

A Parallel Second-Order Adaptive Mesh Algorithm for Incompressible Flow in Porous Media

George S. H. Pau, Ann S. Almgren, John B. Bell,
Michael J. Lijewski

*Lawrence Berkeley National Laboratory, MS 50A-1148, 1 Cyclotron Road,
Berkeley, CA 94720*

Abstract

In this paper we present a second-order accurate adaptive algorithm for solving multiphase, incompressible flows in porous media. We assume a multiphase form of Darcy's law with relative permeabilities given as a function of the phase saturation. The remaining equations express conservation of mass for the fluid constituents. In this setting the total velocity, defined to be the sum of the phase velocities, is divergence-free. The basic integration method is based on a total-velocity splitting approach in which we solve a second-order elliptic pressure equation to obtain a total velocity. This total velocity is then used to recast component conservation equations as nonlinear hyperbolic equations. 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 integration algorithm on the grid hierarchy is a recursive procedure in which coarse grids are advanced in time, fine grids are advanced multiple steps to reach the same time as the coarse grids and the data at different levels are then synchronized. The single grid algorithm is described briefly, but the emphasis here is on the time-stepping procedure for the adaptive hierarchy. Numerical examples are presented to demonstrate the algorithm's accuracy and convergence properties and to illustrate the behavior of the method.

Key words: adaptive mesh refinement, Darcy flow, porous media

PACS: 47.56.+r

1991 MSC: 65M50, 65M55

¹ Support for this work was provided by the LDRD program at Lawrence Berkeley National Laboratory under contract DE-AC02-05CH11231.

1 Introduction

Multicomponent and multiphase flows in the subsurface are often characterized by localized phenomena such as steep concentration gradients or saturation fronts. Accurately resolving these types of phenomena requires high resolution in regions where the solution is changing rapidly. For this reason, the development of some type of dynamic gridding capability has long been of interest in the porous media community.

Heinemann [11] and Ewing *et al.* [10] have considered dynamic local grid refinement approaches. More recent papers by, for example, Sammon [19] and Christensen *et al.* [6], discuss development of adaptive techniques in the context of unstructured grids. An alternative approach to local refinement is based on structured-grid adaptive mesh refinement. This type of approach, based on the strategy introduced for gas dynamics by Berger and Colella [4], was first applied to porous media flow by Hornung and Trangenstein [13] and by Propp [17]. Additional developments are discussed in Trangenstein [20], Trangenstein and Bi [21], and Hoang and Kleppe [12]. In this approach, regions to be refined are uniformly subdivided in both space and time. Related approaches were developed by Nilsson *et al.* [15,16] uses a spatially anisotropic refinement strategy with no temporal refinement and by Edwards [9] who uses a cell-by-cell refinement strategy.

The focus of this paper is on developing a structured adaptive mesh refinement (AMR) algorithm for porous media flow. Our approach is similar to the approach introduced by Hornung and Trangenstein [13] and to the approach discussed by Propp [17]. In this approach, the solution is represented on a hierarchical sequence of nested grids with successively finer spacing in both time and space. Increasingly finer grids are recursively embedded in coarse grids until the solution is sufficiently resolved. An error estimation procedure based on user-specified criteria evaluates where additional refinement is needed and grid generation procedures dynamically create or remove rectangular fine grid patches as resolution requirements change. The method presented here uses subcycling in time so that all levels are advanced at the same CFL number, thus reducing the numerical dissipation of the explicit upwind advection scheme used to advance the solution. The major difference between the approach adopted here and that of Hornung and Trangenstein is that the current method does not require a global, multilevel pressure solve at each fine-grid time step. Instead, when advancing a given level, we solve the pressure on that level with boundary conditions obtained from the coarser levels. This approach avoids the computational expense of the global solve but introduces additional complexity into the synchronization step of the algorithm in which inconsistencies between different levels are corrected. The synchronization approach used here is based on the algorithm developed by Almgren *et al.* [1] for

incompressible flows. The methodology is implemented in parallel using the `BOXLIB` framework discussed in Rendleman *et al.* [18] and Crutchfield [7]. A similar approach was used by Propp; however, his algorithm used a different approach to synchronization and was limited to two dimensions.

Before describing the adaptive algorithm we will briefly review the total velocity splitting approach and briefly discuss our basic fractional step scheme for a single grid. In the third section we describe, in detail, the recursive time-stepping procedure for the adaptive algorithm and other aspects of the adaptive algorithm. The fourth section shows convergence results and presents computational examples illustrating both the performance and parallel scaling of the method.

2 Total velocity-splitting algorithm

Here we consider multiphase, multicomponent incompressible flow in heterogeneous porous media. The basic void fraction of the medium is referred to as the porosity, ϕ , and the fraction of that void occupied by a particular phase is referred to as the phase saturation, u_α . (Here, Greek subscripts refer to mobile phases.) The multicomponent mixture is composed of N components (or lumped pseudo-components). We define $\mathbf{n} \equiv (n_1, \dots, n_N)$ as the vector of component densities per unit pore volume, and \mathbf{n}_α as the subset of \mathbf{n} in phase α .

The equations that describe the flow represent mass conservation, energy conservation, and Darcy's law. Darcy's law expresses the volumetric flow rate, v_α , of each phase in terms of the phase pressure, p_α , namely,

$$v_\alpha = -\lambda_\alpha(\nabla p_\alpha - \rho_\alpha \vec{g}) , \quad (1)$$

where $\lambda_\alpha \equiv K k_{r,\alpha} / \mu_\alpha$ is the phase mobility. Here K is the permeability distribution of the medium, $k_{r,\alpha}$ is the relative permeability, which expresses modification of the flow rate due to multiphase effects, μ_α is the phase viscosity, ρ_α is the phase density, \vec{g} is the gravitational acceleration. The pressure in each phase is related to a reference pressure, p , by a capillary pressure, $p_{c,\alpha} = p_\alpha - p$, which is a function of saturation.

For this system, conservation of mass for each component is given by

$$\phi \frac{\partial \mathbf{n}}{\partial t} + \nabla \cdot \sum_\alpha \frac{\mathbf{n}_\alpha}{u_\alpha} v_\alpha = \nabla \cdot \mathcal{D} + \mathcal{R}, \quad (2)$$

where \mathcal{D} represents diffusive terms that include multiphase molecular diffu-

sion and dispersion, and \mathcal{R} is the reaction source term. The diffusive term associated with capillary pressure is implicitly included in the definition of the phase velocities. The diffusion terms can be quite complex, depending on the particular problem; see [22] for a detailed discussion of these terms.

The component conservation equations specify the change in total mass of each component resulting from the transport and diffusion of that component distributed across the phases. The phase behavior of the system specifies how the components are apportioned into phases and the volume occupied by those phases. In this paper, we consider only simple incompressible systems without mass transfer between phases. We can then define a total velocity

$$v_T \equiv \sum_{\alpha} v_{\alpha} = - \sum_{\alpha} \lambda_{\alpha} (\nabla p_{\alpha} - \rho_{\alpha} \vec{g}), \quad (3)$$

which is divergence-free. This divergence-free condition then leads to a second-order elliptic equation:

$$\nabla \cdot v_T = - \nabla \cdot \sum_{\alpha} \lambda_{\alpha} (\nabla (p + p_{c,\alpha}) - \rho_{\alpha} \vec{g}) = 0, \quad (4)$$

since $p_{\alpha} = p + p_{c,\alpha}$. We can rewrite (3) to express ∇p in terms of v_T , then using (1) we can express v_{α} in terms of the total velocity,

$$v_{\alpha} = \frac{\lambda_{\alpha}}{\lambda_T} v_T + \lambda_{\alpha} (\rho_{\alpha} - \sum_{\beta} \frac{\lambda_{\beta} \rho_{\beta}}{\lambda_T}) \vec{g} - \lambda_{\alpha} (\nabla p_{c,\alpha} - \sum_{\beta} \frac{\lambda_{\beta} \nabla p_{c,\beta}}{\lambda_T}),$$

where $\lambda_T = \sum_{\alpha} \lambda_{\alpha}$ is the total mobility. Writing the conservation equations in terms of the total velocity yields the fractional flow form of the component conservation equations, given by

$$\phi \frac{\partial \mathbf{n}}{\partial t} + \nabla \cdot F(\mathbf{n}, v_T) = \nabla \cdot (H \nabla P_c) + \nabla \cdot \mathcal{D} + \mathcal{R}, \quad (5)$$

where

$$F(\mathbf{n}, v_T) = \sum_{\alpha} \frac{\mathbf{n}_{\alpha}}{u_{\alpha}} \left[\frac{\lambda_{\alpha}}{\lambda_T} v_T + \lambda_{\alpha} (\rho_{\alpha} - \sum_{\beta} \frac{\lambda_{\beta} \rho_{\beta}}{\lambda_T}) \vec{g} \right],$$

and

$$H \nabla P_c = \sum_{\alpha} \frac{\mathbf{n}_{\alpha}}{u_{\alpha}} \lambda_{\alpha} (\nabla p_{c,\alpha} - \sum_{\beta} \frac{\lambda_{\beta} \nabla p_{c,\beta}}{\lambda_T}).$$

