

A Multilevel Multiscale Mimetic (M^3) Method for Two-Phase Flows in Porous Media

K. Lipnikov* J. D. Moulton* D. Svyatskiy*

Abstract

We describe a multilevel multiscale mimetic (M^3) method for solving two-phase flow (water and oil) in a heterogeneous reservoir. The governing equations are the elliptic equation for the reservoir pressure and the hyperbolic equation for the water saturation. On each time step, we first solve the pressure equation and then use the computed flux in an explicit upwind finite volume method to update the saturation. To reduce the computational cost, the pressure equation is solved on a much coarser grid than the saturation equation. The coarse-grid pressure discretization captures the influence of multiple scales via the subgrid modeling technique for single-phase flow recently proposed in [23, 17, 22]. We extend significantly the applicability of this technique by developing a new robust and efficient method for estimating the flux coarsening parameters. Specifically, with this advance the M^3 method can handle full permeability tensors and general coarsening strategies, which may generate polygonal meshes on the coarse grid. These problem dependent coarsening parameters also play a critical role in the interpolation of the flux, and hence, in the advection of saturation for two-phase flow. Numerical experiments for two-phase flow in highly heterogeneous permeability fields, including layer 68 of the SPE Tenth Comparative Solution Project, demonstrate that the M^3 method retains good accuracy for high coarsening factors in both directions, up to 64 for the considered models. Moreover, we demonstrate that with a simple and efficient temporal updating strategy for the coarsening parameters, we achieve accuracy comparable to the fine-scale solution, but at a fraction of the cost.

1 Introduction

The ever increasing power of computers combined with the growing sophistication of numerical algorithms and visualization tools is driving the demand for quantitative simulations of multiscale phenomena. Typical multiscale applications include, contaminant flow in aquifers, enhanced oil recovery, power generation in fission reactors,

*Los Alamos National Laboratory, Theoretical Division, MS B284, Los Alamos, NM, 87545, USA, lipnikov@lanl.gov, moulton@lanl.gov, dasvyat@lanl.gov

and reacting flows in catalyst beds. Here, the term multiscale is intended to convey two important properties of these applications. First, the processes being modeled involve disparate length and time scales. Second, the fine-scale spatial structure and temporal coupling strongly influences coarse-scale properties of the solution, such as average flow rate through an aquifer. Thus, a naive simulation that does not address the influence of the fine-scales is terribly inaccurate, even for coarse-scale properties of the solution, while a fully resolved simulation is computationally intractable.

Thus, the goal of multiscale modeling is to develop methods that balance the competing demands of accuracy and efficiency. Specifically, we are interested in methods that not only capture the influence of fine or unresolved scales, but provide a methodology that facilitates achieving the desired accuracy in the computation of specific quantities of interest with the least possible computational work. Ultimately, these methods need to include error estimation and error control throughout the simulation, but this is beyond the scope of this research.

In modeling single and two-phase flow in porous media, the strongest multiscale influence arises from the heterogeneous structure of the subsurface environment. It is well understood that employing simple averages of the fine-scale parameters in a model of the same form has significant limitations. In particular, only a few heterogeneous structures are such that the coarse-scale flow under a uniform coarse-scale gradient is described by the same mathematical model with an equivalent or upscaled permeability. In general, a rigorous treatment of the fine-scale structure introduces a full tensor permeability, as well as additional closure or non-local terms in the coarse-scale model. For example, even in two-scale periodic media, classical homogenization using asymptotic expansions leads to a full tensor permeability for the coarse-scale model, even if the fine-scale permeability was a scalar [4, 20]. Thus, upscaling the fine-scale model, not just its parameters, is a fundamental challenge in modeling multiscale phenomena.

Recently, research in multiscale methods has focused on this more rigorous approach by upscaling flow models through the creation of multiscale basis functions. Although these concepts have appeared in various forms in a number of applications, such as the shape functions in reactor physics [31], and Generalized Finite Element Methods for conductivity in composite materials [3, 27], they were introduced to the subsurface flow modeling community through the Galerkin Multiscale Finite Element Method (MsFEM) [18]. In this method multiscale basis functions are created on each coarse-scale cell by solving local flow problems for the pressure. Using these basis functions in the weak form of the pressure equation to create a coarse-scale discretization captures the influence of the fine-scale structure directly, and in particular, this discretization is an upscaled model, not simply a model with upscaled parameters. However, this approach does not generate locally conservative velocity fields, which is desirable for two-phase and multiphase flow calculations.

To address this problem researchers turned to the first-order form of the bulk flow equations (i.e., writing mass conservation and Darcy's law as a first order system). Hence, using a Control Volume Finite Element Method (CVFEM) or a Mixed Finite Element Method (MFEM) to discretize the fine-scale model explicitly enforces local

mass conservation. Moreover this approach provided a natural setting to develop multiscale methods that explicitly enforce local mass conservation. For example, the initial work of Chen and Hou [9] introduced the Mixed Multiscale FEM (MMsFEM), which used local problems to develop multiscale basis functions for the velocity field. A modified version of this method was developed by Aarnes [1] to treat wells and to provide a locally conservative flux on both the fine and coarse scales. In addition, Efendiev et al. [14] introduced a Multiscale Finite Volume Element Method (MsFVEM) for two-phase flow calculations.

In all of these methods, the accuracy of the solution depends critically on the accuracy of the local boundary conditions that are used to create the multiscale basis functions. This motivates various approaches to capture the influence of the heterogeneous structure, such as local oscillatory conditions or oversampling [18], as well as special source and boundary functions for near-well regions [1]. However, for channelized flows arising from permeability fields with long correlation lengths, an effective local approach may not be possible and global flow information may be needed. For example, in [14] a highly accurate fine-scale global solution at the initial time is used to construct the boundary conditions for the MsFEVM basis functions. Similarly, in [2] a global solution at the initial time is used to define the local boundary conditions throughout a two-phase flow simulation. In both cases the robustness and accuracy was improved significantly with the use of this global information.

However, this rigid distinction between local and global is largely an artifact of the convenient but artificial partitioning of scales into resolved and unresolved, or fine-scale and coarse-scale. Specifically, this domain-decomposition style partitioning is what places the demands on these boundary conditions, because as the coarsening factor increases these artificial internal boundary conditions become increasingly important. In fact, most upscaling methods achieve a coarsening factor of approximately 10 in each coordinate direction, while the trends in fine-scale realizations of large reservoirs *requires a coarsening factor of 100 or more*.

These observations motivate our interest in developing truly *multilevel* methods that build on the success of robust variational multigrid methods. Here the issue of efficiently generating a hierarchy of coarse-scale models that accurately captures the influence of the fine-scale structure is handled through a minimization principle. Given a fine-scale discrete operator, a coarse-grid, and an interpolation operator, the coarse-scale operator that minimizes the error in the range of the interpolation is readily obtained. This is dubbed the variational or Galerkin coarse-grid operator, and it is the workhorse of robust multigrid solvers such as Black Box Multigrid (BoxMG) [11, 12], Algebraic Multigrid (AMG) [28, 29], and Smoothed Aggregation based AMG [30, 5]. This procedure does not simply average coefficients, but upscales the discrete model algebraically in a manner that captures the low energy modes of the fine-scale operator. The accuracy and efficiency of this coarsening procedure was demonstrated in the Multilevel Upscaling (MLUPS) method introduced by MacLachlan and Moulton [24] for single-phase steady-state flow. In particular, it was demonstrated that even for highly heterogeneous permeability fields with long

correlation lengths accuracy comparable to MsFEM [18] could be obtained much more efficiently. However, like MsFEM this method does not produce locally conservative velocity fields. While regaining locally conservative fields through postprocessing is of interest, here we consider another approach with greater flexibility in terms of the underlying discretization.

In this work we develop a new multilevel multiscale mimetic (M^3) method for finite difference discretizations of the first order system for pressure and bulk fluid velocity. This approach uses the subgrid modeling technique for single-phase flow recently proposed in [23, 17, 22]. We extend significantly the applicability of this technique by developing a new robust and efficient method for estimating the flux coarsening parameters. Specifically, with this advance, the M^3 method can handle full permeability tensors and general coarsening strategies, which may generate polygonal meshes on the coarse grid. The M^3 method is readily extended to polyhedral meshes by leveraging recent developments in mimetic finite difference (MFD) methods [8, 7].

In this subgrid modeling technique the accurate approximation of the flux coarsening parameters is critical for applications involving highly heterogeneous media or unstructured and distorted meshes. For example, using the solution of local problems to approximate these parameters [17, 1, 2] is not sufficiently accurate for channelized flows. Moreover, the cost of solving local problems for every macro-edge may approach or exceed the cost of performing a global solve on the fine grid. The novel technique that we propose for estimating the flux coarsening parameters is based on an approximate and inexpensive solution of this flow problem. To compute this approximate solution, we apply *a small number* of AMG cycles (many less than required for solving the fine-grid flow problem) to the equivalent symmetric positive definite system for the Lagrange multipliers. As noted earlier, a robust variational multigrid algorithm efficiently captures the effects of the fine-scale structure on the flow, and hence, quickly provides accurate estimates of the problem dependent flux coarsening parameters. By design, the structure of the coarse-scale system is identical to the original fine-scale system, hence we can apply the coarsening procedure recursively to obtain a multilevel algorithm. This algorithm is very flexible with a number of free parameters, including the number of coarse levels, the coarsening factor for each level, and the accuracy of the flux coarsening. Numerical experiments show that the cost of the multigrid cycles is a negligible part of the overall simulation cost.

The paper is organized as follows. In Section 2 we introduce a continuum model of two-phase flow in porous media, along with its mimetic finite difference (MFD) discretization. The new multilevel multiscale method for upscaling the pressure equation, which combines the novel subgrid-modeling technique proposed in [23, 17, 22] and a new AMG based methodology for computing the flux coarsening parameters, is developed in in Section 3. Section 4 presents an algorithm for simulating two-phase flow that uses the IMPES (IMplicit Pressure, EXplicit Saturation) scheme with the new multilevel multiscale method for pressure. A discussion of the computational complexity is presented as well. Numerical results of two-phase flow through highly heterogeneous media are presented in Section 5. Finally, in Section 6 we present conclusions and directions for future work.

2 Problem formulation

2.1 Continuum model

We consider immiscible two-phase flow in a two-dimensional porous medium, Ω [10]. The effects of gravity, compressibility, and capillary pressure are neglected. The two immiscible phases will be referred to as oil and water, and will be denoted with the subscripts o and w , respectively. The governing equations are given by,

$$\phi \frac{\partial S_j}{\partial t} + \nabla \cdot \vec{u}_j = -q_j, \quad j = o, w, \quad (2.1)$$

where S_j and \vec{u}_j are the saturation and the velocity of phase j , respectively, and the porosity, ϕ , is to be constant. We assume that the two phases fill the pore volume completely so that the saturation satisfies the constraint,

$$S_o + S_w = 1.$$

The source and sink terms, which represent injection and production wells, are denoted by q_j . The velocity of each phase is given by a generalized form of Darcy's law,

$$\vec{u}_j = -\frac{k_{rj}(S_j)}{\mu_j} \mathbb{K} \nabla p, \quad (2.2)$$

where \mathbb{K} is a symmetric and uniformly positive definite permeability tensor, k_{rj} is the relative permeability of phase j , and μ_j is its viscosity. The relative permeability k_{rj} depends on saturation in such a way that the permeability of one phase is reduced by the presence of other. The total, or bulk fluid, velocity is defined by

$$\vec{u} = \vec{u}_w + \vec{u}_o.$$

The phase mobility $\lambda_j(S_j)$ and the total mobility $\lambda(S_w)$ are defined as follows:

$$\lambda_j(S_j) = \frac{k_{rj}(S_j)}{\mu_j} \quad \text{and} \quad \lambda(S_w) = \lambda_o(1 - S_w) + \lambda_w(S_w). \quad (2.3)$$

Let us introduce the Buckley-Leverett fractional flow functions

$$f_w(S_w) = \frac{\lambda_w(S_w)}{\lambda(S_w)} \quad \text{and} \quad f_o(S_w) = \frac{\lambda_o(1 - S_w)}{\lambda(S_w)}.$$

Then, combining the governing equations (2.1) with Darcy's law (2.2), we obtain the pressure equation,

$$\nabla \cdot (\lambda(S_w) \mathbb{K} \nabla p) = q_o + q_w, \quad (2.4)$$

and the saturation equations,

$$\phi \frac{\partial S_j}{\partial t} + \nabla \cdot (f_j \vec{u}) = -q_j, \quad j = o, w. \quad (2.5)$$

Since the saturations S_j sum to one, we use water saturation and pressure as the primary variables of the model. We will refer to the equations presented above as the fine-scale model. In order to close this model, boundary and initial conditions must be imposed. We consider no-flow (Neumann) boundary conditions and homogeneous, $S_w = \text{const}$, initial condition for water saturation.

