A simulation workflow for large-scale CO_2 storage in the Norwegian North Sea

K.-A. Lie* (SINTEF) H. M. Nilsen (SINTEF) O. Andersen (SINTEF)
O. Møyner (SINTEF)

October 14, 2014

Abstract

Large-scale CO_2 injection problems have revived the interest in simple models, like percolation and vertically-averaged models, for simulating fluid flow in reservoirs and aquifers. A series of such models have been collected and implemented together with standard reservoir simulation capabilities in a high-level scripting language as part of the open-source MATLAB Reservoir Simulation Toolbox (MRST) to give a set of simulation methods of increasing computational complexity. Herein, we outline the methods and discuss how they can be combined to create a flexible tool-chain for investigating CO_2 storage on a scale that would have significant impact on European CO_2 emissions. In particular, we discuss geometrical methods for identifying structural traps, percolation-type methods for identifying potential spill paths, and vertical-equilibrium methods that can efficiently simulate structural, residual, and solubility trapping in a thousand-year perspective. The utility of the overall workflow is demonstrated using real-life depth and thickness maps of two geological formations from the recent CO_2 Storage Atlas of the Norwegian North Sea.

Introduction

The sedimentary basins in the Norwegian North Sea contain a large number of saline aquifers that have small flow rates and offer large volumes of pore space that can be used to store carbon dioxide produced from localized sources. To store carbon geologically, CO_2 will be injected as a supercritical liquid into a high-permeable strata that is limited upward by a low-permeable strata (caprock) that inhibits flow. The injected CO_2 has lower density than the formation fluids and will form a plume that migrates upward by buoyancy forces. Aquifers are typically connected to the surface through permeable strata, and the injected CO_2 may therefore in principle travel in the up-dip direction and eventually leak back to the atmosphere through sedimentary outcrops. In practice, this process will take thousands of years because of the long distances involved. Moreover, as the plume migrates upward, some of the CO_2 will be contained in structural and stratigraphic traps (structural trapping), be trapped as small droplets between rock grains (residual trapping) or dissolve into the formation water (dissolution trapping), or react with rock minerals and become permanently trapped.

The main concern for policy makers and the general public is risk of leakage: how likely is it that the injected CO_2 will leak back to the surface or migrate into active petroleum reservoirs. In other parts of the world, one may also be concerned of leakage of CO_2 or highly saline brine into water resources. In other words, the operator of a potential injection site needs to maximize storage volumes while minimizing leakage risks and undesired effects on areas surrounding the injection points. The operator will obviously also want to ensure operational safety and minimizing financial costs. Similar assessments will be desired by companies, investors, and/or government (agencies) that take an environmental, societal, or financial risk through the operation. The only viable way to make such assessments upfront is through model studies that aim to investigate the likely outcomes of a storage operation. The main controls in a model study are the aquifer geology and the physics of the flow processes.

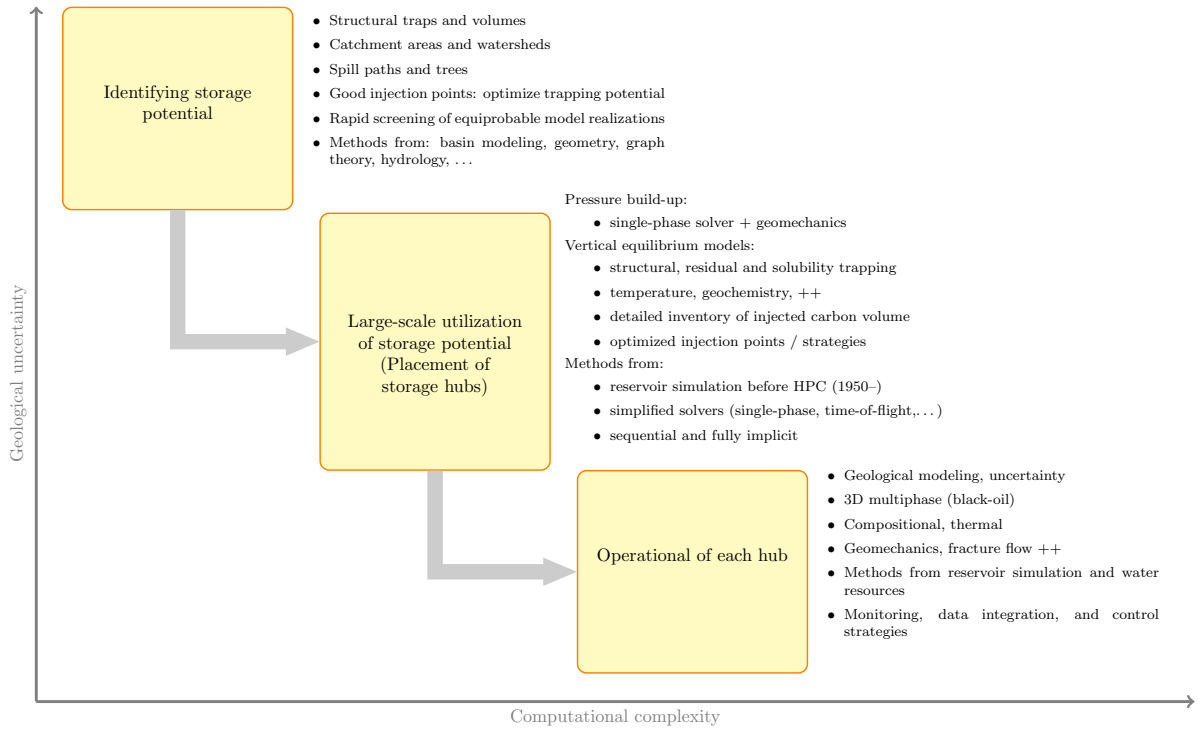

Figure 1 *Conceptual picture of how different modeling tools and approaches could be used in the various phases of the planning of a large-scale storage operation.*

From petroleum exploration and production, one knows that variations in geological structures and rock properties will strongly impact the migration of hydrocarbons on a basin scale as well as the flow on a reservoir scale. Unfortunately, the understanding of most saline aquifers is quite limited: seismic surveys have limited coverage and few core-samples are available since aquifers seldom have been penetrated by more than a few wells, if any. In fact, aquifers with fewer wells present may be better suited for CO_2 storage, as the risk of leakage is reduced. It is therefore very important that modeling tools are able to properly account for the impact of uncertainty in the geological description of the aquifer to correctly span the range of likely outcomes of an injection operation.

The flow dynamics of the injection and migration processes depends on delicate balances between various physical mechanisms. The balances may vary with spatial location and change significantly as time progresses. CO_2 is very mobile and can travel long distances, but the flow is typically confined to thin layers underneath the sealing caprock. Because of the high ratios between the lateral and vertical scales involved, and the large disparity in temporal scales of the different physical processes, modeling CO_2 storage is a very challenging multiscale problem that is best attacked using a range of different tools for flow modeling. Whereas traditional tools for 3D simulation can be used to study the buildup of a plume near the injection point, these tools will in most cases not be able to resolve the long-term, large-scale migration of the CO_2 plume and the associated trapping processes. Indeed, because the time scale of the vertical flow process typically will be much shorter than the time scale of the lateral fluid movement, one will often be better off, both in terms of accuracy and computational efficiency, by describing the migration process in a vertically-averaged sense on a 2D grid that follows the caprock surface, possibly in combination with invasion-percolation computations and similar ideas that have been developed to study migration processes taking place over millions of years.