In the absence of diffusive terms the total velocity defines a splitting that decomposes the dynamics into an elliptic pressure equation and a system of hyperbolic conservation laws. We note that we are assuming here that the

system $\phi \frac{\partial \mathbf{n}}{\partial t} + \nabla \cdot F(\mathbf{n}, v_T) = 0$ is hyperbolic. This is a condition on the model specification; some models for three-phase flow use relative permeability models that lead to non-hyperbolic behavior and the resulting problems are ill-posed in the zero-viscosity limit (see [3]). For purposes of exposition, we will set $\mathcal{D} = 0$ and $\mathcal{R} = 0$ for the remainder of the discussion.

2.1 Single grid algorithm

In this subsection, we discuss the single grid algorithm, giving an overview of the time stepping procedure. The discretization uses a volume of fluid approach in which \mathbf{n}_{ijk}^n denotes the average value of \mathbf{n} over cell- ijk at time t^n ; \mathbf{n} and p are defined on cell centers while F and v_T are defined on cell edges. The temporal discretization fits into the basic framework of IMPES-type methods in which the pressure equation is solved implicitly and the saturation (component) equations are solved explicitly. Here the basic algorithm is modified so that the overall splitting approach is second-order accurate in time and we treat the diffusion terms semi-implicitly so that diffusive terms, $\nabla \cdot \mathcal{D}$ and $H \nabla P_c$, do not limit the time step.

The major steps of the algorithm are:

Step 1: Solve the pressure equation, (4), rewritten in the form

$$D \cdot \sum_{\alpha} \lambda_{\alpha}(Gp) = D \cdot \sum_{\alpha} \lambda_{\alpha} \rho_{\alpha} \vec{g} - D \cdot \sum_{\alpha} \lambda_{\alpha}(Gp_{c,\alpha}) ,$$

for p with properties evaluated using \mathbf{n}^n . We then use equation (3) to define a total velocity v_T^n . Here D and G are second-order accurate discretizations of the divergence and gradient operators, respectively. The divergence operator returns a cell-centered divergence from face-centered values; the gradient operator differences cell-centered values to return normal gradients on faces.

Step 2: Use (5) to advance the solution from time t^n to time t^{n+1} using v_T^n . We use a second-order Godunov scheme to compute the hyperbolic fluxes and a Crank-Nicolson discretization of the diffusive terms, so that

$$\phi \frac{\mathbf{n}^{n+1,*} - \mathbf{n}^n}{\Delta t} + GF^{n+\frac{1}{2}}(\mathbf{n}, v_T^n) = \frac{1}{2} \left(D(H^n GP_c^n) + D(H^{n+1,*} GP_c^{n+1,*}) \right) (6)$$

with a suitable linearization of the coefficients of the diffusion term; here we use $H^{n+1,*}$ and $P_c^{n+1,*}$ to denote H and P_c , respectively, which are functions of \mathbf{n} , evaluated at $\mathbf{n}^{n+1,*}$. We note here that although we have dropped \mathcal{D} and \mathcal{R} terms from the equations for clarity of exposition, additional diffusion terms can be added analogously to the capillary pressure term. If reactions are

included in the system they can be included using the operator-split formalism discussed in Day and Bell [8]. In this step, $F^{n+\frac{1}{2}}(\mathbf{n}, v_T^n)$ denotes time-centered fluxes computed using the Godunov procedure but with the total velocity evaluated at t^n .

Step 3: Using the $\mathbf{n}^{n+1,*}$ data, re-solve the pressure equation (4) to compute a new total velocity v_T^{n+1} from (3). Define $v_T^{n+\frac{1}{2}} = \frac{1}{2}(v_T^n + v_T^{n+1})$.

Step 4: Use (5) to re-advance the solution from time t^n to time t^{n+1} , this time using $v_T^{n+\frac{1}{2}}$ to obtain values of \mathbf{n}^{n+1} ,

$$\phi \frac{\mathbf{n}^{n+1} - \mathbf{n}^n}{\Delta t} + GF^{n+\frac{1}{2}}(\mathbf{n}, v_T^{n+\frac{1}{2}}) = \frac{1}{2} \left(\nabla \cdot H^n GP_c^n + D \cdot H^{n+1} GP_c^{n+1} \right) \quad (7)$$

again with a suitable linearization of the coefficients of the diffusion term.

We postulate that Step 1 may not be critical; it may be acceptable to replace v_T^n with the final v_T computed in Step 3 of the previous time step without significant loss of accuracy, thus allowing us to avoid the computational cost associated with Step 1. In this case, v_T^n is computed with $\mathbf{n}^{n,*}$ instead of \mathbf{n}^n . We shall demonstrate with some examples in Section 4 that this has no noticeable effects on our solution. For clarity, we shall denote a time step using Step 1 as a *full solve* and one without as a *partial solve*.

3 Adaptive Mesh Refinement

In this section we present the extension of the algorithm described above to an adaptive hierarchy of nested rectangular grids. First we describe the creation of the grid hierarchy and the regridding procedure used to adjust the hierarchy during the computation. Next we discuss the adaptive time step algorithm with subcycling in time, focusing on the synchronization between different levels of refinement.

3.1 Creating and Managing the Grid Hierarchy

The grid hierarchy is composed of different levels of refinement ranging from coarsest ($\ell = 0$) to finest ($\ell = \ell_{max}$). Each level is represented as the union of rectangular grid patches of a given resolution. In this implementation, the refinement ratio is always even, with the same factor of refinement in each coordinate direction, i.e. $\Delta x^{\ell+1} = \Delta y^{\ell+1} = \Delta z^{\ell+1} = \frac{1}{r} \Delta x^\ell$, where r is the

refinement ratio. (We note here that neither isotropic refinement nor uniform base grids are requirements of the fundamental algorithm; see the section on future work.) In the actual implementation, the refinement ratio, either 2 or 4, can be a function of level; however, in the exposition we will assume that r is constant. The grids are properly nested, in the sense that the union of grids at level $\ell + 1$ is contained in the union of grids at level ℓ for $0 \leq \ell < \ell_{max}$. Furthermore, the containment is strict in the sense that, except at physical boundaries, the level ℓ grids are large enough to guarantee that there is a border at least one level ℓ cell wide surrounding each level $\ell + 1$ grid. (Grids at all levels are allowed to extend to the physical boundaries so the proper nesting is not strict there.)

The initial creation of the grid hierarchy and the subsequent regridding operations in which the grids are dynamically changed to reflect changing flow conditions use the same procedures as were used by Bell et al. [2] for hyperbolic conservation laws. The construction of the grid hierarchy is based on error estimation criteria specified by the user to indicate where additional resolution is required. The error criteria are currently based on tracking component density gradients for one of the components; however, more sophisticated criteria based on estimating the error can be used (see, e.g., [4]). Given grids at level ℓ we use the error estimation procedure to tag cells where the criteria for further refinement are met. The tagged cells are grouped into rectangular patches using the clustering algorithm given in Berger and Rigoutsos [5]. These rectangular patches are refined to form the grids at the next level. The process is repeated until either the error tolerance criteria are satisfied or a specified maximum level is reached. The proper nesting requirement is imposed at this stage.