2.2 Discretization of the fine-scale model

In this section we define a fine-scale discretization of equations (2.4) and (2.5). Let Ω_h be a polygonal partition of the domain Ω that is referred to as the fine-grid (or fine-scale) partition:

$$\Omega_h = \bigcup_{i=1}^N e_i.$$

In this paper we use the mimetic finite difference (MFD) method [8] to discretize the pressure equation. However, the multiscale method proposed in the next section may be adopted easily to other mixed methods, such as the mixed finite element (MFE) [6] and control volume finite element (CVFE) [15] methods.

For each cell e_i , we define one pressure unknown, p_i , which represents the integral average of p . Let \mathbf{p} be the vector of all pressure unknowns. For each edge ℓ_j of polygon e_i , we define one unknown, u_{i,ℓ_j} , which represents the average normal flux $\bar{\mathbf{u}} \cdot \bar{\mathbf{n}}$ (a scalar) through this edge. Hereafter, we will refer to u_{i,ℓ_j} as simply the flux unknown. Let \mathbf{u} be the vector of flux unknowns. The size of \mathbf{u} is the number of boundary edges plus twice the number of internal edges. When two cells e_i and e_k share a common edge ℓ , the following continuity condition holds:

$$u_{i,\ell} = -u_{k,\ell}. \quad (2.6)$$

We use the IMPES approach (IMplicit Pressure and EXplicit Saturation) to discretize equations (2.4) and (2.5) in time. First, the pressure equation is solved. Second, the hyperbolic saturation equation is integrated explicitly using the single-point upwind finite volume method [16, 25]. Thus, the MFD discretization of the pressure equation may be written in the form,

$$\begin{bmatrix} \mathbf{M}(\mathbf{S}^n) & \mathbf{B}^T & \mathbf{C}^T \\ \mathbf{B} & \mathbf{0} & \mathbf{0} \\ \mathbf{C} & \mathbf{0} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{u}^n \\ \mathbf{p}^n \\ \boldsymbol{\lambda}^n \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ \mathbf{q}^n \\ \mathbf{0} \end{bmatrix}, \quad (2.7)$$

where n denotes the time step, $\mathbf{M}(\mathbf{S}^n)$ is the mass matrix computed using the saturation from the current time step, matrix \mathbf{B} represents the divergence operator multiplied by the matrix of polygon areas, matrix \mathbf{C} represents continuity of the normal flux across mesh edges (2.6), $\boldsymbol{\lambda}^n$ is the vector of Lagrange multipliers, and \mathbf{q}^n is the source term. Due to scaling of the divergence operator, components of \mathbf{q}^n are integrals over fine-grid polygons of the source and sink function $q_o + q_w$.

A saddle-point matrix with the same structure appears in the hybrid MFE method. The weak form of the velocity continuity condition (the last equation in (2.7) or equivalently (2.6)) naturally gives rise to an efficient solution algorithm. Specifically, the mass matrix is block-diagonal with as many blocks as there are cells in Ω_h ; therefore, velocity and pressure unknowns can be easily eliminated, resulting in a sparse symmetric positive definite problem for Lagrange multipliers.

3 Upscaling of the pressure equation

In a direct single-scale implementation of the IMPES algorithm, the solution of the pressure equation (2.7) dominates the computational time. For example, in the direct fine-scale and reference solution computations in Section 5, approximately 90% of the computational time is spent in solving for the pressure. The objective of the proposed method, and of multiscale methods in general, is to mitigate this cost by solving the pressure equation on a much coarser grid. In this section we describe a robust multilevel technique for generating this coarse-grid system. For simplicity, we omit the time superscript from (2.7). The method is truly multilevel and produces locally conservative coarse-grid velocities.

We begin by describing a two-level coarsening method, a building block for the multilevel method. Let $N_0 = N$, $cN_1 \leq N_0$, with $0 < c < 1$, and

$$\Omega_H = \bigcup_{i=1}^{N_1} E_i, \quad E_i = \bigcup_{k \in \mathcal{F}(E_i)} e_k,$$

where $\mathcal{F}(E_i)$ is a set of indices of fine-grid cells and E_i is a *macro-cell* (a polygon). Let $|\mathcal{F}(E_i)|$ be the cardinal number of set $\mathcal{F}(E_i)$, $i = 1, \dots, N_1$. We assume that the coarse-grid partition Ω_H is non-overlapping and conformal, and that each edge of a macro-cell E_i is a segment of a straight line.

The two-level method consists of two steps, which are illustrated in Fig. 1. First, for each macro-cell, we eliminate all internal velocity unknowns, and replace all internal pressure unknowns with a single pressure. This elimination process is performed by *equivalent* modifications of the original system. The unknowns in the resulting system are shown in the middle schematic of Fig. 1. Second, we perform a *conservative* flux coarsening procedure that defines one flux unknown per macro-cell edge from the fine-grid fluxes on the macro-edge. The final form of the reduced system, shown in the right schematic of Fig. 1, is identical to the original system.

3.1 Elimination of internal degrees of freedom

To eliminate the internal degrees of freedom we use the algorithm proposed by Y. Kuznetsov (see [23, 17, 22] where this algorithm is described for the diffusion-reaction problem). Let us consider a macro-cell E formed by $n_p \equiv |\mathcal{F}(E)|$ fine-grid cells e_k , $k \in \mathcal{F}(E)$. Let $\Gamma \equiv \partial E$ be its boundary formed by n_u^{bnd} fine-grid edge γ_k . Furthermore, let n_u^{int} be the number of internal edges of E .

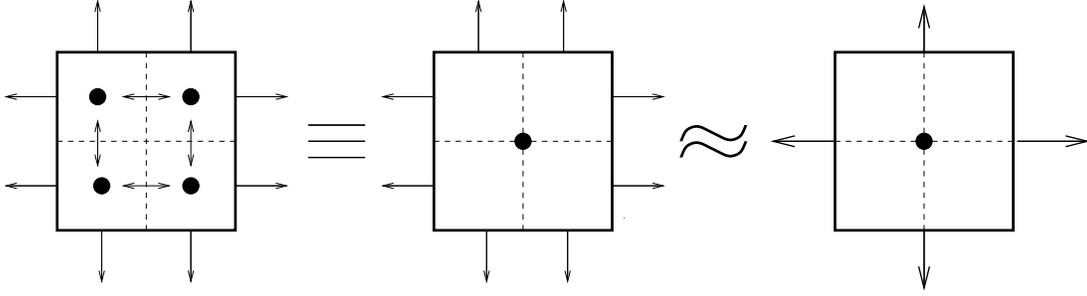


Figure 1: Schematic of the two steps of the two-level upscaling method for a 2×2 square macro-cell. The cell-centered pressure unknowns are represented by circles, and the velocity unknowns are represented by arrows. The first step is an equivalent reduction, while the second is approximate.

Using the continuity condition (2.6), we eliminate duplicated internal fluxes and the corresponding Lagrange multipliers. Equations for the remaining unknowns associated with the macro-cell E may be written in the form,

$$\begin{bmatrix} \mathbf{M}_{\Gamma\Gamma} & \mathbf{M}_{\Gamma i} & \mathbf{B}_{\Gamma}^T \\ \mathbf{M}_{i\Gamma} & \mathbf{M}_{ii} & \mathbf{B}_i^T \\ \mathbf{B}_{\Gamma} & \mathbf{B}_i & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{u}_{\Gamma} \\ \mathbf{u}_i \\ \mathbf{p}_i \end{bmatrix} = \begin{bmatrix} \mathbf{g}_{\Gamma} - \mathbf{C}_{\Gamma}^T \boldsymbol{\lambda}_{\Gamma} \\ \mathbf{g}_i \\ \mathbf{q}_i \end{bmatrix}, \quad (3.1)$$

where \mathbf{u}_i is the vector of size n_u^{int} of internal flux unknowns, \mathbf{u}_{Γ} is the vector of size n_u^{bnd} of boundary flux unknowns, \mathbf{p}_i is the vector of size n_p of pressure unknowns, $\boldsymbol{\lambda}_{\Gamma}$ is the vector of size n_u^{bnd} of the remaining Lagrange multipliers, and \mathbf{q}_i is the vector of size n_p of source data. All the introduced vectors are parts of vectors in the global formulation (2.7). For the fine-grid discretization, $\mathbf{g}_i = 0$ and $\mathbf{g}_{\Gamma} = 0$; however, in a multilevel framework they may be non-zero in macro-cells containing sources.

By properties of the discretization method, the leading 2×2 matrix in (3.1) is symmetric and positive definite. Let us define the following Schur complement matrix:

$$\mathbf{S} = \mathbf{B}_i \mathbf{M}_{ii}^{-1} \mathbf{B}_i^T. \quad (3.2)$$

Note that

$$\mathbf{S} = \mathbf{S}^T \geq 0 \quad \text{and} \quad \ker \mathbf{S} = \ker \mathbf{B}_i^T. \quad (3.3)$$

It can be shown that the null space has dimension one and that it is spanned by the vector $\boldsymbol{\psi} = (1, \dots, 1)^T$ of size n_p . Let us consider the generalized eigenvalue problem

$$\mathbf{S} \mathbf{w}_i = \nu_i \mathbf{D} \mathbf{w}_i, \quad i = 1, \dots, n_p,$$

where \mathbf{D} is the positive definite diagonal matrix with cell areas of the diagonal:

$$\mathbf{D} = \text{diag} \{ |e_{k_1}|, \dots, |e_{k_{n_p}}| \}, \quad k_i \in \mathcal{F}(E).$$

The reason for choosing such a diagonal matrix is explained at the end of this subsection. We assume that the eigenvectors are \mathbf{D} -orthonormal and $\nu_1 = 0$. We define the pseudo-inverse matrix, \mathbf{S}^+ , as follows

$$\mathbf{S}^+ = \sum_{i=2}^{n_p} \frac{1}{\nu_i} \mathbf{w}_i \mathbf{w}_i^T.$$

It is easy to verify that

$$\mathbf{S}^+ \mathbf{S} = \mathbf{I} - \frac{1}{\alpha^2} \boldsymbol{\psi} \boldsymbol{\psi}^T \mathbf{D}, \quad \alpha^2 = \boldsymbol{\psi}^T \mathbf{D} \boldsymbol{\psi}, \quad (3.4)$$

Note that system (3.1) is compatible being a part of the compatible system (2.7). Therefore, it can be solved for \mathbf{u}_i and \mathbf{p}_i :

$$\begin{aligned} \mathbf{p}_i &= \mathbf{S}^+ (\mathbf{B}_\Gamma - \mathbf{B}_i \mathbf{M}_{ii}^{-1} \mathbf{M}_{i\Gamma}) \mathbf{u}_\Gamma + \boldsymbol{\eta}, \\ \mathbf{u}_i &= \mathbf{R}_{i\Gamma} \mathbf{u}_\Gamma + \boldsymbol{\xi} \end{aligned} \quad (3.5)$$

where

$$\begin{aligned} \mathbf{R}_{i\Gamma} &= -\mathbf{M}_{ii}^{-1} \mathbf{B}_i^T \mathbf{S}^+ (\mathbf{B}_\Gamma - \mathbf{B}_i \mathbf{M}_{ii}^{-1} \mathbf{M}_{i\Gamma}) - \mathbf{M}_{ii}^{-1} \mathbf{M}_{i\Gamma}, \\ \boldsymbol{\eta} &= -\mathbf{S}^+ [\mathbf{q}_i - \mathbf{B}_i \mathbf{M}_{ii}^{-1} \mathbf{g}_i] + \beta \boldsymbol{\psi}, \\ \boldsymbol{\xi} &= \mathbf{M}_{ii}^{-1} \mathbf{B}_i^T \mathbf{S}^+ [\mathbf{q}_i - \mathbf{B}_i \mathbf{M}_{ii}^{-1} \mathbf{g}_i] + \mathbf{M}_{ii}^{-1} \mathbf{g}_i, \end{aligned}$$

and β is an arbitrary number.