Figure 1 gives a conceptual illustration of how different modeling tools and approaches could be used in the various phases of the planning of a large-scale storage operation. Herein, we will focus on the two first phases: identification of storage potential and placement of injection hubs for large-scale utilization.

To this end, we will describe a set of simulation tools of increasing computational complexity and discuss how they can be combined to create a flexible tool-chain for investigating CO_2 storage on a scale that would have significant impact on European CO_2 emissions. In doing so, we will consider ideas from computational geometry, basin modeling, hydrology, and reservoir simulation and adapt and combine them in a way that, to the best of our knowledge, is new within CO_2 sequestration modeling. All the tools are implemented in a high-level scripting language as part of the open-source MATLAB Reservoir Simulation Toolbox (MRST; Lie et al., 2012), and made available as a separate CO_2 module, (SINTEF Applied Mathematics, 2013). In addition to the methods discussed herein, MRST also implements standard 3D simulation tools: sequential solvers for (in)compressible, immiscible flow, fully-implicit methods based on automatic differentiation for black-oil models (including certain modeling options for enhanced oil recovery), as well as early prototypes of geochemical, thermal, and geomechanical effects.

Methodology

All computational methods implemented in MRST-co2lab are formulated based on a hybrid 2D grid that represents the 3D aquifer in terms of its top and bottom surfaces, i.e., the 3D surfaces that separate the high-permeable strata of the aquifer from low-permeable strata that bounds it from above and below. These surfaces are represented as depth values at each vertex and edge/cell centroid of a 2D mesh. In addition, each cell in the grid contains information of the petrophysical properties in the volume that lies below it in the 3D representation of the aquifer. The computational methods that will be discussed in the following can roughly be divided into two classes: (i) methods that do not utilize temporal information to identify the potential for structural trapping (Nilsen et al., 2014d), and (ii) methods for estimating the outcomes of injection operations in a long-term, large-scale perspective by simulating the combined effects of structural, residual, and solubility trapping in a vertically averaged sense (Nilsen et al., 2014c,b). The methods are designed to fit together as part of a multi-fidelity tool-chain supporting the flexibility in resolution required of simulations used for decision support. Herein, our focus is on workflow tools that enable interactive inspection of models and rigorous mathematical optimization of injection points and strategies.

Spill-point analysis and structural trapping capacity

In the short term, structural and stratigraphic trapping are the dominant mechanisms for geological storage of CO_2. Structural traps correspond to local maxima in the top surface (see Figure 2), and first-order estimates of the associated storage volumes can be produced quickly by simple geometrical/topological algorithms. In (Nilsen et al., 2014d), we discuss in detail two families of such algorithms that use the depth of the top surface evaluated at the cell centers and cell vertices, respectively, to determine *spill paths* that connect each node (i.e., cell or vertex) to its upslope neighbors and either end up in a local maximum or at the perimeter of the model. The top surface may contain one or more closed regions inside which all spill paths converge to a local maximum. These so-called *spill region* will act as funnels by collecting buoyant fluids within the area covered by the region and channeling them towards the maximum point. As such, spill regions may limit the global sweep efficiency needed to make residual trapping a viable storage mechanism. Each spill region is separated topographically from adjacent spill regions, or regions that spill to the exterior of the model, by a perimeter (hydrology: drainage divide or watershed). All nodes that lie above the highest point on the perimeter, called the *spill point*, are said to belong to a structural trap, which can potentially be used to safely store CO_2, whereas the remaining part of the spill region is said to belong to the trap's catchment area. Based on this analysis, one can provide upper limits on the amount of CO_2 that may be structurally trapped within an aquifer.

To study migration of CO_2, we can use a percolation-type method that assumes that CO_2 is injected at an infinitesimal rate. In the resulting migration model, CO_2 injected at a point within a catchment area will accumulate inside the associated trap and gradually fill it up until the lower surface of the CO_2 extends down to the spill point. When this happens, CO_2 will enter the adjacent spill region and continue its upward migration towards a new local minimum or the model perimeter. A spill region

can obviously be contained within another spill region and/or be linked by spill paths to other non-overlapping spill regions in an upward succession. By nesting the individual spill regions this way, we can define a hierarchical spill system (hydrology: drainage system) that enables us to easily identify potential storage volumes that are upslope of a given point, or vice versa, the catchment areas that are downhill of a point or trap. The upward succession of traps in the spill system can be seen as a set of separate trees (hydrology: drainage basins) that each describe an isolated migration system that starts low in the formation at traps (leaf nodes) that have no downhill connections and ends up at the root, defined as a trap that either has no upslope connections or is connected to the model perimeter. Using these trees, one can quickly suggest injection points that are upward connected to large trap volumes. If structural trapping is the only containment mechanism, the best points for finite-rate injections will lie on the perimeter between two (or more) spill regions that are connected upslope to distinctly different trap trees with large and approximately equal trap volumes.

The primary access in MRST-co2lab to the trapping analysis described above is through a graphical user interface that can be invoked once a proper structure for representing the top-surface grid has been created. However, the software also provides an API to individual functions that implement distinct parts of the analysis, see (Nilsen et al., 2014d) for more details. All functionality is documented through MATLAB's help system and tutorials that accompany the software.

Vertical equilibrium models

Most aquifers that can be considered candidates for carbon storage have spatial extents on the order of tens to hundreds of kilometers in the lateral direction and tens to hundreds of meters in the vertical direction. After the injection has ceased, the driving forces in the horizontal direction will be moderate and fluid redistribution will mainly be caused by buoyant forces caused by high density contrasts between CO_2 and the resident brine. As a result, the lateral movement of fluids will dominate the large-scale dynamics of the CO_2 plume since the vertical flow takes place on a much shorter time scale. As a good approximation, the lateral redistribution of fluids can be considered instantaneous so that the fluid phases are always in vertical equilibrium.

This observation naturally leads to a two-scale modelling approach, in which the three-dimensional flow equations are integrated vertically to form a 2D flow model for the lateral movement of fluids. The effective properties and constitutive relationships that enter the averaged flow equations are in turn determined by hydrostatic phase pressures and analytical expressions that e.g., describe the vertical distribution of fluid phases in vertical equilibrium. This transformation not only reduces the computational cost by reducing the dimension of the problem and increasing the time constants that characterize the dynamics of the model. The analytical expressions provide a subscale resolution that allows us to describe important parts of the plume dynamics that take place on a vertical scale that cannot be accurately resolved by the overly coarse grid resolution that must be imposed to make a 3D model computationally tractable.

