denoting an injector. The two plots on the bottom show the control parameters (BHPs) for injectors and producers. The production curves for the two producers are shown in the top right plot, with solid lines indicating oil production and dashed lines indicating water production.

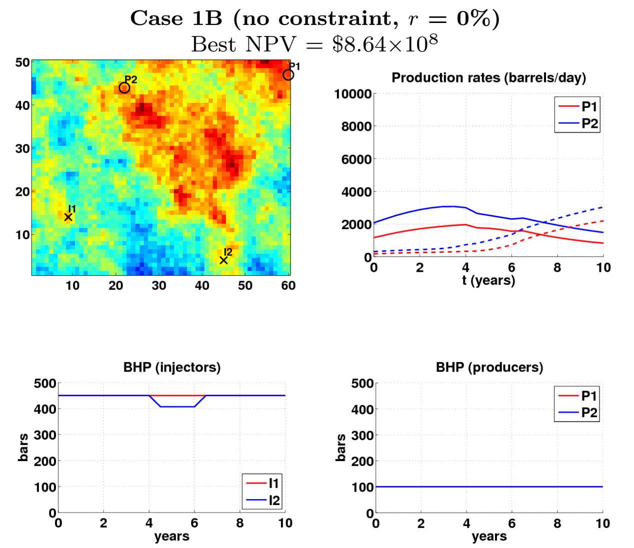

Fig. 6 Best solution found for Case 1B. Symbols used are the same as for Fig. 5.

The best solutions found overall for each test case are shown in Figures 5 to 7. Qualitatively, we can see that while the best solutions found for Cases 1B and 1C are fairly similar to one another, the best solution found for Case 1A is quite different. In Case 1A, since there is cash discounting of 10% and no limit on flow rate,

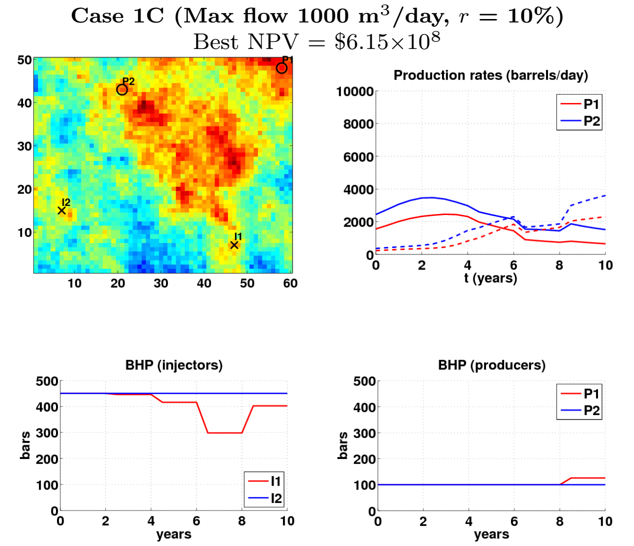

Fig. 7 Best solution found for Case 1C. Symbols used are the same as for Fig. 5.

there is a heavy incentive to produce large amounts of oil quickly. Furthermore, increased water production in later years is not strongly penalized. We therefore obtain an optimal solution in which wells are placed more closely together than in the other two cases, and where water production rates are quite high at later production times. We note as well that the BHP in one of the injectors is lowered to the minimum level after 2 years and held there for much of the production cycle. This occurs because the water front from this injector arrives at the two producers quickly, and so the BHP is subsequently lowered to allow the front from the other injector to arrive before the producers are flooded. One producer is eventually shut in after only 8 years, once the water cut reaches the 78% threshhold.

Since Case 1B does not include any cash discounting, the objective is now to maximize the total amount of oil produced over the time period of 10 years, while minimizing the amount of water produced. As a result, the injectors are placed farther away from the producers, and the production rates are slower and steadier than in Case 1A. The optimal controls in this case are essentially to hold the injector and producer BHPs at their maximum and minimum values, respectively.

For Case 1C, the positioning of the wells is similar, although the first injector is placed in a region of somewhat higher permeability. This allows for a slightly higher overall initial production rate, which is important since Case 1C includes discounting. Since there is an upper limit of 1000 m^3/day on the flow rate at each well, the BHP at the first injector must be eventually lowered in order to maintain a flow rate below the maximum. The flow constraint is also why the wells are placed farther apart than they are in Case 1A. Although both Cases 1A and 1C have the same discounting rate, the flow constraint means that we cannot produce as much oil early on in Case 1C compared to Case 1A. Therefore it is advantageous to place the wells farther apart to delay water production.

