

NUMERICAL STUDIES OF DENSITY-DRIVEN FLOW IN CO₂ STORAGE IN SALINE AQUIFERS

George S. H. Pau, John B. Bell, Karsten Pruess, Ann S. Almgren, Michael J. Lijewski, and Keni Zhang

Lawrence Berkeley National Laboratory
1 Cyclotron Road MS50A-1148
Berkeley, CA 94720, USA
e-mail: gpau@lbl.gov

ABSTRACT

Simulations are routinely used to study the process of carbon dioxide (CO₂) sequestration in saline aquifers. In this paper, we look at some numerical aspects of the accurate modeling and simulation of the dissolution-diffusion-convection process. We perform convergence studies with respect to solver tolerances, grid resolutions, fluctuation strength, and domain size. We show that stringent tolerances and grid resolutions are needed to accurately predict onset time. Domain size must be sufficiently large to contain at least 2 extended fingers to accurately predict the long-term stabilized mass flux of CO₂; otherwise, finite domain effects will adversely change the flow behavior of the system we are modeling.

INTRODUCTION

Carbon dioxide (CO₂) sequestration involves injecting CO₂ into a saline aquifer. While the primary mechanism of securing the CO₂ relies on a leak-proof formation, secondary geochemical mechanisms may play a significant role, especially in a geological time frame. At long time, an immiscible CO₂ gas layer will form on top of the brine in the rock formation. Under ambient temperature and pressure conditions in a typical aquifer, CO₂ will dissolve into the brine and increase the density of the brine at the interface of the layers by 0.1–1%, depending on the salinity of the brine (Pruess and Zhang, 2008). Due to gravitational instability and the heterogeneity in the rock properties of the aquifer, CO₂-rich brine fingers will form, leading to convective flow that transports these CO₂-rich brines downward, while driving brine with low CO₂ concentration upwards. This then accelerates the rate at which CO₂ is dissolved and provides a more secure mechanism by which CO₂ can be stored.

This dissolution-diffusion-convection process has been analyzed in a number of studies. In Ennis-King and Paterson (2003) and Riaz et al. (2006), linear stability analyses yield useful relations for the onset time for convection, dominant wavelength for growth of convective fingers, and the growth rates of these

fingers. Numerical simulations were also performed to further elucidate the process and to validate the linear stability analyses. For example, Riaz et al. (2006) performed numerical simulation of a single-phase two-component model with the Boussinesq assumption and demonstrated that the simulation results are consistent with their analysis. Pruess and Zhang (2008) examined long-term behavior of the CO₂ flux, in addition to the onset of convection. Their simulation uses a full compressible model with very accurate equations of state.

In this work, we examine some of the numerical aspects of studying the dissolution-diffusion-convective process of CO₂ through simulation. We use a second-order accurate adaptive method that is described in the next section. Specifically, we examine how simulation parameters, such as solver tolerances, grid resolution, strength of perturbations, and domain size, affect our solution. The results are compared to those obtained through TOUGH2-MP, a parallelized version of the general-purpose simulator TOUGH2/ECO2N (Zhang et al., 2008; Pruess, 2004; Pruess and Spycher, 2007).

NUMERICAL SCHEME

We assume that a layer of CO₂-saturated brine is formed at the interface of the brine and CO₂ gas, and the dissolution rate is sufficiently high that the layer remains saturated for the length of the simulation. This assumption allows us to use a variable-density single-phase incompressible model to treat the dissolution-diffusion-convection process. The variable-density formulation we use is similar to the formulation of Douglas et al. (2002), and different from that of Riaz et al. (2006), since it does not make the Boussinesq assumption. Here, we provide an overview of the methodology; details of the approach are given in Pau et al. (2009).