Let us multiply the second two block equations in (3.1) by the matrix $\mathbf{R}_{i\Gamma}^T \mathbf{D}$ and add the result to the first block equation. Note that

$$(\mathbf{B}_\Gamma + \mathbf{R}_{i\Gamma}^T \mathbf{B}_i^T) \mathbf{p}_i = \frac{1}{\alpha^2} \mathbf{B}_\Gamma^T \boldsymbol{\psi} \boldsymbol{\psi}^T \mathbf{D} \mathbf{p}_i \equiv \widehat{\mathbf{B}}_\Gamma p_E,$$

where p_E is the coarse-grid pressure unknown given by

$$p_E = \frac{1}{\alpha^2} \boldsymbol{\psi}^T \mathbf{D} \mathbf{p}_i = \frac{\sum_{k \in \mathcal{F}(E)} p_k |e_k|}{\sum_{k \in \mathcal{F}(E)} |e_k|}.$$

This formula is physically sound (p_E is the integral average of fine-grid pressure unknowns) and explains the definition of the pseudo-inverse matrix given in equation (3.4). Now, we multiply the last block equation in (3.1) by $\boldsymbol{\psi}^T$ and use property (3.3). The reduced system may be written in a matrix form:

$$\begin{bmatrix} \widehat{\mathbf{M}}_{\Gamma\Gamma} & \widehat{\mathbf{B}}_\Gamma^T \\ \widehat{\mathbf{B}}_\Gamma & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{u}_\Gamma \\ p_E \end{bmatrix} = \begin{bmatrix} \widehat{\mathbf{g}}_\Gamma - \mathbf{C}_\Gamma^T \boldsymbol{\lambda}_\Gamma \\ q_E \end{bmatrix}, \quad (3.6)$$

where

$$\widehat{\mathbf{M}}_{\Gamma\Gamma} = \begin{bmatrix} \mathbf{I} & \mathbf{R}_{i\Gamma}^T \end{bmatrix} \begin{bmatrix} \mathbf{M}_{\Gamma\Gamma} & \mathbf{M}_{i\Gamma}^T \\ \mathbf{M}_{i\Gamma} & \mathbf{M}_{ii} \end{bmatrix} \begin{bmatrix} \mathbf{I} \\ \mathbf{R}_{i\Gamma} \end{bmatrix}$$

and

$$\widehat{\mathbf{g}}_\Gamma = \mathbf{g}_\Gamma - \mathbf{M}_{i\Gamma}^T \boldsymbol{\xi} + \mathbf{R}_{i\Gamma}^T (\mathbf{g}_i - \mathbf{M}_{ii} \boldsymbol{\xi}), \quad q_E = \boldsymbol{\psi}^T \mathbf{q}_i.$$

3.2 Conservative flux coarsening

After elimination of internal fluxes and reduction to a single pressure, we obtain a discretization with several flux unknowns per macro-edge (see the second picture in Fig. 1). To obtain a discretization with only one flux unknown per macro-edge, we closely follow the approach proposed by Y. Kuznetsov and described in [17], and then introduce our new methodology for approximating the flux coarsening parameters in Section 3.3.

Let L denote a macro-edge of a macro-cell E . By assumption, L is a segment of a straight line. First, we consider the case in which L consists of two fine-grid

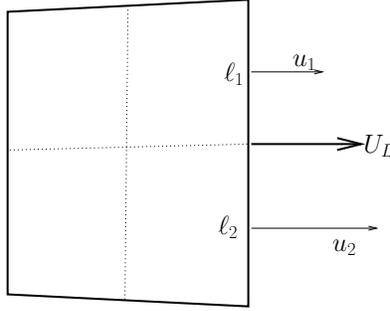


Figure 2: Macro-cell and its edge L (left) consisting of two fine-grid edges ℓ_1 and ℓ_2 .

edges ℓ_1 and ℓ_2 (see Fig. 2). Let u_1 and u_2 be the fluxes corresponding to ℓ_1 and ℓ_2 , respectively. Without loss of generality, we assume that $|u_1| \leq |u_2|$. The goal of the coarsening procedure is to define the flux U_L through the macro-cell edge L in terms of the fine-grid fluxes. Mass conservation dictates that

$$U_L(|\ell_1| + |\ell_2|) = u_1|\ell_1| + u_2|\ell_2|. \quad (3.7)$$

where $|\ell_i|$ denotes the length of edge ℓ_i .

Let α be the ratio of fine-grid fluxes:

$$\alpha = \begin{cases} \frac{u_1}{u_2}, & u_2 \neq 0, \\ 1, & \text{otherwise.} \end{cases} \quad (3.8)$$

This ratio characterizes the first moment of the flux, and will play an essential role in our algorithm. In Section 3.3.1 we show that using the exact ratio leads to an optimal coarse-grid system in the sense that it reproduces fine-scale quantities exactly. However, this calculation is equivalent to solving the fine-scale system, and hence, it is not a practical multiscale method. We introduce new method for approximating the flux ratios, α , in Section 3.3.2 and show its robustness and efficiency in Section 5 with numerical experiments for challenging test problems. Here, we suppose that α is given. Then, condition (3.7) gives

$$u_1 = \frac{\alpha(|\ell_1| + |\ell_2|)}{\alpha|\ell_1| + |\ell_2|}U_L \quad \text{and} \quad u_2 = \frac{|\ell_1| + |\ell_2|}{\alpha|\ell_1| + |\ell_2|}U_L.$$

In other words, α defines the 2×1 interpolation matrix \mathbf{Q}_L such that

$$\begin{bmatrix} u_1 \\ u_2 \end{bmatrix} = \mathbf{Q}_L U_L. \quad (3.9)$$

If the coarse-grid edge L consists of m fine-grid edges, we have $m - 1$ coarsening parameters α_i , $i = 1, \dots, m - 1$. For simplicity we assume again that the fine-grid fluxes satisfy $|u_i| \leq |u_{i+1}|$. Then,

$$\alpha_i = \begin{cases} \frac{u_i}{u_{i+1}}, & u_{i+1} \neq 0, \\ 1, & \text{otherwise.} \end{cases}$$

Using the same arguments as above, we derive that

$$u_i = \frac{r_i^{(1)}}{r_i^{(2)}} \sum_{j=1}^m |\ell_j| U_L, \quad (3.10)$$

where $r_m^{(1)} = 1$,

$$r_i^{(1)} = \prod_{j=i}^{m-1} \alpha_j, \quad i \leq m-1, \quad \text{and} \quad r_i^{(2)} = |\ell_m| + \sum_{j=1}^{m-1} |\ell_j| \prod_{k=j}^{m-1} \alpha_k.$$

Formula (3.10) defines the interpolation matrix \mathbf{Q}_L in this case.

Remark 3.1 *Expression (3.10) can be simplified if we introduce m parameters β_1, \dots, β_m such that $\alpha_j = \beta_j / \beta_{j+1}$. Then*

$$u_i = \frac{\beta_i \sum_{j=1}^m |\ell_j|}{\sum_{j=1}^m \beta_j |\ell_j|} U_L. \quad (3.11)$$

Let \mathbf{Q}_E be the interpolation matrix for the macro-cell E that is assembled from the edge-based matrices \mathbf{Q}_L . Then, given coarse-grid fluxes $\widehat{\mathbf{u}}_\Gamma$, we may calculate the fine-grid fluxes by

$$\mathbf{u}_\Gamma = \mathbf{Q}_E \mathbf{u}_\Gamma^c.$$

Substituting this expression in (3.6), and multiplying the first equation by \mathbf{Q}_E^T , we obtain the coarse-grid discretization for the macro-cell E :

$$\begin{bmatrix} \mathbf{M}_E & \mathbf{B}_E^T \\ \mathbf{B}_E & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{u}_\Gamma^c \\ p_E \end{bmatrix} = \begin{bmatrix} \mathbf{Q}_E^T \widehat{\mathbf{g}}_\Gamma - \mathbf{C}_E \boldsymbol{\mu} \\ q_E \end{bmatrix} \quad (3.12)$$

where

$$\mathbf{M}_E = \mathbf{Q}_E^T \widehat{\mathbf{M}}_\Gamma \mathbf{Q}_E, \quad \mathbf{B}_E = \widehat{\mathbf{B}}_\Gamma \mathbf{Q}_E, \quad \mathbf{C}_E = \text{diag} \{ |L_1|, \dots, |L_m| \},$$

and

$$\boldsymbol{\mu} = \mathbf{C}_E^{-1} \mathbf{Q}_E^T \mathbf{C}_\Gamma \boldsymbol{\lambda}.$$

The coarse-grid discretization (3.12) is illustrated in the right schematic of Fig. 1. The structure and scaling of matrix \mathbf{C}_E has been designed to match the structure and scaling of matrix \mathbf{C}_Γ in (3.1). The global coarse-grid system is obtained by assembling the elemental systems (3.12) for all macro-cells, and is closed by adding flux continuity conditions similar to (2.6). After solving the coarse-grid system, we obtain a vector \mathbf{u}^c of coarse-grid fluxes and a vector \mathbf{p}^c of coarse-grid pressures p_E . The fine-grid velocities needed for the transport equation are calculated using first the interpolation operators \mathbf{Q}_E and then formula (3.5). Moreover, since this global coarse-grid problem has exactly the same sparsity structure as the fine-grid problem (2.7), the multi-level method may be constructed by recursion.

3.3 Flux coarsening parameters

In the M^3 method the accuracy of the flux coarsening parameters directly influences the overall accuracy of the flow calculation, and particularly, the advection of saturation. Specifically, the coarse-grid pressure solve preserves mass locally and provides the zeroth order moment or average flux through each macro-edge. The flux coarsening parameters characterize the distribution, or first moment, of the flux along the macro-edge, and hence, significantly impact the accuracy of interpolation of the flux to the fine-grid.

Although standard approaches in parameter upscaling and in the generation of multiscale basis functions could be used to approximate the flux coarsening parameters, the limitations of these methods, which were highlighted in the introduction, are naturally present in this sub-grid modeling technique as well. Thus, to address these limitations in a robust and efficient manner we introduce a novel approach in Section 3.3.2 that is motivated by the important observation formulated in Lemma 3.1.

3.3.1 Exact flux coarsening

In this subsection we assume that the exact flux ratio, α^{ex} , is known for each macro-edge in the coarse-grid. We prove that the fluxes obtained by the multiscale method are exactly the fine-grid fluxes. For this proof it is sufficient to analyze the two-level method.

Lemma 3.1 *Let α^{ex} be the exact flux ratio for each macro-edge corresponding to the solution of (2.7) on the fine grid. Then interpolation of coarse-scale fluxes, obtained by the solving equations (3.12) on the coarse-grid, yields exactly the fine-grid fluxes.*

Proof. Without loss of generality, we consider a macro-edge L consisting of two fine-grid edges ℓ_1 and ℓ_2 . Let u_1 and u_2 be the fine-grid fluxes corresponding to these edges. Using (3.7), we define the restriction operator \mathbf{R}_L ,

$$\mathbf{R}_L \begin{bmatrix} u_1 \\ u_2 \end{bmatrix} = \frac{|\ell_1|u_1 + |\ell_2|u_2}{|\ell_1| + |\ell_2|} = U_L,$$

where U_L is the flux through macro-edge L . Using the interpolation operator \mathbf{Q}_L defined in (3.9), we write the following equation for the interpolation error:

$$\begin{bmatrix} \varepsilon_1 \\ \varepsilon_2 \end{bmatrix} = \begin{bmatrix} u_1 \\ u_2 \end{bmatrix} - \mathbf{Q}_L \mathbf{R}_L \begin{bmatrix} u_1 \\ u_2 \end{bmatrix}.$$

A straightforward calculation shows that:

$$\varepsilon_1 = (u_1 - \alpha^{ex} u_2) \frac{|\ell_2|}{\alpha^{ex} |\ell_1| + |\ell_2|} \quad \text{and} \quad \varepsilon_2 = (\alpha^{ex} u_2 - u_1) \frac{|\ell_1|}{\alpha^{ex} |\ell_1| + |\ell_2|}.$$

Since $\alpha^{ex} = u_1/u_2$, then $\varepsilon_1 = \varepsilon_2 = 0$. Thus vector $(u_1, u_2)^T$ is the eigenvector of the matrix $\mathbf{Q}_L \mathbf{R}_L$ which corresponds to eigenvalue $\lambda = 1$.