At $t = 0$ the initial data is used to create grids at level 0 through ℓ_{max} . (Grids have a user-specified maximum size, therefore more than one grid may be needed to cover the physical domain.) As the solution advances in time, the regridding algorithm is called every k_ℓ (also user-specified) level ℓ steps to redefine grids at levels $\ell+1$ to ℓ_{max} . Level 0 grids remain unchanged throughout the calculation. Grids at level $\ell + 1$ are only modified at the end of level ℓ time steps, but because we subcycle in time, i.e. $\Delta t^{\ell+1} = \frac{1}{r} \Delta t^\ell$, level $\ell + 2$ grids can be created and/or modified in the middle of a level ℓ time step if $k_{\ell+1} < r$.

When new grids are created at level $\ell + 1$, the data on these new grids are copied from the previous grids at level $\ell + 1$ if possible, or interpolated in space from the underlying level ℓ grids otherwise. After regridding we always recalculate v_T from the new data, thus the first fine time step after regridding is always a *full solve* rather than a *partial solve*. We note here that while there is a user-specified limit to the number of levels allowed, at any given time in the calculation there may not be that many levels in the hierarchy, i.e. ℓ_{max} can change dynamically as the calculation proceeds, as long as it does not

exceed the user-specified limit.

3.2 Overview of Time-Stepping Procedure

The adaptive time stepping approach uses temporal subcycling so that each level is advanced independently at its own time step ($\Delta t^{\ell+1} = \frac{1}{r} \Delta t^\ell$), requiring no interlevel communication other than the supplying of Dirichlet data from a coarse level to be used as boundary conditions at the next finer level. When coarse and fine data reach the same time point, then the data at the different levels are synchronized. The synchronization approaches are based on the same set of algorithmic ideas as those developed in Almgren et al. [1]

The adaptive time-step algorithm can most easily be thought of as a recursive procedure, in which to advance level ℓ , $0 \leq \ell \leq \ell_{max}$ the following steps are taken:

- Advance level ℓ in time as if it is the only level. Supply boundary conditions for component densities and pressure from level $\ell - 1$ if level $\ell > 0$, and from the physical domain boundaries.
- If $\ell < \ell_{max}$
 - Advance level $(\ell + 1)$ r times with time step $\Delta t^{\ell+1} = \frac{1}{r} \Delta t^\ell$. Use boundary conditions for component densities and pressure from level ℓ and from the physical domain boundaries.
 - Synchronize the data between levels ℓ and $\ell + 1$, and interpolate corrections to higher levels if $\ell + 1 < \ell_{max}$.

Before describing the steps of the synchronization in detail, we first discuss, in general terms, how to synchronize the data at different levels so that the solution as computed on each level sequentially can most closely approximate the solution which would be found using composite solves. During the advance of each level, for each operator we supply Dirichlet boundary data for the fine grids from the next coarser grid. This implies that the values at both levels are consistent, but the computed fluxes at the coarse/fine interfaces are not. The synchronization solves correction equations that account for discrepancies in fluxes between levels. The corrections equations reflect the type of operator being corrected. For hyperbolic equations the correction of flux discrepancies is a simple explicit flux correction as discussed in Berger and Colella [4]. For a self-adjoint elliptic operator, the discrepancy in the fluxes represents a discontinuity in normal derivative at the coarse/fine boundary and the correction equation takes the form of an discrepancy layer potential problem (see [1]).

After the level $\ell + 1$ data have been advanced to the same point in time as the level ℓ data, there are three sources of discrepancy in the composite solution

that need to be corrected in the synchronization step.

- (D.1) The data at level ℓ that underlie the level $\ell + 1$ data are not synchronized with the level $\ell + 1$ data.
- (D.2) The composite total velocity computed from the pressure equation, defined as the time-averaged (over a level ℓ time step) level $\ell + 1$ advection velocity on all level $\ell + 1$ faces including the $\ell/(\ell + 1)$ interface, and the level ℓ advection velocity on all other level ℓ faces, does not satisfy the composite divergence constraint at the $\ell/(\ell + 1)$ interface. This mismatch results in spatially constant advected quantities with no source terms not remaining constant.
- (D.3) The advective and diffusive fluxes from the level ℓ faces and the level $\ell + 1$ faces do not agree at the $\ell/(\ell + 1)$ interface, resulting in a loss of conservation.

The aim of the synchronization steps is to correct the effects of each mismatch. Discrepancy (D.1) is easily corrected by averaging the level $\ell + 1$ data onto the level ℓ data beneath. We denote this correction by (S.1). The second discrepancy, (D.2), is discretely manifested as a non-zero difference between the coarse grid total velocity and the effective time-averaged fine grid total velocity at the coarse/fine interface. This difference results from not having satisfied the elliptic matching conditions at the coarse/fine interface during the pressure solve. An elliptic solve is necessary to correct for the discrepancy. We perform a level ℓ “pressure sync solve,” (S.2) with the right-hand-side defined as the divergence of the mismatch between the level ℓ and the time-averaged level $\ell + 1$ total velocity. From this solve we define a total velocity correction field that is used to modify the explicit dependence of the hyperbolic flux terms on the total velocity. These “re-advection corrections,” as well as the interpolation of these corrections to all higher levels, are combined with the refluxing corrections to modify the solution as described below.

In the case of zero viscosity/diffusivity, the re-advection corrections described immediately above, and the correction for (D.3), which is simply a hyperbolic refluxing step, are added directly to the new-time solution. The refluxing corrections modify the solution only on the coarse grid cells immediately outside the fine grids; the re-advection corrections modify the solution at all cells at level ℓ and higher. However, in the case of non-zero viscosity/diffusivity, the modification of the solution by the re-advection and refluxing corrections requires solving additional elliptic equations (S.3).

In the single grid algorithm, hyperbolic and diffusive fluxes are not added directly to the solution update; rather, they form part of the right-hand-sides for the parabolic solves associated with the Crank-Nicolson discretization of the diffusive terms. Similarly, in the synchronization step, the hyperbolic and diffusive flux discrepancies, as well as the re-advection corrections, define the right-hand-sides for the refluxing solves. The solutions to the elliptic refluxing

equations at level ℓ will modify the new-time total velocity and component densities on all grids at level ℓ , and the interpolation of these corrections will modify the new-time data on all grids at all higher levels.

3.3 Details of Time-Stepping Procedure

Assume now that we are advancing level ℓ , $0 \leq \ell \leq \ell_{max}$, one level ℓ time step. We now add the level index, ℓ , to the superscript of each quantity, defining, for example, $\mathbf{n}^{n,\ell}$ as the component density on level ℓ at time t^n . We define Δt^ℓ as the time step of the level ℓ grid, A^ℓ as the area of a face at level ℓ and Vol^ℓ as the volume of a grid cell at level ℓ .

3.3.1 Advancing a single level

To advance the data on level ℓ one level ℓ time step, we follow the time-stepping procedure as described for the single grid algorithm in the previous section. For the level ℓ advancement we advance all the grids at that level simultaneously. Explicit advection work is decoupled except for the exchange of boundary conditions. When a coarse/fine boundary does not coincide with a physical domain boundary for the level ℓ advection step, level $\ell - 1$ data are interpolated linearly in time to specify Dirichlet boundary conditions at the coarse-fine boundary. The procedures used for these interpolations are the same as those discussed in Almgren et al. [1]

3.3.2 Computing the Coarse-Fine Discrepancy

Over the course of a level ℓ time step, we must accumulate several quantities at the $\ell/(\ell + 1)$ interface in order to correctly capture the flux discrepancies at the end of the level ℓ time step. We refer to the face-based data structures that contain these quantities as registers. The total velocity and flux registers accumulate the discrepancies between the level ℓ and level $\ell + 1$ face-based total velocity v_T and fluxes F , respectively.

These registers are defined only on the $\ell/(\ell + 1)$ interface and are indexed by level ℓ indices. Note that in d dimensions, one level ℓ face contains r^{d-1} level $\ell + 1$ faces; the sums over faces below should be interpreted as summing over all level $\ell + 1$ faces which are contained in the level ℓ face. The sums over k should be understood as summing over the r level $\ell + 1$ time steps contained within a single level ℓ time step.

