Efficient adjoint-based well-placement optimization using flow diagnostics proxies

Stein Krogstad¹ · Halvor Møll Nilsen¹

Received: 14 January 2021 / Accepted: 14 October 2021 / Published online: 8 November 2021
© The Author(s) 2021

Abstract

Model-based optimization of placement and trajectories of wells in petroleum reservoirs by the means of reservoir simulation forecasts is computationally demanding due to the high number of simulations typically required to achieve a local optimum. In this work, we develop an efficient flow-diagnostics proxy for net-present-value (NPV) with adjoint capabilities for efficient computation of well control gradients and approximate sensitivities with respect to placement/trajectory parameters. The suggested flow-diagnostic proxy consists of numerically solving a single pressure equation for the given scenario and the solution of a few inter-well time-of-flight and steady-state tracer equations, typically achieved in a few seconds for a reservoir model of medium size. Although the proxy may not be a particularly good approximation for the full reservoir simulation response, we find that for the cases considered, the correlation is very good and hence the proxy is suitable for use in an optimization loop. The adjoint simulation for the proxy model which provides control gradients and placement sensitivities is of similar computational complexity as the forward proxy model (a few seconds). We employ a version of the *generalized reduced gradient* for handling individual well constraints (e.g., bottom-hole-pressures and rates). As a result, the individual well constraints are enforced within the flow-diagnostics computations, and hence every parameter update becomes feasible without sacrificing gradient information. We present two numerical experiments illustrating the efficiency and performance of the approach for well placement problems involving trajectories and simulation models of realistic complexity. The suggested placements are evaluated using full simulations. We conclude by discussing limitations and possible enhancements of the methodology.

Keywords Flow diagnostics · Well placement optimization · Reservoir simulation · MRST

1 Introduction

Optimizing the placement and trajectories of petroleum reservoir wells is a computationally demanding and time-consuming task due to the high number of simulations typically required to achieve a local optimum. A second challenge is the discrete nature of typical reservoir simulation well models, which poses a challenge in employing efficient gradient-based optimization methods. For this reason, the majority of research in model-based

optimization of well placement has used gradient-free methods (see e.g., [5] and references therein). Exceptions do exist, however. Within an adjoint framework, these include, but are not limited to, approximating well placements with smooth Gaussians [14] and the use of *pseudo*-wells [1, 23]. More recently, [21] suggested a more flexible approach in which certain partial derivatives appearing in the adjoint equations are approximated by numerical perturbations. We apply the same approach herein.

Flow diagnostics provide a way to quantify and visually inspect displacement processes in a fraction of the time it would take to carry out a full physics simulation. The prerequisite is a set of representative fluxes, from which we can obtain travel times and passive tracer concentrations. Derived quantities include connection strength among injectors and producers, measures of dynamic heterogeneity (i.e., the distribution of flow path lengths) and simple proxies for recovery and sweep efficiency. Such flow diagnostics quantities have traditionally been associated

✉ Stein Krogstad
Stein.Krogstad@sintef.no

Halvor Møll Nilsen
HalvorMoll.Nilsen@sintef.no

¹ SINTEF Digital, Mathematics and Cybernetics, P.O. Box 124 Blindern, N-0314, Oslo, Norway

with streamline simulation [3, 16, 20], but it has recently been shown that they can be computed equally well using standard finite-volume methods [12, 15].

In this paper, we develop an adjoint model for a set of flow diagnostics equations, and demonstrate the use of a simple proxy to efficiently optimize well controls and positions employing the technique of [21]. The diagnostics proxy is designed to correlate with net-present-value (NPV), and can accordingly be used to quickly optimize a field's well configuration.

2 Flow diagnostics proxies for field performance

We briefly introduce the concept of *flow diagnostics*. For a more detailed introduction, the reader can resort to the textbook [11, Chapter 13] or one of the original articles that introduced these methods in a finite-volume setting [12, 15].

2.1 Time-of-flight and influence regions

A flow diagnostics measure or proxy, as used here, consists of: 1) solving a pressure equation to obtain a flux field; 2) using this flux field to compute time-of-flight and influence regions (numerical tracers); and 3) based on these, compute a scalar measure (objective). If the proxy is to correlate with or approximate, e.g., the long-term displacement performance of a field, the flux field should represent a long-term average (omitting transient effects). Accordingly, we neglect gravitation and capillary effects. In addition, we evaluate phase densities and viscosities at a fixed reference pressure p_{ref} and all relative permeability functions for a fixed saturation field s_0 . In effect, these choices result in an incompressible system, and further assuming immiscible flow, the pressure equation becomes

$$\nabla \cdot \left(\sum_{\alpha} \mathbf{v}_{\alpha} \right) = \sum_{\alpha} q_{\alpha}, \quad \mathbf{v}_{\alpha} = -\lambda_{\alpha} K \nabla p, \quad (1)$$

where p is the reservoir pressure, $\mathbf{v}_{\alpha}$ is the Darcy velocity of phase α and q_{α} is the volumetric phase source term. The phase mobility is given by $\lambda_{\alpha} = k_{r\alpha}(s_0)/\mu_{\alpha}(p_{\text{ref}})$, where $k_{r\alpha}(s_0)$ is the phase relative permeability at saturation s_0 and $\mu_{\alpha}(p_{\text{ref}})$ is the phase viscosity at reference pressure. Finally, K is the (absolute) permeability tensor. If the boundary conditions are given by well rates at surface conditions q_{α}^s , we use the relation $q_{\alpha} = q_{\alpha}^s/b_{\alpha}(p_{\text{ref}})$, where b_{α} is the inverse formation volume factor of phase α . The resulting pressure equation hence becomes linear.