The basic integration scheme is based on the total-velocity splitting approach. Due to the dependence of fluid density on the mass fraction of dissolved CO₂, the velocity divergence constraint is given by

$$\nabla \cdot u = \sum_{i=1}^2 \frac{1}{\rho_i} \nabla \cdot \phi \rho D \nabla X_i,$$

where the summation is over components $i = 1$ (CO_2) and $i = 2$ (H_2O). Here, u is the Darcy velocity, ρ_1 and X_1 are the density and the mass fraction of $\text{CO}_2(\text{aq})$, ρ_2 and X_2 are the density and the mass fraction of H_2O , respectively, ρ is the density of the mixture, ϕ is the porosity and D is the diffusion coefficient. Expressing u in terms of pressure p , this divergence constraint leads to a second-order elliptic pressure equation, given by,

$$-\nabla \cdot \frac{\kappa}{\mu} (\nabla p - \rho \mathbf{g}) = \sum_{i=1}^2 \frac{1}{\rho_i} \nabla \cdot \phi \rho D \nabla X_i,$$

that we can solve to obtain u ; κ is the permeability of the porous media, μ is the viscosity of the mixture and \mathbf{g} is the gravity. This velocity is then used to recast component conservation equations as nonlinear hyperbolic equations.

The discretization procedure we have adopted is similar to the IMPES approach. The pressure equation is solved implicitly using a finite difference method and the mass conservation equations are solved semi-explicitly using an explicit second-order Godunov method for advection and an implicit Crank-Nicholson discretization of diffusion. Unlike the basic IMPES algorithm, however, our method is second-order accurate in both space and time.

The overall time-stepping procedure is integrated into an adaptive mesh refinement (AMR) framework (Almgren et al., 1994) to efficiently accommodate the difference in scale between the diffusive boundary layer and the large-scale convective fingers. Our approach to adaptive refinement uses a nested hierarchy of logically rectangular grids with simultaneous refinement of the grids in both space and time. The grid changes with time based on a set of user-defined refinement criteria. Shown in Figure 1 is a snapshot of the grid; finer grids are placed in region where small features with large concentration gradients are present. The resulting algorithm is parallelized and shows good scaling behavior up to 1024 CPUs.




Figure 1. AMR grid with four levels of refinement. Refinement criterion is based on concentration gradient of all components.

PROBLEM SETUP

The fluid properties, as specified in Table 1, are derived from the ECO2N fluid module of TOUGH2 and correspond to pure water (no salinity) at temperature and pressure conditions of $T = 45^\circ\text{C}$, $P = 100$ bar, as would be encountered in a typical sedimentary basin near 1,000 m depth. A two-dimensional domain was used as shown in Figure 2. In most cases, a domain size of $1 \text{ m} \times 4 \text{ m}$ is used.

Table 1. Fluid and formation properties

Property	Magnitude
Viscosity, μ	0.5947 mPas
Water density, ρ_w	994.56 kg/m^3
Saturated CO_2 mass fraction, X_1	0.049306
Density increase due to CO_2 dissolution, $\Delta\rho$	10.45 kg/m^3
Diffusivity, D	$2 \times 10^{-9} \text{ m}^2/\text{s}$
Mean permeability, κ	10 D
Mean porosity, ϕ	0.3




Figure 2. Schematic of specifications used for numerical simulations. Different heights H and widths W were employed for the simulation domain. Periodic boundary conditions mean that the same conditions were maintained on the left and right boundaries.

The flow system is initialized in gravity equilibrium (hydrostatic pressure gradient). The top boundary is impermeable and only diffusion of dissolved CO_2 into the simulation domain is permitted at the top boundary. Isothermal conditions are assumed throughout. The medium is assumed homogeneous, but small random variations, described by maximum per cent deviation from the mean, in permeability or porosity are imposed to seed the convective activity.

RESULTS

Figure 3 shows the concentration of CO_2 after the initialization of convective flow. Small fingers that form at the top boundary of the domain grow and merge into one or more extended fingers, the main conduits by which aqueous CO_2 is transported efficiently downward. These extended fingers grow wider as they traverse downward, due to diffusion. In addition, the dynamics within each of these extended fingers can be quite complex. Specifically, we observe that blobs of aqueous CO_2 with higher concentration may pinch off from an extended finger, and move downward at a greater speed. These blobs eventually diffuse to a point where they are indistinguishable from the extended finger. The dynamics of convective flow is thus complex and exhibits highly nonlinear behavior.




Figure 3. A snapshot of the concentration of CO_2 after onset of convection. The arrows show the velocity field.