For completeness, let us briefly outline the derivation of the most basic form of a vertical-equilibrium (VE) model, which is based on an assumption of a sharp interface separating CO_2 and resident brine. To this end, we start with a standard incompressible, two-phase model

$$\frac{\partial \phi s_\alpha}{\partial t} + \nabla \cdot \vec{v}_\alpha = 0, \qquad \vec{v}_\alpha = -\mathbf{k}\lambda_\alpha\left(\nabla p - \rho_\alpha \vec{g}\right). \tag{1}$$

Here, ϕ denotes porosity, $\mathbf{k}$ permeability, p pressure, and $\vec{g}$ the gravity vector, whereas s_α, ρ_α, and λ_α denote saturation, density, and mobility for phase $\alpha = \{w,n\}$. Brine is assumed to be the wetting phase and CO_2 the non-wetting phase, and these two fluids fill the pore space completely so that $s_w + s_n = 1$.

We assume that the aquifer is bounded above and below by impermeable strata that are separated a distance H. To keep the presentation as simple as possible, we assume that the top surface is flat and perpendicular to $\vec{g}$ and choose our coordinate system such that the z-axis is (almost) aligned with $\vec{g}$. Moreover, we assume that ϕ and $\mathbf{k}$ are constant in the z-direction. The general case with a non-flat

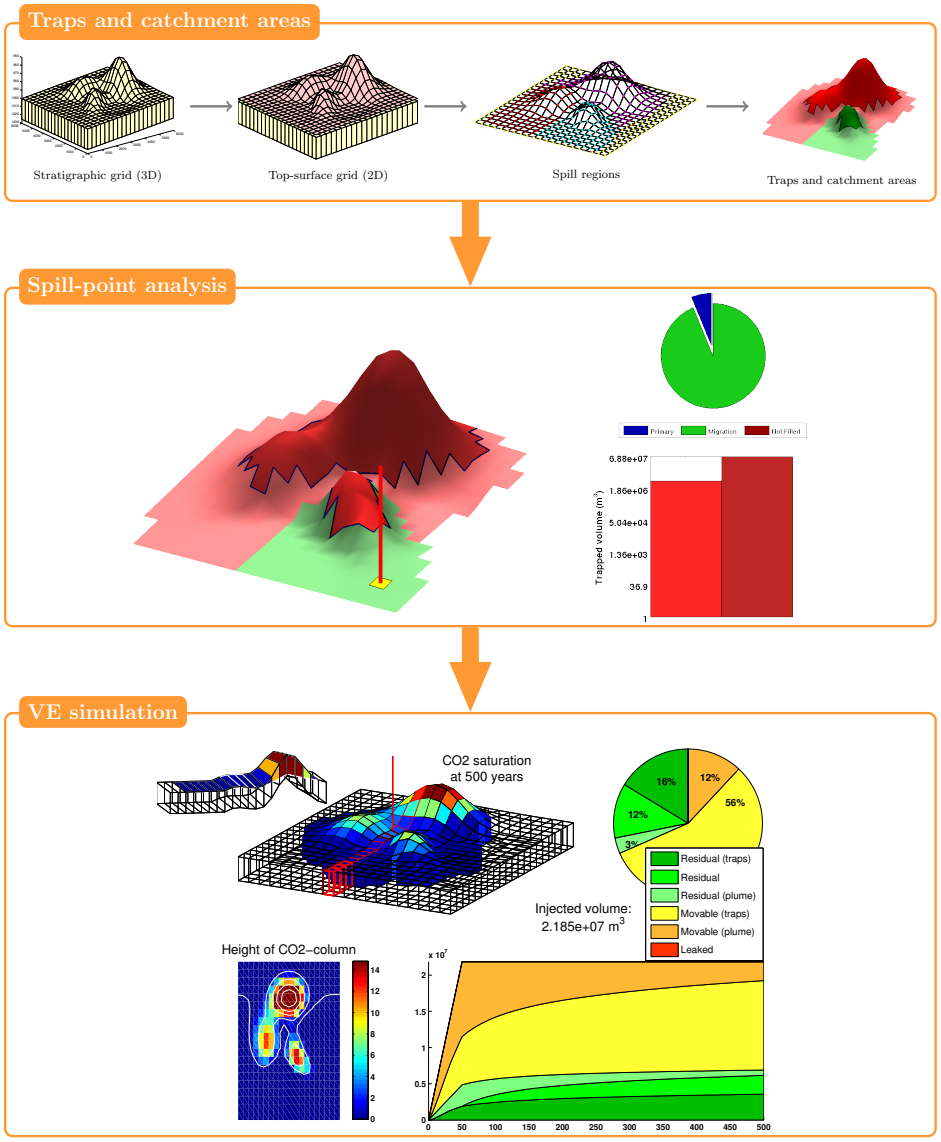

Figure 2 *Illustration of the chain of tools implemented in MRST-co2lab.*

top surface and non-constant petrophysical parameters, and certain geometrical approximations that are tacitly introduced in the following, are discussed in detail in (Nilsen et al., 2014c). By integrating (1) from 0 to H in the z-direction we obtain the coarse-scale flow equations:

$$\frac{\partial \Phi S_\alpha}{\partial t} + \nabla_\| \cdot \vec{V}_\alpha = 0, \qquad \vec{V}_\alpha = -\mathbf{K}\lambda_\alpha(\nabla_\| P - \rho_\alpha \vec{g}_\|), \tag{2}$$

where $\nabla_\|$ denotes a 2D operator, $\vec{g}_\|$ is the gravity component in the (x,y)-plane, and upper-case symbols denote upscaled counterparts of the lower-case symbols used in (1). From (2) one can easily derive a coarse-scale system formulated in terms of P and one of the phase saturations. To determine the constitutive relations entering (2) we need to determine the fluid distribution along the z-axis. First of all, the fluid distribution is assumed to consist of three zones separated by sharp interfaces: CO_2 with residual brine $s_{w,r}$ at the top, brine with residual CO_2 ($s_{n,r}$) below in a zone from which the CO_2 plume has retracted, and pure brine at the bottom:

$$s(z) = \begin{cases} s_{w,r}, & 0 \le z \le h, \\ 1 - s_{n,r}, & h < z \le h_{\max}, \\ 1, & h_{\max} \le z \le H. \end{cases} \tag{3}$$

Moreover, if we take P to be the pressure at the top surface, the hydrostatic pressure reads

$$p(z) = \begin{cases} P + \rho_n g z, & 0 \le z \le h, \\ P + \rho_n g h + \rho_w g (z - h), & h < z \le H. \end{cases} \tag{4}$$

Then, by simple integration we obtain the following relations between h and $h_{\max}$ and the coarse-scale quantities:

$$\begin{aligned} S_n(h, h_{\max}) &= (1 - s_{w,r}) \frac{h}{H} + s_{n,r} \frac{h_{\max} - h}{H} \\ \Lambda_n(h) &= \lambda_n(s_{w,r}) h \\ \Lambda_w(h, h_{\max}) &= \lambda_w(1 - s_{n,r}) \frac{h_{\max} - h}{H} + \lambda_w(1) \frac{H - h_{\max}}{H} \end{aligned} \tag{5}$$

From this expression, we see that the natural variables to use for the upscaled equations are h and the hysteretic parameter $h_{\max}(t) = \max_{s \le}(h(s))$. The model can also be formulated using S as a primary unknown, but then we must always invert the functional form of S to compute h before we can determine the relative mobilities. Finally, we notice that the upscaled model will have hysteretic behavior even if the fine-scale model has not.