Once the pressure equation is solved, we obtain a total flux field $\mathbf{v} = \sum_{\alpha} \mathbf{v}_{\alpha}$, for which we state the time-of-flight (TOF) and steady state tracer equations (see e.g., [11]):

$$\begin{aligned} \mathbf{v} \cdot \nabla \tau &= c_{\phi} \phi, & \tau|_{\text{inflow}} &= 0, \\ \mathbf{v} \cdot \nabla C &= 0, & C|_{\text{inflow}} &= 1. \end{aligned} \quad (2)$$

In the first equation of (2), we solve for the time-of-flight variable τ , i.e., $\tau(\mathbf{x})$ is the time it takes for a passive particle to travel from the inflow boundary (e.g., an injector) to location $\mathbf{x}$. Moreover, ϕ is the porosity and c_{ϕ} is a model-dependent factor that we have included to represent the tractability of the pore volume to flow. For single phase flow, $c_{\phi} = 1$ everywhere, whereas for two-phase flow, as considered here, the choice of c_{ϕ} should be a function of initial saturation and relative permeability relations. This is similar to the retardation factors used in reactive transport [2] and recently proposed for naturally fractured reservoirs [19]. Our particular choice of c_{ϕ} is further discussed in the next subsection and in the numerical experiments. In the second equation of (2), we solve for the tracer variable $C(\mathbf{x})$ which in the continuous setting is just $C(\mathbf{x}) = 1$ if $\mathbf{x}$ is reachable from the inflow boundary, and undefined otherwise. In the discrete setting, a grid cell may be partially flooded and hence obtain any tracer value between zero and one. For the numerical tracer computations, we typically subdivide the inflow boundary (or source terms) into multiple parts and solve a separate tracer equation for each of them, e.g., to get influence regions associated with each well or well segment. Analogously to Eq. 2, we may define equations for *backward* time-of-flight and stationary tracer distributions associated with producers simply by reversing the velocity field and setting boundary conditions at the outflow (see e.g., [11]). After discretizing (2), we get a family of linear systems that all contain the same system matrix, and hence multiple equations of the form Eq. 2 can be solved simultaneously with multiple right-hand-sides. Since we use upstream weighting, the linear systems of the backward equations get system matrices that are transposes of those of the forward systems.

Typical heterogeneity measures derived from the forward and backward time-of-flight and tracer equations, include the Lorenz coefficient and (phase-weighted) sweep efficiency. These have shown good correlation with field performance measures such as recovery [22].

2.2 A diagnostics proxy for net-present-value

For the numerical tests performed herein, we utilize an extended version of the net-present-value (NPV) proxy suggested by [12], which is based on discounting reservoir phase volumes by their travel-time to a producer, i.e.,

backward time-of-flight (τ_b). We express the diagnostics NPV (DNPV) at time T as follows:

$$\begin{aligned} \text{DNPV}(T) = & \int_{t=0}^T b_w r_{wi} q_{w,\text{tot}}^+ (1+d)^{-t} dt \\ & + \int_{\Omega} \phi \mathbf{1}_{\tau_b \leq T} (b_w r_{wp} w_w + b_o r_o w_o) (1+d)^{-\tau_b} d\mathbf{x} \\ & + \int_{\Omega} b_w r_{wp} q_w^+ \int_{\min(\tau_b, T)}^T (1+d)^{-t} dt d\mathbf{x} + r_l l. \quad (3) \end{aligned}$$

The *first term* is simply the total discounted cost of water injection, which we calculate analytically since we have constant rates. Here $q_{w,\text{tot}}^+$ is the total rate (at reservoir conditions) of injected water (where $q^+ \equiv \max(q, 0)$), b_w is the inverse water formation factor, which in this setting is constant, r_{wi} is the expected cost for injecting a unit volume of water, and d is the (time-unit) discount rate.

The *second term* attempts to estimate the revenue of the produced oil volumes and the cost of produced in situ water. This can be expressed as an integral over the domain Ω (reservoir), where b_o is the inverse oil formation factor, r_o is the expected (volume-unit) oil production revenue, and r_{wp} is the expected water production cost. The indicator function $\mathbf{1}_{\tau_b \leq T}$ equals one for the region where $\tau_b \leq T$ and zero otherwise. When gradients are required (in an optimization setting), it is necessary to approximate $\mathbf{1}_{\tau_b \leq T}$ by a smooth transition from one to zero. For the numerical experiments, this was done by the approximation

$$\mathbf{1}_{\tau_b \leq T}(\mathbf{x}) \approx \frac{T - \tau_b(\mathbf{x})}{2(\beta^2 T^2 + (T - \tau_b(\mathbf{x}))^2)^{\frac{1}{2}}} + \frac{1}{2}, \quad (4)$$

where we used $\beta = 0.05$. This value for β was chosen based on trial and error observing that a too small value could lead to extreme gradient entries due to the steep transition. Approximations based on the exponential function were also considered, but would require some additional updates to the automatic differential library to prevent overflow/rounding error issues. We also note that for cases with relatively high discount, using a smooth cut-off as Eq. 4 has limited impact on the computed gradient and hence on the resulting optimization performance.. Finally, w_w and w_o are weights to reflect the mobile pore-volume fraction of water and oil. For a unit mobility displacement, a natural choice is to set these weights equal to the saturations, but in general, as with c_ϕ in Eq. 2, a good choice will depend on the relative permeability functions and mobility ratio. For the two-phase oil/water cases considered herein, we set

$$w_o = s_o \frac{\bar{s}_w - s_{w,\text{con}}}{1 - s_{w,\text{con}}}, \quad w_w = s_w - s_{w,\text{con}}, \quad (5)$$

where s_o and s_w are initial oil and water saturations, $s_{w,\text{con}}$ is the connate water saturation, and $\bar{s}_w$ is the average water saturation of the corresponding 1D Buckley–Leverett

profile over the 1D flooded section. Since the profile is self similar, this average is a constant and can equally be defined through the relative speed v_f of the displacement front by $\bar{s}_w = 1/v_f + s_{w,\text{con}}$. Accordingly, in a 1D problem with initial connate water saturation, w_o will be the average amount of displaced oil at water breakthrough. We could use other weights, of course, but for the examples considered herein, we found that (5) results in good correlations with two-phase simulations. For the numerical experiments, we also include the weights w_w and w_o in the time-of-flight (2) by setting $c_\phi = w_w + w_o$. For the selection (5), this means that τ will be approximating the travel time of particles following the displacing water front. The resulting mobile volumes are finally discounted according to their travel time to a producer.

