

ADJOINT BASED OPTIMAL CONTROL OF PARTIALLY MISCIBLE TWO-PHASE FLOW IN POROUS MEDIA WITH APPLICATIONS TO CO₂ SEQUESTRATION IN UNDERGROUND RESERVOIRS

MORITZ SIMON* AND MICHAEL ULBRICH†

Abstract. With the target of optimizing CO₂ sequestration in underground reservoirs, we investigate constrained optimal control problems with partially miscible two-phase flow in porous media. Our objective is to maximize the amount of trapped CO₂ in an underground reservoir after a fixed period of CO₂ injection, while time-dependent injection rates in multiple wells are used as control parameters. We describe the governing two-phase two-component Darcy flow PDE system, formulate the optimal control problem and derive the continuous adjoint equations. For the discretization we apply a variant of the so-called BOX method, a locally conservative control-volume FE method that we further stabilize by a periodic averaging feature to reduce oscillations. The timestep-wise Lagrange function of the control problem is implemented as a variational form in Sundance, a toolbox for rapid development of parallel FE simulations, which is part of the HPC software Trilinos. We discuss the BOX method and our implementation in Sundance. The MPI parallelized Sundance state and adjoint solvers are linked to the interior point optimization package IPOPT, using limited-memory BFGS updates for approximating second derivatives. Finally, we present and discuss different types of optimal control results.

Key words. optimal control, partially miscible two-phase flow, adjoint approach, complementarity condition, control-volume FE method, CO₂ sequestration

AMS subject classifications. primary 35Q35, 90C90; secondary 65M08.

1. Introduction. This paper studies an adjoint based optimization approach [24] for optimal control problems governed by multiphase multicomponent flow in porous media. Our work is motivated by optimal control problems in the context of subsurface CO₂ sequestration, where we consider the problem of maximizing the amount of trapped CO₂ after a fixed interval of injection into a 2D reservoir Ω . To this end, we apply a partially miscible two-phase two-component flow model [7]. We introduce the flow PDE model, derive its adjoint equation, discuss a suitable discretization of the state and adjoint systems, then present some numerical results. The discretization (currently in 2D) is achieved with the aid of the PDE toolbox Sundance [32], a package of the Trilinos framework [23] that we extended by finite volume (FV) and upwinding capabilities. Furthermore, we developed an interface that links the state and adjoint solvers to the interior point optimization package IPOPT [45]. Extensions to 3D and more complex flows are possible. The current paper elaborates on several aspects of the authors' paper [40] and provides significant extensions of this earlier work. For the reader's convenience we maintain a self-contained presentation to avoid extensive cross-referencing.

For the discretization of the partially miscible flow PDE system we apply a variant of the BOX method [4], a locally conservative control-volume FE method that we extend by a mechanism to reduce oscillations of the CO₂ mass fraction in regions where it exhibits steep slopes. For additional information on the BOX method we refer to [22, 26]. Further possible numerical methods for multiphase flow include, among others, streamline diffusion techniques [48] and conservative mixed FE methods [38]. For a survey of suitable FE/FV methods for multiphase multicomponent flows in porous media we refer to [14] and the references therein.

Concerning PDE-constrained optimization in reservoir modeling, the research focus has

*Chair of Mathematical Optimization, Department of Mathematics, Technische Universität München, Boltzmannstr. 3, 85748 Garching b. München, Germany (simon@ma.tum.de).

†Chair of Mathematical Optimization, Department of Mathematics, Technische Universität München, Boltzmannstr. 3, 85748 Garching b. München, Germany (mulbrich@ma.tum.de).

so far not been on the CO₂ trapping mechanisms, but more on inverse history matching, i.e., parameter estimation in petroleum reservoirs (see [37] for an overview), or on the economic aspect to maximize oil production. For the latter economic objective, optimal control of mixed CO₂-water injection was considered, e.g., in the early work [34], using IMPES (implicit-pressure–explicit saturation) simulation of a modified black oil model and iterative gradient methods for optimization. An overview on recent research in optimal control of enhanced oil recovery (EOR) via water-flooding can be found in [11, 27], based on the original papers [3, 12, 15, 39, 41, 42, 44, 46]. Most current EOR publications deal with multiphase flow, especially immiscible two-phase flow and variants of the black oil model, as far as water-flooding of oil reservoirs is concerned. Recently, adjoint based methods for multiphase multicomponent flow have been evaluated [28], including a comparison between discretize-then-optimize and optimize-then-discretize.

To consider a suitable two-phase flow model for the mentioned optimal control problem in the realm of subsurface CO₂ sequestration, we restrict our attention to the fluid phases CO₂ and water/brine, neglecting oil and other substances, which leads to partially miscible two-phase flow. Note that the aforementioned control approaches to optimal water-flooding usually apply immiscible two-phase or black oil type three-phase flow models. Miscibility of CO₂ in brine introduces a further degree of complexity in our context: Switching between saturation and undersaturation of the wetting phase with CO₂ is modeled by a (smoothed) complementarity condition in the state system, as suggested in [5, 29]. Our implementation

- is flexibly extendable to more complex flow/reservoir models,
- makes use of a high-end MPI parallelized PDE library (Trilinos/Sundance),
- links to the state-of-the-art interior point optimization software (IPOPT).

The toolbox Sundance was chosen as software platform mainly due to its flexibility. Among other advantages it supports a variational problem formulation [33] that, in particular, enables the direct derivation of the exact discrete adjoint. To this end, the timestep-wise Lagrange function of the optimal control problem is formulated as a variational form. We developed an object oriented C++ interface that links the simulations and adjoint computations to the optimization software IPOPT [45], giving access to a state-of-the-art nonlinear optimization code that can handle general constraints and provides L-BFGS [49] approximations to the Hessian of the Lagrangian. In [18] several test problems demonstrate that limited-memory BFGS shows good performance for EOR history matching with multiphase flow.

The paper is organized as follows: Section 2 introduces the state equations for Darcy based partially miscible two-phase flow, together with primary state variables and control objectives. Section 3 contains a formal derivation of the continuous system of adjoint equations from the state system. In Section 4 we present the adjoint based optimal control framework, while locally conservative numerical methods for the state equation are sketched in Section 5. Our implementation of the adjoint based framework within the PDE toolbox Sundance will be discussed in Section 6. Some modeling and implementation details are presented in Section 7. Finally, Section 8 presents and discusses several optimization results.

2. State equations, variables and control objectives. CO₂ injection takes place at multiple well locations in a subregion $\Omega_c \subset \Omega$ of the underground reservoir Ω , being purely saturated with saline water (called “brine”) at the beginning. The thermodynamical reservoir conditions are chosen to have the injected CO₂ in supercritical state, so that a continuous underground flow model with two phases $\alpha \in \{w, n\}$ and two components $i \in \{1, 2\}$ is justified. Partial miscibility is understood as follows: On the one hand, the CO₂ component $i = 2$ can dissolve in the water-rich wetting phase $\alpha = w$. On the other hand, the “evaporation” of the water component $i = 1$ into the CO₂-rich nonwetting phase $\alpha = n$ is neglected.

The reservoir medium is assumed rigid, i.e., with time-invariant functions $\phi = \phi(x)$ and

$\mathbf{K} = \mathbf{K}(x)$ for porosity and absolute permeability. Moreover, temporal and spatial variations of the reservoir temperature are neglected and chemical equilibrium is assumed. Here $x \in \Omega$ denotes the spatial position within the reservoir. The numerical studies in Sections 5 and 8 are done for 2D reservoirs $\Omega \subset \mathbb{R}^2$ (x^1 horizontal displacement, x^2 vertical displacement), while the other sections cover 2D and 3D reservoirs.

As we apply multiphase Darcy's law in the flow PDE system, the mass conservation equations for the components $i \in \{1, 2\}$ (1 = water, 2 = CO₂) in the phases $\alpha \in \{w, n\}$ (wetting and nonwetting) [7] are given by

$$\phi \partial_t \sum_{\alpha} \rho_{\alpha} X_{\alpha}^i S_{\alpha} - \sum_{\alpha} \operatorname{div} [\lambda_{\alpha} \rho_{\alpha} X_{\alpha}^i \mathbf{K} (\nabla p_{\alpha} + \rho_{\alpha} \mathbf{g} e_2)] - \operatorname{div} (D^i \rho_w \nabla X_w^i) = q^i. \quad (2.1)$$

Here ρ_{α} , S_{α} and p_{α} denote the phase densities, saturations and pressures, while X_{α}^i stands for the mass fraction of component i in phase α . Especially, we have $X_n^1 = 0$ and $X_n^2 = 1$ due to partial miscibility, while we set $X := X_w^2$ so that $X_w^1 = 1 - X$. Advection and buoyancy via Darcy's law with phase mobilities λ_{α} and absolute permeability matrix \mathbf{K} are expressed in the first divergence term of (2.1), where \mathbf{g} denotes gravitational constant and e_2 is the vertical unit vector (both in \mathbb{R}^2 and in \mathbb{R}^3); slow diffusion in the wetting phase is modeled in the second divergence term. Since there is no rigorous function space theory for this challenging class of equations, we follow the common approach in the respective literature on numerical multi-phase flow models in porous media (e.g., [4, 5, 14, 29, 36]) and do not rigorously state function spaces. Modeling and implementation details will be discussed in Section 7.

Within our partially miscible setting, the CO₂ component is allowed to dissolve in the wetting phase $\alpha = w$ to some extent. Therefore we need the following three primary state variables $y := (p_n, S_w, X)$ in the flow model:

CO ₂ pressure	$p_n = p_n(x, t)$	$p_w = p_n - p_c$ with $p_c = p_c(S_w)$
brine saturation	$S_w = S_w(x, t)$	$S_n = 1 - S_w$ (natural constraint)
CO ₂ mass fraction	$X = X(x, t)$	$X = \varphi(p_n)$ when $S_w(x, t) < 1$

As an initial condition we state $S_w(\cdot, 0) = 1$ (no CO₂ injected); in such single-phase regimes the nonwetting phase pressure $p_n = p_w$ exists as a mutual parameter. Any value $X < \varphi(p_n)$ corresponds to $S_w = 1$, while the nonwetting phase $\alpha = n$ can only appear (that means $S_w < 1$) if X reaches the threshold $\varphi(p_n)$. In the presence of the nonwetting phase we always have the relation $X = \varphi(p_n)$. Thus an adaptive switching between the state variables X and S_w is required; we highlight here that chemical equilibrium is assumed, so phase transitions are neglected on a finer time scale. We model this by a complementarity condition:

$$a := 1 - S_w \geq 0, \quad b := \varphi(p_n) - X \geq 0 \quad \text{and} \quad ab = 0.$$

This condition is equivalent to $\tilde{g}(a, b) = 0$, where \tilde{g} is a complementarity function, e.g., the Fischer-Burmeister [17] function $\tilde{g}(a, b) = a + b - \sqrt{a^2 + b^2}$. To maintain differentiability, we smooth this algebraic equation in our implementation to

$$g(y) := \tilde{g}_{\varepsilon}(a_{\varepsilon}, b_{\varepsilon}) = a_{\varepsilon} + b_{\varepsilon} - \sqrt{a_{\varepsilon}^2 + b_{\varepsilon}^2 + 2\varepsilon^2} = 0, \quad (2.2)$$

so that the sharp switching condition $ab = 0$ mollifies to $a_{\varepsilon} b_{\varepsilon} = \varepsilon^2 \ll 1$. The smoothed complementarity variables are chosen as

$$a_{\varepsilon} := 1 + \varepsilon^2 - S_w, \quad b_{\varepsilon} := \chi[\varphi(p_n) - X], \quad (2.3)$$

where the factor χ brings a_ε and b_ε to a similar magnitude in the off-threshold realm. The application of complementarity conditions for the modeling of phase transitions in miscible multiphase flow has also been suggested in [5, 29].