In MRST-co2lab we have implemented several types of two-phase VE models (Nilsen et al., 2014c,b), ranging from the simple sharp-interface model outlined above, which can only account for the basic dynamics of structural and residual trapping, to quite sophisticated models that account for compressibility, fine-scale capillary forces, dissolution, hysteretic effects, subscale caprock trapping, etc. The simplest models are formulated both in terms of h and S and are discretized and solved using a sequential splitting method to give high computational efficiency. The more sophisticated methods are, unlike most other VE models reported in recent literature, formulated in terms of S using the black-oil framework that is standard in the petroleum industry. To ensure maximum robustness, the models are discretized and solved using a fully-implicit method as implemented in leading commercial reservoir simulators (e.g., including standard techniques to safeguard the time steps). The fully-implicit method is implemented using automatic differentiation, which enables simple computation of parameter sensitivities and gradients that can be used for mathematical optimization as will be explained later.

Analysis of data from the North Sea CO_2 Storage Atlas

To explore the possibility of large-scale CO_2 storage offshore Norway, the Norwegian Petroleum Directorate (NPD) has produced two volumes of a CO_2 Storage Atlas, one for the Norwegian part of the North Sea (Halland et al., 2011) and one for the Norwegian Sea (Halland et al., 2012). In the combined publication, twenty-seven formations have been grouped into aquifers whose qualities have been assessed with regard to CO_2 storage potential. Similar atlases have been compiled in other parts of the world, as well; see e.g., (Lewis et al., 2008, 2009; NACSA; NATCARB/Atlas; Cloete, 2010; Bradshaw et al., 2011).

As part of MRST-co2lab, we have implemented functionality that provides simple access to the data set published along with the North Sea CO_2 Storage Atlas (Halland et al., 2011). The data set includes formation thickness and depth maps that cover vast scales and hence have a spatial resolution of 500–1000 m. Despite the coarse resolution, the data can still be used to provide indicative estimates of the capacity for structural trapping and simulate likely outcomes of specific injection scenarios. To establish a volumetric grid of an aquifer, the minimal required information is a depth map of the top surface and a map of the formation thickness, or equivalents thereof. Not all formations have both a depth and a thickness map, and when both are present, the data sets are not necessarily fully consistent: coordinates of the depth and thickness maps do not always coincide and sections of the depth map may not be included in the thickness map and vice versa. Using linear interpolation in regions where the scattered

data overlap, we were able to construct volumetric models of fourteen different sand volumes (Nilsen et al., 2014d). In the following, we will use the various tools from MRST-co2lab to analyze two of the formations: the Utsira Formation and the Sandnes Formation.

The Utsira Formation

The world's first commercial CO_2 storage project started in 1996 at the Sleipner West field which is part of the Utsira Formation on the Norwegian Continental Shelf. The injection rate has been up to one metric megaton per year and so far there is no evidence of CO_2 leaking into other horizons or back to the atmosphere. Estimates of the total storage capacity for the Utsira Formation range from 0.3 to 60 billion tonnes (Halland et al., 2011; Thibeau and Mucha, 2011).

The Utsira model derived from the CO_2 Storage Atlas data set has a spatial resolution of 500 m and a total bulk volume of 3800 billion cubic meters, out of which 0.44% or approximately 17 Gm^3 fall inside what can be characterized as structural traps. The annual emissions of CO_2 in countries within the European Union is approximately 3.7 billion tonnes (Olivier et al., 2013). If carbon storage in the Utsira Formation is going to have significant impact on European CO_2 emissions, one must consider storage operations with injection rates in the range of hundreds to thousands of megatonnes per year for the next 30–50 years. It is obviously highly questionable whether one will be able to capture and transport such amounts of CO_2, but for the sake of our analysis we will simply assume that injectable CO_2 volumes are in abundance and try to determine how to best distribute injection hubs throughout the formation to maximize the storage resulting from a 50-year injection period.

As a starting point, we use a simple volumetric argument to compute the upper limit on the total mass that theoretically can be contained by basic trapping mechanisms. First, we approximate the aquifer geometry as a grid of vertical pillars and evaluate an equation-of-state (Span and Wagner, 1996) inside each pillar to determine the corresponding CO_2 density. The full trapping capacity is reached inside each pillar if the pore space in any part of the pillar that lies within a structural trap is fully saturated with CO_2 and the remaining pore space contains the maximal residual CO_2 saturation. Likewise, any brine should contain the maximum amount of dissolved CO_2. Summing over all pillars shows that 1.1 Gt of CO_2 can be contained in structural traps, 77 Gt can be residually trapped, while 34 Gt can be dissolved in the formation brine; see (Andersen et al., 2014) for more details.

To estimate the amounts that can be contained in an actual, large-scale storage operation, we will first look at the spill system and use a simple greedy algorithm to maximize the utilization of structural traps. For each spill tree that has not yet been filled, we pick the leaf node that lies lowest in the formation and place a well as low as possible in the associated catchment area. (In doing so, areas near the boundaries are excluded to avoid pushing injected volumes out of the aquifer during injection). The well is assigned a rate that will fill all the upslope-connected traps in 50 years. If this gives a rate larger than 5 Mt/year, the well is split into multiple wells that are scattered within the same catchment area. If the computed capacity is less than 0.1 Mt/year, the rate is set to 1 Mt/year. All branches of the tree that will not be filled by injecting in the current catchment are cut off and assigned to the list of trees, and the procedure is continued until there are no more trees available.

After the optimization, we have identified a storage potential for 2.5 Gt of CO_2, based on a distribution of wells as shown to the left in Figure 3. To investigate the feasibility of this injection plan, we use a VE simulation to study the pressure buildup and the effect of the various trapping mechanisms. Both the rock and brine are compressible: rock has constant compressibility of $10^{-5}\,\mathrm{bar}^{-1}$, brine has constant compressibility equal to $4.3\times10^{-5}\,\mathrm{bar}^{-1}$ and a reference density of 1020 kg/m^3, whereas the density of CO_2 is assumed to follow a model taken from (Span and Wagner, 1996; Bell et al., 2014) with reference density 760 kg/m^3 and a constant temperature of 35° C. Finally, the viscosities of brine and CO_2 are $8\cdot10^{-4}$ and $6\cdot10^{-4}$ mPa/s, respectively, and the end-point saturations are 0.11 and 0.21.