The *third term* represents the cost of producing injected water that breaks through in a producer. Here, q_w^+ represents positive water sources over the domain, from which the breakthrough time to a producer is estimated by τ_b . If this value is less than T , we calculate the discounted cost of producing the corresponding water volume from time τ_b to T . In the discrete setting, this term amounts to summing over all grid-cells connected to a water injector for which $\tau_b < T$. Finally, the *fourth term* represents the cost of drilling and completing the wells, where l is simply the total length of well-trajectories considered and r_l is the length unit cost of drilling (assumed constant). The total cost is here actuated at time zero.

We now set $c_\phi = w_w + w_o$. It then follows from conservation of mass in the continuous setting that

$$\int_{\Omega} \mathbf{1}_{\tau_b \leq T} \phi c_\phi d\mathbf{x} = T q_{w,\text{tot}}^+ - \int_{\Omega} q_w^+ (T - \tau_b)^+ d\mathbf{x}. \quad (6)$$

In the discrete finite-volume setting, however, the time-of-flight in each cell is computed as a weighted mean of all upstream cells, which can introduce large errors because of averaging (smearing) for highly heterogeneous cases. As a result, the numerical approximations of both the integrals in Eq. 6 will suffer from a bias towards under-estimation. A consequence of this is that the expression (3) might give a poor approximation of the simulated NPV. However, our claim here is that the diagnostics proxy *correlates* with the full simulation version, and this is sufficient for optimization. When introducing absolute costs in the objective as, e.g., well costs, it nonetheless becomes also important that the proxy is sufficiently close to the simulated objective. For this reason, we introduce scaling factors for the approximation of these integrals in our final example and observe both good correlations and approximations.

We finally note that the observed volume mismatch just discussed can be remedied by computing residence-time-distributions (RTDs) instead as in, e.g., [8], but as observed by [22], this does not necessarily improve correlation with

full simulations enough to justify the extra computational cost. The objective (3) could equally well be approximated using streamlines, but herein we focus on a finite-volume based, simple and fast proxy with readily available adjoint formulations for gradient computations, as discussed next.

3 Obtaining gradients for well controls and well trajectories

We employ an adjoint implementation of the flow-diagnostics methodology just described to obtain gradients for well controls and positions. In general, a flow-diagnostics measure (or proxy) will be a function of n forward and m backward diagnostics variables such that

$$J(d_f^1, \dots, d_f^n, d_b^1, \dots, d_b^m, u) \rightarrow \mathbb{R}, \quad (7)$$

where d_f^i is a basic diagnostic quantity (time-of-flight or tracer fields) computed from Eq. 2, d_b^j is an analogous backward quantity computed with the reverse flow field, and u represents control inputs and/or parameters. Let E_p , E_f^i and E_b^j denote the discretized pressure equation, forward diagnostics equation i , and backward diagnostics equation j , respectively. Hence, to compute objective J , we first need to solve

$$E_p(p, u) = 0, \quad E_f^i(p, d_f^i) = 0, \quad E_b^j(p, d_b^j) = 0, \quad (8)$$

for $i = 1, 2, \dots, n$ and $j = 1, 2, \dots, m$. In the above, we have assumed that controls and parameters u only explicitly affect the pressure equation. Let J_u denote the gradient vector of J with respect to u , i.e., $J_u^T = dJ/du$. Analogously to [12] we may now compute J_u by an adjoint procedure as

$$J_u = \left(\frac{\partial E_p}{\partial u} \right)^T \lambda_p + \left(\frac{\partial J}{\partial u} \right)^T, \quad (9)$$

where λ_p are the Lagrange multipliers resulting from solving the following set of adjoint linear systems

$$\begin{aligned} \left(\frac{\partial E_f}{\partial d_f} \right)^T \lambda_f^i &= - \left(\frac{\partial J}{\partial d_f^i} \right)^T, \\ \left(\frac{\partial E_b}{\partial d_b} \right)^T \lambda_b^j &= - \left(\frac{\partial J}{\partial d_b^j} \right)^T, \end{aligned} \quad (10)$$

for $i = 1, 2, \dots, n$ and $j = 1, 2, \dots, m$, and finally

$$\left(\frac{\partial E_p}{\partial p} \right)^T \lambda_p = - \sum_{i=1}^n \left(\frac{\partial E_f^i}{\partial p} \right)^T \lambda_f^i - \sum_{j=1}^m \left(\frac{\partial E_b^j}{\partial p} \right)^T \lambda_b^j. \quad (11)$$

Note that in Eq. 10 we have deliberately skipped the i, j indices in the partial derivatives of E_f and E_b to point out that these are the same for all $i = 1, \dots, n$ and

$j = 1, \dots, m$, respectively. Accordingly, as in the forward systems, both the forward and backward adjoint diagnostics Eq. 10 may be solved as single linear systems with multiple right-hand-sides.

3.1 Well controls

In a reservoir simulator, each well is typically assigned multiple potential control modes (e.g., bottom-hole-pressure, component rate, etc.) and corresponding control values. A potential control value serves either as a target value if the control is/becomes active, or as a limit otherwise. For example, a producer may be assigned an oil rate target or upper limit and a bottom-hole-pressure (BHP) target or lower limit. The simulator will always select the control mode for which all other potential control modes stay within their limits. If no such mode exist (e.g., well pressure above reservoir pressure), the well is shut in. In essence, this control strategy can be viewed as a local optimization problem: maximize production within the given limits. We can only calculate the gradient (e.g., by adjoints) with respect to a particular control mode while that control is active. However, prior to a simulation, we can generally not know what control modes a well will operate under. One approach to deal with this situation is to extend the control vector to the full set of potential well controls as suggested by [7] and used for ensemble optimization in [9]. This methodology is based on the Generalized Reduced Gradient (GRG) method as described in e.g., [10]. In this way, a well that switches from, e.g., bottom-hole-pressure (BHP) control to rate control will contribute to both the corresponding entries in the gradient vector and ultimately lead to updates for both entries in the schedule (control policy). A possible concern with this approach is that some potential controls never become active, and hence result in a zero entry in the gradient vector. For a Newton-method this would be detrimental since it would result in a singular Hessian. For a quasi-Newton method, however, the inverse Hessian is typically approximated by low rank updates that ensure non-singularity.