At the end of the level ℓ time step, the velocity register (δv_T^ℓ) holds the area-weighted difference between the total velocity at level ℓ and the time average

over one level ℓ time step of the space average over the area of the level ℓ face of the total velocity at level $\ell + 1$:

$$A^\ell \delta v_T^\ell = -A^\ell v_T^{n+\frac{1}{2},\ell} + \frac{1}{r} \sum_{k=1}^r \sum_{faces} (A^{\ell+1} v_T^{k+\frac{1}{2},\ell+1}).$$

The flux registers $\delta F^{hyp,\ell}$ and $\delta F^{visc,\ell}$ contain the area-weighted differences in the hyperbolic and diffusive fluxes calculated at level ℓ and the time average over the level ℓ time step of the space average over the area of the level ℓ face of the advective fluxes at level $\ell + 1$:

$$A^\ell \delta \mathcal{F}^{hyp,\ell} = -A^\ell \mathcal{F}^{hyp,\ell} + \frac{1}{r} \sum_{k=1}^r \sum_{faces} (A^{\ell+1} \mathcal{F}^{hyp,k,\ell+1}),$$

$$A^\ell \delta \mathcal{F}^{visc,\ell} = -A^\ell \mathcal{F}^{visc,\ell} + \frac{1}{r} \sum_{k=1}^r \sum_{faces} (A^{\ell+1} \mathcal{F}^{visc,k,\ell+1}),$$

where

$$\begin{aligned} \mathcal{F}^{hyp,\ell} &= F^{n+\frac{1}{2},\ell}(\mathbf{n}^{n+\frac{1}{2},\ell}, v_T^{n+\frac{1}{2},\ell}), \\ \mathcal{F}^{visc,\ell} &= \frac{1}{2} \left(H^{n,\ell} GP_c^{n,\ell} + H^{n+1,\ell} GP_c^{n+1,\ell} \right), \end{aligned}$$

as computed in Step 4 of the algorithm discussed above. In these formulae, the k index indicates the subcycled time steps at the finer level.

We note here that the signs of the quantities added to the flux registers actually depend on the orientation of the normal facing away from the fine grid. We follow the convention below that the signs are given for the faces at which the fine grid is in the direction of the lower coordinate indices.

3.3.3 Synchronization of data

The first synchronization step, (S.1), has already been described in Section 3.2. Here we give the details of the other two steps of the synchronization.

The discrepancy in the total velocity, (D.2), is captured in δv_T^ℓ ; the divergence of δv_T^ℓ defines the right-hand-side for the level ℓ elliptic sync solve (S.2). We solve

$$-D(\lambda_T G(\delta e^\ell)) = \widetilde{D}(\delta v_T^\ell)$$

on all grids at level ℓ for the correction δe^ℓ . Recall δv_T^ℓ is defined only at the

coarse/fine interface; here \widetilde{D} is defined to be the discrete divergence operator evaluated only on the level ℓ cells adjacent to the interface but not underlying any level $\ell + 1$ grids. (At level ℓ cells underlying level $\ell + 1$ grids the right-hand-side is zero.) Boundary conditions on physical no-flow boundaries are homogeneous Neumann ($\frac{\partial(\delta e)^\ell}{\partial n} = 0$); on outflow $\delta e^\ell = 0$. If $\ell > 0$, the boundary conditions for δe^ℓ are given as homogeneous Dirichlet conditions on the level $\ell - 1$ cells outside the level ℓ grids. We then define the correction velocity field from δe^ℓ :

$$v_T^{\ell,corr} = -\lambda_T G(\delta e^\ell).$$

We now use the correction velocity field to define flux corrections at all level ℓ faces. Because of memory considerations we do not store all the time-centered face states, so we must redefine these on all level ℓ faces: we recreate $\mathcal{F}^{hyp,\ell}$ using $v_T^{n+\frac{1}{2},\ell}$ for upwinding, identically to the procedure used in forming the update, and define $\mathcal{F}_{v_T}^{corr,\ell}$ as

$$\mathcal{F}_{v_T}^{corr,\ell} = F^{hyp,n+\frac{1}{2},\ell}(\mathbf{n}^n, v_T^{n+\frac{1}{2},\ell}) - F^{hyp,n+\frac{1}{2},\ell}(\mathbf{n}^n, v_T^{n+\frac{1}{2},\ell} + v_T^{\ell,corr}) .$$

Because we must diffuse the re-advection and refluxing corrections before adding them to the new-time solution, we do not add them directly to the solution. Rather, the divergence of the re-advection flux corrections is added to the advective and viscous/diffusive flux mismatches to define the cell-centered right-hand-sides for the refluxing solves (S.3):

$$RHS_{sync}^\ell = -D\mathcal{F}_{v_T}^{corr,\ell} - \frac{A^\ell}{Vol^\ell}(\delta\mathcal{F}^{hyp,\ell} + \delta\mathcal{F}^{visc,\ell}). \quad (8)$$

Then, we solve for the correction to the solution, \mathbf{n}_{sync}^ℓ

$$\phi\mathbf{n}_{sync}^\ell - \frac{\Delta t}{2}D(H^sGP_c^s) = RHS_{sync}^\ell . \quad (9)$$

where $H^s = H(\mathbf{n}^{n+1} + \mathbf{n}_{sync}^\ell)$ and $P_c^s = P_c(\mathbf{n}^{n+1} + \mathbf{n}_{sync}^\ell)$.

If $\ell > 0$, we must now modify the level $\ell - 1$ velocity registers and flux registers to account for the corrections to the solution due to the re-advection corrections as well as the diffused corrections. This is analogous to the accumulation of advective and diffusive fluxes during the advance of a single level. To do this, we set

$$\begin{aligned}
A^{\ell-1} \delta v_T^{\ell-1} &:= A^{\ell-1} \delta v_T^{\ell-1} + \frac{1}{r} \sum_{faces} A^\ell v_T^{\ell,corr}, \\
A^{\ell-1} \delta \mathcal{F}^{hyp,\ell-1} &:= A^{\ell-1} \delta F^{hyp,\ell-1} + \frac{1}{r} \sum_{faces} A^\ell \mathcal{F}^{corr,\ell}, \\
A^{\ell-1} \delta \mathcal{F}^{visc,\ell-1} &:= A^{\ell-1} \delta \mathcal{F}^{visc,\ell-1} + \frac{1}{r} \sum_{faces} \left\{ \frac{1}{2} A^\ell H^s G P_c^s \right\}.
\end{aligned}$$

The corrected solution at level ℓ is then given by

$$\mathbf{n}^{n+1,\ell} := \mathbf{n}^{n+1,\ell} + \mathbf{n}_{sync}^\ell.$$

If $\ell < \ell_{max}$, we interpolate the correction onto the fine grids at all finer levels, $q, \ell < q \leq \ell_{max}$ using conservative interpolation:

$$\mathbf{n}^{n+1,q} := \mathbf{n}^{n+1,q} + \text{Interp}_{cons}(\mathbf{n}_{sync}^\ell).$$

This completes the synchronization steps for scalar quantities.

4 Numerical Results

In this section, we examine the numerical properties of the method described in the previous section. In the first subsection we examine the convergence behavior of the method. While the basic discretization presented here is similar to a traditional IMPES fractional step discretization, we have modified the method such that the total velocity is centered in time. This makes the method formally second-order accurate, in contrast to the standard IMPES approach. In our first example, we use a stable one-dimensional problem based on a two-component single-phase system to demonstrate the method's second-order rate of convergence and to illustrate the loss of accuracy associated with the standard IMPES-type fractional step scheme. We then demonstrate the second-order convergence rate of the algorithm with a two-dimensional problem and show that the *partial solve* algorithm maintains the overall accuracy of the method relative to the *full solve* version. We also examine the effects of centering the total velocity and compare the current approach with the standard IMPES discretization when the flow is unstable. In the second subsection, we compare solutions obtained on a uniform grid with those computed with AMR. For this purpose, we use a more complicated example involving a three-component two-phase system. In the final two subsections, we simulate a large three-dimensional problem using the parallelized AMR algorithm, and present the scaling behavior. We note that for the numerical results presented here, we neglect capillary pressure.