The differences between these solutions give some insight into the relative performances of our optimization approaches. For Case 1A, we found that although it was more reliable, the decoupled approach was not able to find solutions that were as good as the best ones found by the hybrid and PSO algorithms. This may be attributable to the fact that the optimal solutions for this test case typically involved placing the wells close together and in regions of high permeability, which required varying the BHPs of injectors and producers in order to prevent premature water flooding. Thus, the positions found during the placement phase of the decoupled algorithm, which assumes that wells are held at the extreme BHP values during the entire production period, were not the best positions for this particular problem. In Cases 1B and 1C, where the optimal solutions did not require varying the control parameters to such a great degree, the performance of the decoupled approach was better.

Experiment 2

The results presented in Table 4 indicate that for the unconstrained problem (Case 2A), the performance of the `hybrid-5S` and `decoupled` algorithms was comparable in terms of the solutions to which they converged. The `decoupled` algorithm did score slightly better in terms of reliability, however, and the convergence plot for this test case (Fig. 4) indicates that its convergence was quicker as well. The overall best solution found is shown in Fig. 8. We see again that when a discounting rate of 10% is imposed and there are no constraints on production, the best solution favours producing large amounts of oil early on, at the cost of increased water production in later years.

The results for Case 2B, which featured a maximum flow constraint of 5000 m^3/day on the four injectors and 3000 m^3/day on the seven producers, clearly indicate that the `decoupled-M` approach used for this test case outperformed the `hybrid-5S` algorithm. The convergence plot for this test case indicates that during the well placement phase of the `decoupled-M` algorithm, its performance lags slightly behind that of `hybrid-5S`, because the BHP values for injectors and producers are held fixed at 450 and 200 bars, respectively. The `hybrid-5S` algorithm is able to explore the full range of BHP values. During the control optimization phase, however, the `decoupled-M` algorithm is able to explore the full range of BHP values as well, and finds significantly better solutions than those found by `hybrid-5S`.

The overall best solution found for Case 2B is shown in Fig. 9. In this case, the incentive to produce more oil early on is counterbalanced by the fact that one can only produce a certain amount of oil per day, due to the maximum flow contraints. Thus we see that in the best solution for Case 2B, the producers tend to be placed farther away from injectors than in the best solution for Case 2A. It is also apparent that the optimal control scheme in this case involves holding injectors fairly close to the maximum BHP value, and producers close to the minimum BHP value. This indicates why the `decoupled-M` approach was more successful than `hybrid-5S` for this test case; since the optimal control scheme is fairly simple, maximizing production for this test case is driven primarily by finding good well locations. Thus, the `decoupled-M` approach benefits by reducing the size of the problem initially (from 77 to 22 variables) and focusing on well placement. For Case 2A, we see that the optimal solution involves raising the BHP at several producers to values more towards the middle and even upper values of the permitted range. Thus, well control plays more of an impact in this test case, and the benefit that the decoupled approach gains by focusing on well placement is offset somewhat by the fact that the best solutions may not have simple control schemes.

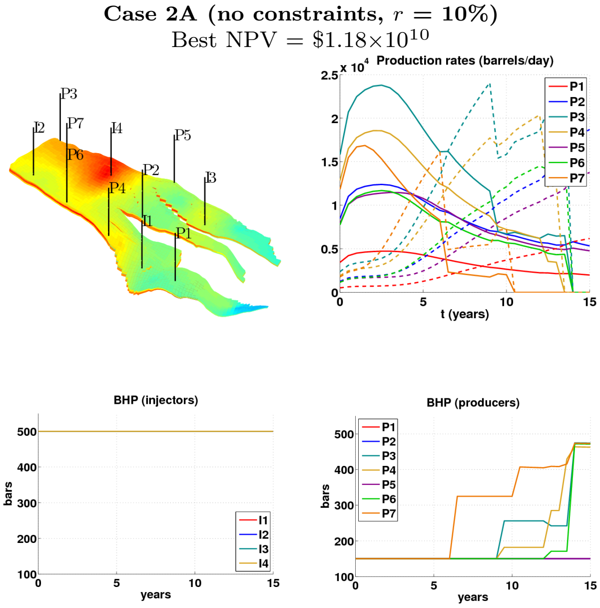

Fig. 8 Best solution found for Case 2A. Plot meanings are the same as for Fig. 5. Note that several lines overlap in the plots of the BHP values.