In the approach considered herein, we solve just a single pressure Eq. 1 that attempts to represent the average flow over the full time horizon considered. Accordingly, controlling or limiting individual component rates does not make sense, since these are very dynamic in time (and our pressure equation is in essence single phase). For this reason, our potential well controls are total rate (at reservoir conditions) and bottom-hole pressure. Although the BHP-values from Eq. 1 are not necessarily directly transferable to a full reservoir simulation, it is important to impose BHP-limits to, e.g., avoid placing wells in infeasible low-permeability regions. In this way, our control vector will have one rate and one BHP entry per well. Since each well

only can have a single control mode, this means that for a single model scenario the corresponding gradient vector will have zeros corresponding to nonactive controls. In Section 5.1 we consider optimization of mean NPV over an ensemble of models. In such a setting, different models may result in different control modes and hence result in a *fuller* gradient vector.

To implement the gradient Eq. 9 and the adjoint Eqs. 10–11, we used the automatic differentiation capabilities of MRST [11] to obtain all the Jacobians/partial derivatives of Eqs. 10–11. Some extra effort was nonetheless required to assemble the adjoint equations to enable an efficient sequential solution strategy and allow for multiple right-hand sides. We validated our implementation of the approach just described by comparing the adjoint-based gradients to approximate gradients obtained by numerical perturbations.

3.2 Trajectory parameterization

For computing the gradient with respect to trajectory parameters, the adjoint Eqs. 10–11 remain the same as for the well control case. For the gradient expression (9), however, partial derivatives with respect to the parameters are required, and this constitutes a challenge for parameters that do not enter the pressure equations explicitly. Well positions, in particular, only appear in a discrete sense in a reservoir simulator, as a list of connected grid cells with corresponding well connection transmissibility factors (well indices). A clever way to overcome this difficulty, whilst still maintaining the efficiency of the adjoint approach, was suggested by [21]. Rather than perturbing positions to obtain approximate gradients, they suggested to perturb positions to obtain approximate partial derivatives $\partial E_p / \partial u$ in Eq. 9. This makes a huge difference, since rather than solving the equations for each perturbed position, one just has to evaluate the equation residuals for each perturbed position. Accordingly, the extra effort required to compute an (approximate) adjoint-based gradient for position parameters compared to well controls, boils down

to a few extra evaluations of the equation residuals. In this work, we follow the approach of [21].

We parametrize a well trajectory by a set of 3D coordinates and obtain a continuous curve by employing a smooth interpolation between the points (see Fig. 1). Here, we use a spline interpolation, but any technique could be applied for this purpose. By using a triangulation of the grid cell-faces, we compute all trajectory-grid intersections and end up with a list of traversed grid cells and corresponding 3D line segments (vectors from entry to exit points). For a traversed grid cell with segment lengths (l_x, l_y, l_z) , we compute the corresponding connection transmissibility factor as

$$WI = \sqrt{\sum_{d=x,y,z} (l_d / L_d)^2 WI_d^2}, \quad (12)$$

where WI_d is the Peaceman well index [13] for a well parallel to the d -axis and L_d is the corresponding grid cell dimension along the d -axis. Since (12) considers projections of the trajectory segment along the coordinate directions, this approach is referred to as the *projection well index* [17]. Figure 1 depicts the interpolation point perturbation approach taken herein. The two end-points are assigned three perturbation directions each, one tangential to the trajectory and two orthogonal. For interior points, however, we skip the tangential direction, since this (at least infinitesimally) will have zero effect. For objectives that include trajectory length, we also need length derivatives with respect to the perturbation directions. We obtain these approximately, by only considering the tangential perturbation directions of the end-points. For curved trajectories, also the length derivative with respect to other perturbations/directions will be non-zero, but these are neglected for simplicity. This simplification is in part supported by the fact that the trajectory curvatures (dogleg) have quite strict upper limits. As in [21], we use a two-sided perturbation, such that the directional partial derivative of the pressure equation with

Fig. 1 Left plot shows well trajectory with intersected cells. Right plot depicts the same trajectory (black) and perturbed trajectories (gray) for one of the four interpolation points

respect to a point $\mathbf{c}$ along direction $\mathbf{v}$ is approximated by

$$\frac{\partial E_p}{\partial \mathbf{c}} \cdot \mathbf{v} \approx \frac{E_p(\mathbf{c} + \epsilon \mathbf{v}) - E_p(\mathbf{c} - \epsilon \mathbf{v})}{2\epsilon}. \quad (13)$$

Here, we have omitted other input to the pressure equation for brevity. As noted by [21], the quality of the obtained gradients depends largely on the perturbation size. For a well centered in a vertical grid column, for instance, a perturbation smaller than half a grid cell will have no effect and produce a zero gradient. As an example, in Fig. 2 we consider perturbation of the top-point of PROD1 along the unit direction $(2, 1, 0)/\sqrt{5}$. The DNPV as a function of perturbation size (sampled at a resolution of 1 meter) results in the black curve (right plot). We clearly observe a flat objective response around zero. The grid-cell dimensions for this case are 8 m in the x - and y -directions and 4 m in the z -direction, and the flat region corresponds to the perturbations that are sufficiently small for the well trajectory to be confined within the original grid cells (absolute perturbation of the top-point less than $2\sqrt{5} \text{ m} \approx 4.47 \text{ m}$). Using the adjoint approach, we next compute directional derivatives along $(2, 1, 0)/\sqrt{5}$ at the point of zero perturbation, hence approximating the slope of the black curve around zero. For the adjoint computations, we consider several perturbation sizes $\epsilon \mathbf{v}$ for the approximation of partial derivatives (13). For instance, the blue curve of Fig. 2 is the slope corresponding to the adjoint-based gradient using $\pm 6 \text{ m}$ for the x - and y -perturbations, and $\pm 3 \text{ m}$ for the z -perturbation. We observe that the gradients resulting from perturbations corresponding to one and two grid-cells match quite nicely with the average slope of the black curve. We note that through this adjoint approach we are attempting to calculate the gradient at a specific point, so

using larger increments in the partial derivative estimation does not necessarily result in any further *averaging* of the gradient. This is illustrated by Fig. 2, where we observe that even for the (24 m, 24 m, 12 m)-perturbation, our gradient approximation does not *see* the downward trend in the objective for perturbations larger than 10 meters. We note that the selection of perturbation size for the approximation of partial derivatives is further examined in [21].