4.1 Convergence

For our first study we consider a single-phase two-component system where components 1 and 2 are completely miscible, and form a mixture with viscosity given by

$$\mu = \frac{1}{(1 - c + M^{1/4}c)^4} , \quad (10)$$

where $c = n_1/(n_1 + n_2)$ and $M = 10$. In our algorithm as presented, both the pressure discretization and the component conservation equations are discretized to second-order accuracy. In the standard IMPES formulation, the total velocity is not centered in time, which introduces a small first-order temporal error. In most realistic problems, this error is typically dominated by spatial truncation errors, so we examine the problem in the context of one dimension where we can use fine spatial resolution. We take as initial data $c(x) = 0.5(1 + \tanh((0.3 - x)/0.05))$ on the interval $[0, 2]$ and impose a constant pressure drop. The time step, which is fixed for each resolution, is set so that the CFL is approximately 0.35. In addition, we turn off the flux limiting in the advection scheme.

The L^1 norm of the errors for the standard IMPES discretization and the *full solve* version of the new algorithm are presented in Table 1. The new algorithm shows clear second-order convergence whereas the standard IMPES algorithm is only first-order accurate asymptotically. At the coarsest resolution, the two approaches have nearly the same error and the IMPES discretization shows nearly second-order convergence between the two coarsest discretizations. At these coarser resolutions the temporal discretization error is dominated by the spatial truncation error, which is second-order accurate in both cases.

Next, we examine the convergence properties of the single grid algorithm in two dimensions. For this problem we consider a domain $[0, l_x] \times [0, l_y]$ where $l_x = 16$ and $l_y = 4$. The domain is initially filled with component 2; component 1 is injected from the left edge of the domain at a rate given by $n_1(0, y; t) = \tau \rho_1$ and $n_2(0, y; t) = (1 - \tau) \rho_2$ where $\tau(t) = 0.5(1 + \tanh(t/10^6))$. The densities of component 1 and 2 are 1000 and 800, respectively. We impose no-flow conditions on the top and bottom edges and a pressure difference of 2 between the inlet and the outlet. We define $M = 2$. The porosity is uniformly 1 while the permeability function $\kappa(x, y)$ is given by

$$\kappa(x, y) = 10^{-5}(1 + \sin(\pi y/l_y)^4). \quad (11)$$

We note that $\tau(t)$ and $\kappa(x, y)$ are chosen so that the solution will be smooth. Gravity points in the negative y direction. The calculation is run for 1.2×10^7 s

with $\Delta x = \Delta y = 1/2^n$ for $n = 3, 4, 5, 6$. Fixed time step and uniform grids are used for the purpose of evaluating the convergence of the solution.

Table 2 shows the discrete L^∞ , L^1 and L^2 norms of the difference between the solution \mathbf{n} obtained on each grid and that obtained on the next finer grid, and the resulting convergence rates. The rate between the two columns of error norms is defined as $\log_2(\varepsilon_l/\varepsilon_r)$ where ε_l and ε_r are the errors shown in the left and right columns. Table 2 clearly demonstrates the second-order convergence property of the algorithm when measured in the L^1 and L^2 norms. The convergence rate of the L^∞ error is less than 2; this can be attributed to the discontinuity in the solution resulting from limited smoothness of the κ function. Table 3 shows the same data but for calculations done using the *full solve* approach. We see by comparing the two tables that the errors obtained from the different approaches are very similar. We note that the omission of Step 1 does not adversely affect the convergence rate of our method.

The next set of calculations compares the solution obtained using our method to that obtained using a first-order method when the solution is no longer smooth. The goal is to show that even though the flow is no longer in the asymptotically second-order regime of the algorithm, the accuracy properties of the formally second-order method still have important consequences for the behavior of the flow. We construct the first-order method by eliminating Steps 3 and 4 from the method presented in Section 2.1, and setting $\mathbf{n}^{n+1} = \mathbf{n}^{n+1,*}$.

However, the error in the solution cannot be directly measured by taking the norm of the difference between two solutions. We instead compare the mixing length given by the length of the fingered zone. Following [14], we define the mixing length, L_δ , as

$$L_\delta = x|_{\bar{c}=\delta} - x|_{\bar{c}=1-\delta} \quad (12)$$

where

$$\bar{c}(x) = \frac{\int_0^{l_y} n_1 dy}{\int_0^{l_y} n_1 dy + \int_0^{l_y} n_2 dy} \quad (13)$$

and $\delta = 1 \times 10^{-5}$. We first look at a stable fingering profile that allows us to establish a hypothesis on how mixing length is related to accuracy. We then examine an example where the flow is highly unstable.

For the first example, we again look at the two-component single-phase system described in the previous section. We consider here a uniform $\kappa(x, y) = 2.0 \times 10^{-5}$ and a pressure difference of 0.5 between the inlet and the outlet. We see in Figure 1 that a stable finger extends from the inlet to the outlet at the bottom of the simulation domain. We then measure the mixing length for three different grid sizes and three different approaches. The grids we consider are 512×128 , 1024×256 and 2048×512 . The approaches are first-order, second-order with the *partial solve* algorithm, and second-order with the *full solve* approach. The solution obtained from the 2048×512 grid and the *full solve* approach is used as the reference solution. The simulation is run to a final time of 8×10^6 s. As shown in Table 4, the mixing length increases with accuracy, suggesting that mixing length can be used as a measure of accuracy. We thus hypothesize that for a given resolution, the solution with a larger mixing length at a fixed time is the more accurate solution. We note that this metric does not extend across different resolutions, since instabilities in the flow may behave differently at different resolutions, leading to potentially different flow profiles — the mixing length is then no longer an appropriate measure of accuracy.

To amplify the difference between the first-order method and the second-order methods, we change the permeability distribution function $\kappa(x, y)$ to

$$\kappa(x, y) = \begin{cases} 1.5 \times 10^{-5}, & y/l_y < .25; \\ 3.0 \times 10^{-5}, & 0.25 \leq y/l_y < .5; \\ 1.0 \times 10^{-5}, & 0.5 \leq y/l_y < .75; \\ 2.0 \times 10^{-5}, & y/l_y \geq .75. \end{cases} \quad (14)$$

In addition, M is increased to 10. Figure 2 shows that the viscous fingers resulting from the first-order method and the second-order methods have some notable differences. Slight variations in the solution can develop into distinct fingering profiles due to instabilities in the solution. To further illustrate this point, we examine a two-component single-phase system in a domain where $\kappa = 10^5 (1 + 10^{-8}\epsilon)$ and ϵ is a random number between 0 and 1. Figure 3 shows that two different sets of ϵ lead to two different fingering profiles. This suggests that any minor perturbation can lead to variation in the details of the fingering. However, we again see that the omission of Step 1 does not have significant effects on the solution. In Table 5, we list the mixing lengths for grid sizes 512×128 and 1024×256 computed with the three approaches. We clearly see that mixing length is larger for the second-order methods, in keeping with the hypothesis proposed earlier.

Results of the subsequent sections are based on the second-order method with the *partial solve* algorithm.

4.2 Comparison between AMR and Uniform Grid

The next set of calculations compare the composite grid solution with the uniform fine grid solution for a two-phase, three-component system. Phase α , consisting of component 1 and 2, is injected into the domain filled with phase β consisting of component 3. We impose the same boundary conditions as in the previous two subsections. For this problem, we have $\rho_1 = \rho_2 = 1000$, $\mu_\alpha = 0.175 + 0.35(\frac{n_2}{n_1+n_2})$, $\rho_3 = 800$, and $\mu_\beta = 0.35$. The porosity is again uniformly 1. The permeability function $\kappa(x, y)$ is given by (14) and the injection rate is given by

$$n_1 = \begin{cases} 100, & 0 < t < 5 \times 10^5 \\ 900, & t \geq 5 \times 10^5 \end{cases} ;$$

$n_2 = 1000 - n_1$, $t > 0$; and $n_3 = 0$, $t > 0$. Again, the flow is not in the asymptotically second-order regime of the algorithm. The goal here is to show that with the adaptive algorithm one can achieve accuracy comparable to that on a uniform fine grid. The regridding criterion, which flags coarse grid cells for refinement, is based on the magnitude of $|\nabla n_1|$.

We note that although the composition for $t < 5 \times 10^5$ will lead to a stable front, the composition for $t \geq 5 \times 10^5$ will lead to a highly unstable front. A direct comparison between the solution after $t = 5 \times 10^5$ is then not meaningful for reasons elucidated in the previous section. It nevertheless allows us to examine the behavior of both propagation modes of the aqueous phase.