Using the state variables $y = (p_n, S_w, X)$, connected through the algebraic complementarity condition $g(y) = 0$, we can rewrite the flow system (2.1) in the schematic form of an algebro-differential system

$$\phi \partial_t b^i(y) - \operatorname{div} B^i(y, \nabla y, x) = q^i(u) \quad (i = 1, 2) \quad \text{and} \quad g(y) = 0$$

in the space-time interior domain $\Omega \times (0, T)$; for this compact notation we abbreviate

$$b^1 := \rho_w(1 - X)S_w, \quad b^2 := \rho_n(1 - S_w) + \rho_w X S_w,$$

$$B^1 := \lambda_w \rho_w (1 - X) \mathbf{K} [\nabla(p_n - p_c) + \rho_w \mathbf{g} e_2] - D^1 \rho_w \nabla X,$$

$$B^2 := \lambda_w \rho_w X \mathbf{K} [\nabla(p_n - p_c) + \rho_w \mathbf{g} e_2] + \lambda_n \rho_n \mathbf{K} (\nabla p_n + \rho_n \mathbf{g} e_2) + D^2 \rho_w \nabla X.$$

The spatial boundary $\partial\Omega = \Gamma_n \cup \Gamma_d$ is decomposed into Neumann and Dirichlet parts: A no-flux condition $n^T B^i(y, \nabla y, x) = 0$ for $i = 1, 2$ is considered on the Neumann boundary Γ_n (impermeable caprock). The Dirichlet conditions $y = y_d$ on Γ_d impose standard values without CO₂ injection, i.e., hydrostatic pressure $p_n = p_n(x)$, full brine saturation $S_w = 1$ and a small CO₂ mass fraction $X = X_0 \ll \varphi(p_n)$. Here $X_0 = X_0(x)$ is chosen such that $b_\varepsilon = 1$ holds on the Dirichlet boundary, which by virtue of (2.2), (2.3) necessitates $g(y) = 0$. The initial values inside the reservoir reproduce the Dirichlet conditions.

Control parameters $u = u(t)$ purely appear in the source terms $q^i = q^i(u)$:

$$q^1 = 0 \quad \text{and} \quad q^2 = \sum_{n=1}^N q_n(t) \omega(x - x_n).$$

So no water is to be injected, while q^2 models the injection of supercritical CO₂ at $N \in \mathbb{N}$ injection wells with locations x_n in a certain control area $\Omega_c \subset \Omega$; here ω models the injection distribution, given, e.g., by a narrow Gaussian normal distribution. In the following, the N injection rates $q_n(t)$ are used as control parameters:

$$u(t) = (q_1(t), \dots, q_N(t)), \quad t \in [0, T]. \quad (2.4)$$

Imposing box constraints on the above rates and making sure that the total amount of injected CO₂ does not exceed a certain threshold, we result in control constraints

$$0 \leq q_n(t) \leq a, \quad t \in [0, T], \quad 1 \leq n \leq N \quad \text{and} \quad \frac{1}{T} \sum_{n=1}^N \int_0^T q_n(t) dt \leq b \quad (2.5)$$

for fixed numbers $a, b > 0$ with $b < Na$ (otherwise the latter constraint is redundant). The second average flow constraint would also be particularly pertinent for the infinite time limit $T \rightarrow \infty$. However, very large time horizons $T \gg 1$ years lie beyond our current computational capacities and are therefore not investigated here.

For future investigation, it will also be promising to introduce additional state constraints, such as an upper bound on the pressure p_n to preserve caprock integrity. State constraints are known to be challenging from a theoretical point-of-view, see for instance Section 2.7.2 in [24] or [43]. A suitable approach to tackle such constraints is by regularization, e.g., of

Moreau-Yosida type [24, 35, 43]. In this paper we only consider control constraints, but highlight that our optimization interface is capable of handling state constraints by a regularization approach.

The optimization objective here is to maximize the amount of trapped CO_2 after a finite time $T > 0$ of injection. On a modest time scale (months–years), such “trapping” can happen in several ways: The injected CO_2 can be

- immobilized below a certain residual saturation S_n^r (residual trapping),
- dissolved into the brine and thereby trapped for the moment (solubility trapping),
- controlled to stay in a certain subregion $V \subset \Omega$ (weak area trapping).

The residual saturation $S_n^r = S_n^r(\phi)$ depends on the reservoir porosity. The nonwetting phase mobility λ_n drops to zero below this threshold, i.e., for saturations $S_n \leq S_n^r$. For an online illustration of CO_2 trapping mechanisms we refer the reader to [50].

A general optimization objective concerns the amount of trapped CO_2 in either of the above variants, which leads to an objective functional of the general form

$$J(y, u) = \int_0^T \int_{\Omega} d(y, x) dx dt + \int_{\Omega} f(y, x)|_{t=T} dx - \gamma R(u), \quad (2.6)$$

where a control-dependent penalty term $R(u)$ allows for regularization of the optimization algorithm in terms of a user-defined parameter $\gamma \geq 0$. The terminal payoff f is allowed to include all three different kinds of trapping, whereas the distributed term d is restrained to contain only the weak area trapping. Hence we obtain

$$\begin{aligned} d &= \frac{\beta_0}{T} \phi \rho_n S_n \mathbf{1}_{\{x \in V\}}, \\ f &= \beta_1 \phi \rho_n S_n \mathbf{1}_{\{x \in V\}} + \beta_2 \phi \rho_n S_n \mathbf{1}_{\{S_n \leq S_n^r\}} + \beta_3 \phi \rho_w X S_w, \end{aligned} \quad (2.7)$$

with respective weighting coefficients $\beta_j \in [0, 1]$. The penalty term for our regularization purposes – necessary especially due to the composite indicator function $\mathbf{1}_{\{S_n \leq S_n^r\}}$ in the case of residual trapping – is chosen in the weighted H^1 -type manner

$$R(u) := \frac{q_*^2}{T} \sum_{n=1}^N \int_0^T \left[T^2 q_n'(t)^2 + \theta(t) q_n(t)^2 \right] dt, \quad (2.8)$$

including a scaling factor q_* ($q_* = 200$ in SI units) and a weight function $\theta : [0, T] \rightarrow [0, \infty[$ that introduces certain time preferences. This strictly convex regularization smooths out the optimization algorithm for reasonable choices of the regularization parameter γ and thereby enhances the flexibility and robustness of the optimization interface.

We highlight that most of the concrete choices made in this section could be replaced by a wide range of other settings without leaving our overall framework. This modeling freedom stems from suitably chosen locally conservative numerical simulation methods and a flexible implementation of the state and adjoint PDE solvers that makes use of variational problem features in Sundance (see Sections 5 and 6 for details). While the above optimal control formulation as well as our numerical experiments address the open loop case, the presented approach can be embedded into a nonlinear model predictive control (NMPC) framework. NMPC [2, 20] is a closed loop control strategy, where subproblems of the above form are solved on a sequence of overlapping, moving prediction time horizons $[t_k, t_k + T_p]$, $k = 0, 1, 2, \dots$, where $t_k = k\Delta T$ and $0 < \Delta T < T_p$. The NMPC feedback control on $(t_k, t_{k+1}]$ is obtained from the corresponding part of the optimal control of the k -th subproblem. Thus, a main building block of NMPC is to solve optimal control problems of the form considered

in this paper. We mention that there are additional aspects to be addressed in NMPC, in particular state estimation as well as the fact that the solution of the k -th subproblem is not instantaneously available at time $k\Delta t$.

3. Formal derivation of the continuous adjoint system. As seen above, the system of state equations for partially miscible two-phase flow with state variables $y = (p_n, S_w, X)$ can be written in the compact form

$$\phi \partial_t b^i(y) - \operatorname{div} B^i(y, \nabla y, x) = q^i(u) \quad (i = 1, 2) \quad \text{and} \quad g(y) = 0 \quad (3.1)$$

in the space-time interior domain $\Omega \times (0, T)$, subject to the boundary conditions

$$\begin{aligned} n^T B^i(y, \nabla y, x) &= 0 \quad \text{on} \quad \Gamma_n \times (0, T) \quad (i = 1, 2), \\ y &= y_d \quad \text{on} \quad \Gamma_d \times (0, T), \end{aligned} \quad (3.2)$$

as well as initial and final time conditions $y = y_d$ for $t = 0$ and $g(y) = 0$ for $t = T$. Furthermore, we have to ensure the complementarity condition $g(y) = 0$ on the spatial boundary $\partial\Omega \times (0, T)$. This condition automatically holds on the Dirichlet boundary and at $t = 0$, since we choose the Dirichlet values y_d such that $g(y_d) = 0$. A weak formulation of the state equation can be found in (5.2).

The following derivations in this section are formal in the sense that a rigorous function space setting is not available for the nonlinear system of PDEs (3.1) under consideration. When changing (3.1) into its weak variational formulation $\langle \lambda, E(y, u) \rangle = 0$ with

$$\lambda := (\mu, \sigma, \nu) = (\mu_1, \mu_2, \mu_3, \sigma_1, \sigma_2, \sigma_3, \nu_1, \nu_2),$$

the primary test functions μ refer to the PDEs and the complementarity $g(y) = 0$; the three spatial functions σ are included to test the initial condition, while the terms ν test the complementarity at the final time $t = T$ and on the Neumann boundary. Here, the dual pairing $\langle v, w \rangle = \sum_i \langle v_i, w_i \rangle$ is the sum over the dual pairings for the respective function spaces of the components.

The Dirichlet conditions are built into the state and test space: $y = y_d$ and $\mu = 0$ on $\Gamma_d \times (0, T)$, respectively. Hence, the weak variational formulation $\langle \lambda, E(y, u) \rangle = 0$ of the state system can be written as

$$\begin{aligned} & \int_0^T \int_{\Omega} \left\{ \sum_{i=1}^2 \left[\mu_i \phi \partial_t b^i(y) + \nabla \mu_i^T B^i(y, \nabla y, x) - \mu_i q^i(u) \right] + \mu_3 g(y) \right\} dx dt \\ & + \int_{\Omega} \left\{ \sum_{j=1}^3 \sigma_j \left[y_j|_{t=0} - (y_d)_j \right] + \nu_1 g(y)|_{t=T} \right\} dx \\ & + \int_0^T \int_{\Gamma_n} \nu_2 g(y) dS dt = 0 \end{aligned} \quad (3.3)$$

for all test functions $\lambda = (\mu, \sigma, \nu)$ as described above.