The global interpolation and restriction matrices are block diagonal with as many blocks as there are coarse-grid edges, therefore the interpolation of coarse-grid fluxes coincides with the fine-grid fluxes on all fine-grid edges. This proves the assertion of the lemma. \square

This lemma proves that if the exact flux ratios, α^{ex} , are known for a specific problem, then the second step in the scheme shown in Figure 1 is an equivalent reduction as well as the first step.

3.3.2 Flux coarsening with algebraic multigrid

The overarching goal of this research is to develop a multilevel algorithm that provides a flexible hierarchical approach to balancing the competing demands of efficiency and accuracy. This is particularly important for problems with highly heterogeneous permeability fields that lack scale separation or have long correlation lengths, as these are not effectively treated with local approximation methods [14]. Indeed, it is critical that global information be incorporated in an efficient and hierarchical manner that does not require a highly accurate solution of the fine-scale pressure equation.

We propose to *define the coarsening parameters α from an approximate solution of (2.7) calculated with an efficient multigrid method.*

Algorithm 1 New method for calculating flux coarsening parameters

- 1: Form the Schur complement of system (2.7) for the Lagrange multipliers, $\boldsymbol{\lambda}$.
 - 2: Perform a few preconditioned conjugate gradient (PCG) iterations with a single V(1,1)-cycle of the Ruge-Stüben algebraic multigrid [26, 28] as the preconditioner, until a *weak* convergence tolerance ε_r is reached.
 - 3: Calculate fluxes from elemental systems (3.1) and the flux coarsening parameters α from (3.8).
-

For problems involving the second-order form of the pressure equation, the algebraic multigrid method (AMG) is known to be a good candidate. In particular, AMG uses the fine-scale discrete model to create a hierarchy of coarse-scale models through

a variational principle that minimizes the error in the range of the interpolation. This methodology is critical to the robustness of AMG for problems with discontinuous and tensor permeabilities. Moreover, this is a very efficient way to incorporate tightly coupled scales ranging from the cell size to the domain size. Note that a robust multi-grid method that deals directly with the saddle point problem (2.7) *does not* exist; therefore Step 1 is required. The Schur complement matrix is symmetric and positive definite. It is formed element-by-element ensuring the *linear complexity* of this algorithm.

The AMG method provides optimal algorithmic scaling with respect to mesh resolution, and achieves a typical convergence rate of 0.1 – 0.3 for the permeability fields considered in Section 5. The conjugate gradient method is used to accelerate convergence of the AMG method. In the numerical experiments, we use *a small number* of PCG iterations, typically 3–5, to solve the Lagrange multiplier system to a relative residual tolerance ε_r , in the L_2 norm. In following discussions the accuracy of the coarsening parameters α will be described by the convergence criteria ε_r .

Although the multilevel formulation that we presented above defines key components that could be used in a multilevel solution algorithm for the saddle point problem, there are still many unresolved issues, such as smoothers, that are beyond the scope of this article.

4 Two phase flow simulation

4.1 The M^3 method

In this implementation of the multilevel multiscale mimetic (M^3) method for two-phase flow we use the IMPES time discretization scheme (IMplicit Pressure and EXplicit Saturation). As noted in the previous section, the solution of the pressure equation is the most computationally demanding part of each time-integration step. To solve it more efficiently, we use the new multilevel method described in the previous section. Having solved the pressure equation on the coarsest grid, we interpolate the fluxes onto the fine-scale mesh and use them to update the saturation explicitly. By construction this interpolation is conservative.

The accuracy of the multiscale solution depends on the various facets of the M^3 method that impact how accurately α approximates α^{ex} (Section 3.3.1) throughout the simulation. For example, the exact flux ratio α^{ex} varies with time, and hence, updating α in time will improve the overall accuracy of the multiscale solution. In fact, updating α at each time-step to high precision is equivalent to solving the fine-grid problem. However, in many multiscale methods, such as the MsMFEM [2], MsFVEM [14], and Multiscale Finite Volume (MsFV) [19] method, the fine-grid pressure equation is solved to high precision only once, at the initial time step. Then, the corresponding multiscale basis functions are used throughout the entire simulation. Although we implement this approach for comparison purposes (see Section 5.1.1), we are more interested in the gains in overall efficiency offered by occasional, compu-

tationally inexpensive corrections of α . These gains in efficiency are possible because of our multigrid approximation technique (see Section 3.3.2), which is not only computationally efficient, due to the fast multigrid convergence, but provides estimates that include valuable information about global flow behavior.

Another important facet of the M^3 method is the coarsening strategy. Specifically, in a multilevel algorithm the coarsening factor between levels, as well as the total coarsening factor, affects the trade-off between accuracy and efficiency. In the numerical examples, we consider rectangular meshes and generally coarsen uniformly in each coordinate. Hence, using c_l to denote the coarsening factor between levels l and $l + 1$ in each coordinate, we may describe a coarsening strategy in a hierarchy of \mathcal{L} levels with the compact notation $c_1 : c_2 : \dots : c_{\mathcal{L}-1}$. Moreover, the total coarsening factor in each coordinate may be written

$$c_T = \prod_{l=1}^{\mathcal{L}-1} c_l.$$

while the total coarsening factor is simply $C_T = c_T^2$. For example, a two-level method with $c_1 = 4$, and a three-level method with $c_1 : c_2 = 2 : 2$, both have a total coarsening factor of $C_T = 16$. However, as we will see in the next subsection, the computational cost of forming the coarse-scale pressure system favors hierarchies that use coarsening factors of two.

These facets of the M^3 method are summarized in the following list:

- the coarsening strategy and the total coarsening factor (see Sec. 5.1.1);
- the total number and frequency of flux ratio updates (see Sec. 5.1.2);
- the convergence tolerance, ε_r , of the PCG(AMG) iteration used to determine the flux ratios (see 5.1.2).

The role of the corresponding parameters are highlighted in pseudo-code description below, Algorithm 2. In the numerical tests, we investigate the impact of these parameters on the quality of the multiscale solution.

4.2 Computational complexity

The most computationally demanding part of the M^3 method is derivation of matrices \mathbf{M}_{ii}^{-1} and \mathbf{S}^+ in (3.2) – (3.5). To analyze the method complexity, we consider rectangular meshes and neglect all algorithmic operations except matrix inversions. Let us assume that the coarsening factors in each direction are $c_l = 2^k$. Then, each macro-cell consists of $n_c(k) = 2^{2k}$ sub-cells and has $n_i(k) = 2^{k+1}(2^k - 1)$ internal edges. Matrices \mathbf{M}_{ii}^{-1} and \mathbf{S}^+ have sizes $n_i(k)$ and $n_c(k)$, respectively. Recall that inversion of a dense symmetric matrix of size n requires $O(n^3)$ arithmetic operations. These inversions are independent of each other and can be easily parallelized.

Algorithm 2 The M³ Method for Two-Phase Flow

- 1: Estimate parameters α using Algorithm 1 (see Section 3.3.2) with a convergence tolerance ε_r .
 - 2: Set the number of levels \mathcal{L} , and the coarsening factors $c_1 : c_2 : \dots : c_{\mathcal{L}-1}$.
 - 3: Form hierarchy of coarse-scale discretizations of the pressure equation (3.12).
 - 4: **while** $t^n < T_f$ **do**
 - 5: **if** (*update of parameter α is scheduled*) **then**
 - 6: compute new parameters α using Algorithm 1 with the convergence tolerance ε_r .
 - 7: **else**
 - 8: use the parameters α from the previous time step
 - 9: **end if**
 - 10: Update coarse-scale discretizations using the criteria described in Section 5.1.3 and the upscaling algorithm described in Section 3.
 - 11: Solve the pressure equation on the coarsest level.
 - 12: **for** $\mathcal{L}, \mathcal{L} - 1, \dots, 2$ **do**
 - 13: Interpolate fluxes from level l to level $l - 1$ using formulas (3.9), (3.5).
 - 14: **end for**
 - 15: Define time step Δt^n using the saturation based criterion [10, Section 7.3].
 - 16: Update the saturation.
 - 17: $n \leftarrow n + 1; t^n \leftarrow t^{n-1} + \Delta t^n$.
 - 18: **end while**
-

If we perform $\mathcal{L} - 1$ levels of coarsening with the factor 2^k on each level, the total number of numerical operations N_{oper} per a coarsest macro-cell is

$$N_{oper} \sim \sum_{l=1}^{\mathcal{L}-1} N_c(l)(n_c^3 + n_i^3) = \sum_{l=1}^{\mathcal{L}-1} 2^{2k(\mathcal{L}-l)}(2^{6k} + 2^{(3k+3)}(2^k - 1)^3),$$

where $N_c(l) = 2^{2k(\mathcal{L}-1-l)}$ is the number of macro-cells on the l -th level.

Now we fix the total coarsening factor. Note that one level of coarsening with the factor 2^k corresponds to two levels of coarsening with the factor $2^{k/2}$ or to four levels of coarsening with the factor $2^{k/4}$ and so on.

The estimation of N_{oper} for different numbers of coarsening levels shown in Figure 3 demonstrates that the coarsening strategy based on the factor 2 in each direction is the most efficient from computational point of view. In this case \mathbf{M}_{ii}^{-1} and \mathbf{S}^+ are symmetric 4×4 matrices. A similar conclusion regarding memory savings can be made. The multilevel strategy based on the coarsening factor 2 in each direction requires much less memory.

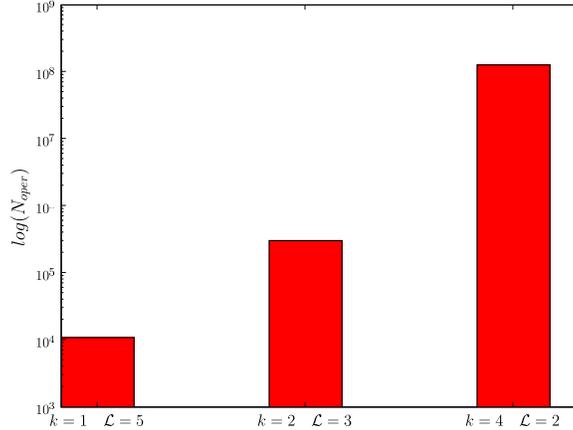


Figure 3: Log of a computational work requires to generate local matrices for a macro-cell consisting of 16×16 fine-scale cells as a function of a number of coarsening levels ($\mathcal{L} - 1$) and a coarsening factor on each level (2^k in each direction).

5 Numerical results

To demonstrate the effectiveness of the proposed multiscale method, we consider two models of the permeability. In both models the permeability field is assumed to be a highly heterogeneous scalar function constant on each fine-grid cell. The first model uses realizations of layered media generated using the GSLIB software package [13]. The second model is taken from the Tenth SPE Comparative Solution Project. Specifically, we used the two-dimensional field defined by layer 68 of the three-dimensional reservoir in model 2. As noted in [21], this is a fluvial layer that provides a challenging test for multiscale methods.

We have three research goals. First, we analyze different coarsening strategies and different total coarsening factors (Sec. 5.1.1 and 5.2). Second, we analyze how frequency of updates of the coarsening parameters α 's impacts solution accuracy (Sec. 5.1.1 and 5.2). Third, we analyze how solution accuracy is affected by accuracy of α 's (Sec. 5.1.1).