4 Optimization approach and constraint handling

Having developed a parametrized optimization problem, as in the previous sections, we could in theory feed the objective function and its derivative to any gradient-based optimizer. In our initial tests, however, we observed that running the problem in an external optimizer was not very efficient and resulted in a large number of iterations. One reason for this might be that the position gradient approximation does not very well predict the objective change for very small perturbations as just discussed. Accordingly, we resort to a simple gradient-search approach, in which the gradient is projected in every step onto the imposed constraints. We note, however, that [21] successfully utilized a Sequential Quadratic Programming (SQP) solver for the well placement problem in which they also provided derivatives for the placement constraints considered. Here we simply consider a gradient projection approach for the gradient entries corresponding to trajectory parameters (entries corresponding to well controls remain unchanged through this projection step). Consider a vector of parameters and controls u with corresponding gradient vector J_u . We obtain a new

Fig. 2 The left plot depicts the grid and wells for the Egg ensemble problem, where INJECT i and PROD j are the i th injector and j th producer. Producer PROD1 is depicted with several perturbations of its top point in direction $(2, 1, 0)/\sqrt{5}$. Objective values along this direction sampled at a resolution of 1 meter is shown by the black curve

in the right plot. Colored lines (right plot) show directional derivatives obtained from adjoint computations where the legend indicates coordinate perturbation magnitudes in meters used in the approximation of the reservoir equation partial derivatives

candidate vector of controls/parameters $u + d = u + \Pi J_u$ where d is a projection of the gradient vector J_u . The projection Π consists of the following steps (initially setting $d \leftarrow J_u$):

- We clip the magnitude of d according to a maximal allowed update in each coordinate direction. The maximal allowed update is a global setting selected initially.
- For each internal (if any) trajectory interpolation point resulting from the updated parameters $u + d$, we check if a constraint on well trajectory curvature (dogleg) is violated. For each violating point, we shift the point towards the straight line between its neighbors until the curvature/dogleg constraint is fulfilled. This procedure is performed in sequence from a well's heel to toe.
- Let u_c and d_c denote the control/parameter vector and update vector, respectively, restricted to a trajectory interpolation point c . The feasible region for each interpolation point is defined by the boundary of a set of grid cells. Since such a set might not be simply connected (due to inactive cells or layers), we consider all boundary intersections along the parametrized line $u_c + s d_c$. Let $s_{\max}$ represent the largest s -value for which there is an intersection. If $s_{\max} \geq 1$, then $u_c + d_c$ is feasible and we leave d_c unchanged. Otherwise, we locate the closest boundary point to $u_c + d_c$, say c_γ , and update $d_c \leftarrow c_\gamma - u_c$. In this way, interpolation points are allowed to move across separated layers but also on the outside of concave boundaries. The procedure nonetheless ensures that the updated $u_c + d_c$ will always stay within the convex hull of the boundary.
- To prevent the updated well trajectories from pointing upwards, we next enforce monotonicity for the z -values for the interpolation points. This is performed simply by updating a point's z -value (depth) to the maximal of its current value and nearest interpolation point in the direction of the heel.
- Finally, we assure that trajectories do not collapse into a single point. This situation may very well indicate that the well is not profitable, but we ensure that the length remains above a minimal value in order not to lose gradient information. In particular, if $u + d$ results in a trajectory of length less than the minimal value, we simply stretch the well trajectory along its heel-to-toe direction so its heel-to-toe distance matches the minimal length, and update d accordingly.

Each of these bullet points represents on its own a projection of the gradient vector J_u . The combination of all the projections in sequence, however, does not necessarily do so (e.g., enforcing a feasible region may lead to violation on curvature etc.). Here, we resort to a fixed-point procedure and iterate these steps until the update

vector does not change between iterations. It is important that the resulting projected gradient $d = \Pi J_u$ indeed is an increasing direction, i.e., that $J_u^T d > 0$. We do not claim that the procedure just described ensures this, but merely remark that it seems difficult to construct a counterexample.

Since we are using a gradient-search approach, it is important to scale the problem properly. Herein, we scale the problem in such a way that all controls are between zero and one; for well controls, zero and one correspond to lower and upper limits, respectively, whereas for positions, zero and one correspond to the boundaries of the bounding box of the feasible regions. This does not, of course, completely prevent scaling issues, and we do observe a quite slow convergence of the combined well control/position optimization compared to the only well control case for which a BFGS-algorithm is utilized to produce approximate Hessians.

5 Numerical experiments

Next, we present two numerical experiments to validate our new flow-diagnostics methodology.

5.1 The egg ensemble

In this numerical experiment we consider optimization of expected (mean) NPV over the 100 realizations of the Egg ensemble [6]. The ensemble is a synthetic model of a channelized oil reservoir produced under water flooding. The model has eight injectors and four producers and each model realization consists of 18 553 grid cells (see Fig. 2). The ensemble realizations only differ in permeability, i.e., the locations of the high permeable channels. All injectors are injecting water at a rate of 79.5 m³/day with a BHP upper limit of 420 bar. Producers are set to produce at 395 bar (no limits on rates). To calculate net-present-value (NPV) we assume an oil revenue of 50 \$/stb, water injection cost 5 \$/stb, and water production cost 3 \$/stb. The yearly discount rate is set to 10%. With these settings, all the 100 ensemble members reach their maximal NPV before 2.5 years. For our optimizations, we seek to optimize NPV at 3 years of production. Simulation of a single realization from this ensemble takes roughly 90 seconds on a laptop using MRST [18], while the evaluation of the proxy and its adjoint takes about one second.