To derive the corresponding adjoint system $E_y^* \lambda + J_y = 0$ in strong formulation, where the subscript y denotes the partial Fréchet derivative of the respective operator in direction y , we linearize the variational form $\langle \lambda, E(y, u) \rangle$ in the state variables and integrate by parts.

The linearized variational form $\langle \lambda, E_y \delta y \rangle$ in direction δy with $\delta y = 0$ on $\Gamma_d \times (0, T)$ reads

$$\begin{aligned} & \int_0^T \int_{\Omega} \left\{ \sum_{i=1}^2 \left[\mu_i \phi \partial_t (b_y^i \delta y) + \nabla \mu_i^T \left(B_y^i \delta y + \sum_{j=1}^3 B_{\nabla y_j}^i \nabla \delta y_j \right) \right] + \mu_3 g_y \delta y \right\} dx dt \\ & + \int_{\Omega} \left[\sum_{j=1}^3 \sigma_j \delta y_j \Big|_{t=0} + \nu_1 (g_y \delta y) \Big|_{t=T} \right] dx + \int_0^T \int_{\Gamma_n} \nu_2 g_y \delta y dS dt, \end{aligned}$$

while the Fréchet derivative of the general objective functional (2.6) results in

$$\langle J_y, \delta y \rangle = \int_0^T \int_{\Omega} d_y \delta y dx dt + \int_{\Omega} (f_y \delta y) \Big|_{t=T} dx.$$

Upon integration by parts both in time and space, this leads to the adjoint equation

$$\begin{aligned} \langle E_y^* \lambda + J_y, \delta y \rangle &= \int_0^T \int_{\Omega} (\mu_3 g_y \delta y + d_y \delta y) dx dt \\ &+ \sum_{i=1}^2 \int_0^T \int_{\Omega} \left[-\phi b_y^i \delta y \partial_t \mu_i + \nabla \mu_i^T B_y^i \delta y - \sum_{j=1}^3 \delta y_j \operatorname{div} \left((B_{\nabla y_j}^i)^T \nabla \mu_j \right) \right] dx dt \\ &+ \int_{\Omega} \left[\left(\sigma^T \delta y - \sum_{i=1}^2 \phi \mu_i b_y^i \delta y \right) \Big|_{t=0} + \left(\nu_1 g_y \delta y + \sum_{i=1}^2 \phi \mu_i b_y^i \delta y + f_y \delta y \right) \Big|_{t=T} \right] dx \\ &+ \int_0^T \int_{\Gamma_n} \left[\nu_2 g_y \delta y + \sum_{i=1}^2 \sum_{j=1}^3 \delta y_j n^T (B_{\nabla y_j}^i)^T \nabla \mu_i \right] dS dt = 0, \end{aligned}$$

here already presented in a partially strong writing; note that directions δy are considered as test functions with $\delta y = 0$ on the space-time Dirichlet boundary $\Gamma_d \times (0, T)$. In the following, we derive the strong form of the adjoint equation system.

Choosing for δy all functions that have compact support in $\Omega \times (0, T)$, we obtain the following adjoint system of PDEs with primary variables μ :

$$\sum_{i=1}^2 \left[-\phi b_{y_k}^i \partial_t \mu_i + (B_{y_k}^i)^T \nabla \mu_i - \operatorname{div} \left((B_{\nabla y_k}^i)^T \nabla \mu_k \right) \right] + g_{y_k} \mu_3 + d_{y_k} = 0 \quad (3.4)$$

for $k = 1, 2, 3$ in the space-time interior domain $\Omega \times (0, T)$. The adjoint PDE system (3.4) evolves backwards in time, in contrast to the state system (3.1), and the corresponding final time ‘‘initial conditions’’ are derived below.

Considering δy that vanish outside a small neighborhood of $\Omega \times \{0\}$, we can eliminate the adjoint variables σ as an expression of the primary variables μ :

$$\sigma = \phi \sum_{i=1}^2 (\mu_i b_{y_k}^i) \Big|_{t=0}$$

In a similar manner we can consider δy that vanish outside a neighborhood of $\Omega \times \{T\}$, which results in the final time condition

$$\nu_1 g_{y_k} + \phi \sum_{i=1}^2 \mu_i b_{y_k}^i + f_{y_k} = 0 \quad \text{at the final time } t = T \quad (k = 1, 2, 3) \quad (3.5)$$

for the time-reversed adjoint PDE system (3.4). Here the linear system (3.5) provides three equations for three unknowns ν_1 and $\mu_{1,2}|_{t=T}$, so ν_1 serves as an auxiliary variable.

Boundary conditions on the space-time Dirichlet boundary $\Gamma_d \times (0, T)$ are $\mu = 0$ by construction. Finally, taking test functions on the Neumann boundary into account, the respective boundary integral results in adjoint Neumann conditions

$$\nu_2 g_{y_k} + n^T \sum_{i=1}^2 (B_{\nabla y_k}^i)^T \nabla \mu_i = 0 \quad \text{on } \Gamma_n \times (0, T) \quad (k = 1, 2, 3). \quad (3.6)$$

Notice once more that ν_2 provides the third variable for the linear system at hand.

To sum up, the adjoint system for the primary adjoint states μ is given by (3.4), (3.5), and (3.6) for $k = 1, 2, 3$ with Dirichlet conditions $\mu = 0$ on $\Gamma_d \times (0, T)$, while $\nu_{1,2}$ provide auxiliary algebraic variables on $\Omega \times \{T\}$ and $\Gamma_n \times (0, T)$. For immiscible flow the variables μ_3 and $\nu_{1,2}$ are not needed since $y_3 = 0$ and $g(y) = 0$ may be eliminated.

4. Adjoint based optimal control framework. Provided that the functionals $E(y, u)$ and $J(y, u)$ are sufficiently smooth and that the weak state equation $E(y, u) = 0$ has a unique solution $y = y(u)$ for every control u , the adjoint approach [24] can be applied for efficient gradient computation of the reduced objective functional $j(u) := J(y(u), u)$:

$$j'(u) = E_u(y, u)^* \lambda + J_u(y, u) \quad \text{where} \quad E_y(y, u)^* \lambda + J_y(y, u) = 0. \quad (4.1)$$

The following numerical investigations are done under the assumption that the adjoint approach is feasible and that the reduced functional is at least continuously Fréchet differentiable. As already inferred at the beginning of Section 3, rigorous functional analytic proofs of these assumptions are very challenging and would, in particular, require a rigorous function space analysis (existence and uniqueness, etc.) of the multiphase flow equation, which is not available to the best of our knowledge. Thus, a rigorous function space analysis is beyond the scope of this article.

Our implementation follows the “discretize-then-optimize” approach, i.e., the discretized adjoint system coincides with the exact discrete adjoint of the discretized state equation. The time horizon $t \in [0, T]$ is discretized into n_t equidistant time steps t_k with step size $\Delta t := t_k - t_{k-1}$. We do not apply adaptive time stepping currently, but it would be possible to integrate this: While the spatial discretization and the computation of adjoints in space are obtained automatically from the features provided by Sundance, the time stepping for forward and adjoint solves is implemented manually as a for-loop and thus can be adjusted to include adaptivity. It is known that adaptive time stepping requires care to ensure sufficient accuracy of the adjoint based gradient. To achieve this, a local (in terms of the control) freezing of the time grid during adjoint based derivative computations would be an appropriate way to include adaptivity in time [19, 25, 30]. A selection of further work on adjoint based optimal control with adaptive time stepping can be found in [1, 16, 19, 25, 30].

Within our implementation of equidistant time stepping, we only take scaled injection rates at subsequent time steps t_{Kj} and $t_{K(j+1)}$ into account, which reduces the number of control parameters. Here $K \in \mathbb{N}$ is a divisor of n_t and linear interpolation is applied to obtain the injection rates in between. Abbreviating $k_t := \frac{n_t}{K}$, the corresponding vector

$$u \in \mathbb{R}^{Nk_t} \quad \text{with} \quad u_{(n-1)k_t+j} := q_* q_n(t_{Kj}), \quad 1 \leq j \leq k_t, \quad 1 \leq n \leq N, \quad (4.2)$$

discretizes the control variables (2.4) with the scaling factor q_* , while $q_n(0) = 0$ is chosen for the initial interpolation steps.

The control constraints (2.5) in discrete form read

$$0 \leq u_i \leq q_* a, \quad 1 \leq i \leq Nk_t \quad \text{and} \quad c(u) := \frac{1}{k_t} \sum_{i=1}^{Nk_t} u_i \leq q_* b,$$

while the penalty term (2.8) can be discretized as

$$\begin{aligned} R(u) &\approx \sum_{n=1}^N k_t \sum_{j=1}^{k_t-1} [u_{(n-1)k_t+j+1} - u_{(n-1)k_t+j}]^2 \\ &+ \sum_{n=1}^N \frac{1}{k_t} \sum_{j=1}^{k_t} \theta_j u_{(n-1)k_t+j}^2. \end{aligned} \quad (4.3)$$

Here we set $\theta_j := \theta(t_{Kj})$ for the user-defined weighting coefficients.

To solve the resulting constrained nonconvex optimization problem, we choose the interior point based software package IPOPT [45]. Since our state equation is highly nonlinear and we use the reduced control problem, which only depends on the time-dependent control, we have to solve a constrained, highly nonconvex optimization problem of moderate dimension. Evaluating the objective function is expensive since it involves a solution of the multi-phase flow equations. Obtaining the gradient of the reduced objective function requires the solution of the adjoint equation. Therefore, an efficient and robust gradient-based optimization code that can handle constraints is the best choice for our problem. IPOPT is among the most efficient and robust open source nonlinear optimization codes and thus meets our requirements very well. Certainly, other comparably good choices would be possible. An object oriented C++ interface was developed that calls the Sundance state and adjoint solvers to compute $j(u)$ and its derivative $j'(u)$. We configured the interior point algorithm in IPOPT to use limited-memory BFGS updates [49] for approximating the Hessian matrix of the Lagrange function.

For a given control vector u , $j(u)$ and $j'(u)$ can be computed from the state and adjoint systems, using the adjoint derivative representation (4.1). As the control parameters explicitly just appear in the space-time interior part of the weak state equation, we can restrict our attention to the primary variables y and μ in the respective weak implementation.

5. Locally conservative methods for the state system. Let y^k denote the time-discrete state variables at the time levels $t = t_k$. Applying an implicit Euler scheme to the state system (3.1), we obtain the following semi-discrete state equation in strong form:

$$\phi \frac{b^i(y^k) - b^i(y^{k-1})}{\Delta t} - \operatorname{div} B^i(y^k, \nabla y^k, x) - q^i|_{t=t_k} = 0 \quad (i = 1, 2) \quad (5.1)$$

for $1 \leq k \leq n_t$, together with the complementarity condition $g(y^k) = 0$ and the boundary conditions. The respective initial condition $y^0 = y_d$ reproduces the Dirichlet values.

For the space discretization we apply the BOX method [4], a control-volume FE method. It has two main advantages: It is locally conservative due to its control-volume nature and it possesses sufficient stability due to involved full upwinding [21] and mass lumping [26]. Therefore it is very suitable for simulation of the highly nonlinear multiphase flow system under consideration (see [4] for a thorough discussion). To make the PDE toolbox Sundance [32] within the Trilinos framework accessible for such stabilized locally conservative methods, we integrated both the control-volume and upwinding features into its object oriented C++ core structure.