To complete the definition of the two-phase flow model given in Section 2, we define the the relative permeability curves as,

$$k_{rw}(S) = (S^*)^2 \quad k_{ro}(S) = (1 - S^*)^2 \quad S^* = \frac{S - S_{wc}}{1 - S_{wc} - S_{or}}, \quad (5.1)$$

where S_{wc} is the critical saturation, and S_{or} is the residual saturation. We set $S_{wc} = S_{or} = 0.1$. In addition, the initial saturation is set to a constant, $S(t = 0) = S_{wc} = 0.1$; the phase viscosities are $\mu_w = 1$ and $\mu_o = 4$; and the porosity of the medium is assumed to be constant, $\phi = 0.2$. For all simulations we use no-flow boundary conditions. The flow is driven by source and sink terms, which are considered as

injector and producer wells with constant rates. The time is expressed in PVI (Pore Volume Injected) units, which is a common practice in reservoir simulations.

To estimate the accuracy of the multiscale solution obtained using the proposed M^3 method (Algorithm 2), we define a *reference* solution. The reference solution is computed by forming the MFD discretization (Section 2.2) on a mesh that is twice finer, in each direction, than the fine-grid. We also define a *fine-scale* (or fine-grid) solution which uses the MFD discretization on the fine-grid. For both the fine-scale and reference solutions the pressure equation is solved with PCG to a tolerance of $\varepsilon_r = 10^{-8}$, and the IMPES time discretization approach is used to advance saturation. The PCG solver is preconditioned with a single V(1,1) cycle of standard Ruge-Stüben AMG [28] per iteration, and hereafter, will be referred to as PCG(AMG).

We focus our investigation of accuracy on three quantities that are important in two phase flow simulations: water-cut, saturation in the production wells, and the breakthrough time. The relative error in the saturation at the production wells, $\delta_S(t)$, is defined as follows:

$$\delta_S(t) = \frac{|S_r(t) - S(t)|}{S_r(t)}, \quad (5.2)$$

where $S_r(t)$ denotes the reference saturation. For strongly heterogeneous permeability fields, *the reference solution may be much closer to the continuum solution than the fine-grid solution*. Therefore, using it in (5.2) results in a more reliable estimate of the simulation error.

We also choose several macro-edges in the domain and monitor the behavior of the flux ratios on these edges as functions of time. In particular, we study how well the computed flux ratio α approximates the exact α^{ex} . In addition, we study how different parameters of the M^3 method affect the quantities of interest.

5.1 GSLIB model

The fine grid is a 128x128 uniform partition of a square domain $\Omega = [0, 1000\text{m}]^2$. We consider two heterogeneity scenarios for the permeability field generated such that the $\log_{10}(\mathbb{K})$ is normally distributed with mean zero and variance four. The correlation lengths along the principle axes of statistical anisotropy are 0.8 and 0.04, leading to a strongly layered structure with a dynamic range of approximately 6 orders of magnitude. The average direction of high permeability layers is vertical in Scenario 1 and is rotated by 45° clockwise in Scenario 2 (see Fig. 4). Both scenarios represent structures that are difficult to upscale. The source is located at the south-west corner of Ω and the sink is in the north-east corner (see Fig. 5).

5.1.1 Simulations with fixed coarsening parameters α 's

The goal of the first group of experiments is to demonstrate that the accuracy of the M^3 method decreases very modestly with the number of levels, \mathcal{L} . This is important because using smaller coarsening factors, and hence more levels, improves the efficiency of the method (see Section 4.2). Here we consider three coarsening strategies

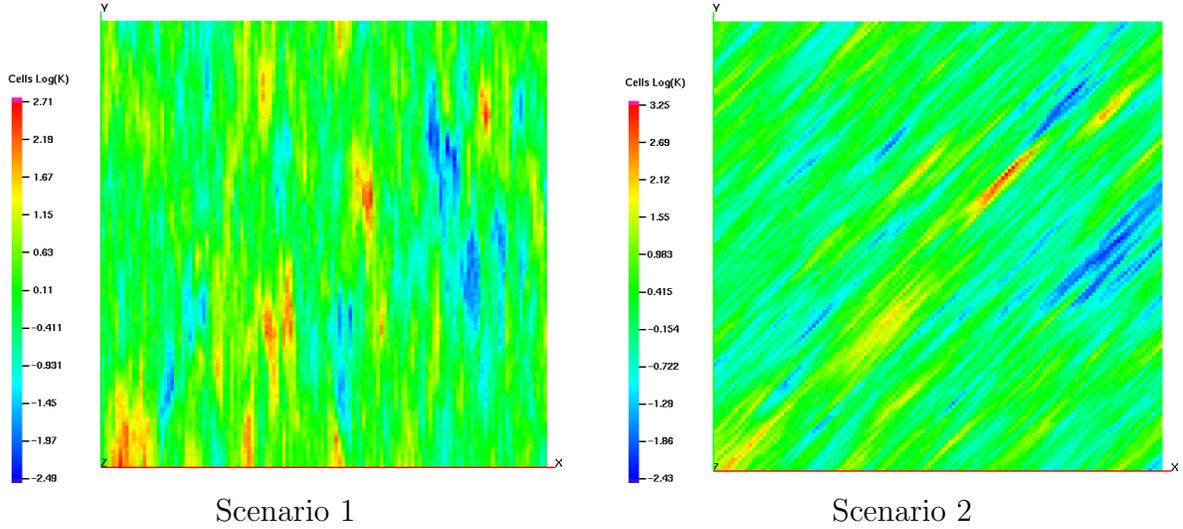


Figure 4: Two permeability fields generated by the GSLIB package: Scenario 1 (left) and Scenario 2 (right).

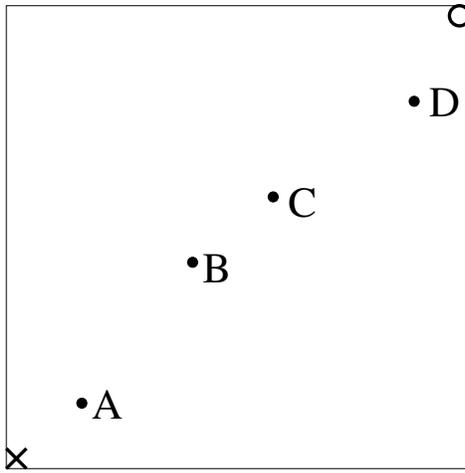


Figure 5: Locations of the injector (\times) and production (\circ) wells and four monitored macro-edges A, B, C, and D.

with a total coarsening factor of 8, and a 16×16 coarsest mesh. Specifically we consider a two-level method, $c_1 = 8$, a three-level method, $c_1 : c_2 = 4 : 2$, and a four-level method, $c_1 : c_2 : c_3 = 2 : 2 : 2$.

At the initial time the fine-scale system (2.7) is solved with PCG(AMG) to a tolerance $\varepsilon_r = 10^{-8}$. This solution is used to determine the flux ratio α for each macro-edge. The flux ratio is fixed for the rest of the simulation. The water-cut curves and relative errors of saturation in the producer well are shown in Fig. 6.

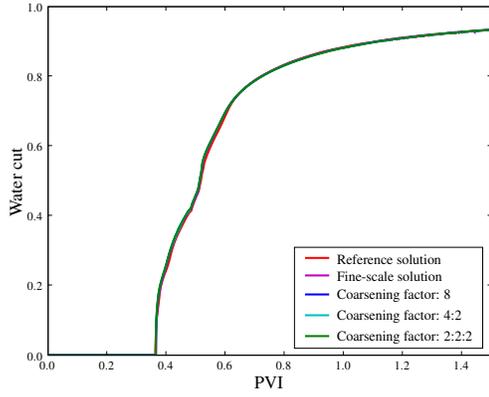
The full scale water-cut curves shown in Fig. 6, (a1) and (a2), are in excellent agreement with the reference solution for all three coarsening strategies. However, differences are apparent in the zoomed water cuts shown in Fig. 6 (b1) and (b2), and

in $\delta_S(t)$ shown in Fig. 6 (c1) and (c2). Specifically, the water cut for Scenario 1 (vertical streaks) shows that as the number of coarsening levels increases, the breakthrough time decreases. Moreover, for all coarsening strategies the shape of the water cut is very similar to the fine-scale solution, while the reference solution has even earlier but less abrupt breakthrough. In the simulation of Scenario 2 (streaks at 45 degrees), the water cut for different coarsening strategies is still in excellent agreement with the fine-scale solution. However, once again, the reference solution exhibits an earlier breakthrough than the fine-scale solution. Nevertheless, the difference in the breakthrough time between any of these coarsening strategies and the fine-grid solution is less than the difference in the breakthrough time between the fine-grid solution and the reference solution. This indicates that new M^3 method has captured the multi-scale nature of the fine-grid model very well, while the fine-grid model has not fully resolved the continuum model.

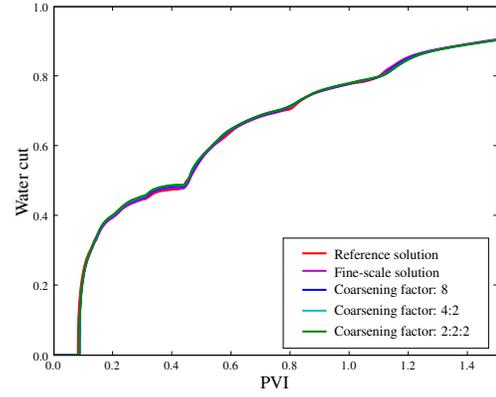
Similarly, the relative error in the saturation, $\delta_S(t)$ is plotted in Fig. 6, (c1) and (c2), and shows that over the majority of the simulation the error is less than 2%. Note that the large peaks in $\delta_S(t)$ are due to small errors in the breakthrough times. By definition (5.2) values of these peaks are bounded by 8.

In the next group of tests, we demonstrate that even with a large total coarsening factor, we obtain reasonable accuracy in the multiscale solution. Specifically, we consider from four to seven levels, with $c_i = 2$ on each level. This produces a total coarsening factor that varies from 8 to 64, in each direction. The coarsening parameters, α , are determined as in the first group of tests; to a high tolerance, only at the initial time. From the viewpoint of computational efficiency, there is no need to use more than four to six levels for such small two-dimensional problems. However, the experiments demonstrate that the M^3 method produces a reasonable solution even for a rather extreme total coarsening factor, and this may be necessary in real large-scale models.

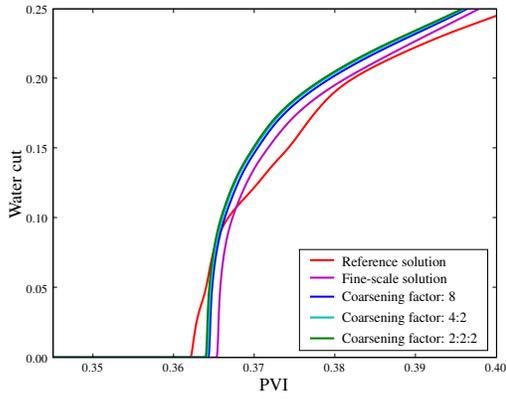
The water-cut curves and relative errors of saturation in the producer wells are shown in Fig. 7. Here the full scale water-cut curves, (a1) and (a2), are in good agreement with the reference solutions. In contrast to the previous tests, discrepancies are noticeable in these plots for both scenarios with the 2×2 coarsest mesh, and for Scenario 2 with the 4×4 coarsest mesh. The zoomed water cuts shown in Fig. 7 (b1) and (b2) indicate that the overall trend of earlier breakthrough with increased coarsening has continued. However, there are interesting differences in the performance for these scenarios, particularly for the coarsening strategy with a 2×2 coarsest mesh. First, with Scenario 1 (vertical layers), the overall shape of the water-cut is well approximated, despite its early breakthrough. In contrast, for Scenario 2 the breakthrough is not significantly earlier than with the other coarsening strategies, and is in fact later than the reference solution. However, the shape of the water-cut is more complex, and is not tracked well beyond approximately 0.5 PVI. This problem is apparent in the plots of $\delta_S(t)$ shown in Fig. 7 as well. Specifically, we see that for Scenario 1 the error is well below 5% for most of the simulation, while for Scenario 2, despite achieving a good estimate of the breakthrough time, $\delta_S(t) \approx 10\%$ for 0.2–0.4 PVI time range. Thus, in order for the M^3 method to achieve the desired robustness and