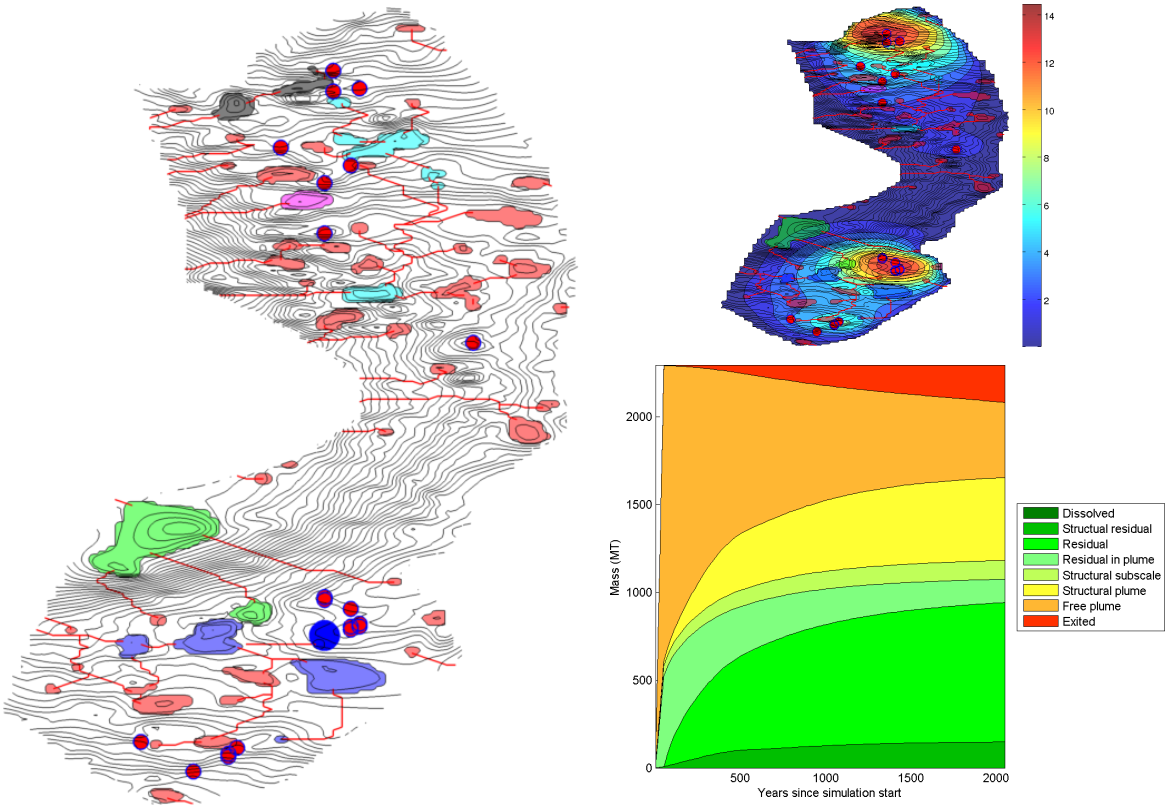

Figure 3 Left: *Distribution of potential injection hubs in the Utsira Formation suggested by greed optimization of the spill system. Structural traps are shown as colored regions (unique color for the five largest trees, red for the other), whereas the spill paths are indicated by red lines.* ***Upper right:*** *Pressure increase in units bar at the end of the injection period.* ***Lower right:*** *Detailed trapping inventory showing the volumetric distribution of CO_2 as function of time over a period of two thousand years after the start of injection. The subdivision into different trapping categories follows the suggestions in (Nilsen et al., 2014c), but with slightly revised labels.*

At the boundary of the domain, Dirichlet boundary conditions corresponding to a hydrostatic pressure distribution are prescribed to simulate an open domain. To compensate for the low spatial resolution in the model, we use an effective (accretion layer) model that accounts for subscale caprock variations, as discussed in (Nilsen et al., 2014c). The input data to this model are estimated by comparing the atlas model to a high-resolution model of the Sleipner injection site, see (Nilsen et al., 2014a) for details. However, solubility trapping is disregarded since there is not yet consensus how this effect should be modeled on Utsira. From Figure 3 we see that pressure buildup will not be a problem for the proposed injection strategy; at the end of the injection period the maximum pressure increase is less than 15 bar. On the other hand, the simulation cannot account for 9% of the injected mass, which has exited the model domain over a period of two thousand years. Likewise, at the end of the simulation 430 Mt of CO_2, which is almost one fifth of injected mass, is still free to move around. Including dissolution will generally reduce the height and retard the movement of the migrating plume (Nilsen et al., 2014c). To provide an upper bound on this effect, we introduce a rate-dependent dissolution model that has been developed for homogeneous media. Specifically, referring to (Nilsen et al., 2014c, Eq. (25)) we set $C/c_{\max} = 0.05$ m/year and $c_{\max} = 0.03$. However, even when large dissolution effects are included, a significant amount of CO_2 will exit the aquifer, as shown in Figure 4.

To reduce the potential leakage out of the aquifer, we obviously need to improve our injection strategy. Before we start looking for a better well-placement strategy, we will first try to rigorously optimize the

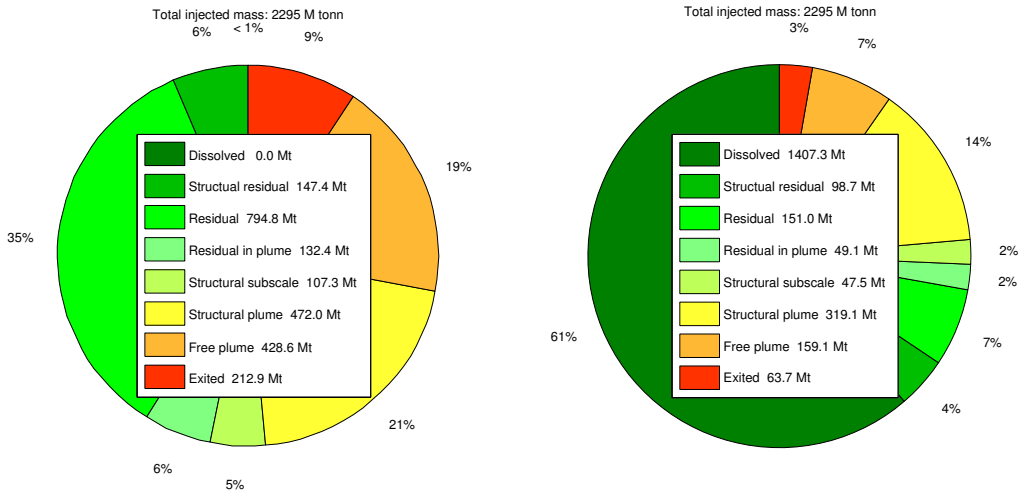

Figure 4 *Detailed inventory of the CO_2 trapping in the Utsira Formation two thousand years after the start of the injection.* ***Left:*** *VE simulation without dissolution effects.* ***Right:*** *VE simulation with significant dissolution effects.*