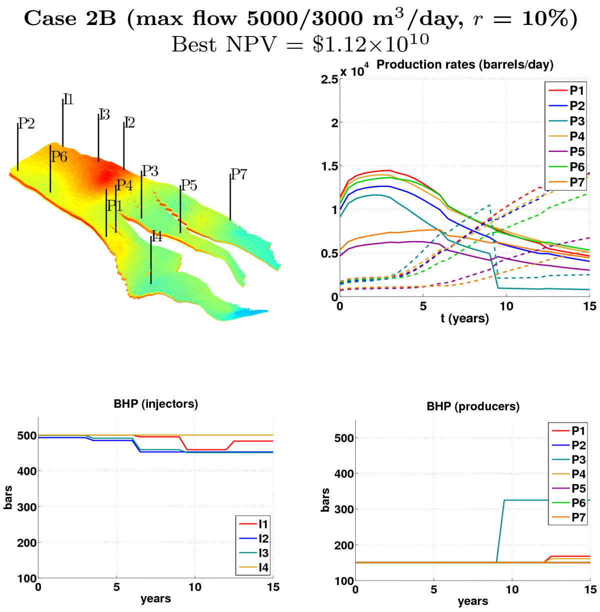

Fig. 9 Best solution found for Case 2B. Plot meanings are the same as for Fig. 5. Note that several lines overlap in the plots of the BHP values.

7 Conclusions

We have examined several approaches to simultaneous optimization of well placement and control, which combine particle swarm optimization (PSO) with pattern search (GPS). We focused on two general approaches: a hybrid algorithm combining PSO and GPS (based on the previously proposed PSwarm algorithm [33]), which we applied to all variables simultanously, and a decoupled method where PSO was applied initially to a well placement problem (assuming a fixed control scheme), and GPS was applied to the controls afterwards. These approaches were applied to a total of five test cases, with some variants of the different approaches being tested as well.

Overall, we find that there may be benefits to decoupling the well placement and control aspects of the problem. In three out of five experiments, the decoupled algorithm found better solutions, on average, than any of the approaches that attempted to optimize over all variables simultaneously. In one case (denoted Case 2B), every one of the five solutions found by the decoupled algorithm was better than those found by the hybrid algorithm. We hypothesize that this is due to the fact that during the well placement phase, the fixed control scheme assumed by the decoupled approach is one that is conducive to finding good solutions. Thus, by reducing the size of the problem and focusing on well placement, the size of the solution space is reduced, and a more thorough exploration of that space is possible. Although the solution space of the problem when all variables are considered simultaneously contains any solution that our decoupled approach could possibly find, it is harder to find those solutions in the larger space. This finding is similar to results published in some papers on optimal placement of large numbers of wells, where it was possible to find better solutions by initially placing wells according to some pattern, rather than allowing all well positions to vary freely [28, 29].

One caveat is that the decoupled approach is sensitive to the control scheme that is assumed during the initial well placement phase. In our experiments, we found in one case (1A) that best solution found by the decoupled approach was not as good as those found by any of the simultaneous approaches; the optimal solutions found for this case tended to require raising and lowering BHPs significantly in order to avoid premature water flooding. In another test case (1C), we compared two variants of the decoupled approach and found that the performance of the algorithm could be improved by holding BHP values closer to the middle of the permitted range, rather than at the extremes. Thus, if one employs a decoupled approach, some thought should be given to the assumed control scheme during the well placement phase. It is also likely that the effectiveness of the decoupled approach will suffer for cases where the control scheme is expected to require varying the control parameters significantly.

There are many avenues for further exploration of the joint placement and control problem. These include modeling more complicated well types such as horizontal, deviated or multilateral wells; incorporating other decision variables in addition to well location and control, such as the number and type of wells to drill, and scheduling of drilling operations; and finally, modeling of geological uncertainty. Taking this considerations into account will also likely require investigating new optimization techniques to account for the increasing complexity of the problem.

Acknowledgments

The authors acknowledge funding from the Natural Sciences and Engineering Research Council of Canada (NSERC), the Atlantic Canada Opportunities Agency (ACOA) and Research & Development Corporation of Newfoundland and Labrador (RDC). The authors thank the developers of MRST at SINTEF Applied Mathematics for making their software freely available and for providing support. The authors also thank StatoilHydro (operator of the Norne field) and its license partners ENI and Petoro for the release of the Norne data used in Experiment 2. Further, the authors acknowledge the Center for Integrated Operations at NTNU for cooperation and coordination of the Norne cases. The views expressed in this paper are the views of the authors and do not necessarily reflect the views of StatoilHydro and the Norne license partners. Finally, we thank Obi Isebor (Stanford University) for several productive conversations on this topic.