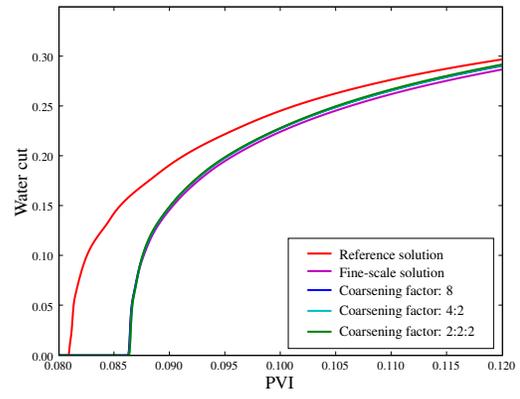
(a1)



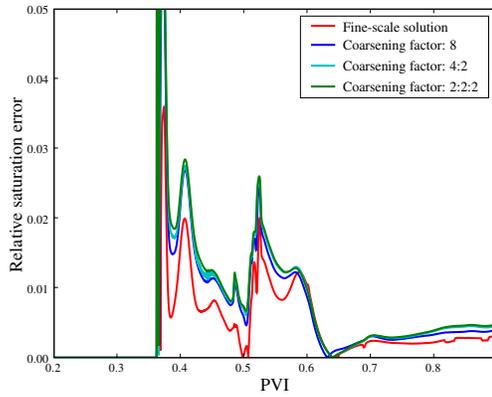
(a2)



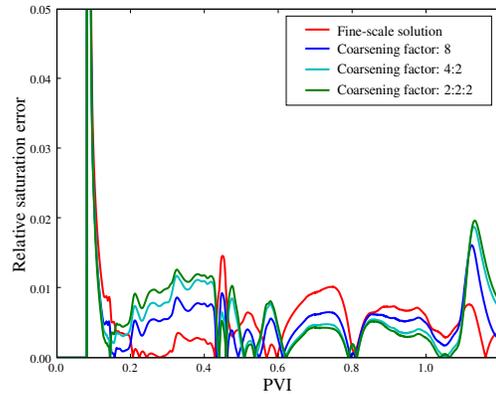
(b1)



(b2)



(c1)



(c2)

Figure 6: Comparison of three coarsening strategies for Scenario 1 (left column) and Scenario 2 (right column): water-cut curves (a1, a2), zoom of water-cut curves near the breakthrough time (b1, b2), and relative saturation error $\delta_S(t)$ in the producer well (c1, c2).

accuracy with a large total coarsening factor and over a broad class of problems, the temporal dependence of the coarsening parameters must be addressed.

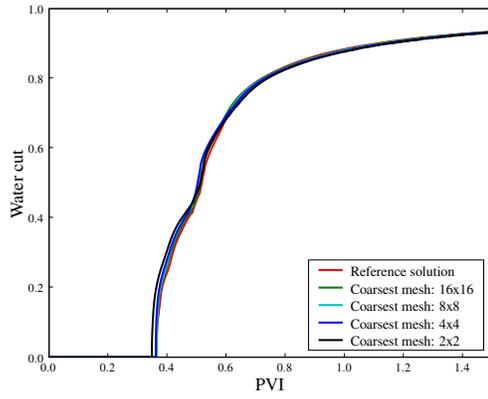
5.1.2 Simulations with updated coarsening parameters α 's

In the previous section, the coarsening parameters α were computed very accurately but only once, at the initial time step. In this section, we propose to perform a similar amount of the computational work but to distribute it more uniformly across the whole simulation. More precisely, instead of solving the fine-scale problem very accurately ($\varepsilon_r = 10^{-8}$) at the initial time step, we solve it approximately ($\varepsilon_r = 10^{-1} - 10^{-2}$) several times during the simulation. This approach provides a flexible strategy for controlling the accuracy of α 's and, in turn, the accuracy of the multilevel solution.

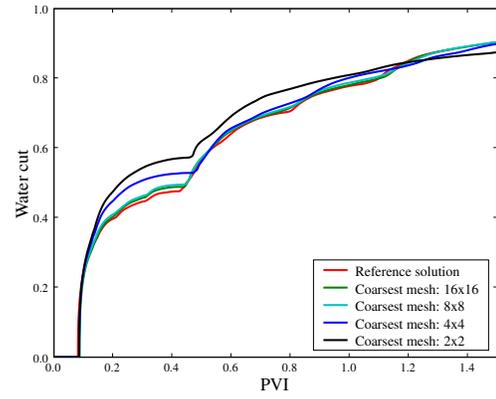
The total computational work required to compute the α 's depends on the frequency of the updates and the convergence tolerance, ε_r . The convergence rate of the PCG(AMG) iteration is bounded well below one independent of the mesh size, and is dominated by the heterogeneous structure of the absolute permeability. Consequently, the number of PCG(AMG) iterations required to achieve the tolerance ε_r , is essentially fixed independent of time. In addition, the total amount of work is always controllable, and is significantly less than solving the fine-scale model. Thus, the trade-off between accuracy and efficiency in the M³ method has significantly greater potential to be monitored and controlled than typical upscaling methods.

Different update strategies for α can be realized. In this section we demonstrate the importance of these updates using the simplest strategy: the flux ratios are updated 10 times in equal time intervals of 0.15 PVI. We consider four macro-edges located at points A, B, C, and D (see Fig. 5) and study the time behavior of the flux ratios for the fine-grid and multiscale solutions. The results are shown in Fig. 8, where it is clear from the fine-scale solution that rapid variations in the flux ratio of approximately 10% occur as the saturation front passes through these macro-edges. Moreover, the extent to which the flux ratios return to their initial values (the level line marking the non-updated case) depends on the location as well as the local flow characteristics. The stair-step approximation that we obtain with 10 equally spaced updates is shown for two different values of the convergence tolerance, $\varepsilon = 0.1$ and $\varepsilon = 0.01$.

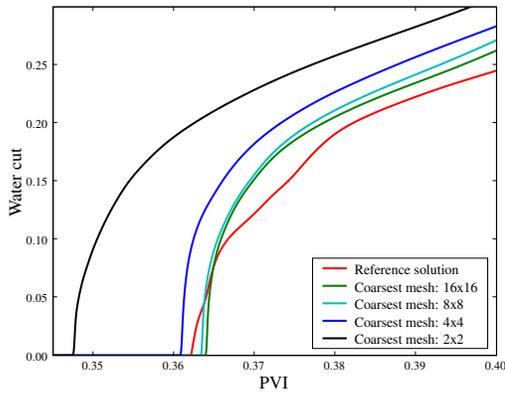
To demonstrate the weak influence of the convergence tolerance we perform two simulations for each scenario that differ only in the value of ε_r , namely 0.1 and 0.01. In particular, we consider a six-level method, which has a 4×4 coarsest mesh. The resulting water cuts and relative errors in saturation are shown in Fig. 9. Here we see that even this simple update strategy has significantly improved the multiscale solution for Scenario 2 (45° layering). Specifically, with updates $\delta_S(t)$ has been reduced by about a factor of two, to approximately 2.5%, in the 0.2 – 0.4 PVI time range. In addition, we see that solution with updates tracks the water cut much more closely, than the solution without updates. The differences for Scenario 1 are less pronounced but still apparent in the left column of Fig. 9. Most important, the



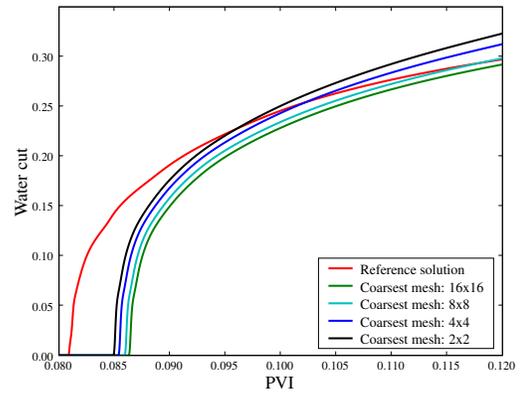
(a1)



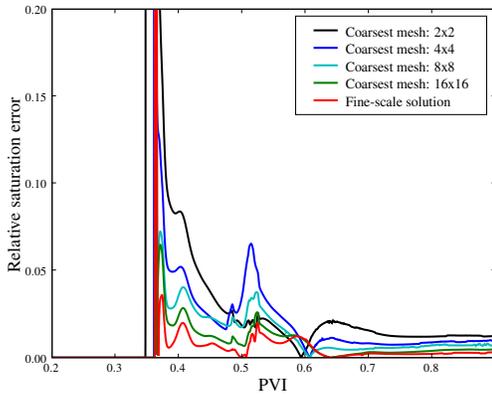
(a2)



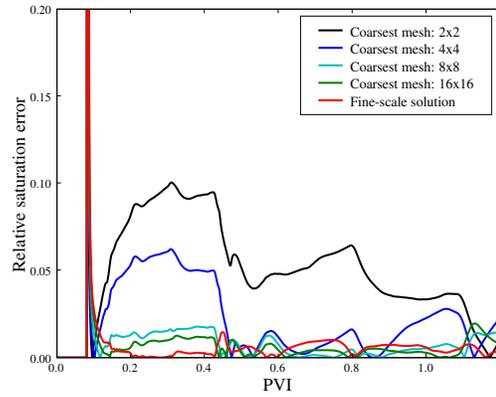
(b1)



(b2)

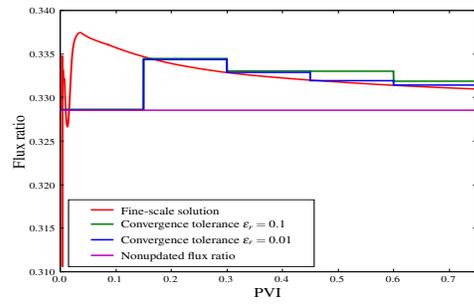
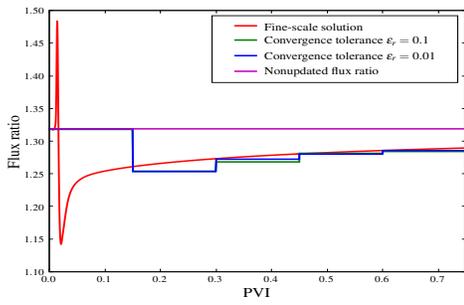


(c1)

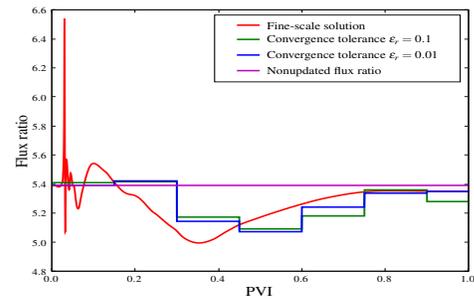
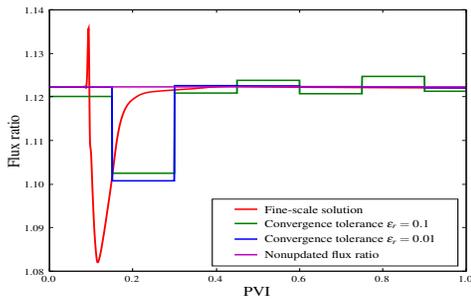


(c2)

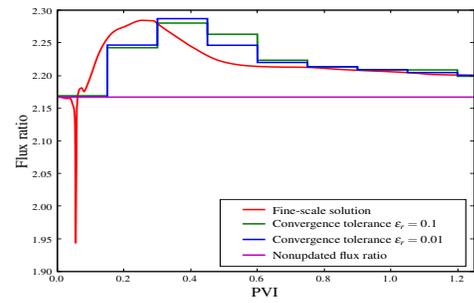
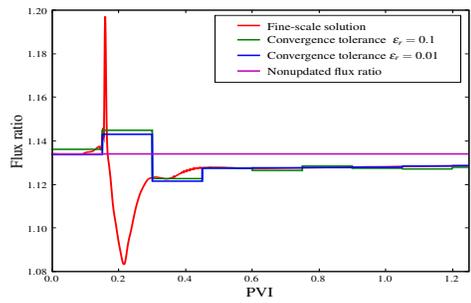
Figure 7: Dependence on the size of the coarsest mesh for Scenario 1 (left column) and Scenario 2 (right column): water-cut curves (a1, a2), zoom of water-cut curves near the breakthrough time (b1, b2), and the relative saturation error $\delta_S(t)$ in the producer well (c1, c2).