On the one hand, the BOX method discretizes the state variables y^k to piecewise linear continuous finite elements on a primal (here triangular) mesh. On the other hand, similar to finite volume methods, constant functions on vertex-centered cells (“boxes”) of a dual mesh are used as test functions; see Figure 5.1 for the construction of a box in the dual mesh.

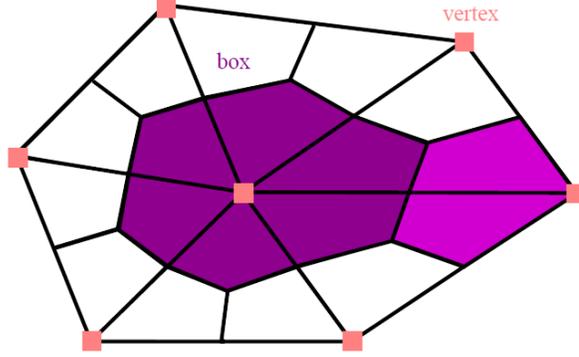


FIG. 5.1. Construction of boxes in the dual mesh.

Given a primal triangulation of the domain $\Omega \subset \mathbb{R}^2$, we denote the collection of boxes in the dual mesh by B_j ($j = 1, \dots, m$) and their respective box volumes by ω_j . Following the derivation of the BOX method in [4], the semi-discrete state system (5.1) is first integrated over each box B_j – note that indicator functions $\mathbf{1}_{B_j}$ constitute a basis of our test space – and Gauß’ divergence theorem is applied to yield a boundary integral:

$$\begin{aligned} & \int_{B_j} \left[\phi \frac{b^i(y^k) - b^i(y^{k-1})}{\Delta t} - \operatorname{div} B^i(y^k, \nabla y^k, x) - q^i|_{t=t_k} \right] dx \\ &= \int_{B_j} \left[\phi \frac{b^i(y^k) - b^i(y^{k-1})}{\Delta t} - q^i|_{t=t_k} \right] dx - \int_{\partial B_j} n^T B^i(y^k, \nabla y^k, x) dS = 0 \end{aligned} \quad (5.2)$$

Inserting piecewise linear Lagrange elements for the state variables and applying upwinding to the vector-valued terms B^i in the boxwise weak form (5.2), we arrive at

$$\int_{B_j} \left[\phi \frac{b^i(y^k) - b^i(y^{k-1})}{\Delta t} - q^i|_{t=t_k} \right] dx - \int_{\partial B_j} n^T B_{\text{up}}^i dS = 0 \quad (i = 1, 2). \quad (5.3)$$

The fully implicit upwinding procedure that produces B_{up}^i is based on the decomposition

$$B^i = - \sum_{\alpha} C_{\alpha}^i V_{\alpha} + D^i \rho_w X_w^i \quad \text{with} \quad C_{\alpha}^i := \lambda_{\alpha} \rho_{\alpha} X_{\alpha}^i, \quad V_{\alpha} := -\mathbf{K}(\nabla p_{\alpha} + \rho_{\alpha} \mathbf{g} e_2).$$

Upon implicit evaluation of the Darcy velocities V_{α} at the current discretized state y^k , the above advective coefficients C_{α}^i are implicitly evaluated

- at the vertex v_j of the box B_j for outflow $n^T V_{\alpha} \geq 0$,
- at the vertex of the opposite box for inflow $n^T V_{\alpha} < 0$.

Here n denotes the outer dual cell normal of the box B_j and “opposite” in 2D is subject to the respective line segment of the polygonal curve ∂B_j (see Figure 5.1). The flux terms B_{up}^i are upwinded boxwise, using this edge based decision strategy.

Box quadrature rules are chosen as follows: To evaluate the boundary integrals in (5.3), a midpoint rule is applied on each line segment of ∂B_j in triangular 2D grids. Volume integrals are evaluated vertex-centered, i.e.,

$$\int_{B_j} f(x) dx \approx \omega_j f(v_j), \quad 1 \leq j \leq m,$$

where v_j denotes the unique primal vertex inside the box B_j (see Figure 5.1). For the control-volume FE method this results in a mass lumping, where the entries of the FE mass matrix are assigned to its main diagonal [26]. Such lumped mass matrices counteract the appearance of unphysical oscillations in the solution [13]. For further details we refer the reader to Chapter 3 in [4] for two-phase flow or Section 4.2 in [7] for miscible flow.

Finally, since at discontinuity-like fronts of the CO_2 mass fraction X sometimes local undershooting that results in unphysical oscillations could be observed (see Subsection 8.1 for an illustration), we augmented the BOX method by a simple strategy that reduces this effect (and which could be replaced by more sophisticated techniques): dual cell averaging of X every $M \in \mathbb{N}$ time steps, where M is a divisor of n_t and $m_t := \frac{n_t}{M}$. In detail, after application of the BOX scheme at the time levels t_{Mj} ($1 \leq j < m_t$), the variable X^{Mj} is substituted via a linear integral transformation by its average value \bar{X}^{Mj} over the respective boxes in the dual mesh. This transforms the discretized state y^{Mj} into an intermediate state \bar{y}^{Mj} , which is taken as a basis for the respective next implicit time step t_{Mj+1} in the BOX scheme. We highlight that this averaging procedure is locally conservative.

6. Implementation strategies in Trilinos/Sundance. As a software platform for state and adjoint simulations, we use the FEM package Sundance [32] within the scientific computing framework Trilinos [23]. This software is also used and further developed in other projects of our research group, thus generating significant synergies [6, 8, 9, 31]. In particular, the C++ interface of the Sundance simulation code to IPOPT stems from our earlier work in a project on shape optimization with Navier-Stokes flow [8, 9, 31], partially funded by the DFG.

Sundance has a wide range of benefits, for instance efficient treatment of nonlinearities, automatic generation of linearized equations and thereby support for implementing adjoints within gradient based optimization algorithms [33]. We represent the whole optimal control problem by its Lagrange function $\mathcal{L}(y, \bar{y}, u, \mu, \bar{\mu})$, implemented as a timestep-wise variational form in Sundance. Here μ denotes the adjoint state, i.e., the Lagrange multiplier of the state equation, while the additional variables \bar{y} and $\bar{\mu}$ account for the periodic averaging steps. Especially, Sundance includes routines that automatically derive nonlinear and linear variational problems from this Lagrangian corresponding to $\mathcal{L}_{\mu^k} = 0$, which is the state equation, and $\mathcal{L}_{y^k} = 0$, which is the adjoint equation. Furthermore, Sundance is MPI based and provides parallel system assembly as well as interfaces to a whole library of direct and iterative parallel solvers (e.g., NOX, Amesos, AztecOO).

To make these convenient features available, we had to implement additional class structures and quadrature rules within the PDE toolbox, so far without direct support for control volumes and upwinding. As those are essential ingredients of the applied variant of the BOX method (see Section 5), a 2D implementation was built into the object oriented C++ framework of Sundance.

Our overall approach is designed such that it could be extended to 3D mainly by adjusting the BOX method for the PDE discretization. We briefly discuss this here before we continue with the details of our 2D approach. As illustrated below, our top-level coding is very flexible and thus mainly requires to adjust the Lagrange function of the optimal control problem to the 3D setting. The higher level state and adjoint computations and the optimization interface

are essentially independent of the space dimension. Since the discrete problem dimension would be significantly higher in 3D, efficient preconditioning of the Krylov subspace solvers would be more important than in 2D. The main work in extending our code to 3D would be required for the extension of several geometric low-level subroutines (dual elements etc.) from 2D to 3D. For 3D quad grids this is possible with reasonable effort, whereas unstructured tetrahedral meshes in 3D would need more low-level implementation work. Alternatively, one could stack copies of the 2D grid to extrude it along the third dimension. This latter approach, however, would require adjustments in the top-level problem coding.

Returning to our actual implementation, we now illustrate our quite compact, Lagrange function based implementation of the state and adjoint equation systems: The state equation for y^k and its adjoint for μ^k at time step t_k as well as the periodic averaging procedure after the steps t_{Mj} can be written in terms of the time-discretized Lagrange function:

$$\begin{aligned} \mathcal{L}_{\mu^k}(y, \bar{y}, u, \mu, \bar{\mu}) &= 0 \quad (\text{time-discretized state equation}), \\ \mathcal{L}_{y^k}(y, \bar{y}, u, \mu, \bar{\mu}) &= 0 \quad (\text{corresponding adjoint equation}), \\ \mathcal{L}_{\bar{\mu}^{Mj}}(y, \bar{y}, u, \mu, \bar{\mu}) &= 0 \quad (\text{state equation for averaging}), \\ \mathcal{L}_{\bar{y}^{Mj}}(y, \bar{y}, u, \mu, \bar{\mu}) &= 0 \quad (\text{respective adjoint equation}), \\ \text{where } y &:= (y^1, \dots, y^{n_t}), \quad \mu := (\mu^1, \dots, \mu^{n_t}) \\ \text{and } \bar{y} &:= (\bar{y}^M, \dots, \bar{y}^{M(m_t-1)}), \quad \bar{\mu} := (\bar{\mu}^M, \dots, \bar{\mu}^{M(m_t-1)}). \end{aligned}$$

Note that the linear integral transformation for each averaging step can be written in weak form as $A(\bar{y}^{Mj}, y^{Mj}, \bar{\mu}^{Mj}) = 0$ and may be interpreted as an intermediate time step. The Lagrangian \mathcal{L} contains the stepwise weak state equation functionals $\langle \mu^k, E_k \rangle$, the stepwise evaluations of the objective functional (2.6) and the averaging functional as follows:

$$\begin{aligned} \mathcal{L}(y, \bar{y}, u, \mu, \bar{\mu}) &= \sum_{\substack{k=1 \\ k \neq Mj+1}}^{n_t} L_k(y^k, y^{k-1}, u, \mu^k) \\ &+ \sum_{j=1}^{m_t-1} \left[A(\bar{y}^{Mj}, y^{Mj}, \bar{\mu}^{Mj}) + L_{Mj+1}(y^{Mj+1}, \bar{y}^{Mj}, u, \mu^{Mj+1}) \right] \end{aligned}$$

with the respective timestep-wise Lagrangian defined by

$$L_k := \Delta t \left[\langle \mu^k, E_k(y^k, y^{k-1}, u) \rangle + \int_{\Omega} d(y^k) dx \right] + \delta_{kn_t} \int_{\Omega} f(y^{n_t}) dx - \gamma r_k(u)$$

for $1 \leq k \leq n_t$ (with \bar{y}^{k-1} instead of y^{k-1} for $k = Mj + 1$), where the terms r_k denote the corresponding stepwise components of the time-discretized penalty term (4.3).