For the uniform grid, we use a grid size of 1024×256 . For the AMR grid, we have three levels of refinement, and a factor of two refinement between each level. Thus the base grid of the AMR example has size 256×64 . A particular cell is tagged for refinement if $|\nabla n_1| > 500$. For both simulations, the time step is computed with a CFL number of 0.4. Figure 4 shows that stable solution fronts for the uniform and the adaptive simulations are identical. However, while the unstable fronts exhibit similar features, there are notable differences in the fingering structure. Unstable fronts can thus develop differently even though captured at the same effective resolution, as shown by the uniform and adaptive solutions shown in Figure 5.

In Table 6, we compare the computational time needed to solve the current problem on a uniform grid and an adaptive grid. We note that at 7×10^5 s, AMR approach reduces the computational time by 88%. However, as the % of domain refined increases, the performance degrades. At 10^6 s, the computational saving achieved by the AMR approach is reduced to 80%, when the % of domain refined at level 2 increases to 22.5%. Thus, the efficiency gains from using AMR will be problem-dependent.

We note that the implementation of the AMR algorithm involves a flux correction strategy, as described in Section 3.3.3. However, the velocity v_T is not corrected as its procedure will be considerably more involved. If the correction required is large, this can lead to large error in subsequent flux calculations and degrade the solutions. In the above example, the required maximum correction required is only 0.2% and 0.6% of v_T in the x and y -direction. This is sufficiently small that a correction is not needed. This certainly does not rule out the need of velocity correction for other problems, and it may be necessary to incorporate a velocity correction procedure in the future.

4.3 A three-dimensional example

We consider the three-component two-phase system for a three-dimensional problem. The domain is now given by $[0, 16] \times [0, 4] \times [0, 2]$. We impose no-flow boundary conditions along the faces with normals in the y and z direction, and a pressure difference of 2 in the x -direction. The properties of the components are the same as in Section 4.2, and porosity is again uniformly 1. The permeability function $\kappa(x, y, z)$ here is a log-random distribution with a mean of 1.5×10^{-5} , and a variance of 0.1. In addition, κ is highly correlated in the x and y directions but not in the z direction, as shown in Figure 6.

We perform the simulation on a Cray XT4 supercomputer, with up to 4 levels of refinement. At the coarsest level, the grid size is $128 \times 64 \times 16$, and the refinement ratio is 2 between all adjacent levels. The resulting grid and component density of component 1 at time 1.6×10^6 is shown in Figure 7.

4.4 Parallel Performance

Here, we examine the scaling of the proposed algorithm on parallel architectures. The weak-scaling study we use here is based on the two-component single-phase system with layered permeability function. By replicating the problem in the y and z direction, we are able to scale the problem size without modifying the problem characteristics. We first note that % of domain refined, as shown in the right figure of Figure 8, holds approximately constant at level 1, but increases only slightly at level 2 as the domain size increases. This implies that the computational load on each CPU should be approximately constant as number of processors increases. In left figure of Figure 8, we, however, see a slight deviation from the ideal scaling, resulting in a modest increase in execution time as we increase the number of processors from 4 to 256. We can attribute this increase to the elliptic solver (which does not scale well with number of processors as the explicit advection components of the

algorithm) and the increase in the communication cost (a four-fold increase going from 4 processors to 256 processors).

5 Conclusions and Future Work

In this paper, we have presented a structured adaptive mesh refinement algorithm for incompressible flows in porous media. The method is based on a total velocity splitting of the equations into an elliptic pressure equation and a system of hyperbolic conservation laws for the component densities. The resulting system is discretized using cell-centered differencing for the pressure equation and a second-order Godunov scheme for the component conservation equations. The adaptive algorithm uses subcycling in time, i.e., the ratio of the time step to the mesh spacing is constant across levels of refinement. Rather than solve a global composite equation for the pressure at every time step at the finest level, we perform single-level solves at each fine step then use an elliptic correction solve to synchronize the solution across levels. This leads to a considerable improvement in computational efficiency. The method has been implemented in parallel and shows excellent scalability up to 256 processors.

Our goal in future work will be to extend this approach to more realistic problems. As a first step in that direction, we will extend the current approach to include capillary pressure, diffusion and reactions. Beyond that, our target will be to extend the approach to more realistic fluid models that include compressibility, interphase mass transfer and thermal effects.

Table 1

L^1 errors and convergence rates of the method for a 1D two-component single-phase system.

h	IMPES	rate	NEW	rate
2/50	1.06e-2	–	9.92e-3	–
2/100	2.86e-3	1.89	2.44e-3	2.02
2/200	1.25e-3	1.19	5.97e-4	2.03
2/400	6.07e-4	1.04	1.50e-4	1.99
2/800	3.00e-4	1.01	3.58e-5	2.07

Table 2

L^∞ , L^1 and L^2 errors and convergence rates for n_1 in a 2-D two-component single-phase calculation with the *partial solve* algorithm.

n	3	rate	4	rate	5	rate	6
L^∞	7.980	1.38	3.0612	1.57	1.29	1.11	0.5960
L^1	0.752	1.91	2.005e-1	1.99	5.047e-2	1.95	1.306e-2
L^2	1.438	1.83	4.036e-1	1.88	1.097e-1	1.81	3.137e-2

Table 3

L^∞ , L^1 and L^2 errors and convergence rates for n_1 in a 2-D two-component single-phase calculation, using the *full solve* approach.

n	3	rate	4	rate	5	rate	6
L^∞	7.978	1.38	3.061	1.29	1.250	1.08	0.5913
L^1	0.752	1.91	2.006e-1	1.99	5.050e-2	1.98	1.282e-2
L^2	1.439	1.83	4.039e-1	1.88	1.098e-1	1.81	3.128e-2

Table 4

Comparing the mixing length L_δ at time 8×10^6 ; ϵ_r is the relative error when compared to solution obtained for grid size 1024×256 with the *full solve* second-order method.

grid size	512×128		1024×256		2048×512	
	L_δ	ϵ_r	L_δ	ϵ_r	L_δ	ϵ_r
1st-order	11.0313	7.23e-2	11.5000	3.29e-2	11.8828	6.57e-4
2nd-order, partial	11.0625	6.96e-2	11.5156	3.16e-2	11.8906	0
2nd-order, full	11.0625	6.96e-2	11.5156	3.16e-2	11.8906	—

Table 5

Comparing the mixing length L_δ at time 5×10^6 for an unstable two-component single-phase system.

grid size	L_δ	
	512×128	1024×256
1st-order	7.219	7.906
2nd-order, partial	7.500	8.219
2nd-order, full	7.531	8.250

Table 6
Comparison of the computational cost

	$t = 7 \times 10^5$	$t = 1 \times 10^6$
time taken for uniform grid	7233.5s	10461.9s
time taken for AMR grid	878.4s	2106.2s
% domain refined at level 1	26.4	45.6
% domain refined at level 2	11.9	22.5

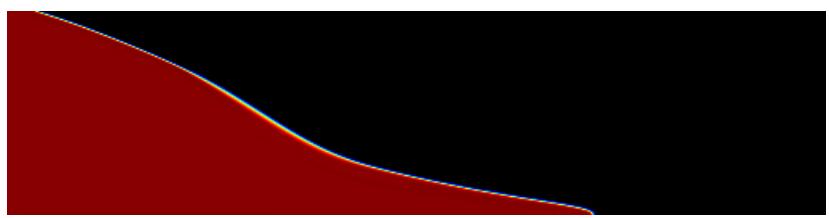


Fig. 1.