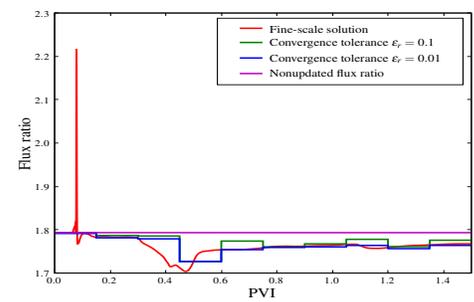
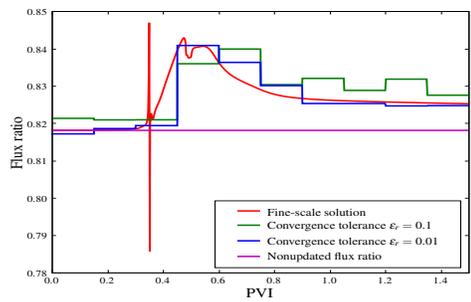
macro-edge A



macro-edge B



macro-edge C



macro-edge D

Figure 8: The flux ratios for the fine-scale solution and the multilevel solutions with flux-ratio updates on macro-edges A, B, C, and D. Left column corresponds to Scenario 1, Right column corresponds to Scenario 2.

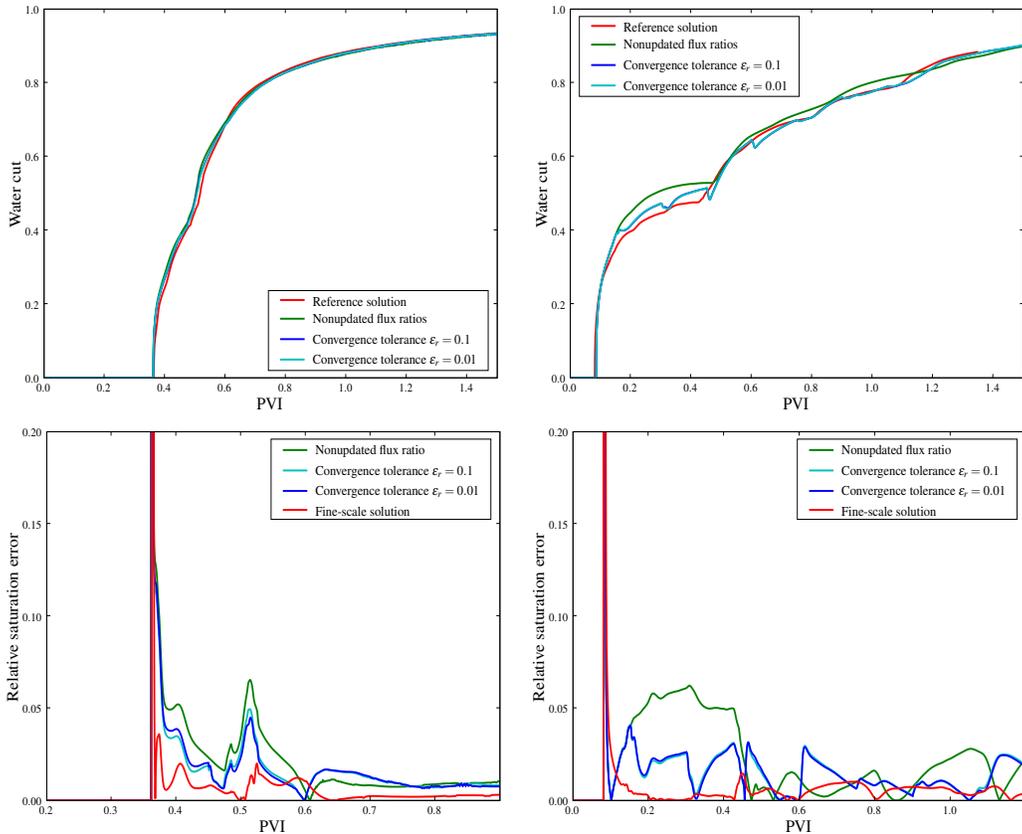


Figure 9: Dependence on the residual convergence factor ϵ_r . The flux ratios are updated every 0.15 PVI. The top row shows water-cuts for Scenario 1 (left) and Scenario 2 (right). The bottom row shows the relative saturation error in the producer well.

multiscale solutions generated with $\epsilon_r = 0.1$ and $\epsilon_r = 0.01$ are virtually indistinguishable. Thus, the design of the M^3 method, along with the robust AMG preconditioner underpinning the solver, allows a weak convergence tolerance to be used. Moreover, after the initial time there is some level of error in the saturation. Consequently, even if α is updated with $\epsilon_r = 10^{-8}$, the result would differ from those associated with the fine-scale solution.

In Fig. 10 we set $\epsilon_r = 0.1$, which corresponds to approximately 3 PCG(AMG) iterations, and consider several different levels of coarsening. Although, improvements in the solution are obtained for all coarsening strategies, it is apparent these updates increase in importance with the number of levels. Specifically, in Scenario 2 the temporal updates drive the water cut of the multiscale solutions back toward the fine-scale solution. The magnitude of these temporal bumps increases with the coarsening level, suggesting that we should consider increasing the frequency of these updates in conjunction with more extreme coarsening.

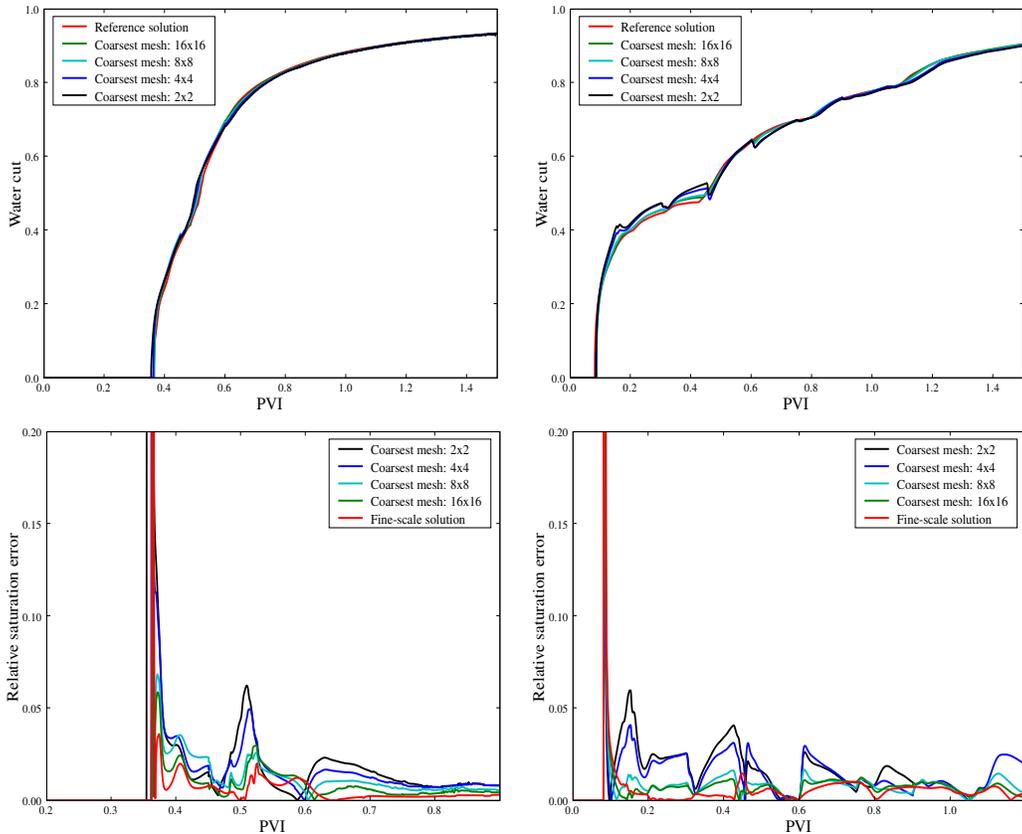


Figure 10: Dependence on the number of coarsening levels. The flux ratios are updated every 0.15 PVI. The residual convergence factor is $\varepsilon = 0.1$ (3 PCG iterations). The top row shows water-cuts for Scenario 1 (left) and Scenario 2 (right). The bottom row shows the relative saturation error in the producer well.

5.1.3 Adaptive calculation of subgrid matrices

The computational complexity of the M^3 method was analyzed in Section 4.2 and focused on the construction of the matrices \mathbf{M}_{ii}^{-1} and \mathbf{S}^+ for the macro-cells. These matrices depend on the geometry of the fine-scale cells, the media properties (total mobility and absolute permeability tensor), and the coarsening parameters. Since, the geometry and the absolute permeability tensor are assumed to be constant in time, temporal adaptation of the M^3 method is focused on the total mobility and the coarsening parameters. The coarsening parameters are updated very infrequently, approximately every 5,000 time steps in the previous GSLIB Scenarios (Section 5.1.2). At these isolated time steps, all matrices are recomputed. Therefore, the total mobility is the primary parameter that facilitates adaptive computations of local matrices.

The total mobility varies in time and space; however, it changes significantly in the vicinity of a sharp water front. The essential idea of the adaptive approach is to update local matrices only for those macro-cells where a change of the total mobility is relatively large. We implemented the update criteria proposed in [19]. Specifically

we use the condition,

$$\frac{1}{1 + \varepsilon_\lambda} < \frac{\lambda^n}{\lambda^{n-1}} < 1 + \varepsilon_\lambda, \quad (5.3)$$

where λ^n is total mobility at time steps t^n and ε_λ is a user-defined threshold. If this condition fails for any sub-cell of a particular macro-cell then the corresponding local matrices are updated. Thus, the smaller ε_λ , the larger the number of matrices that will be updated. This rule is implemented on all levels of the hierarchy.

Numerical experiments with the M³ method suggest that if $\varepsilon_\lambda \approx 0.1$, then the adaptive strategy does not introduce significant errors in the multiscale solution. These results are consistent with the results presented in [19]. Furthermore, with $\varepsilon_\lambda = 0.1$ a significant savings in computational work is achieved. To characterize this savings, let T_{pr} and T_{tr} be the times required to solve the pressure and transport equations, respectively, in computing the fine-scale solution. For the problems described in this section,

$$T_{pr} \approx 5T_{tr}.$$

Next, let T_{ms} be the time required to solve the pressure equation with the M³ method using four levels and $c_1 : c_2 : c_3 = 2 : 2 : 2$. Note that due to the small size of the fine-scale model, increasing the number of levels further does not significantly reduce the cost of the pressure solve. The numerical tests show that

$$T_{pr} \approx 2.5T_{ms}.$$

The effectiveness of the adaptive strategy depends on the time step. For larger time steps, the mobility changes more significantly and hence the larger the number of local matrices that must be updated. In our experiments, the adaptive strategy speeds up the pressure solver approximately 20 times with respect to the full multiscale algorithm:

$$T_{pr} \approx 2.5T_{ms} \approx 50T_{adp},$$

where T_{adp} is the time required to solve the pressure equation with the adaptive M³ method. For larger problems the speed up is even larger.

In general, the time required to update the coarsening parameters α should be considered as well. However, in our experiments these updates are very infrequent, and as a result, have no appreciable impact on the total simulation time.

5.2 SPE 10th Model

In this section, we consider a more realistic model from the upscaling benchmark tests in the Tenth SPE Comparative Solution project. Here, we focus on layer 68 of the three-dimensional reservoir. The absolute permeability field is shown in Fig. 11, where the high permeability channels of this fluvial layer are clearly visible. The dynamic range of the absolute permeability is approximately 6 orders of magnitude, as it was for the GSLIB examples.

We consider the classic five-spot well configuration with the injector in the middle of the domain and the producer wells at the corners (see Fig. 11). The fine-scale

geological model is a Cartesian grid with 60×220 cells, and the two-phase flow parameters are defined at the beginning of Section 5. We assume that the rates of the injector and producer wells are constant, with each producer accounting for one quarter of the injected fluid. Although, this assumption is not physically realistic, it allows us to evaluate the new multiscale method without the additional complication of well modeling, which is beyond the scope of this paper. Moreover, the algebraic nature of the M^3 method provides a natural mechanism to incorporate different well models, and this will be explored in the future.