References

1. Artus, V., Durlofsky, L., Onwunalu, J., Aziz, K.: Optimization of nonconventional wells under uncertainty using statistical proxies. Comput. Geosci. **10**(4), 389–404 (2006)
2. Audet, C., Dennis, J.: A pattern search filter method for nonlinear programming without derivatives. SIAM J. Optim **14**(4), 980–1010 (2004)
3. Bangerth, W., Klie, H., Wheeler, M., Stoffa, P., Sen, M.: On optimization algorithms for the reservoir oil well placement problem. Comput. Geosci. **10**, 303–319 (2006)
4. Bellout, M., Echeverría Ciaurri, D., Durlofsky, L., Foss, B., Kleppe, J.: Joint optimization of oil well placement and controls. Comput. Geosci. **16**(4), 1061–1079 (2012)
5. Bittencourt, A., Horne, R.: Reservoir development and design optimization. In: Paper SPE 38895 presented at the SPE Annual Technical Conference and Exhibition, San Antonio, Texas, 5-8 October (1997)
6. Bouzarkouna, Z., Ding, D., Auger, A.: Using evolution strategy with meta-models for well placement optimization. In: 12th European Conference on Mathematics of Oil Recovery (ECMOR XII), Oxford, UK, 6-9 September (2010)
7. Brouwer, D., Jansen, J.: Dynamic optimization of water flooding with smart wells using optimal control theory. In: Paper SPE 78278 presented at the SPE 13th European Petroleum Conference, Aberdeen, Scotland, UK, 29-31 October (2002)
8. Bukhamsin, A., Farshi, M., Aziz, K.: Optimization of multilateral well design and location in a real field using a continuous genetic algorithm. In: Paper SPE 136944 presented at the SPE/DGS Annual Technical Symposium, Al-Khobar, Saudi Arabia, 4-7 April (2010)
9. Clerc, M.: Particle Swarm Optimization. iSTE, London (2006)
10. Echeverría Ciaurri, D., Isebor, O., Durlofsky, L.: Application of derivative-free methodologies to generally constrained oil production optimization problems. Procedia Computer Science **1**(1), 1301 – 1310 (2010)
11. Echeverría Ciaurri, D., Mukerji, T., Durlofsky, L.: Derivative-free optimization for oil field operations. In: X.S. Yang, S. Koziel (eds.) Computational Optimization and Applications in Engineering and Industry, *Studies in Computational Intelligence*, vol. 359, pp. 19–55. Springer Berlin / Heidelberg (2011)
12. Emerick, A., Silva, E., Messer, B., Almeida, L., Szwarcman, D., Pacheco, M., Vellasco, M.: Well placement optimization using a genetic algorithm with nonlinear constraints. In: Paper SPE 118808 presented at the SPE Reservoir Simulation Symposium, The Woodlands, Texas, 2-4 February (2009)