Since μ^k only appears in the functional L_k , whereas y^k appears both in L_k and L_{k+1} for $k < n_t$ with $k \neq Mj$ and only in L_{n_t} for $k = n_t$, this means that, defining two further functionals F and G at a respective fixed time step t_k without averaging via

$$F := L_k \quad \text{and} \quad G := L_k + (1 - \delta_{kn_t})L_{k+1},$$

one can write the weak time-discretized state and adjoint equations as $F_{\mu^k} = 0$ and $G_{y^k} = 0$. At the discrete times t_{Mj} ($1 \leq j < m_t$) with averaging the time step splits into two separate parts that set up the respective equations for the stepwise computation of

$$y^{Mj} \xrightarrow{F_1} \bar{y}^{Mj} \xrightarrow{F_2} y^{Mj+1} \quad \text{and} \quad \mu^{Mj} \xleftarrow{G_2} \bar{\mu}^{Mj} \xleftarrow{G_1} \mu^{Mj+1},$$

where the forward settings are $F_1 = A$ and $F_2 = L_{Mj+1}$, while the corresponding adjoint functionals $G_{1,2}$ combine the timestep-wise Lagrangians with the averaging functional. This variational setup provides the basis for our flexible implementation of the discrete state system on the one hand and its exact discrete adjoint system on the other hand. For an in-depth view on these implementations and for an illustration of the modular capabilities of Sundance, the interested reader is referred to Section 5 of the authors' recent article [40].

For runtime acceleration, we implemented an IMPES-like strategy to generate good initial iterates for implicit time steps that turn out to require too many NOX iterations: When solving the time-discretized state system, it is natural to use the respective state at time t_{k-1} as an initial value \underline{y}_S for NOX' inexact Newton method. However, using a uniform step size Δt , sometimes too many NOX iterations occur in time steps where comparably large state changes take place. Note that we apply a fully implicit time stepping scheme. If this situation occurs, we generate better estimates as follows: Given the state iterate \underline{y} on the time level t_{k-1} , we decompose the time interval $[t_{k-1}, t_k]$ into n_0 equidistant subintervals and perform simplified IMPES-like nonlinear solves on each interval (here IMPES abbreviates implicit-pressure-explicit-saturation), each of which is called from a C++ function that does not contain upwinding and is as explicit as possible (no details here). The thus obtained estimate for the current iterate is then taken as initial value \underline{y}_S in the fully implicit state system, reducing the number of NOX solver iterations significantly. Since the IMPES-like preiterates are quite cheap from a computational point-of-view, this accelerates the state simulation by up to a factor of three.

Furthermore, the state and adjoint simulation was parallelized by domain decomposition. Tests were run on up to $n_p = 32$ processor units with good strong scaling properties even on relatively coarse grids. The parallelization uses in-built features of Sundance with MPI based communication between processors.

7. Model and implementation details. All following formulas are written in standard SI units, neglecting the exact physical nomenclature of the units for the sake of clarity. First of all, we restrict our attention to reservoirs with constant porosity $\phi = 0.2$ and therefore also impose constant residual saturations $S_w^r = S_n^r = 0.1$ for water and CO_2 . Our choices for the absolute permeability matrix $\mathbf{K} = \mathbf{K}(x)$ will be discussed in Section 8.

Standard Brooks-Corey [10] models are taken for the capillary pressure p_c and the phase mobilities λ_α : Defining the effective saturations

$$\mathcal{S}_w := \frac{S_w - S_w^r}{1 - S_w^r - S_n^r} \quad \text{and} \quad \mathcal{S}_n := 1 - \mathcal{S}_w,$$

the Brooks-Corey models may be represented in the form

$$p_c = p_0 \mathcal{S}_w^{-\frac{1}{\lambda}} \quad \text{and} \quad \lambda_w = \frac{1}{\mu_w} \mathcal{S}_w^{\frac{2+3\lambda}{\lambda}}, \quad \lambda_n = \frac{1}{\mu_n} \mathcal{S}_n^2 \left(1 - \mathcal{S}_w^{\frac{2+\lambda}{\lambda}}\right). \quad (7.1)$$

Here the dynamic viscosities $\mu_w = 10^{-3}$ and $\mu_n = 10^{-4}$ are chosen constant, while we set $\lambda = 2$ and $p_0 = 2 \cdot 10^5$ for the entry level capillary pressure. To establish smoothness, cubic splines were used to mollify the functions (7.1) at the boundaries $\mathcal{S}_w = 1$ and $\mathcal{S}_w = 0$.

The brine phase is assumed incompressible, but a linear density increase with the CO_2 mass fraction X is taken into account. Slight compressibility of the CO_2 phase is expressed in a linearized density $\rho_n(p_n)$. We implemented the linearized models

$$\rho_n = 4 \cdot 10^2 + c_n(p_n - 8 \cdot 10^6), \quad \rho_w = 1.1 \cdot 10^3 + 400X,$$

where $c_n = 10^{-4}$ is the approximated compressibility coefficient about $p_n = 8 \cdot 10^6$. The values of the brine density ρ_w are chosen to fit a moderately saline reservoir with salt mass

fraction $s \approx 0.15$ in the brine phase; see Chapter 3 in [7] for a thorough discussion of the underlying physical properties, including more complicated modeling efforts.

In analogy to the CO_2 density, the solubility threshold $X = \varphi(p_n)$ is again linearized about $p_n = 8 \cdot 10^6$, which yields the approximate implementation

$$\varphi(p_n) = 3.3 \cdot 10^{-2} + c_X (p_n - 8 \cdot 10^6) \quad \text{with} \quad c_X = 10^{-9}.$$

The corresponding smoothing parameter in the complementarity condition (2.2)(2.3) is chosen as $\varepsilon = 3 \cdot 10^{-4}$. A model for the diffusion coefficients $D^i = D^i(\phi, S_w)$ is taken from Section 2.5 of [7], namely

$$D^1 = 0, \quad D^2 = D_0 (\phi^4 S_w^{10})^{\frac{1}{3}} \quad \text{with} \quad D_0 = 2 \cdot 10^{-9}.$$

In comparison to advection this process happens rather slowly. Note that $D^1 = 0$ means we neglect the self-diffusion of water in the brine phase during our implementation.

In order to receive a sufficiently smooth objective functional, the indicator functions in (2.7) were implemented using an exponential C^∞ -approximation of the Heaviside function. A similar strategy has been applied to smooth out the edge based upwinding procedure.

8. Discussion of some optimization results. SI units are suppressed as above. Two case studies were performed in rectangular reservoirs $\Omega = [0, a_1] \times [0, a_2]$ of width $a_1 = 400$ and height $a_2 = 100$, having the Neumann boundary $\Gamma_n = [0, a_1] \times \{0, a_2\}$ on top and bottom and the Dirichlet boundary $\Gamma_d = \{0, a_1\} \times [0, a_2]$ on the side strips. The Dirichlet conditions are chosen as $S_w = 1, p_n = 8 \cdot 10^6$ on top (otherwise hydrostatic) and $X \approx 10^{-4}$, so that complementarity $g(y) = 0$ holds on this boundary.

The Sundance simulation is done on a triangular grid with $n_1 \times n_2$ uniform rectangles in Ω (two triangles per rectangle). Let us however stress that our implementation supports unstructured grids. For both case studies we applied a horizontal domain decomposition with $n_p = 16$ parallel MPI processes on a mesh with $n_1 = 128$ and $n_2 = 32$.

8.1. Case study for $N = 5$ wells with residual-solubility trapping objective. The model reservoir is supplied with a heterogeneous isotropic absolute permeability distribution $\mathbf{K}(x) = k_0(x)\mathbf{id}$ and with $N = 5$ injection wells. The function $k_0(x)$ is illustrated in Figure 8.1: The permeability increases from left to right and has a horizontal layer of significantly lowered permeability. The five wells are ordered in a trapezoidal pattern, while the two upper wells are located within the blue layer of low permeability. The optimization was run for

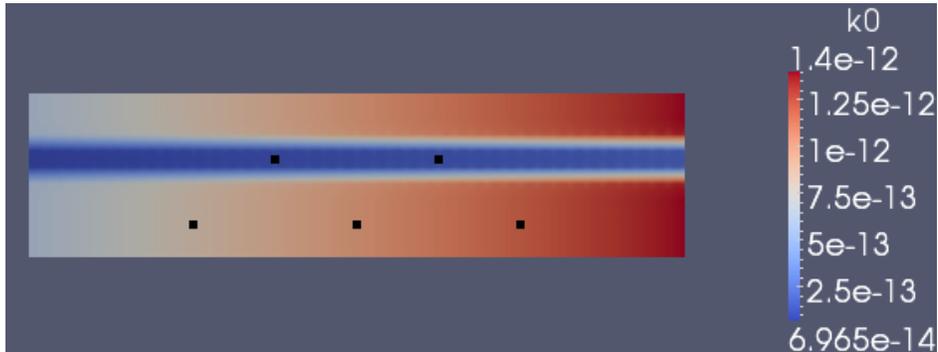


FIG. 8.1. Permeability distribution k_0 and injection well locations.

$T = 4$ months with $n_t = 200$ time steps, while interpolation and averaging periods were set

to $K = 5$ and $M = 10$, and has converged up to an IPOPT tolerance of 10^{-5} , in our case corresponding to a dual infeasibility norm (see [45] for precise definitions); the objective function was scaled by -10^{-3} in the interface, since IPOPT minimizes and expects problem-adjusted scaling. We set $\beta_2 = \beta_3 = 1$ and $\beta_0 = \beta_1 = 0$ in (2.7) to combine residual and solubility trapping in the objective, while the regularization parameter is $\gamma = 7.5 \cdot 10^{-4}$ and the further weight function $\theta(t) = 1.6 \cdot 10^3 \left(1 + \frac{4t}{T}\right)$ in (2.8) penalizes late injection. Given a scaling of $q_* = 200$ in (4.2), the control constraints

$$0 \leq u_i \leq 9 \quad \text{and} \quad c(u) = \frac{1}{k_t} \sum_{i=1}^{Nk_t} u_i \leq 18$$

make sure that at least two wells must be active for maximal injection $c(u) = 18$.