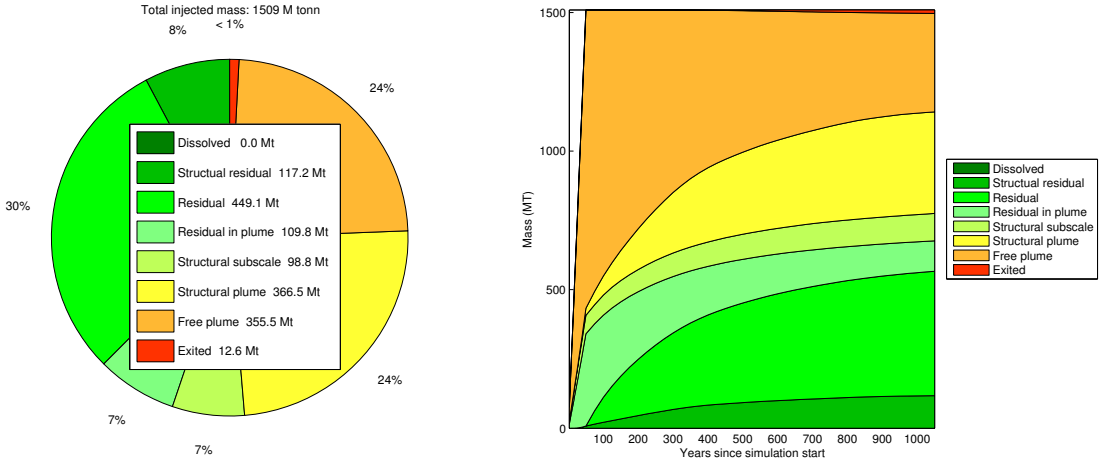

Figure 5 *Detailed inventory of the CO_2 trapping in the Utsira Formation for an optimized injection strategy.* ***Left:*** *Carbon inventory at the end of simulation.* ***Right:*** *Volumetric distribution of CO_2 over the first two thousand years after start of injection.*

injection rates given the current well placement. As our objective function, we will use

$$J = Mass_{injected} - C\big(Mass_{present} - Mass_{injected}\big), \tag{6}$$

which means that we seek to maximize the injected mass and minimize the mass that exits across the boundaries of the aquifer. As our optimization method, we use steepest descent method with gradients computed numerically by an adjoint method as part of the simulation thanks to the use of automatic differentiation. For robustness, dynamic time-stepping is implemented and constraints are handled by projections, see (Raynaud et al., 2014) for details.

The optimized strategy reduces the total injected amount from 2.5 to 1.5 Gt and is thereby able to delay the time the plume reaches the boundary to approximately 450 years after the injection start, which in turn eliminates 94% of the leakage over the 2000-year simulation period, see Figure 5. On the other hand, the carbon mass inside the free plume is almost the same after two thousand years in both strategies. For a more in-depth study of the Utsira Formation, we refer the reader to (Andersen et al., 2014).

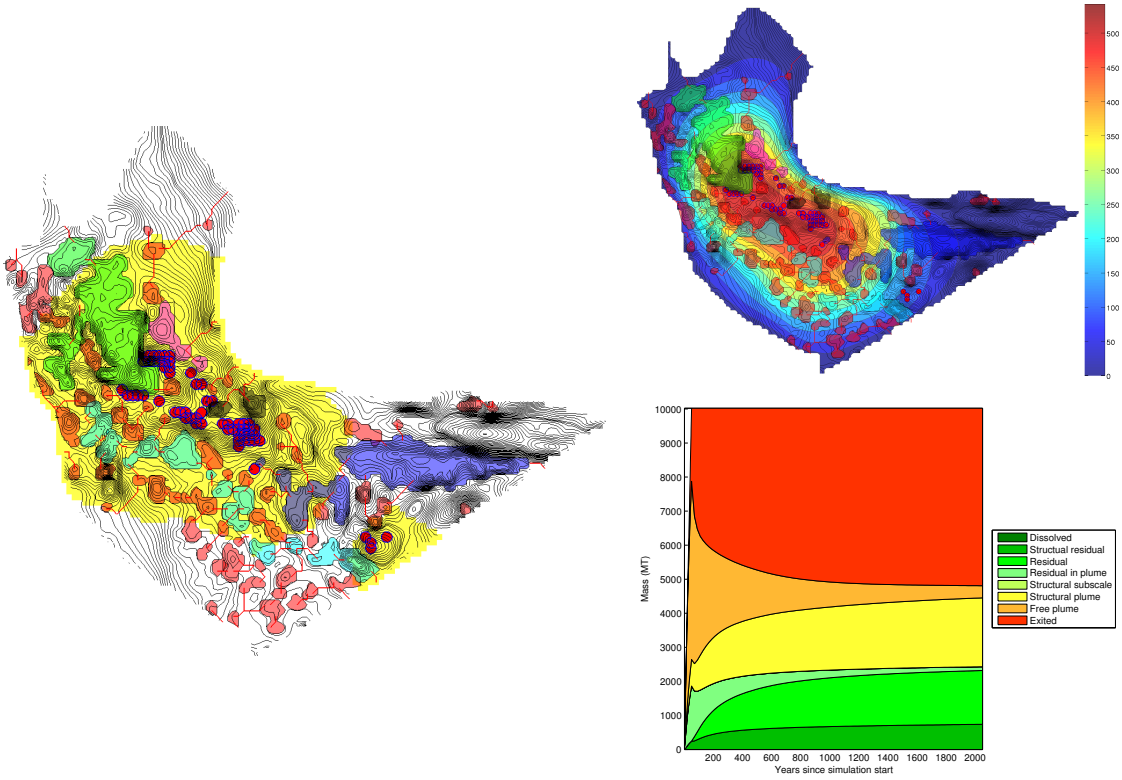

Figure 6 Left: *Distribution of potential injection hubs in the Sandnes Formation suggested by greedy optimization of the spill system. Circles indicate wells, the five largest spill trees have a unique color, the remaining are shown in red, and yellow color is the plume after fifty years.* ***Upper right:*** *Pressure increase in units bar at the end of the injection period.* ***Lower right:*** *Detailed inventory of the volumetric CO_2 distribution as a function of time.*

The Sandnes Formation

Among all the data sets extracted from the North Sea CO_2 Storage Atlas (Halland et al., 2011), the Sandnes Formation is the aquifer that has the highest potential of structural trapping because of a combination of many medium-sized traps and several huge domes that may potentially store large amounts of CO_2. Out of a total bulk volume of 1550 Gm^3, approximately 213 Gm^3 (or 14%) fall inside what can be characterized as structural traps. Repeating the same volumetric analysis as for the Utsira formation indicates that the upper bound on the storage capacity is 24 Gt of CO_2, distributed as 9 Gt in structural traps, 14 Gt residually trapped, and 1 Gt dissolved into the formation water.

Applying the greedy algorithm introduced above suggests an injection plan that will seek to fill 10 Gt within a period of fifty years, which obviously is an overly optimistic plan. To fill the two largest traps, the greedy algorithm distributes a high number of wells within limited areas inside the two traps. Altogether, these wells are set to inject a very large volume within a limited region and will cause an unacceptably high pressure buildup that will force large volumes of CO_2 out across the perimeter of the aquifer, as shown in Figure 6. (For the simulation, we used the same fluid parameters as for Utsira, but with a temperature of 85° C.)

We will therefore try a different strategy for placing the injection hubs. We start by distributing forty-seven injection hubs relatively uniformly throughout the aquifer, as illustrated to the left in Figure 7. Second, we modify the objective function (6) so that the present mass is only computed inside catchment areas that do not spill to the boundary (i.e., cells that are within the green region in the right plot of Figure 7).