In subsequent analyses, we will focus on two quantities of interest. First is the onset time of convection, t_{onset} , which determines the time at which convection becomes an important transport mechanism. We define t_{onset} as the time at which the average mass flux at the top boundary, F_0 , has a relative deviation of 1% from a pure diffusive mass flux.

Second is the stabilized mass flux, $\langle F_0 \rangle$, at the top boundary. The existence and the quantification of $\langle F_0 \rangle$ provide a simple model for larger-scale simulations of CO_2 injection. To determine $\langle F_0 \rangle$, we first compute the moving average $F_m(t)$ defined as

$$F_m(t) = \frac{1}{\Delta t} \int_{t-\Delta t/2}^{t+\Delta t/2} F_0(s) ds,$$

where Δt is chosen such that F_m does not fluctuate more than 5% over a time interval Δt . The stabilized mass flux is then given by

$$\langle F_0 \rangle = \frac{1}{\Delta T} \int_{T-\Delta T}^T F_m(t) dt,$$

where $T+\Delta t/2$ is time at end of simulation. We ensure that the fingers have not reached the bottom of the simulation domain at time $T+\Delta t/2$.

Both the definitions we adopted above are admittedly arbitrary, but they provide consistent measures for

performing our convergence studies. We note that F_0 is numerically equal to the CO₂ dissolution rate. Thus, F_0 denotes the rate at which CO₂ is removed from the highly mobile and buoyant gas phase, and put into a less mobile and negatively buoyant aqueous phase. This rate has important ramifications for storage security.

Effects of solver tolerances on the onset time

We begin this section with an analysis of the influence of numerical errors on the onset time. For homogeneous permeability and porosity, the two-dimensional problem at hand can be essentially reduced to a one-dimensional diffusion problem. Convective transport is only induced when a non-uniform flow field is generated by the heterogeneity in the permeability or the porosity function. However, linear solvers with finite tolerances may introduce small non-uniform errors that can also eventually lead to instability and convective flow. It is then important that we ensure that any convective transport observed in our simulation is induced by heterogeneities in formation's properties, and not due to finite tolerances of the linear solver we use. In particular, the tolerances must be sufficiently small that it has limited influence on the dynamics of the flow, and thus the onset time.

Since the Darcy velocity is computed from the pressure, we shall look at tolerances of the linear solver used to solve the pressure equation in our numerical scheme. Our multigrid linear solver uses two tolerances to control the accuracy of the linear solve: the relative error tolerance, ϵ_{rel} and the absolute error tolerance, ϵ_{abs} . Figure 4 shows that for 1% fluctuation in κ and $\epsilon_{rel} = 10^{-12}$, the onset time converges to a value of 2.3×10^5 s when $\epsilon_{abs} < 10^{-14}$. We note that the numerically-induced onset time is larger by a factor greater than 2. In addition, as we decrease ϵ_{abs} , the numerically induced onset time increases, but the true onset time remains unchanged. Decreasing ϵ_{rel} to 10^{-14} also does not change the results. Thus, we can conclude that with $\epsilon_{rel} = 10^{-12}$ and $\epsilon_{abs} = 10^{-14}$, the convective behavior that we observed is induced only by the heterogeneity in the permeability, and not by effects of finite tolerances in the linear solver.

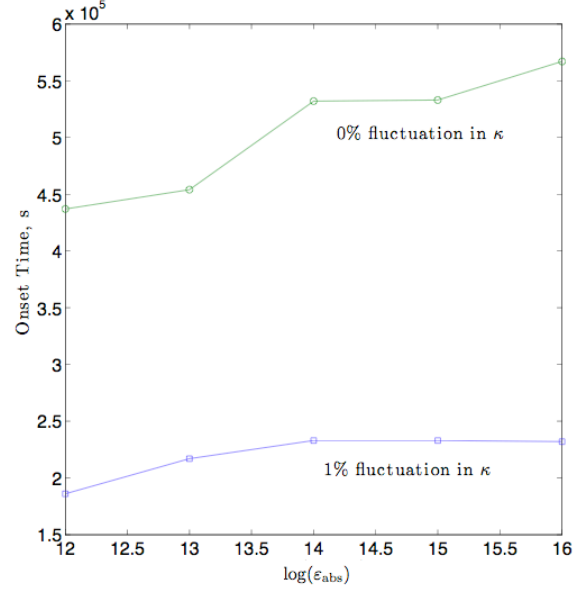


Figure 4. Effects of numerical errors on the determination of onset time. The onset time has converged when the absolute tolerance is 10^{-14} .