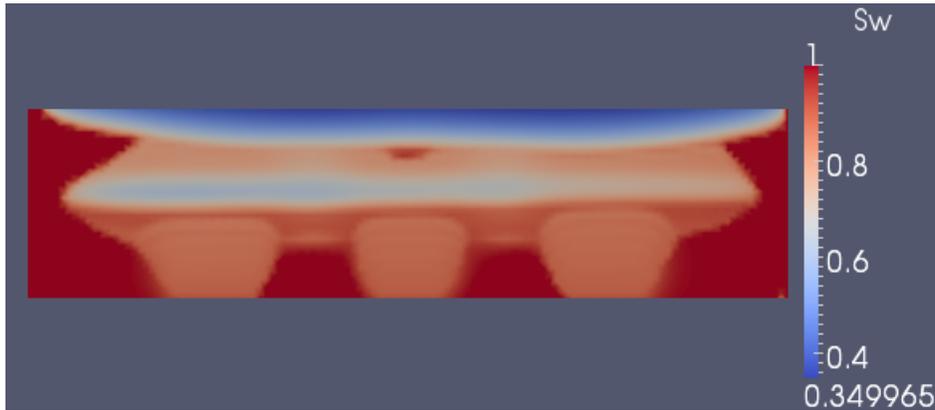
In Figures 8.2–8.4 we show the final saturation profiles for the computed optimal injection strategy in comparison to an early bang-bang strategy – i.e., inject the maximal amount in all wells for the first 80 time steps, then switch injection off – and to uniform injection with $c(u) = 18$. The optimal control has $c(u) \approx 17.318 < 18$, so the upper bound is not active in this situation. The objective functional values are shown in Table 8.1. According to these

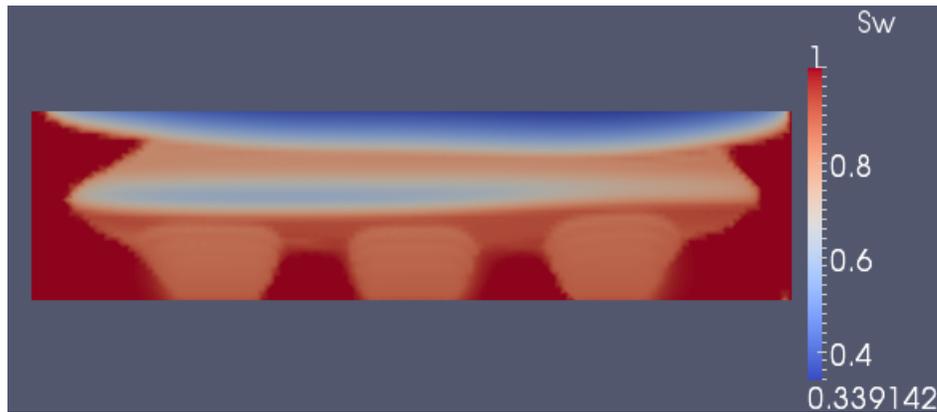
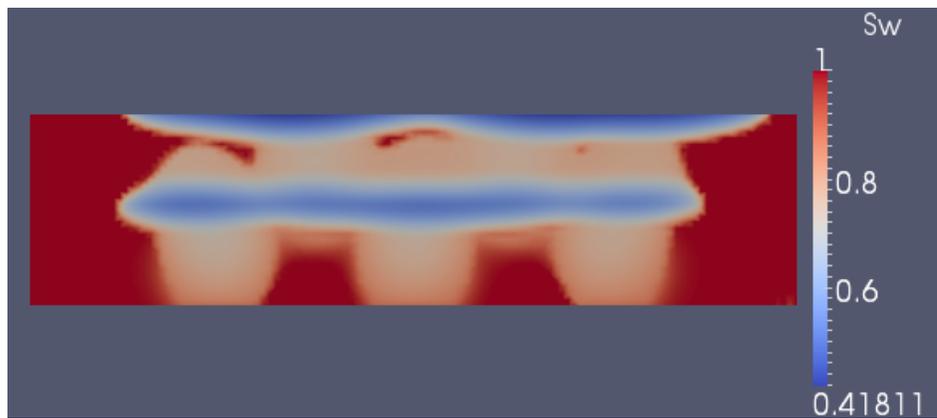
TABLE 8.1

Objective function values for the computed optimal control of subsection 8.1 and for two alternative suboptimal strategies (higher values are better). For the optimal control we give the objective function value for both with regularization term (J) and without regularization term ($J|_{\gamma=0}$).

strategy ($N = 5$)	J	$J _{\gamma=0}$
optimal control	$2.4709 \cdot 10^5$	$2.4747 \cdot 10^5$
early bang-bang	—	$2.4441 \cdot 10^5$
uniform ($u_i = 3.6$)	—	$1.9797 \cdot 10^5$

values, the early bang-bang strategy comes close to the optimal solution and is significantly superior to uniform injection. Also the Figures 8.2 and 8.3 look rather similar, except for the red brine hole below the upper CO_2 plume in Figure 8.2. Notice that, in contrast to the uniform strategy from Figure 8.4, the three lower “pillars” of nonwetting supercritical CO_2 in Figures 8.2 and 8.3 are cut off from the upper portion of the supercritical plume – here the CO_2 is immobilized below its residual saturation $S_n^r = 0.1$.

FIG. 8.2. Final saturation S_w for optimal control.

FIG. 8.3. Final saturation S_w for early bang-bang.FIG. 8.4. Final saturation S_w with uniform $u_i = 3.6$.

The optimal injection rates in the upper and lower wells are graphed in the respective Figures 8.5 and 8.6. Here we recognize the following tendencies: All rates come rather close to “on-off” bang-bang strategies, where the upper wells switch back to maximal injection shortly before the end. This can be understood as a consequence of their location in the blue layer of low permeability: Parts of the final portion of injected CO_2 can thus remain below residual saturation, essentially due to a slower CO_2 plume evolution.

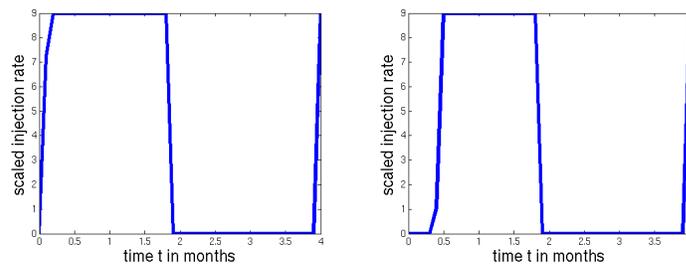
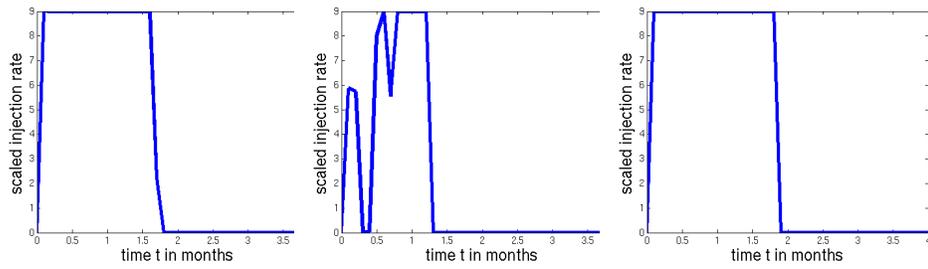


FIG. 8.5. Optimal rates in the upper two wells.

FIG. 8.6. *Optimal rates in the lower three wells.*

Eventually, we illustrate the effect of averaging (here applied every $M = 10$ time steps) in the given scenario by comparison of the early bang-bang simulation and a corresponding simulation without averaging. Figures 8.7 and 8.8 show the respective profiles of the CO_2 mass fraction X in the brine phase at the final time $T = 4$ months.

FIG. 8.7. *Final mass fraction X for early bang-bang.*FIG. 8.8. *Final mass fraction X without averaging.*

Notice that X drops below zero in both simulations, which should certainly not be the case from a physical point-of-view. But the dual cell averaging significantly ameliorates this

problem: The minimal value $X_{\min} = -0.00283$ in Figure 8.7 is rather close to zero, whereas one has $X_{\min} \approx -0.74 X_{\max}$ in Figure 8.8 without averaging. Graphically, several blue negative fluctuations of X appear close to the left and right interface regions between CO_2 plume and brine in Figure 8.8, while the CO_2 -brine interface in Figure 8.7 is sharper and not overly surrounded by such unphysical fluctuations.

We conclude this subsection by some aspects concerning robustness. In fact, several parameters in the model — such as porosity or permeability — are not accessible by direct measurements. In practice, the model thus would contain errors in the parameters. To study the stability of our computed optimal control with respect to model perturbations, we randomly generated 10 realizations of perturbed ϕ - and k_0 -fields in the model. To this end we add a uniformly distributed random perturbation in each grid point of the discrete ϕ - and k_0 -field. More precisely, these coefficients were multiplied in each grid point by iid random variables that were uniformly distributed on $[1 - \eta/100, 1 + \eta/100]$. We considered the cases $\eta = 10$ and $\eta = 20$ (%). For the 10 realizations with 20% perturbation of ϕ and k_0 , the objective function values without regularization term, evaluated at the computed optimal control u^* of the unperturbed problem, result in the sample mean value

$$\text{mean}(J(u^*)|_{\gamma=0}^{\text{pert}=20\%}) = 2.4692 \cdot 10^5 = 0.9978 \cdot J(u^*)|_{\gamma=0}^{\text{unpert}}.$$

Here, $J(u^*)|_{\gamma=0}^{\text{unpert}} = 2.4747 \cdot 10^5$ is the optimal objective value for the unperturbed ϕ - and k_0 -field (see Table 8.1). The sample standard deviation is

$$\text{std}(J(u^*)|_{\gamma=0}^{\text{pert}=20\%}) = 397.09 = 0.0016 \cdot J(u^*)|_{\gamma=0}^{\text{unpert}}.$$

The worst case objective function value among all 10 perturbations is

$$J(u^*)|_{\gamma=0}^{\text{wc.pert}=20\%} = 2.4640 \cdot 10^5 = 0.9957 \cdot J(u^*)|_{\gamma=0}^{\text{unpert}}.$$

We see that the optimal control is quite stable with respect to larger variations. The objective function under perturbation is still much closer to the unperturbed optimal value than to the unperturbed objective value for the other strategies listed in Table 8.1.

For $\eta = 10$ % we obtain the following smaller deviations from $J(u^*)|_{\gamma=0}^{\text{unpert}} = 2.4747 \cdot 10^5$:

$$\text{mean}(J(u^*)|_{\gamma=0}^{\text{pert}=10\%}) = 2.4721 \cdot 10^5 = 0.9989 \cdot J(u^*)|_{\gamma=0}^{\text{unpert}}.$$

The sample standard deviation is

$$\text{std}(J(u^*)|_{\gamma=0}^{\text{pert}=10\%}) = 204.85 = 0.00083 \cdot J(u^*)|_{\gamma=0}^{\text{unpert}}.$$

The worst case objective function value among all 10 perturbations is

$$J(u^*)|_{\gamma=0}^{\text{wc.pert}=10\%} = 2.4695 \cdot 10^5 = 0.9979 \cdot J(u^*)|_{\gamma=0}^{\text{unpert}}.$$

In a further test, we considered the sensitivity of the optimal value with respect to the smoothing parameter ε for the complementarity condition. Our choice $\varepsilon = 3 \cdot 10^{-4}$ is already quite small and we expect almost no change of the objective function value if we further decrease ε . In fact for $\varepsilon = 10^{-4}$ (one third of our default value) we obtain $J(u^*)|_{\gamma=0}^{\varepsilon=10^{-4}} = 2.4747 \cdot 10^5$ and, thus, there is no change in the first 5 digits. For $\varepsilon = 6 \cdot 10^{-4}$ (twice our default value), we also obtain $J(u^*)|_{\gamma=0}^{\varepsilon=6 \cdot 10^{-4}} = 2.4747 \cdot 10^5$. For $\varepsilon = 3 \cdot 10^{-3}$ (10 times our default value), the corresponding objective function value is $J(u^*)|_{\gamma=0}^{\varepsilon=3 \cdot 10^{-3}} = 2.4744 \cdot 10^5$. The deviations from the value $2.4747 \cdot 10^5$ in Table 8.1 are thus only visible in the last digit for the

10 times larger ε -value. This shows that the smoothing of the complementarity function does not introduce an error of any relevant size.