To select the phase-weights w_w and w_o in the diagnostics proxy (3), we consider the two-phase Buckley–Leverett solution for the given oil/water properties. In particular, we set $w_o = 1/2.7$ (i.e., the inverse of the frontal speed), whereas $w_w = 0$, since the model is initially everywhere at connate water saturation. In this example, we do not use

Fig. 3 The left plot shows simulated (solid line) and diagnostics (dashed line) evolution of NPV for realizations 60, 5, and 4 of the Egg ensemble. The right plot shows cross-correlation of simulation

and diagnostics evaluations of NPV at 3 years for all 100 realizations where the linear trend (dashed line) is obtained by a least-squares fit

additional scaling of the proxy, and hence use $c_\phi = w_o$ in Eq. 2.

We initially compute the proxy NPV and run full simulations for the entire ensemble. Figure 3 (right plot) compares the diagnostic and simulation values. Cross-correlation for the base case is 0.89, which we consider quite good, especially taking into account that the NPV-spread is quite modest (the standard deviation is approximately 3.25% of the mean). In Fig. 3 (left plot), we have plotted the NPV-evolution with time for the worst, median, and best performing realization with respect to the proxy NPV.

We next consider three optimization problems:

1. Optimize all well controls only (keep well placements fixed).
2. Optimize well placements of the producers (keep well controls fixed).

3. Optimize both well placements of producers and all well controls.

We only consider straight line well trajectories, so the trajectory of each producer is parameterized with two coordinates (six parameters). The eight injectors have two potential well controls each (rate and bhp), while the four producers all have a single well control each (bhp). Accordingly, the three optimization problems contain $2 \times 8 + 4 = 20$, $4 \times 6 = 24$, and $20 + 24 = 44$ parameters, respectively. The well placement optimization uses a simple line-search logic along the projected gradient, reducing the step length whenever encountering a non-increasing control vector. The optimization halts after five consecutive step reductions. After each optimization, we fed the resulting optimized well controls/positions into a reservoir simulator for evaluations over the entire ensemble. Figure 4 shows cross-correlation of simulated NPV versus

Fig. 4 Cross-correlation of simulation and diagnostics evaluations of NPV at 3 years after optimization for all 100 realizations of the Egg ensemble. The three cases are well control optimization only (left), well position optimization only (middle), and combined control and

position optimization (right). Linear trends (dashed lines) are obtained by least-squares fits. For comparison, realizations 60, 5 and 4 are highlighted as in the base-case plot (Fig. 3)

the diagnostics NPV after well controls/well placements have been optimized. Compared to the base-case shown in Fig. 3, we observe a deterioration in correlation for the well control only case (left plot) and combined control/placement case (right plot) mainly due to a few outliers. Figure 5 reports the NPV-histograms (simulated NPV) for the base case and the three optimized cases. The improvement in mean NPV over the base case is 4.7%, 3.2%, and 9.5% for the control, position, and combined optimizations, respectively. As expected, the improvement achieved by optimizing both control and position is greater than the sum of the individual improvements of the other two approaches. We also observe that the spread of the ensemble is reduced substantially, from 3.25% in the base case to 1.15% in the combined optimization case.

Figure 6 reports producer positions resulting from the placement only optimization (left) and the combined optimization of position and control (right). When only optimizing placement, the deviations from the original positions are quite modest. For the combined optimization, however, the producer positions have changed drastically. Also, three injectors and producer PROD1 are operating at minimal rates and are hence suggested shut down. The resulting sweep pattern for this case is almost entirely from left to right, as seen in Fig. 6.

5.2 The olympus ensemble

Next, we take a more detailed look at the performance of the suggested methodology for well-trajectory optimization

Fig. 5 Histograms and ensemble mean values (dashed lines) for simulated NPV at 3 years over the 100 realizations of the Egg ensemble. The four plots show the base case (top left), optimized controls

only (top right), optimized positions only (bottom left) and combined optimization of controls and positions (bottom right)

Fig. 6 Well placements for the Egg ensemble resulting from placement optimization only (left) and combined control and placement optimization (right). Transparent vertical lines show location of producers for

base case, while blue lines show injector locations (these were not optimized). Wells with a cross operate at minimal rate, and hence are suggested removed

by optimizing the trajectory of a single producer (PROD-1) of the Olympus benchmark ensemble [4]. This ensemble contains 50 realizations, and each model realization consists of 192 750 grid cells. Permeability, porosity and fluid regions vary across the ensemble realizations. Simulation of a single realization from this ensemble takes about 30 minutes in MRST, while the evaluation of the proxy and its adjoint takes about 10 seconds with the current implementation. Rather than optimizing over the mean (as in the previous experiment), we here treat each realization individually to gain some performance statistics. We study realization 1 in detail, and repeat the experiment for the remaining 49 realizations. Figure 7 shows the first realization. The upper layers of the model show apparent channelized features and high permeability contrasts. The initial oil zone encapsulates the region where the wells are situated, while the remaining part of the model below this cap contains water. Furthermore, the model contains four saturation regions with distinctly different relative permeability functions. Hence, in the selection of the oil/water weights (5) for the proxy, these both depend on initial saturation and saturation regions. Table 1 reports the

resulting approximate values for these weights, as computed from the Buckley–Leverett profiles of the corresponding relative permeability functions. We note the distinct values for the fourth region due to a very narrow band of saturations for which both the water and oil are mobile.