Effects of grid resolution on the onset time

Riaz et al. (2006) derived a critical wavelength λ_c from the linear stability analysis for examining the onset of convective flow. To fully capture the dynamics during the onset of convection, the dimensions of the simulation domain must be much larger than λ_c . For the parameters given in Table 1, $\lambda_c = 0.1$ m. Thus, our simulation domain, which has a width that is 10 times larger, is adequate.

However, the grid resolution must still be sufficiently high to resolve λ_c so that we can capture the initiation of convective flow accurately. We perform a grid convergence study to determine the appropriate grid size. Initial experiments show that we need to adjust the fluctuations value with grid size so that the underlying statistics are consistent for an uncorrelated random distribution. It reflects the notion that for a truly random medium, as gridblocks become larger, the variance of the fluctuations will be lower because of averaging. For example, an initial 1% fluctuation for a given resolution has to be reduced by half when the grid size in each direction is doubled.

We examine four different grid sizes: $\Delta x = 1/256$, $1/512$, $1/1024$, and $1/2048$. The corresponding fluctuations are given by 0.25%, 0.5%, 1% and 2%. Figure 5 shows that the onset time converges to a single value. From Table 2, we can conclude that the

onset time converges to 2.23×10^5 s. An effective resolution of $\Delta x = 1/1024$, which is a factor of 100 smaller than λ_c , is thus sufficient.




Figure 5. Effects of numerical errors on onset time. The onset time has converged when the absolute tolerance is 10^{-14} .

Table 2. Onset time, t_{onset} , at different resolutions.

h, m	$t_{onset} \times 10^{-5}$ s
1/256	2.44
1/512	2.33
1/1024	2.26
1/2048	2.23

Effects of fluctuation strength on the onset time

As the fingering phenomenon is driven by the heterogeneity in the permeability and porosity, we would like to study what influence fluctuations in permeability and porosity have on the onset time. Figure 6 shows that for similar level of relative per cent fluctuation, fluctuation in ϕ leads to faster onset time than fluctuation in κ . The variations with respect to per cent fluctuation are, however, similar; the onset time decreases with increasing fluctuation strength.

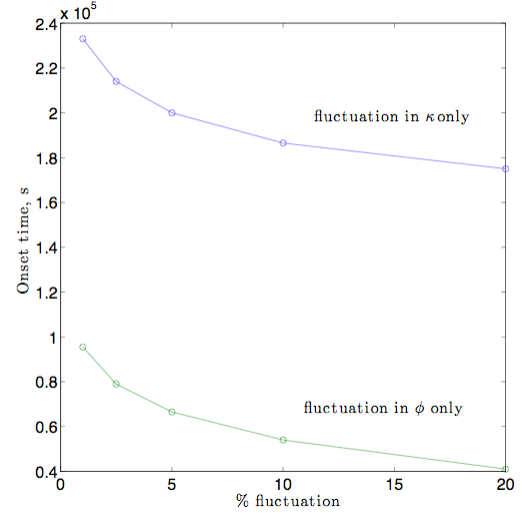


Figure 6. Variation of onset time with % fluctuation in permeability κ and porosity ϕ

Effects of domain size on the stabilized mass flux

At long time, the flow is characterized by fingers that extend along the length of the domain. Initial experiments show that boundary effects will adversely affect the solutions if the simulation domain cannot accommodate at least two extended fingers. To provide a good averaging to F_0 , here should be a modest number of these extended fingers. In Table 3, we show the computed stabilized mass fluxes, the maximum relative deviations from the $\langle F_0 \rangle$ in the sample set used to compute $\langle F_0 \rangle$, and the number of fingers for domains with different widths.