Summarizing our perturbation and sensitivity studies, we observe that the computed optimal control is, in terms of the resulting objective value, remarkably robust with respect to model parameter variations. Furthermore, modifications of the smoothing parameter ε show that our choice affects the complementarity condition in the model so little that only the largest variation affects the 5th digit of the objective value.

8.2. Case study for $N = 3$ wells with weak area trapping objective. Again the reservoir is supplied with heterogeneous permeability $\mathbf{K}(x) = k_0(x)\mathbf{id}$, now containing $N = 3$ injection wells below an arc-shaped region with lowered permeability. The optimization was run for $T = 6$ months with $n_t = 300$ time steps, interpolation and averaging set to $K = 6$ and $M = 15$. Concerning the trapping objective, we have $\beta_1 = 1$ and $\beta_0 = \beta_2 = \beta_3 = 0$ in (2.7) with regularization settings $\gamma = 2 \cdot 10^{-4}$ and $\theta(t) = 2.5 \cdot 10^3 \left(1 + \frac{4t}{T}\right)$. So we maximize the amount of supercritical CO_2 that is trapped below the blue arc at $t = T$ (inside a half-disc region V , see Figure 8.9 for details). Constraints for the injection wells are

$$0 \leq u_i \leq 9 \quad \text{and} \quad c(u) \leq 9,$$

so at least maximal injection in one well is required for the upper bound $c(u) = 9$.

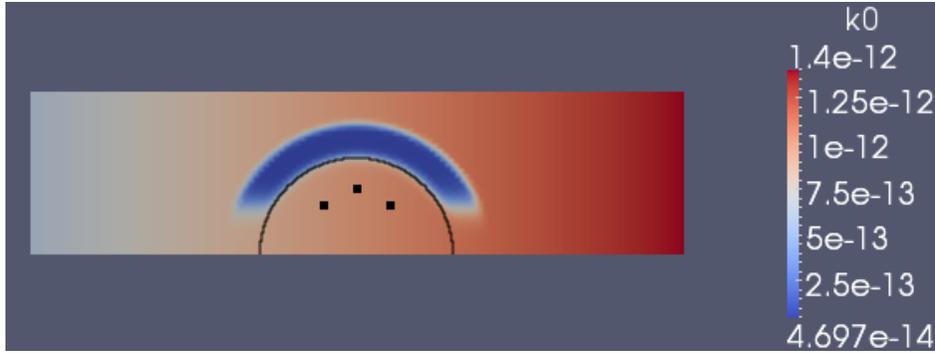
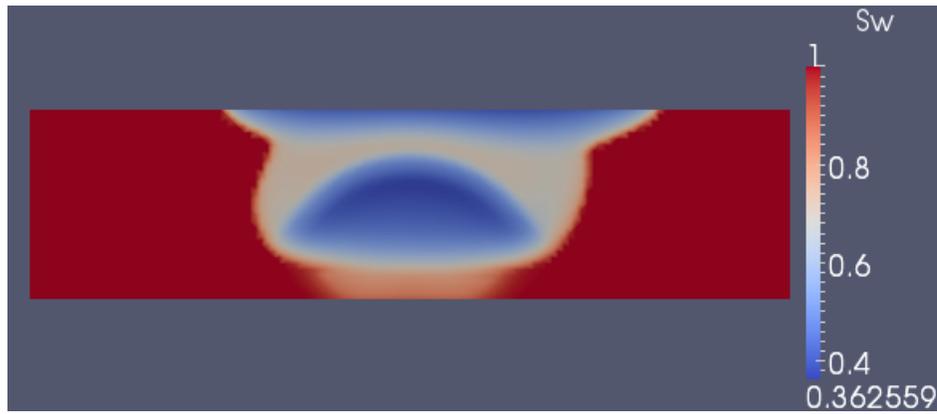
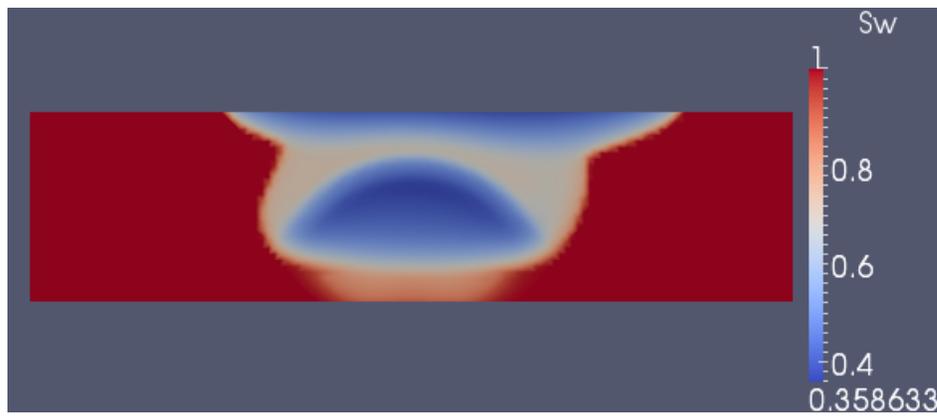
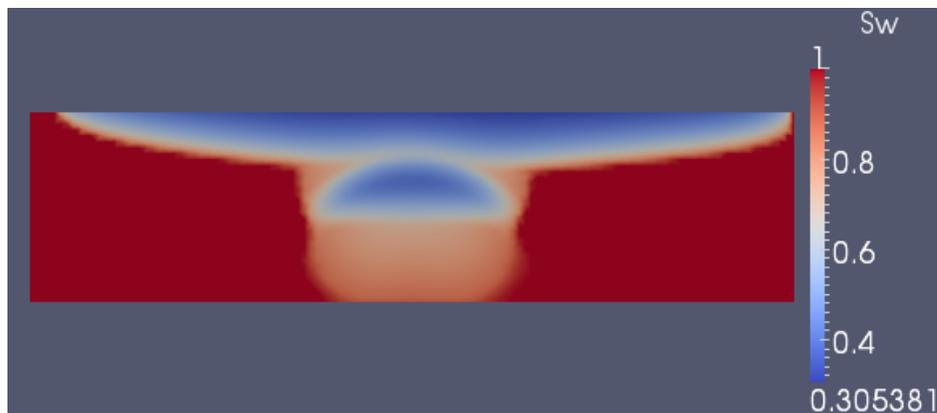


FIG. 8.9. Permeability distribution k_0 , well locations and trapping area V .

The IPOPT problem was again scaled by -10^{-3} , whereas its tolerance was sharpened to 10^{-6} . The IPOPT iteration was started from a late bang-bang control and has converged in a decent way, requiring less than 100 iterates. This acceptable convergence behavior is mainly due to the better conditioning of the weak area trapping objective, in comparison to residual trapping. Therefore we could also impose a slightly weaker penalization than in Subsection 8.1. In Figures 8.10–8.13 we illustrate the final saturation profiles for optimal injection, late bang-bang, uniform injection and maximal injection purely in the central well – the upper bound $c(u) = 9$ being active in all four strategies.

The saturation profiles for uniform and pure central injection (Figures 8.12 and 8.13) differ significantly from the other two pictures: Here most of the initially injected CO_2 has already passed through the blue arc and risen to the caprock. The late bang-bang strategy comes quite close to being optimal, but the linear increase in permeability from left to right makes way for slightly better strategies. The objective functional values are listed in Table 8.2. This shows that the late bang-bang approach almost doubles the objective from the uniform strategy, while the optimal strategy tops it by another 1.2 %. Anyway, uniform distributed injection is superior to pure central injection. The difference between optimal

FIG. 8.10. Final saturation S_w for optimal control.FIG. 8.11. Final saturation S_w for late bang-bang.FIG. 8.12. Final saturation S_w with uniform $u_i = 3$.

(Figure 8.10) and late bang-bang (Figure 8.11) control in the half-disc region V is hard to recognize from the saturation profiles, but the CO_2 plume above the trapping arc is slightly

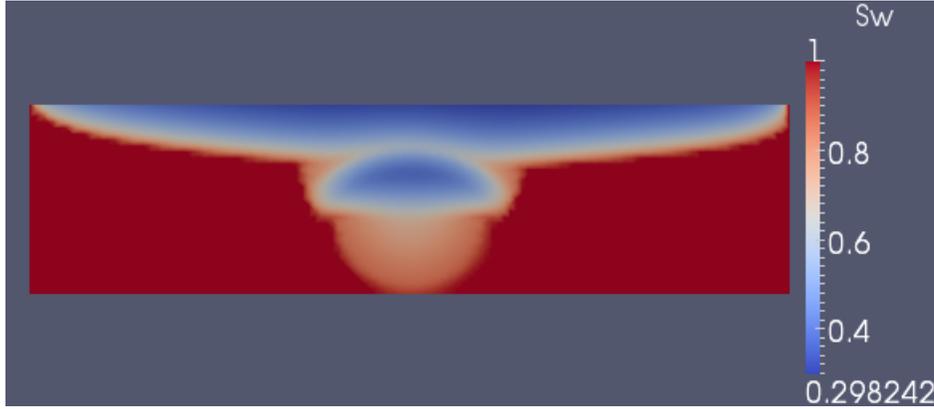


FIG. 8.13. Final saturation S_w for pure central injection.

TABLE 8.2

Objective function values for the computed optimal control of subsection 8.2 and for three alternative suboptimal strategies (higher values are better). For the optimal control we give the objective function value for both with regularization term (J) and without regularization term ($J|_{\gamma=0}$).

strategy ($N = 3$)	J	$J _{\gamma=0}$
optimal control	$2.1006 \cdot 10^5$	$2.1023 \cdot 10^5$
late bang-bang	—	$2.0769 \cdot 10^5$
uniform ($u_i = 3$)	—	$1.1702 \cdot 10^5$
purely central	—	$8.8938 \cdot 10^4$

narrower in the optimal picture. In the optimal rates (Figure 8.14) this is reflected by a short time of maximal initial injection in the left well, together with a delayed start of injection in the right well in comparison to late bang-bang (i.e., after $t = 4$ months). Overall, the lower permeable left of the reservoir is favored to the right, which makes sense as the injected CO_2 is to be kept below the trapping arc as long as possible (at least until $t = T$).

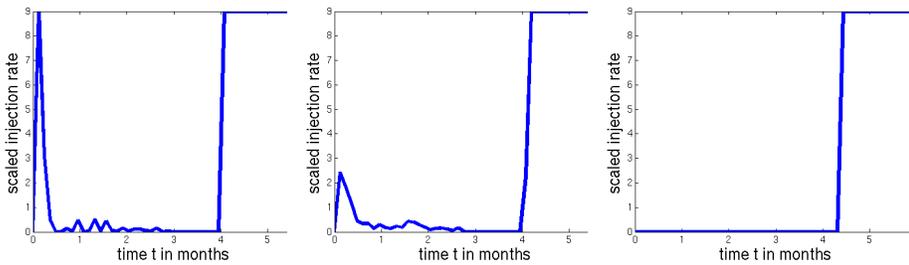


FIG. 8.14. Optimal injection rates (from left to right).