13. van Essen, G., Van den Hof, P., Jansen, J.: Hierarchical long-term and short-term production optimization. SPE Journal **16**(1), 191–199 (2011)
14. Forouzanfar, F., Li, G., Reynolds, A.: A two-stage well placement optimization method based on adjoint gradient. In: Paper SPE 135304 presented at the SPE Annual Technical Conference and Exhibition, Florence, Italy, 20–22 September (2010)
15. Güyagüler, B., Horne, R.: Uncertainty assessment of well-placement optimization. SPE Reservoir Evaluation and Engineering **7**(1), 24–32 (2004)
16. Hu, X., Eberhart, R.: Solving constrained nonlinear optimization problems with particle swarm optimization. In: Proceedings of the Sixth World Multiconference on Systemics, Cybernetics and Informatics (2002)
17. Humphries, T., Haynes, R., James, L.: Simultaneous optimization of well placement and control using a hybrid global-local strategy. In: Proceedings of the 13th European Conference on the Mathematics of Oil Recovery (ECMOR XIII), Biarritz, France, 10–13 September (2012)
18. Isebor, O., Durlofsky, L., Echeverría Ciaurri, D.: A derivative-free methodology with local and global search for the joint optimization of well location and control. In: Proceedings of the 13th European Conference on the Mathematics of Oil Recovery (ECMOR XIII), Biarritz, France, 10–13 September (2012)
19. Jansen, J., Douma, S., Brouwer, D., Van den Hof, P., Bosgra, O., Heemink, A.: Closed-loop reservoir management. In: Paper SPE 119098 presented at the SPE Reservoir Simulation Symposium, The Woodlands, Texas, USA, 2-4 February (2009)
20. Kennedy, J., Eberhart, R.: Particle swarm optimization. In: Proceedings of the IEEE International Conference on Neural Networks, vol. 4, pp. 1942 –1948 (1995)
21. Lewis, R., Torczon, V.: Pattern search algorithms for bound constrained minimization. SIAM J. Optim **9**(4), 1082–1099 (1999)
22. Li, L., Jafarpour, B.: Generalized field development optimization: Coupled well-placement and control under geologic uncertainty. In: Proceedings of the 13th European Conference on the Mathematics of Oil Recovery (ECMOR XIII), Biarritz, France, 10–13 September (2012)
23. Li, L., Jafarpour, B.: A variable-control well placement optimization for improved reservoir development. Comput. Geosci. **16**(4), 871–889 (2012)
24. Lie, K.A., Krogstad, S., Ligaarden, I., Natvig, J., Nilsen, H., Skaflestad, B.: Open-source MATLAB implementation of consistent discretisations on complex grids. Computational Geosciences pp. 1–26 (2011)
25. Nogueira, P.d.B., Schiozer, D.: An efficient methodology of production strategy optimization based on genetic algorithms. In: Paper SPE 122031 presented at the SPE Latin American and Caribbean Petroleum Engineering Conference, Cartagena, Colombia, 31 May - 3 June (2009)
26. NTNU IO Centre: Norne benchmark case. `http://www.ipt.ntnu.no/~norne/` (2012)
27. Onwunalu, J., Durlofsky, L.: Application of a particle swarm optimization algorithm for determining optimum well location and type. Comput. Geosci. **14**, 183–198 (2010)
28. Onwunalu, J., Durlofsky, L.: A new well-pattern-optimization procedure for large-scale field development. SPE Journal **16**(3), 594–607 (2011)
29. Ozdogan, U., Sahni, A., Yeten, B., Guyaguler, B., Chen, W.: Efficient assessment and optimization of a deepwater asset development using fixed pattern approach. In: Paper SPE 95792 presented at the SPE Annual Technical Conference and Exhibition, Dallas, Texas, 9-12 October (2005)
30. Peaceman, D.: Interpretation of well-block pressures in numerical reservoir simulation. SPE Journal **18**(3), 183–194 (1978)
31. Sarma, P., Chen, W., Durlofsky, L., Aziz, K.: Production optimization with adjoint models under nonlinear control-state path inequality constraints. In: Paper SPE 99959 presented at the SPE Intelligent Energy Conference and Exhibition, Amsterdam, The Netherlands, 11-13 April (2006)
32. SINTEF Applied Mathematics: Matlab reservoir simulator toolbox v. 2012a. `http://www.sintef.no/Projectweb/MRST/` (2012)
33. Vaz, A., Vicente, L.: Pswarm: a hybrid solver for linearly constrained global derivative-free optimization. Optimization Methods and Software **24**(4-5), 669–685 (2009)
34. Vaz, A., Vicente, L.N.: A particle swarm pattern search method for bound constrained global optimization. J. Glob. Optim. **39**(2), 197–219 (2007)
35. Wang, C., Li, G., Reynolds, A.: Production optimization in closed-loop reservoir management. SPE Journal **14**(3), 506–523 (2009)
36. Yang, D., Zhang, Q., Gu, Y.: Integrated optimization and control of the production-injection operation systems for hydrocarbon reservoirs. Journal of Petroleum Science and Engineering **37**, 69–81 (2003)
37. Yeten, B., Durlofsky, L., Aziz, K.: Optimization of nonconventional well type, location and trajectory. SPE Journal **8**(3), 200–210 (2003)
38. Zandvliet, M., Bosgra, O., Jansen, J., Van den Hof, P., Kraaijevanger, J.: Bang-bang control and singular arcs in reservoir flooding. Journal of Petroleum Science and Engineering **58**, 186–200 (2007)
39. Zandvliet, M., Handels, M., van Essen, G., Brouwer, D., Jansen, J.: Adjoint-based well-placement optimization under production constraints. SPE Journal **13**(4), 392–399 (2008)