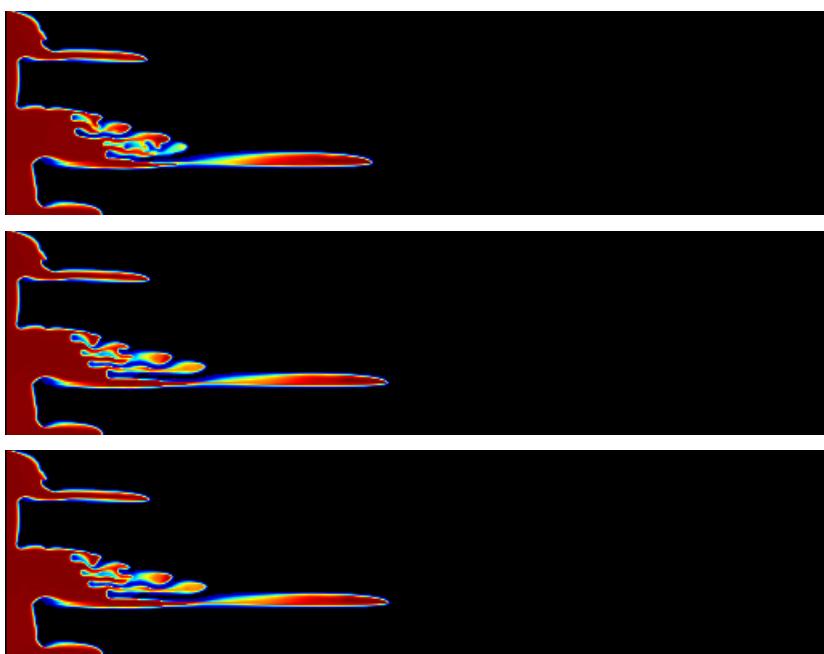


Fig. 2.

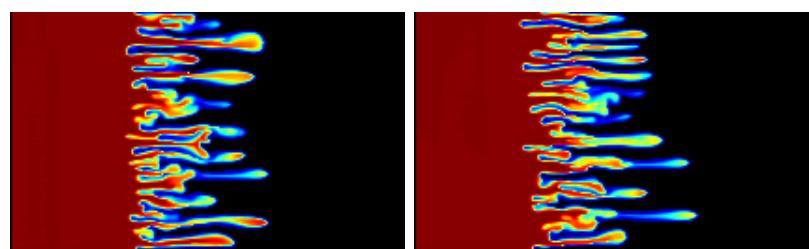


Fig. 3.

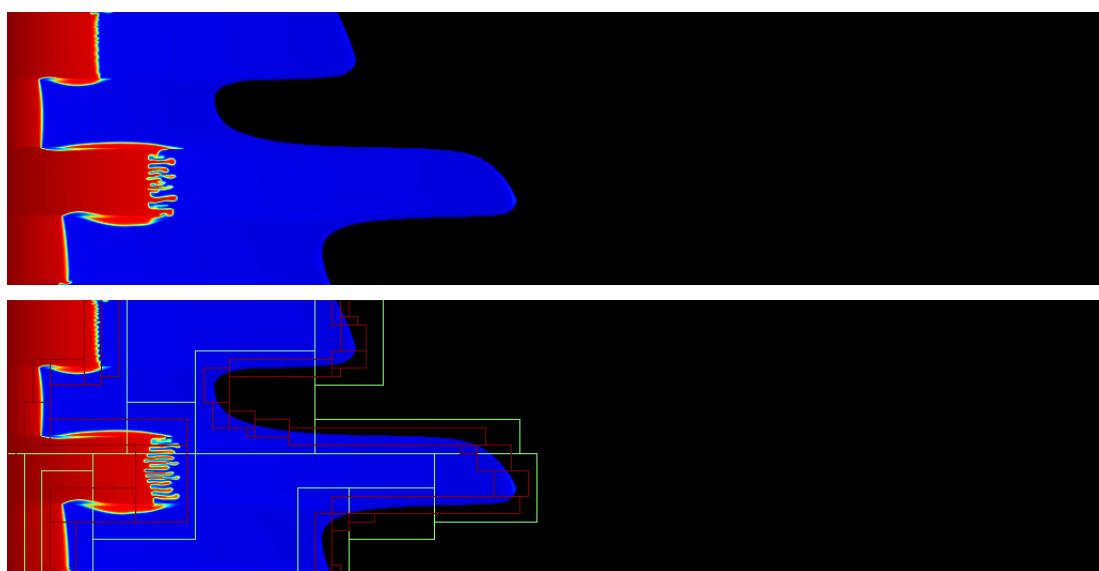


Fig. 4.

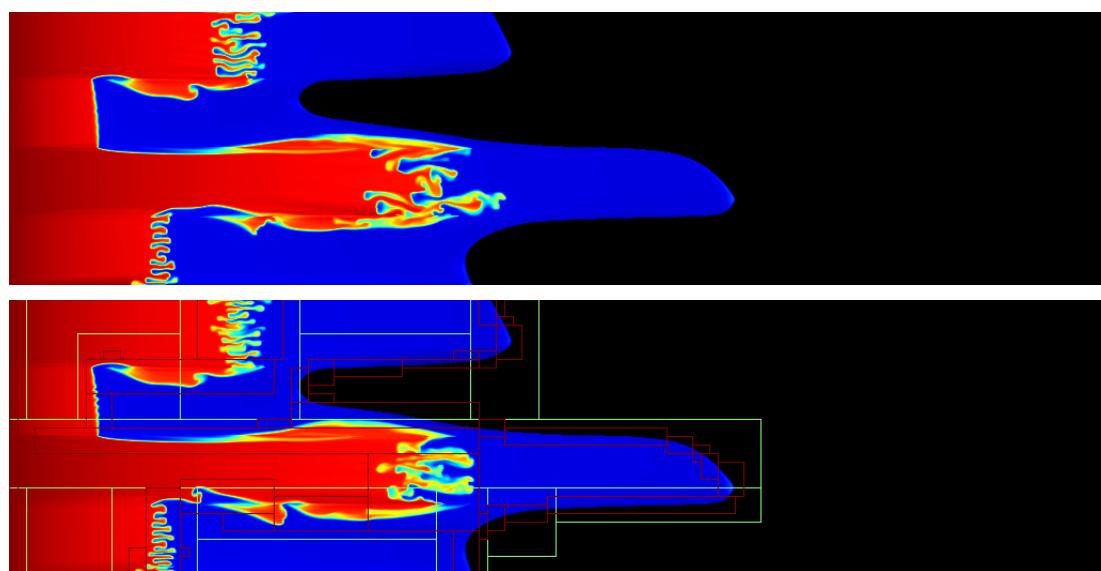


Fig. 5.

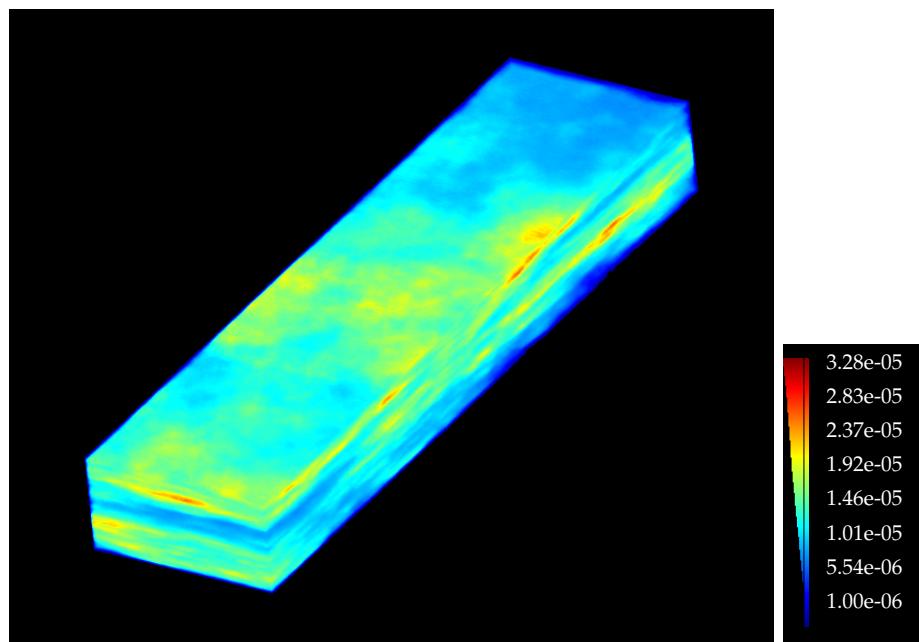


Fig. 6.