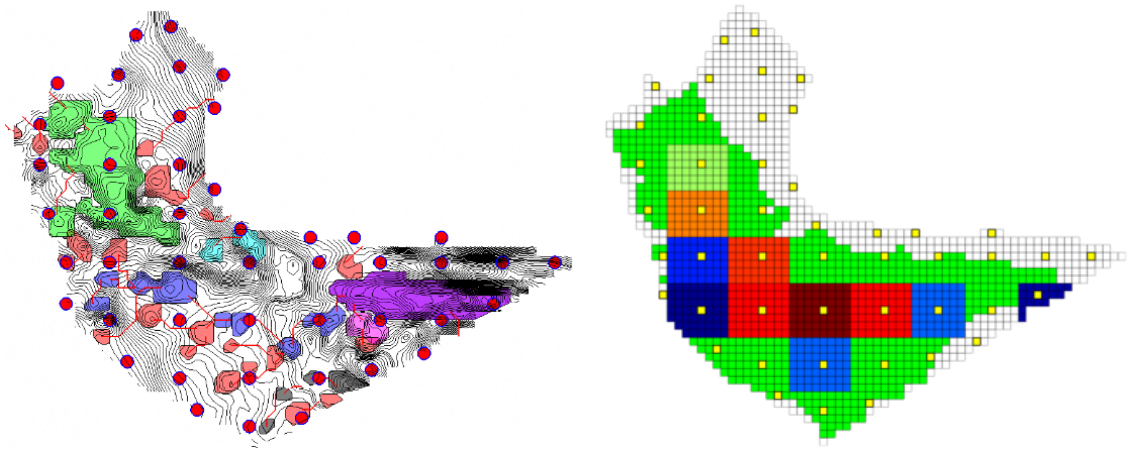

Figure 7 *Constrained optimization of injection strategy for the Sandnes Formation.* ***Left:*** *A regular coarse partition (from the agglom module of MRST) is imposed on the top surface grid and a well is place at the center of each coarse block.* ***Right:*** *Green color signify catchment areas that do not spill to the boundary of the aquifer. The colors of the blocks indicate the rate of the corresponding well after optimization with constant rate.*

Constraining the objective function by the catchment areas as explained above suggests a constant-rate strategy that will inject a total of 2.5 Gt. The outcome of this strategy is that more than 50% of the injected volume is inside a plume that covers a large fraction of the top surface and is expanding outward, as shown to the left in Figure 8. The fact that the plume is expanding can be seen since no residual trapping has taken place outside of the structural traps ('free residual' denotes the fraction of the movable plume that eventually will become structurally trapped at the tail of the plume). This outcome is still not feasible in a long-term perspective because the height of the plume and the pressure inside will cause it to continue to expand, which will move significant amounts of CO_2 into the spill region associated with the boundary.

To overcome this problem, we allow the optimization algorithm to specify a time-varying rate and maximize the objective function (6) one hundred years after the injection has stopped. By then, the dynamics of the plume has switched from being viscous dominated to being driven by buoyancy forces, and the objective function will have more time to penalize leakage. In the resulting strategy, 1.9 Gt is injected. The right plot in Figure 8 shows that the with the new strategy the movable plume is significantly reduced and that unlike in the previous case, residual trapping has taken place already during the injection period. Figure 9 shows the objective function for the two different injection strategies. Here we see that the constant-rate strategy attains a high value at the end of simulation by pushing a lot of CO_2 into the domain. However, as the migration phase, the objective function decreases because of leakage across the boundary. The second strategy injects less volumes and therefore attains a lower value at the end of injection, but causes less leakage across the boundary and will therefore experience less drop-off in the long perspective.

Concluding remarks

Accurate modeling and optimization of CO_2 storage operations is a challenging multiscale problem. The choice of modeling scale and computational approach will to a large extent depend on the type of questions asked and the characteristics of the aquifer in question, and to be able to provide answers that are trustworthy, the successful modeler should have a large toolbox of models and computational methods of varying fidelity and computational complexity

Whereas traditional 3D simulation tools are indispensable to study problems on a mesoscopic scale

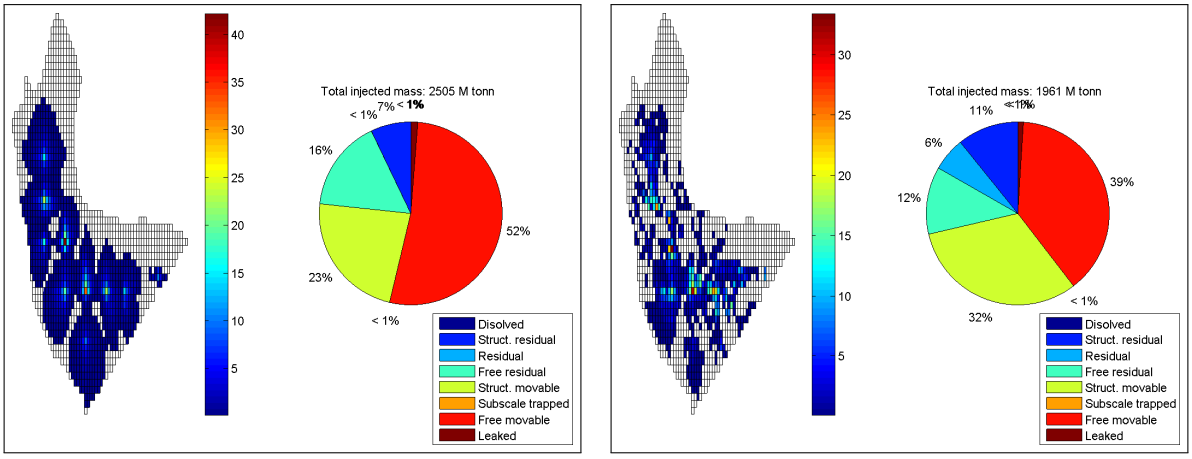

Figure 8 *Distribution of CO_2 in the Sandnes Formation at the end of injection. The pie charts show the inventory of the CO_2, while the grid plots show the height of the plume.* ***Left:*** *injection strategy derived by maximizing the value of the objective function* (6) *at the end of the injection period.* ***Right:*** *alternative strategy suggested by maximizing the value of the objective function 100 years after the injection has stopped.*

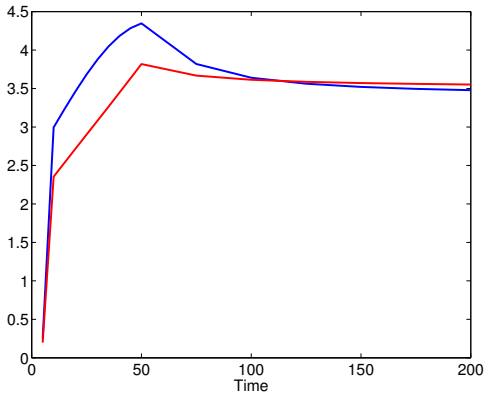

Figure 9 *Objective function* (6) *as function of time for two different injection strategies: blue line denotes a constant rate strategy that has been optimized over the injection period of fifty years, whereas the red line denotes a variable rate strategy that has been optimized up to one hundred years after the injection has ceased.*

(injectivity, early formation of the 3D plume, pressure buildup and thermo-mechanical effects in the near-well zone, etc), they generally fall short if the focus is on the large-scale, long-term distribution of injected volumes. For this type of problem, combinations of simple geometrical methods borrowing ideas from hydrology and basin modeling with vertically integrated models formulated on a 2D grid that follows the top surface of an aquifer, constitute a powerful approach that can be used to efficiently delineate structural traps and migration paths and simulate the likely outcomes of injection scenarios to provide estimates of structural, residual, and solubility trapping. As a general rule, we strongly advice that this type of models and methods are used early in the modeling workflow to efficiently investigate alternative hypotheses and explore as much as possible of parameter space to accurately span the range of likely outcomes.