Table 3. Stabilized mass fluxes for different domain size. The deviation shown is the maximum absolute relative deviation of F_0 from $\langle F_0 \rangle$.

width, m	$\langle F_0 \rangle \times 10^6$ kg/m ² /s	deviation	number of fingers
1	1.46	0.27	2
2	1.48	0.12	4
8	1.52	0.07	8
16	1.50	0.06	17

The fluctuation decreases with increasing width; the relative deviation of F_0 from $\langle F_0 \rangle$ for $W = 16$ m is about a factor of 5 smaller than the relative deviation for $W = 1$ m. The mean also appears to converge to 1.5×10^{-6} kg/m²/s, although the variation is less drastic compared to the variation in the relative deviation. This shows that having a width that can

accommodate two fingers may be sufficient, but larger width will lead to better averaging.

COMPARISON TO TOUGH2-MP

The same problem of convective instability discussed above was simulated with a parallelized version TOUGH2-MP (Zhang et al., 2008) of the general-purpose reservoir simulator TOUGH2. Initial calculations used the ECO2N fluid property module (Pruess and Spycher, 2007). However, since the conditions in the present problem are limited to a single aqueous phase, a more efficient simulation can be obtained by using EOS7, in which the fluid is represented as a two-component mixture of H₂O and brine. We use the brine component to represent density changes in the aqueous phase from CO₂ dissolution. Numerical work is approximately proportional to NEQ², where NEQ is the number of equations per gridblock. Accordingly, a given problem can be solved with EOS7 (NEQ = 2) in less than half the time required with ECO2N (NEQ = 3). Our test calculations have confirmed excellent agreement between simulations using EOS7 and ECO2N.

We experimented with different domain sizes and grid resolutions, to achieve results with “small” space discretization errors. Most calculations were performed for a domain of width $W = 1$ m, height $H = 5$ m, and a vertical grid resolution of 1 mm near the top boundary, which gradually was coarsened going downward. Horizontal grid resolution was 10 mm, and the total number of gridblocks was 52,300. Simulations presented here were performed on a Dell T5400 dual quad core computer with a total of 8 cores, and in most cases 16 processes were run (two per processor), because this was found to reduce total execution time compared to running eight processes. Depending on grid resolution and simulation time, individual runs typically took from 1/2 to 4 hours.

Results

The convective activity shows similar features as seen in the adaptive grid simulations discussed previously, including fingering convection, merging and pinch-off of fingers, and continuous generation of new fingers as older ones grow. As an example, Figure 7 shows convective patterns for three different random number seeds after a time of 101.6 days. The resolution of features is somewhat inferior in comparison to the adaptive gridding (Figure 3). Specifics of the convective activity are very sensitive to small problem variations, but we observe that integral measures of the process, such as onset time of convection and long-term behavior of the CO₂

mass flux carried by the convection, are quite robust to modest changes in problem parameters. Indeed, for the three cases with different random permeability fields shown in Figure 7, onset times of convection are identical, and long-term stabilized fluxes show random fluctuations of $\pm 15\%$ about the same mean of approximately 1.3×10^{-6} kg/s/m².