In the original setup of the Olympus ensemble [4], producers are set to operate at a target oil rate of 900 m³/day and injectors at a water rate of 1600 m³/day. The well BHP-limits are set to 150 bar for producers and 235 bar for injectors. Nearly all BHP-limits are active during the first part of the simulations but switch to rate control as the water starts to flood the oil cap. To make the comparison with the proxy more clear, we introduce constant liquid rates for all injectors corresponding to the average simulation rate for the base case. Producers are set to operate at 150 bar. An oil price of 45 \$/bbl, cost of produced water 6 \$/bbl, cost of injected water 2 \$/bbl, and a discount rate of 8% per year, gives the simulated NPV-function depicted in Fig. 9 for ensemble realization number one. Due to the high degree of heterogeneity of this model, the approximation of the integrals in Eq. 3 suffers from smearing/averaging of the time-of-flight solution, which results in a delayed

Fig. 7 Horizontal permeability and wells for realization 1 of the Olympus ensemble. Colors indicating low permeability in blue and high permeability in yellow

Table 1 Oil/water weights used in the NPV proxy for the Olympus benchmark case

	Region 1	Region 2	Region 3	Region 4
w_w	$s_w - 0.14$	$s_w - 0.22$	$s_w - 0.08$	$s_w - 0.57$
w_o	$0.45 s_o$	$0.43 s_o$	$0.46 s_o$	$0.15 s_o$

peak of the NPV-proxy. To achieve a closer match with the NPV-function, we here use a manual tuning of c_ϕ in Eq. 2 and in the lower limit in the integral approximating the production of injected water. In particular, we found that setting $c_\phi = (w_w + w_o)/1.3$ gave a quite good match between the simulation-based and diagnostics-based NPVs, as can be observed in Fig. 9. We note that this tuning was performed only for the first realization, but gave nice results for all of the 50 ensemble members. The left plot of Fig. 8 shows diagnostics versus simulation NPV at 10 years using the base-case setup for all the 50 realizations. The correlation for this plot is 0.96.

We next consider well-trajectory optimization for PROD-1, keeping injection rates fixed and setting all producers at 150 bar. Our objective is NPV at 10 years. The base case trajectory is shown in Fig. 9 (blue trajectory), whereas a vertical cross-section is shown in Fig. 10. Since the model does not provide well trajectories, the shown base-trajectory is based on fitting a quadratic curve to the centers of the connected cells. Hence the base trajectory used here produces slightly different results than the connected cell-list provided by the model. We parametrize the trajectory by using four equally distributed interpolation-points. For the

first ensemble realization, we considered two optimization cases; the first with zero well cost, and the second with cost 10 000 \$/m. The resulting (local) optimal trajectories are shown in Figs. 9 and 10. Both trajectories have approximately the same location, whereas the zero-cost trajectory is significantly longer than the other. Figure 10 also shows that whereas the base trajectory only penetrates the lower layers, both the optimized trajectories cross the inactive (sealing) layer to achieve better drainage. We also note that the S-shape of the resulting trajectories would probably not make a driller happy, although both were confined with a dogleg constraint of $1/3^\circ/\text{m}$.

Finally, we repeat the well-trajectory optimization for the remaining 49 ensemble members, setting the well cost to 10 000 \$/m. Figure 11 reports the NPV improvements resulting from trajectory optimization of well PROD-1. The predicted improvements from the NPV proxy agree quite well with those observed in simulations, although the diagnostic predictions are a bit more optimistic on average. The correlation between the NPVs from diagnostics and simulation remains good also for the optimized placement cases, as can be observed in Fig. 8 (right plot). The correlation coefficient for this case is 0.97.

Fig. 8 NPV proxy versus simulated NPV at 10 years for all 50 realizations of the Olympus ensemble. The left plot shows results for the base case, whereas the right plot is after optimizing the placement of well PROD-1. Dashed lines show linear trends obtained by least-squares fits

Fig. 9 Left figure shows estimated net-present-values for the first realization of the Olympus ensemble. Estimates are based on simulations (solid lines) and diagnostics (dashed lines) for base case, and

optimized placement of PROD-1 for the cases with zero well cost and a well cost of 10000 \$/m. Right figure shows top view of the resulting trajectories within a section of the model grid

Fig. 10 Perforated section of well trajectories for producer PROD-1 from Fig. 9 shown along vertical grid sections colored by horizontal permeability for the first realization of the Olympus ensemble. Low permeability shown in blue, while high permeability is shown in bright yellow

Fig. 11 Relative improvements in proxy NPV/simulated NPV resulting from placement optimization of producer PROD-1 for each of the 50 realizations of the Olympus ensemble. For comparison, also the improvement in maximal (over time horizon) simulated NPV is included

6 Concluding remarks and future work

In this work, we develop and test a methodology for well control and well-trajectory optimization using a proxy based on flow diagnostics with gradients obtained by adjoints. Using full reservoir simulations for model-based optimization of well trajectories is computationally intensive, and hence the number of work-flow iterations one can afford to run on problem setup, optimization, and analysis of results will typically be limited. Rougher, but much faster estimates of well performance for different placements and controls can speed up this loop and enable trying out a much larger number of scenarios before a few are selected for more thorough analysis with full simulations. For the well placement problem in particular, it is advantageous that the proxy operates on the original grid. Hence, model reduction and upscaling techniques are not necessarily suited for this application.

For the two model ensembles considered in this work, we observed good correlations between diagnostics-based and simulation-based NPVs. For the Olympus ensemble, however, good correlation and approximation properties relied on manual tuning/scaling of a few terms in the proxy. In particular, we introduced scaling of the pore volumes in the time-of-flight equation, phase weighting functions in the diagnostics proxy, and scaling of the time-of-flight values for arrival times of the injected water. Herein, we scaled these factors manually by comparing results to the full simulation output of the first realization. The same scaling was used for all realizations, however. Automation of this tuning process is a natural next development step for our new methodology. That is, we should optimize the scaling factors by applying the already implemented adjoint formulation to calibrate the proxy to reservoir simulator results.

One current shortcoming of the proposed approach is that an optimal approach as proposed by a diagnostics proxy will have constant rates or BHPs over the full time horizon. Enabling dynamic controls does not necessarily lead to a significant improvement in objective, but we expect that shutting in wells may, for instance, impair the correlation with the current proxy. Future efforts will be made to further investigate the application of diagnostics proxies to more complex schedules such as reactive strategies. In particular, we will look into different proxy parametrizations to better represent and hopefully approximate such situations.