In this paper, we have briefly outlined a set of geometrical tools and VE simulation models that have been implemented as a separate module, MRST-co2lab, in the open-source MATLAB Reservoir Simulation Toolbox (MRST). Apart from a rapid development cycle in MATLAB compared with compiled languages, the main purpose of using MRST is to make the new methods interoperable with the wide

range of traditional 3D modeling tools that are already implemented in the software. To this end, the methods in MRST-co2lab have mainly been developed using unified data models and modeling frameworks inherited from commercial reservoir simulators.

In the main part of the paper we have demonstrated how the various functionality in MRST-co2lab can be combined in two-level workflow for developing optimized injection plans: In the first phase, estimates of structural traps and spill paths are used to determine injection points that are optimal in terms of residual trapping and reduced risk for leakage through the boundary. Then, these injection points are used as an initial guess in a rigorous mathematical optimization of injection rates that uses and adjoint formulation that can be efficiently realized using automatic differentiation. In two cases considered herein, all wells have been constrained by rate only. In a more realistic setting, wells will be constrained by pressure, or by a combination of rate and pressure constraints. While this may make the optimization problem harder to solve, in particular if wells switch between rate and pressure constraints, our forward simulators will provide accurate gradients that can be fed into existing optimization methods developed for petroleum applications. We also believe that analysis of the spill system will still give good initial guesses in many cases. In a future work, our two-level optimization workflow should be expanded to include simulation-based well-placement optimization, e.g., using a combination of single-phase pressure solves and time-of-flight computation (Møyner et al., 2014).

Acknowledgments

The work was funded in part by Statoil ASA and the Research Council of Norway through grants no. 199878 (Numerical CO_2 laboratory) and 21564 (MatMoRA-II).

References

Andersen, O., Nilsen, H.M. and Lie, K.A. [2014] Reexamining CO_2 storage capacity and utilization of the Utsira Formation. *ECMOR XIV – 14th European Conference on the Mathematics of Oil Recovery, Catania, Sicily, Italy, 8-11 September 2014*, EAGE, doi:10.3997/2214-4609.20141809.

Bell, I.H., Wronski, J., Quoilin, S. and Lemort, V. [2014] Pure and Pseudo-pure Fluid Thermophysical Property Evaluation and the Open-Source Thermophysical Property Library CoolProp. *Ind. Eng. Chem. Res.*, **53**(6), 2498–2508, doi:10.1021/ie4033999.

Bradshaw, B.E. et al. [2011] An assessment of Queensland's CO2 geological storage prospectivity–The Queensland CO2 geological storage atlas. *Energy Procedia*, **4**(0), 4583–4590, doi:10.1016/j.egypro.2011.02.417, 10th International Conference on Greenhouse Gas Control Technologies.

Cloete, M. [2010] Atlas on geological storage of carbon dioxide in South Africa. Tech. rep., Council for Geoscience, Johannesburg, South Africa.

Halland, E.K., Johansen, W.T. and Riis, F. (Eds.) [2011] *CO_2 Storage Atlas: Norwegian North Sea*. Norwegian Petroleum Directorate, P. O. Box 600, NO–4003 Stavanger, Norway.

Halland, E.K., Johansen, W.T. and Riis, F. (Eds.) [2012] *CO_2 Storage Atlas: Norwegian Sea*. Norwegian Petroleum Directorate, P.O. Box 600, NO-4003 Stavanger, Norway.

Lewis, D. et al. [2008] Assessment of the potential for geological storage of carbon dioxide for the island of Ireland. Tech. rep., Sustainable Energy Ireland, Environmental Protection Agency, Geological Survey of Northern Ireland, and Geological Survey of Ireland.

Lewis, D. et al. [2009] Assessment of the potential for geological storage of carbon dioxide in Ireland and Northern Ireland. *Energy Procedia*, **1**(1), 2655–2662, doi:10.1016/j.egypro.2009.02.033, Greenhouse Gas Control Technologies 9.

Lie, K.A., Krogstad, S., Ligaarden, I.S., Natvig, J.R., Nilsen, H.M. and Skaflestad, B. [2012] Open source MATLAB implementation of consistent discretisations on complex grids. *Comput. Geosci.*, **16**, 297–322, ISSN 1420-0597, doi:10.1007/s10596-011-9244-4.

Møyner, O., Krogstad, S. and Lie, K.A. [2014] The application of flow diagnostics for reservoir management. *SPE J.*, accepted.

MRST [2014] The MATLAB Reservoir Simulation Toolbox, version 2014a. `http://www.sintef.no/MRST/`.

NACSA [2012] *The North American Carbon Storage Atlas*. Natural Resources Canada, Mexican Ministry of Energy, and U.S. Department of Energy.

NATCARB/Atlas [2012] *The 2012 United States Carbon Utilization and Storage Atlas*. U. S. Department of Energy, Office of Fossil Energy, 4th edn.

Nilsen, H.M., Lie, K.A. and Andersen, O. [2014a] Analysis of trapping capacities in the Norwegian North Sea

using `mrst-co2lab`. *submitted*.

Nilsen, H.M., Lie, K.A. and Andersen, O. [2014b] Fully implicit simulation of vertical-equilibrium models with hysteresis and capillary fringe. *submitted*.

Nilsen, H.M., Lie, K.A. and Andersen, O. [2014c] Robust simulation of sharp-interface models for fast estimation of CO_2 trapping capacity. *submitted*.

Nilsen, H.M., Lie, K.A., Møyner, O. and Andersen, O. [2014d] Spill-point analysis and structural trapping capacity in saline aquifers using `mrst-co2lab`. *submitted*.

Olivier, J.G.J., Janssens-Maenhout, G., Muntean, M. and Peters, J.A.H.W. [2013] Trends in global CO2 emissions: 2013 report. Tech. rep., PBL Netherlands Environmental Assessment Agency, The Hague, Netherlands.

Raynaud, X., Krogstad, S. and Nilsen, H.M. [2014] Reservoir management optimization using calibrated transmissibility upscaling. *ECMOR XIV – 14th European Conference on the Mathematics of Oil Recovery, Catania, Sicily, Italy, 8-11 September 2014*, EAGE, doi:10.3997/2214-4609.20141864.

SINTEF Applied Mathematics [2013] The MATLAB Reservoir Simulation Toolbox: Numerical CO_2 laboratory.

Span, R. and Wagner, W. [1996] A new equation of state for carbon dioxide covering the fluid region from triple-point temperature to 1100 K at pressures up to 800 MPa. *J. Phys. Chem. Ref. Data*, **25**(6), 1509–1597.

Thibeau, S. and Mucha, V. [2011] Have we overestimated saline aquifer CO2 storage capacities? *Oil Gas Sci. Technol. – Rev. IFP Energies nouvelles*, **66**(1), 81–92, doi:10.2516/ogst/2011004.