The bang-bang like structure – i.e., either maximal or minimal injection in every well at all times (“on-off”) - of the optimal strategies in Subsections 8.1 and 8.2 is supported by the recent article [47]: There it is shown that, given optimal control of EOR water-flooding with immiscible two-phase flow, such bang-bang solutions may occur if the control constraints on the injection rates are linear. Despite some deviations in the rates of the respective central wells, our optimization results for miscible two-phase flow agree with this observation. We

expect that, given more exact convergence of the IPOPT iterations, the bang-bang nature would also be apparent in the respective central wells.

Acknowledgements. The support from Award No. UK-C0020, made by King Abdullah University of Science and Technology (KAUST) is gratefully acknowledged. This work was conducted as part of the MAC-KAUST project K1 “Simulating CO₂ Sequestration” within the Munich Centre of Advanced Computing (MAC) at TUM. The computations were performed on a compute cluster that was partially funded by DFG INST 95/919-1 FUGG. Finally, we thank the referees for their valuable suggestions that helped us to improve the quality of the article.

REFERENCES

- [1] M. Alexe, A. Sandu: *On the discrete adjoints of adaptive time stepping algorithms*; J. Comput. Appl. Math. 233(4), 1005–1020, 2009.
- [2] F. Allgöwer, R. Findeisen, C. Ebenbauer: *Nonlinear model predictive control*; Encyclopedia for Life Support Systems (EOLSS) article contribution 6.43.16.2, 2003.
- [3] H. Asheim: *Maximization of water sweep efficiency by controlling production and injection rates*; SPE paper 18365, presented at “SPE European Petroleum Conference”, London, 1988.
- [4] P. Bastian: *Numerical Computation of Multiphase Flows in Porous Media*; Habilitation thesis, Christian-Albrechts-Universität Kiel, 1999.
- [5] I. Ben Gharbia, J. Jaffre: *Gas phase appearance and disappearance as a problem with complementarity conditions*; Mathematics and Computers in Simulation 99, 28–36, 2014.
- [6] J. Benk, M. Mehl, M. Ulbrich: *Sundance PDE solvers on Cartesian fixed grids in complex and variable geometries*; in Proceedings of “ECCOMAS Thematic Conference on CFD and Optimization”, Antalya, 2011.
- [7] A. Bielinski: *Numerical Simulation of CO₂ Sequestration in Geological Formations*; Dissertation, Universität Stuttgart, 2006.
- [8] C. Brandenburg, F. Lindemann, M. Ulbrich, S. Ulbrich: *A continuous adjoint approach to shape optimization for Navier Stokes flow*; in K. Kunisch et al. (eds.), *Optimal Control of Coupled Systems of Partial Differential Equations*, Int. Ser. Numer. Math. 158, Birkhäuser, Basel, 35–56, 2009.
- [9] C. Brandenburg, F. Lindemann, M. Ulbrich, S. Ulbrich: *Advanced numerical methods for PDE constrained optimization with application to optimal design in Navier Stokes flow*; in S. Engell et al. (eds.), *Constrained Optimization and Optimal Control for Partial Differential Equations*, Birkhäuser, Basel, 257–275, 2011.
- [10] R.H. Brooks, A.T. Corey: *Hydraulic properties of porous media*; Hydrology Papers 3, Colorado State University, Fort Collins, 1964.
- [11] D.R. Brouwer: *Dynamic Water Flood Optimization with Smart Wells Using Optimal Control Theory*; PhD thesis, Delft University of Technology, 2004.
- [12] D.R. Brouwer, J.D. Jansen: *Dynamic optimization of water flooding with smart wells using optimal control theory*; SPE J. 9(4), 391–402, 2004.
- [13] M.A. Celia, P. Binning: *A mass-conservative numerical solution for two-phase flow in porous media with application to unsaturated flow*; Water Resour. Res. 28(10), 2819–2828, 1992.
- [14] Z. Chen, G. Huan, Y. Ma: *Computational Methods for Multiphase Flows in Porous Media*; Computational Science and Engineering Series, vol. 2, SIAM, 2006.
- [15] N. Dolle, D.R. Brouwer, J.D. Jansen: *Dynamical optimization of water flooding with multiple injectors and producers using optimal control theory*; in Proceedings of “14th International Conference on Computational Methods in Water Resources”, Delft, 2002.
- [16] M.U. Enriquez: *The Effects of Coupling Adaptive Time-Stepping and Adjoint-State Methods for Optimal Control Problems*; PhD thesis, Rice University, Houston, Texas, USA.
- [17] A. Fischer: *Solution of monotone complementarity problems with locally Lipschitzian functions*; Math. Program. 76(3B), 513–532, 1997.
- [18] G. Gao, A.C. Reynolds: *An improved implementation of the LBFGS algorithm for automatic history matching*; SPE J. 11(1), 5–17, 2006.
- [19] M.S. Gockenbach, W.W. Symes: *Adaptive simulation, the adjoint state method, and optimization*; in L.T. Biegler et al. (eds.), *Large Scale PDE-Constrained Optimization*, Springer, 281–297, 2003.
- [20] L. Grüne, J. Panneck: *Nonlinear Model Predictive Control*; Communications and Control Engineering Series, Springer, 2011.
- [21] R. Helmig: *Multiphase Flow and Transport Processes in the Subsurface*; Springer, 1997.
- [22] R. Helmig, J. Niessner, H. Class: *Recent advances in finite element methods for multiphase flow processes in*

- porous media*; Int. J. Comput. Fluid Dyn. 20(3), 245–252, 2006.
- [23] M.A. Heroux, J.M. Willenbring, R. Heaphy: *Trilinos developers guide*; Sandia National Laboratories, SAND 2003–1898, 2003.
- [24] M. Hinze, R. Pinnau, M. Ulbrich, S. Ulbrich: *Optimization with PDE Constraints*; Mathematical Modelling: Theory and Applications, vol. 23, Springer, 2008.
- [25] M. Hinze, J. Sternberg: *A-revolve: an adaptive memory-reduced procedure for calculating adjoints; with an application to computing adjoints of the stationary Navier-Stokes system*; Optim. Methods Softw. 20(6), 645–663, 2005.
- [26] R. Huber, R. Helmig: *Multiphase flow in heterogeneous porous media: A classical finite element method versus an implicit pressure–explicit saturation-based mixed finite element–finite volume approach*; Int. J. Numer. Methods Fluids 29(8), 899–920, 1999.
- [27] J.D. Jansen: *Adjoint-based optimization of multi-phase flow through porous media – a review*; Comput. Fluids 46(1), 40–51, 2011.
- [28] D. Korounis, D. Voskov, K. Aziz: *Adjoint methods for multicomponent flow simulations*; in Proceedings of “ECMOR XII: European Conference on Mathematics of Oil Recovery”, Oxford, 2010.
- [29] A. Lauser, C. Hager, R. Helmig, B. Wohlmuth: *A new approach for phase transitions in miscible multi-phase flow in porous media*; Adv. Water Resour. 34(8), 957–966, 2011.
- [30] S. Li, L. Petzold: *Adjoint sensitivity analysis for time-dependent partial differential equations with adaptive mesh refinement*; J. Comput. Phys. 198(1), 310–325, 2004.
- [31] F. Lindemann: *Theoretical and Numerical Aspects of Shape Optimization with Navier-Stokes Flows*; Dissertation, Technische Universität München, 2012.
- [32] K.R. Long: *Sundance rapid prototyping tool for parallel PDE optimization*; in L.T. Biegler et al. (eds.), Large scale PDE-constrained optimization, Springer, 331–342, 2003.
- [33] K.R. Long, P.T. Boggs, B.G. van Bloemen Wanders: *Sundance: High-level software for PDE-constrained optimization*; SAND report, preprint, 2012.
- [34] G.J. Mehos, W.F. Ramirez: *Use of optimal control theory to optimize carbon dioxide miscible-flooding enhanced oil recovery*; J. Petroleum Sci. Eng. 2(4), 247–260, 1989.
- [35] I. Neitzel, F. Tröltzsch: *On convergence of regularization methods for nonlinear parabolic optimal control problems with control and state constraints*; Control and Cybernetics 37(4), 1013–1043, 2008.
- [36] R. Neumann, P. Bastian, O. Ippisch: *Modeling and simulation of two-phase two-component flow with disappearing nonwetting phase*; Computational Geosciences 17(1), 139–149, 2013.
- [37] D.S. Oliver, A.C. Reynolds, N. Liu: *Inverse Theory for Petroleum Reservoir Characterization and History Matching*; Cambridge University Press, 2008.
- [38] M. Peszynska, E.W. Jenkins, M.F. Wheeler: *Boundary conditions for fully implicit two-phase flow models*; in X. Feng, T.P. Schulze (eds.), Recent Advances in Numerical Methods for Partial Differential Equations and Applications, Contemporary Mathematics Series 306, AMS, 85–106, 2002.
- [39] P. Sarma, K. Aziz, L.J. Durlofsky: *Implementation of adjoint solution for optimal control of smart wells*; SPE paper 92864, presented at “SPE Reservoir Simulation Symposium”, Houston, 2005.
- [40] M. Simon, M. Ulbrich: *Optimal control of partially miscible two-phase flow with applications to subsurface CO₂ sequestration*; in M. Bader et al. (eds.), Advanced Computing, Lecture Notes in Computational Science and Engineering 93, Springer, 81–98, 2013.
- [41] B. Sudaryanto, Y.C. Yortsos: *Optimization of fluid front dynamics in porous media using rate control*; Phys. Fluids 12(7), 1656–1670, 2000.
- [42] B. Sudaryanto, Y.C. Yortsos: *Optimization of displacements in porous media using rate control*; SPE paper 71509, presented at “SPE Annual Technical Conference and Exhibition”, New Orleans, 2001.
- [43] M. Ulbrich: *Semismooth Newton Methods for Variational Inequalities and Constrained Optimization Problems in Function Spaces*; MOS-SIAM Series on Optimization, vol. 11, SIAM, 2011.
- [44] G.A. Virmovski: *Water flooding strategy design using optimal control theory*; in Proceedings of “6th European Symposium on IOR”, Stavanger, 1991.
- [45] A. Wächter, L.T. Biegler: *On the implementation of a primal-dual interior point filter line search algorithm for large-scale nonlinear programming*; Math. Program. 106, 25–57, 2006.
- [46] I.S. Zakirov, S.I. Aanonsen, E.S. Zakirov, B.M. Palatnik: *Optimization of reservoir performance by automatic allocation of well rates*; in Proceedings of “ECMOR V: European Conference on Mathematics of Oil Recovery”, Leoben, 1996.
- [47] M.J. Zandvliet, O.H. Bosgra, J.D. Jansen, P.M.J. van den Hof, J.F.B.M. Kraaijewanger: *Bang-bang control and singular arcs in reservoir flooding*; J. Petroleum Sci. Eng. 58(1), 186–200, 2007.
- [48] Q. Zhang: *A finite difference–streamline diffusion (FSD) method and its error estimates for two-phase incompressible miscible flow in porous media*; Acta Math. Appl. Sin. 26(2), 318–327, 2003.
- [49] C. Zhu, R.H. Byrd, J. Nocedal: *L-BFGS-B: Algorithm 778: L-BFGS-B, FORTRAN routines for large scale bound constrained optimization*; ACM Transactions on Mathematical Software 23(4), 550–560, 1997.
- [50] CO₂ trapping mechanisms, www.co2captureproject.org/co2_trapping.html.