Acknowledgements The research reported in this paper was funded in part by the Research Council of Norway through grant no. 280950 and in part by Equinor Energy AS, Total E&P Norge AS, and Wintershall DEA Norge AS. We wish to thank the anonymous reviewers for their very valuable comments and suggestions for improving the clarity and readability of the presented material.

Funding Open access funding provided by SINTEF AS.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Castineira, D., Alpak, F.O.: Automatic well placement optimization in a channelized turbidite reservoir using adjoint based sensitivities. In: SPE Reservoir Simulation Symposium. Society of Petroleum Engineers, The Woodlands, Texas (2009). <https://doi.org/10.2118/119156-MS>
2. Cvetkovic, V., Dagan, G.: Transport of kinetically sorbing solute by steady random velocity in heterogeneous porous formations. *J. Fluid Mech.* **265**, 189–215 (1994). <https://doi.org/10.1017/S0022112094000807>
3. Datta-Gupta, A., King, M.J.: Streamline simulation: theory and practice, vol. 11 Society of Petroleum Engineers Richardson (2007)
4. Fonseca, R., Della Rossa, E., Emerick, A., Hanea, R., Jansen, J.: Overview of the olympus field development optimization challenge. In: EAGE/TNO Workshop on OLYMPUS Field Development Optimization, Barcelona, Spain (2018). <https://doi.org/10.3997/2214-4609.201802246>
5. Forouzanfar, F., Reynolds, A.: Well-placement optimization using a derivative-free method. *J. Pet. Sci. Eng.* **109**, 96–116 (2013). <https://doi.org/10.1016/j.petrol.2013.07.009>
6. Jansen, J.D., Fonseca, R.M., Kahrobaei, S., Siraj, M.M., Van Essen, G.M., Van den Hof, P.M.J.: The egg model—a geological ensemble for reservoir simulation. *Geosci. Data J.* **1**(2), 192–195 (2014). <https://doi.org/10.1002/gdj3.21>
7. Kraaijevanger, J.F.B.M., Egberts, P.J.P., Valstar, J.R., Buurman, H.W.: Optimal waterflood design using the adjoint method. In: SPE Reservoir Simulation Symposium. Society of Petroleum Engineers, Houston, Texas, U.S.A (2007). <https://doi.org/10.2118/105764-MS>
8. Krogstad, S., Lie, K.A., Nilsen, H.M., Berg, C.F., Kippe, V.: Efficient flow diagnostics proxies for polymer flooding. *Comput. Geosci.* **21**(5–6), 1203–1218 (2017). <https://doi.org/10.1007/s10096-017-9681-9>
9. Krogstad, S., Nilsen, H., Møyner, O., Rasmussen, A.: Well control optimization of the olympus case using MRST and OPM. In: EAGE/TNO Workshop on OLYMPUS Field Development Optimization, Barcelona, Spain (2018). <https://doi.org/10.3997/2214-4609.201802287>
10. Lasdon, L.S., Waren, A.D., Jain, A., Ratner, M.: Design and testing of a generalized reduced gradient code for nonlinear programming. *ACM Trans. Math. Softw.* **4**(1), 34–50 (1978). <https://doi.org/10.1145/355769.355773>
11. Lie, K.A.: An Introduction to Reservoir Simulation Using MATLAB/GNU Octave: User Guide for the MATLAB Reservoir

- Simulation Toolbox (MRST). Cambridge University Press, Cambridge (2019)
12. Møyner, O., Krogstad, S., Lie, K.A.: The application of flow diagnostics for reservoir management. *SPE J.* **20**(2), 306–323 (2014). <https://doi.org/10.2118/171557-PA>
 13. Peaceman, D.: Interpretation of well-block pressures in numerical reservoir simulation with nonsquare grid blocks and anisotropic permeability. *Soc. Pet. Eng. J.* **23**(03), 531–543 (1983). <https://doi.org/10.2118/10528-PA>
 14. Sarma, P., Chen, W.: Efficient well placement optimization with gradient-based algorithms and adjoint models. *Soc. Pet. Eng. - Intell. Energy Conf. Exhibition Intell. Energy* **2**, 2008 (2008). <https://doi.org/10.2118/112257-MS>
 15. Shahvali, M., Mallison, B., Wei, K., Gross, H.: An alternative to streamlines for flow diagnostics on structured and unstructured grids. *SPE J.* **17**(3), 768–778 (2012). <https://doi.org/10.2118/146446-PA>
 16. Shook, G., Mitchell, K.: A robust measure of heterogeneity for ranking earth models: The F-Phi curve and dynamic Lorenz coefficient. In: *SPE Annual Technical Conference and Exhibition* (2009). <https://doi.org/10.2118/124625-MS>
 17. Shu, J.: Comparison of various techniques for computing well index. Master's thesis, Stanford University (2005)
 18. SINTEF: Mrst 2020a. <http://www.mrst.no>, <https://doi.org/10.5281/zenodo.3891372> (2020)
 19. Spooner, V.E.: Flow diagnostics for naturally fractured reservoirs. Ph.D. thesis, Heriot-Watt University (2020)
 20. Thiele, M., Batycky, R., Fenwick, D.: Streamline simulation for modern reservoir-engineering workflows. *J. Petrol. Tech.* **62**(01), 64–70 (2010). <https://doi.org/10.2118/118608-jpt>
 21. Volkov, O., Bellout, M.C.: Gradient-based constrained well placement optimization. *J. Pet. Sci. Eng.* **171**, 1052–1066 (2018). <https://doi.org/10.1016/j.petrol.2018.08.033>
 22. Watson, F., Krogstad, S., Lie, K.: Flow diagnostics for model ensembles. In: *ECMOR XVII – 17th European Conference on the Mathematics of Oil Recovery*. EAGE (2020). <https://doi.org/10.3997/2214-4609.202035133>
 23. Zandvliet, M., Handels, M., van Essen, G., Brouwer, R., Jansen, J.D.: Adjoint-based well-placement optimization under production constraints. *SPE J.* **13**(04), 392–399 (2008). <https://doi.org/10.2118/105797-PA>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.