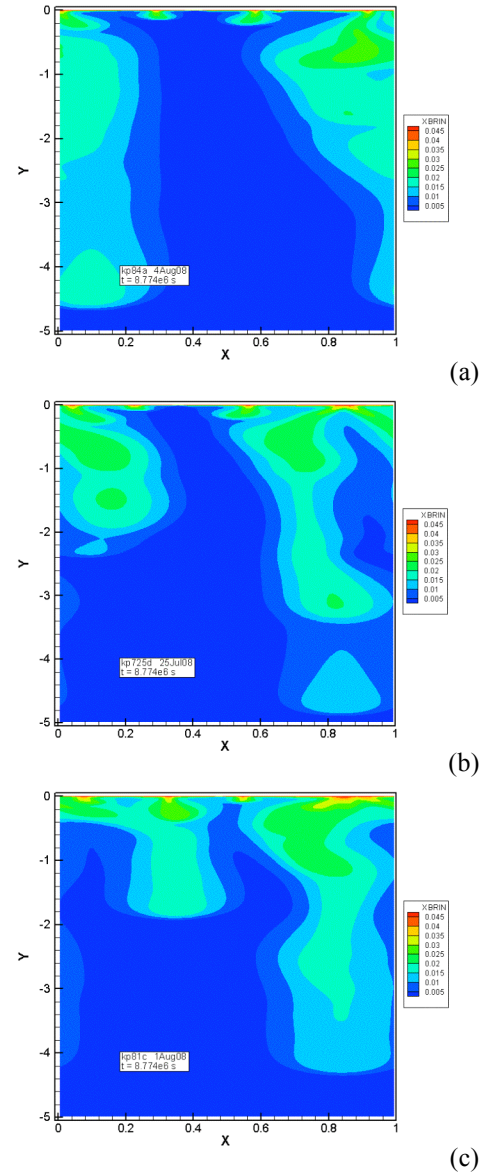


Figure 7. Simulated distribution of dissolved CO₂ after 101.6 days for three different random number seeds to generate permeability heterogeneity (a-0.7; b-0.8; c-0.9). Note the different horizontal and vertical scales.

We explored the sensitivity of the onset time for convection to the random perturbations applied to the medium. Figure 8 shows that the onset time for convection decreases with increasing strength of the applied perturbation of the medium. Porosity perturbations are seen to be more effective in triggering convective instability than permeability perturbations. These results are consistent with the results in the previous section, as shown in Figure 6.

A comparison calculation with a perfectly homogeneous medium yields a substantially larger onset time, as in this case convective instability arises only from numerical roundoff. The onset times obtained from TOUGH2-MP simulations are around $6\text{-}9 \times 10^4$ s, about a factor 3 smaller than obtained with the adaptive algorithm. The reason for this significant discrepancy is unknown. At a nominal onset time of 7.5×10^3 s, the thickness of the diffusive boundary layer is = 12.2 mm, which is well resolved with our 1 mm grid resolution.

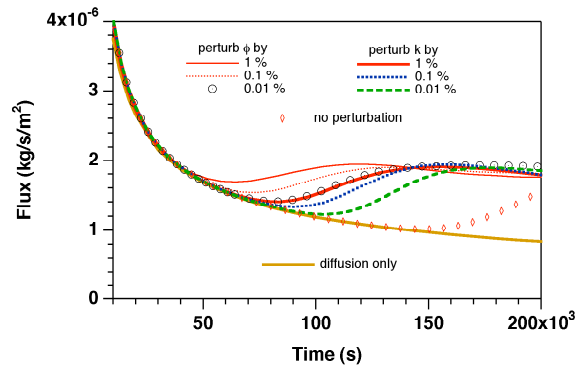


Figure 8. Simulated CO_2 fluxes at the top boundary at early times, for different random perturbations of the porosity or permeability field.

The convective CO_2 flux at the top is equal to the CO_2 dissolution rate per unit area. Due to the partially chaotic nature of the convection process, this rate fluctuates, but the fluctuations are modest in size, and fluxes stabilize at $1.3 \times 10^{-6} \text{ kg/s/m}^2$ ($\pm 15\%$), regardless of how the instability was triggered (Figure 9). This is approximately 13% smaller than the value obtained from adaptive gridding simulations.

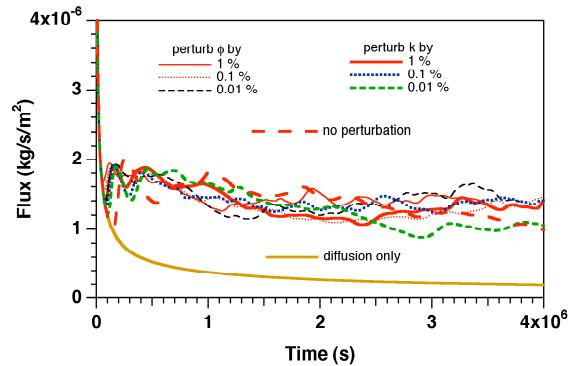


Figure 9. Longer-term behavior of simulated CO_2 fluxes at the top boundary.

We also performed simulations with different boundary conditions at the sides and bottom of the domain, such as no-flow conditions on the sides, and constant pressure conditions on the bottom. Onset times for convection and long-term stabilized fluxes were found to be insensitive to boundary conditions. A more extensive account of TOUGH2 analyses of the dissolution-diffusion-convection process is available in a laboratory report (Pruess and Zhang, 2008).