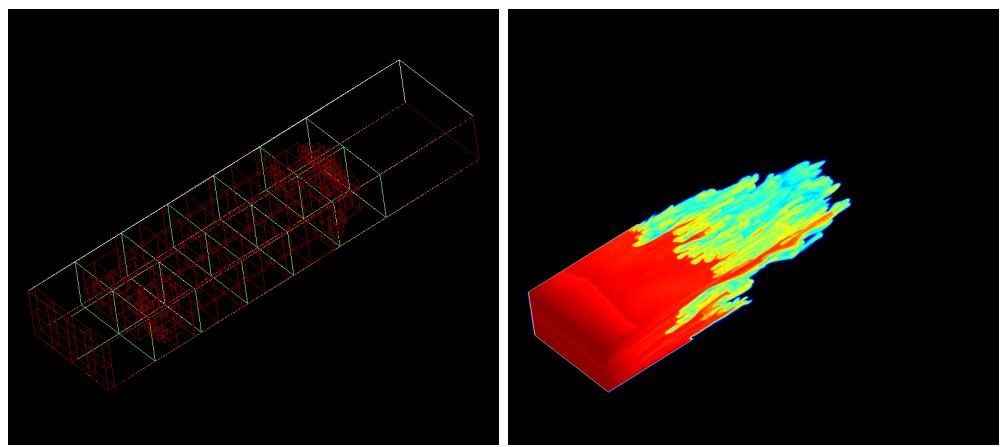


Fig. 7.

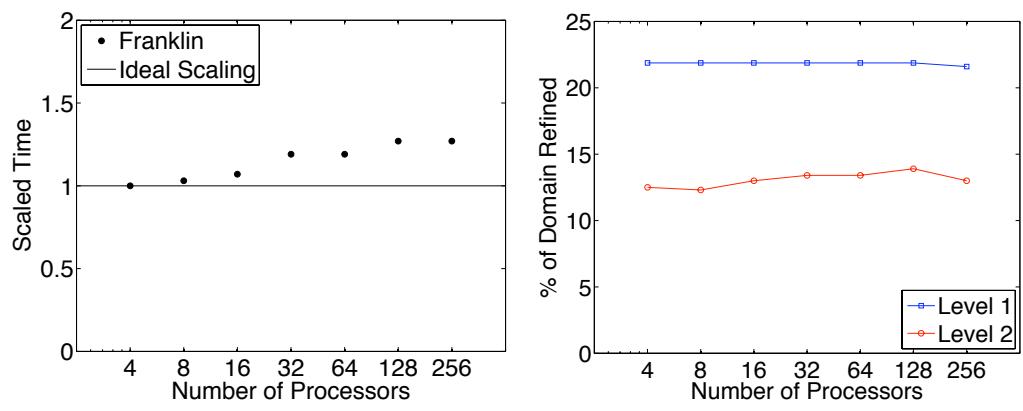


Fig. 8.

Figures:

- Fig. 1 Solution of n_1 at 8×10^6 .
- Fig. 2 Concentration of component 1 after time 5×10^6 based on first order method (top), second order method (middle), and second order method with Step 1 (bottom).
- Fig. 3 Fingering details of two different perturbation profiles of κ .
- Fig. 4 Concentration of water at time 7×10^5 , solved on an uniform grid (top) and adaptive grid (bottom).
- Fig. 5 Concentration of water at time 1×10^6 , solved on an uniform grid (top) and adaptive grid (bottom).
- Fig. 6 The permeability function used in the 3D example.
- Fig. 7 The adaptive grid (left) and solution at time 1.6×10^6 .
- Fig. 8 Parallel performance of the adaptive algorithm: Weak scaling behavior (Left) and % of domain refined (Right).

References

- [1] A. S. Almgren, J. B. Bell, P. Colella, L. H. Howell, and M. Welcome. A conservative adaptive projection method for the variable density incompressible Navier-Stokes equations. *J. Comp. Phys.*, 142:1–46, 1998.
- [2] J. B. Bell, M. J. Berger, J. S. Saltzman, and M. Welcome. Three-dimensional adaptive mesh refinement for hyperbolic conservation laws. *SIAM J. Sci. Statist. Comput.*, 15:127–138, 1994.
- [3] J. B. Bell, J. A. Trangenstein, and G. R. Shubin. Conservation laws of mixed type describing three-phase flow in porous media. *SIAM J. Appl. Math.*, 46(6):1000–1017, 1986.
- [4] M. J. Berger and P. Colella. Local adaptive mesh refinement for shock hydrodynamics. *J. Comp. Phys.*, 82(1):64–84, May 1989.
- [5] M. J. Berger and J. Rigoutsos. An algorithm for point clustering and grid generation. *IEEE Transactions on Systems, Man, and Cybernetics*, 21:1278–1286, 1991.
- [6] J.R. Christensen, G. Darche, B. Dechelette, H. Ma, and P.H. Sammon. Applications of dynamic gridding to thermal simulations. Paper SPE 86969. In *SPE International Thermal Operations and Heavy Oil Symposium and Western Regional Meeting*, Bakersfield, California, March 2004.
- [7] W. Y. Crutchfield. Load balancing irregular algorithms. Technical Report UCRL-JC-107679, Lawrence Livermore National Laboratory, July 1991.
- [8] M. S. Day and J. B. Bell. Numerical simulation of laminar reacting flows with complex chemistry. *Combust. Theory Modelling*, 4(4):535–556, 2000.
- [9] M. G. Edwards. A higher-order Godunov scheme coupled with local grid refinement for flow in a porous medium. *Comput. Methods Appl. Mech. Engrg.*, 131:287–308, 1996.
- [10] R. E. Ewing, B. A. Boyett, D. K. Babu, and R. F. Heinemann. Efficient use of refined grids for multiphase reservoir simulation. Paper SPE 18413. In *SPE Reservoir Simulation Symposium*, Houston, Texas, February 1989.
- [11] Z. E. Heinemann. Using local grid refinement in a multiple-application reservoir simulator. Paper SPE 12255. In *SPE Reservoir Simulation Symposium*, San Francisco, California, November 1983.
- [12] H. M. Hoang and J. Kleppe. A parallel adaptive finite difference algorithm for reservoir simulation. Paper SPE 99572. In *SPE Europe/EAGE Annual Conference and Exhibition*, Vienna, Austria, June 2006.
- [13] R. D. Hornung and J. A. Trangenstein. Adaptive mesh refinement and multilevel iterations for flow in a porous media. *J. Comp. Phys.*, 136:522–545, 1997.

- [14] O. Manickam and G. M. Homsy. Simulation of viscous fingering in miscible displacements with nonmonotonic viscosity profiles. *Physics of Fluids*, 6(1):95–107, 1994.
- [15] J. Nilsson, M. Gerritsen, and R. Younis. An adaptive, high-resolution simulation for steam-injection processes. Paper SPE 93881. In *SPE Western Regional Meeting*, Irvine, California, March 2005.
- [16] J. Nilsson, M. Gerritsen, and R. Younis. A novel adaptive anisotropic grid framework for efficient reservoir simulation. Paper SPE 93243. In *SPE Reservoir Simulation Symposium*, Woodlands, Texas, January 2005.
- [17] R. M. Propp. *Numerical Modeling of a Trickle Bed Reactor*. Ph.D., Dept. of Mechanical Engineering, Univ. of California, Berkeley, Sept. 1998.
- [18] C. A. Rendleman, V. E. Beckner, M. J. Lijewski, W. Y. Crutchfield, and John B. Bell. Parallelization of structured, hierarchical adaptive mesh refinement algorithms. *Computing and Visualization in Science*, 3(3):147–157, 2000.
- [19] P. H. Sammon. Dynamic grid refinement and amalgamation for compositional simulation. Paper SPE 79683. In *SPE Reservoir Simulation Symposium*, Houston, Texas, February 2003.
- [20] J. A. Trangenstein. Multi-scale iterative techniques and adaptive mesh refinement for flow in porous media. *Advances in Water Resources*, 25:1175–1213, 2002.
- [21] J. A. Trangenstein and Z. Bi. Large multi-scale iterative techniques and adaptive mesh refinement for miscible displacement simulation. Paper SPE 75232. In *SPE/DOE Improved Oil Recovery Symposium*, Tulsa, Oklahoma, April 2002.
- [22] Tianfu Xu, Eric Sonnenthal, Nicolas Spycher, and Karsten Pruess. TOUGHREACT user's guide: A simulation program for non-isothermal multiphase reactive geochemical transport in variably saturated geologic media. Technical Report LBNL-55460, Lawrence Berkeley National Laboratory, September 2004.