CONCLUSION

High-resolution simulations of the diffusion-convection process in CO_2 sequestration were performed using a block structured adaptive mesh refinement method. Some numerical aspects of the simulations were examined, specifically variations in the onset time and stabilized mass flux with respect to solver tolerances, grid resolution, fluctuation strength, and domain size. Our main findings are as follows.

The diffusion-convection process involves complex fluid dynamics on multiple scales, including creation, growth, movement, merging, and pinch-off of convective fingers. While details of the convection process are chaotic in nature, integral measures, such as onset time for convection, and long-term CO_2 mass flux associated with the convective activity, are robust and insensitive to modest problem variations. By employing an adaptive gridding method combined with semi-implicit time stepping, we were able to control discretization errors and demonstrate convergence of onset times for convection. High-resolution simulations with TOUGH2-MP, using fixed gridding with resolution down to 1 mm at the CO_2 dissolution boundary, showed similar features as the simulations with adaptive gridding. Long-term stabilized CO_2 fluxes obtained from adaptive

gridding simulations and TOUGH2-MP agree to within 13%. Onset time for convection with TOUGH2-MP was only about 1/3 of what was obtained with the adaptive gridding approach. The reason for these differences is unknown at present, but may have to do with different approximations made for the flow physics.

ACKNOWLEDGMENT

Support for this work was provided by the LDRD program at Lawrence Berkeley National Laboratory (LBNL) and by the Office of Basic Energy Sciences under Contract No. DE-AC02-05CH11231 with the U.S. Department of Energy. This research used resources of the National Energy Research Scientific Computing Center and the Lawrence Livermore computational cluster resource provided by the LBNL's IT Division, both of which are supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

REFERENCES

- Almgren, A. S., J. B. Bell, L. H. Howell, and P. Colella, An Adaptive Projection Method for the Incompressible Navier-Stokes Equations. In *Proceedings of the IMACS 14th World Conference*, Atlanta, GA, July 11-15, 1994.
- Douglas, J. Jr., D. Frias, N. Henderson, and F. Pereira, Simulation of Single-Phase Multi-Component Flow Problems in Gas Reservoirs by Eulerian-Lagrangian Techniques, *Transport in Porous Media*, 50:307–342, 2002.
- Ennis-King, J., and L. Paterson, Role of Convective Mixing in the Long-Term Storage of Carbon Dioxide in Deep Saline Formations, paper SPE-84344, presented at *Society of Petroleum Engineers Annual Fall Technical Conference and Exhibition*, Denver, CO, October 2003.
- Pau, G. S. H., A. S., Almgren, J. B. Bell, and M. J. Lijewski, A Parallel Second-Order Adaptive Mesh Algorithm for Incompressible Flow in Porous Media, *Phil. Trans. R. Soc. A*, 2009. To be published.
- Pruess, K. The TOUGH Codes—A Family of Simulation Tools for Multiphase Flow and Transport Processes in Permeable Media, *Vadose Zone J.*, Vol. 3, pp. 738 - 746, 2004.
- Pruess K. and N. Spycher. ECO2N – A Fluid Property Module for the TOUGH2 Code for Studies of CO₂ Storage in Saline Aquifers, *Energy Conversion and Management*, 48(6), 1761-1767, 2007.
- Pruess, K., and K. Zhang, Numerical Modeling Studies of the Dissolution-Diffusion-Convective Process during CO₂ Storage in Saline Aquifers,

- Paper LBNL-1243E, Lawrence Berkeley National Laboratory, Berkeley, Calif., 2008.
- Riaz, A., H. A. Tchelepi, and F. M. Orr, Onset of Convection in a Gravitationally Unstable Diffusive Boundary Layer in Porous Media, *Journal of Fluid Mechanics*, 548:87–111, 2006.
- Zhang, K., Y. S. Wu and K. Pruess. *User's Guide for TOUGH2-MP - A Massively Parallel Version of the TOUGH2 Code*, LBNL-315E, Lawrence Berkeley National Laboratory, Berkeley, Calif., 2008.