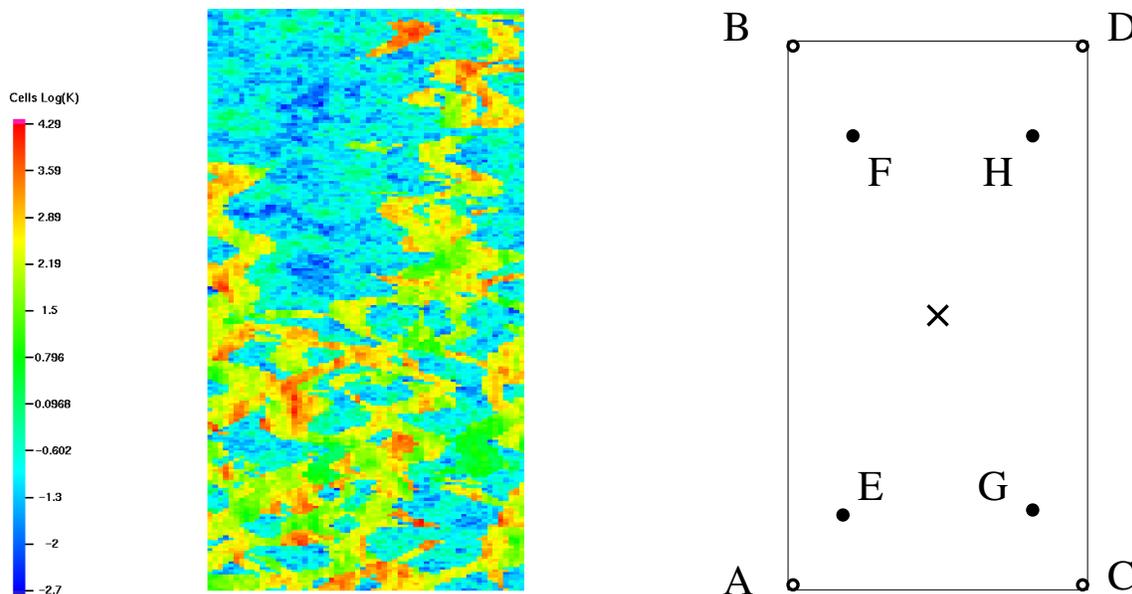


Figure 11: the left picture shows the permeability field of the layer 68 in the SPE 10th model. the right picture shows locations of the injector well (\times) and producer wells A, B, C, and D (\circ). The monitored macro-edges (\bullet) are located at points E, F, G, and H.

In the first group of tests, as in Section 5.1.1, we demonstrate that the M^3 method produces reasonably accurate solutions even with a large total coarsening factor and many levels. Specifically, we consider the total coarsening factors 8, 16, and 32 in each direction, which correspond to methods with 4, 5, and 6 levels, respectively. In fact, in the case of 6 levels, the coarsest mesh in the M^3 method is 2×7 macro-cells. This is approximately a factor of 3 coarser in each direction than most upscaling studies of this SPE benchmark. We note, that since the fine-grid dimensions are not powers two, rectangular macro-cells are introduced as needed along the boundaries. The coarsening parameters, α , are determined as in Section 5.1.1; to a high tolerance, only at the initial time.

The full scale water-cut curves for the four producer wells are shown in Fig. 12. Overall the water-cut obtained with the different coarsening strategies is in good agreement with the fine-scale solution. However, small visible discrepancies suggest similarities with the previous GSLIB based Scenario 2 case (45° layering). Specifically,

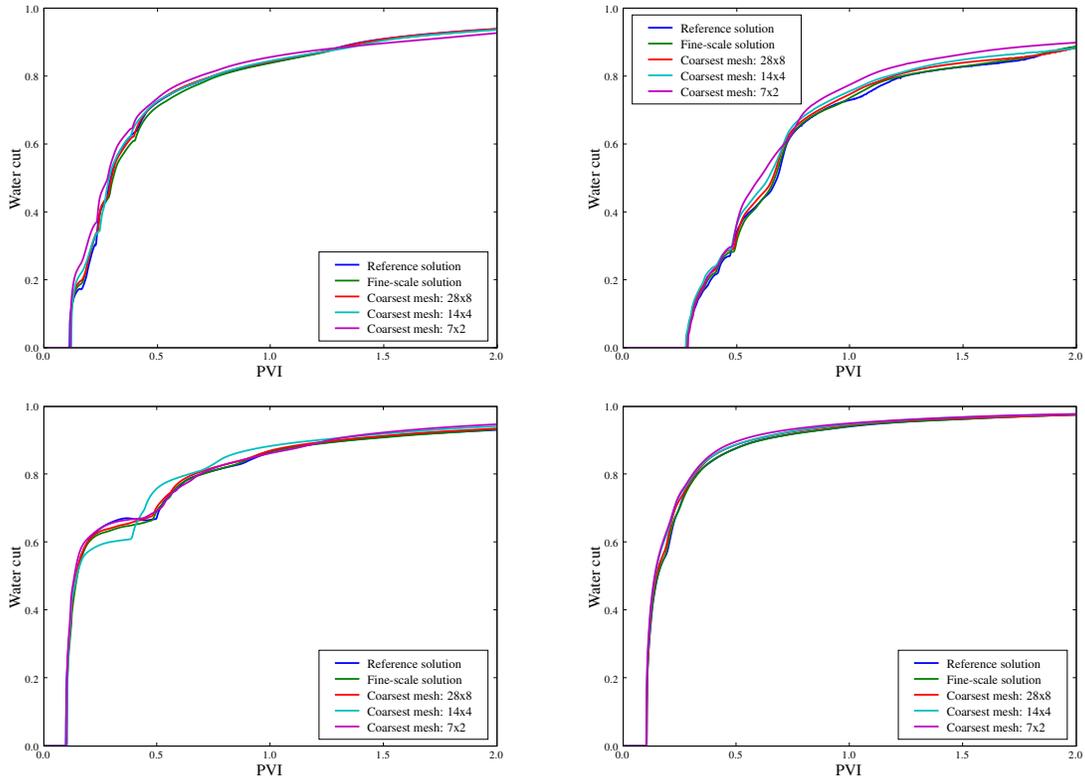


Figure 12: Water-cut curves for total coarsening factors 8, 16, and 32 at producers A (top-left), B (top-right), C (bottom-left), and D (bottom-right).

the water-cuts for producers A, B, and C have noticeable temporal structure. Once again, since the spatial structure of the highly permeable flow channels is not aligned with the grid, the single-point upwind advection scheme is stressed. Moreover, the larger total coarsening factor is the less accurate multiscale solution is obtained. Together these factors produce a noticeable drift and damping of the temporal structure, which is most apparent in the six level multiscale solution.

To highlight these features of the multiscale solution, Fig. 13 shows zoomed relative errors of saturation at producer A, which has the largest $\delta_S(t)$, and producer D, which has the most uniform and smallest $\delta_S(t)$. As expected, the errors increase modestly with the number of levels, but are most pronounced shortly after breakthrough in producer A. In fact, for the six-level case, the breakthrough time is well approximated, yet the relative error in saturation is approximately 10% in the 0.1 PVI following breakthrough, and slowly drops below to less than 5% around 0.3 PVI. Thus, to achieve the desired accuracy and robustness we must address the temporal dependence of the coarsening parameters.

It is important to note that, in general, the multiscale solution cannot be uniformly more accurate than the fine-scale solution over the entire domain and throughout the entire simulation time. However, localized anomalies will naturally arise through the fortuitous cancellation of errors. For example, for producer A, between approximately

0.3 and 1.2 PVI the multiscale solutions based on 4 and 5 levels appear to be more accurate than the fine-scale solution. While errors elsewhere, at producer D for example, increase modestly and monotonically with the number of levels.

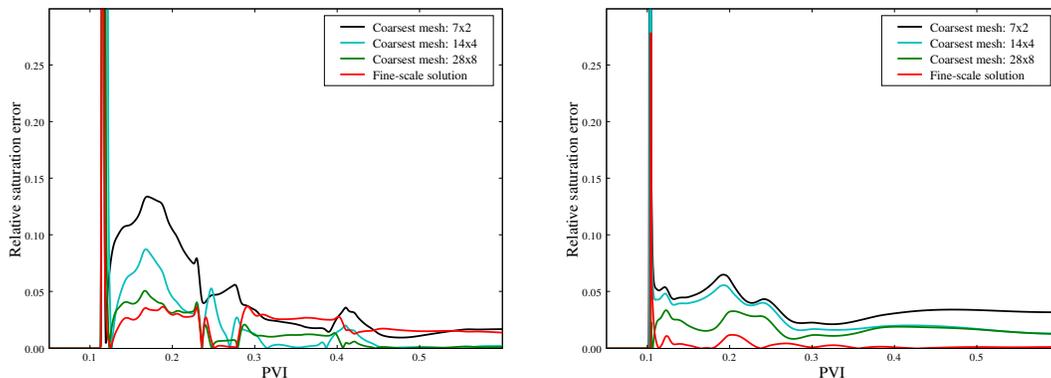


Figure 13: Relative error of saturation at producer A (largest relative error) and at producer D (smallest relative error) for total coarsening factors 8, 16, and 32.

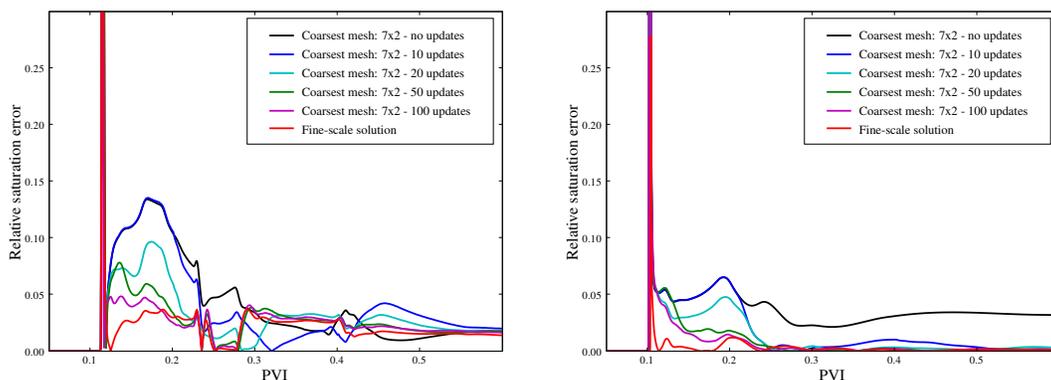


Figure 14: Relative error in saturation at producer A (largest relative error) and at producer D (smallest relative error) for the different number of updates. The coarsest mesh consists of 7×2 macro-cells. The convergence tolerance is $\varepsilon_r = 0.01$.

To reduce the errors and improve robustness, particularly for the simulations with the large total coarsening factors, we update the coarsening parameters in time. We focus this study on the the six-level case, with a total coarsening factor of 32. As in Section 5.1.2 we use the simple strategy of uniform updates in time, and study the performance of the M^3 method with 10, 20, 50 and 100 updates. This corresponds to very infrequent updates of the coarsening parameters, approximately ranging from every 5,000 to every 500 time steps. The convergence tolerance, $\varepsilon = 0.01$ is used and corresponds to approximately 5 PCG(AMG) iterations per update.

The relative error in saturation for producers A and D with different numbers of updates, are shown in Fig. 14. The general trend is as expected, the more frequent

the updates of α , the more reliably the multiscale solution approximates the fine-scale solution. A problem area worth noting is the relative error in saturation at producer A in the first 0.1 PVI after breakthrough. Here, 10 updates does not have a significant impact on the error, but 20 updates reduces the peak error (at about 0.18PVI) down below 10%. With 100 updates not only this peak is reduced but the multiscale solution exhibits a nearly uniform error of approximately 5%, which is comparable to the fine-scale solution. Moreover, the error at later times is converging to the fine-scale error with fewer anomalous regions over which the multiscale solutions appear more accurate. Similar improvements are observed in producer D. Here, the error was well behaved even without updates, but was steadily around 5% at later times. Once again, 10 updates does not have a significant impact in the first 0.1 PVI after breakthrough. However, once again, with 100 updates, the error after breakthrough quickly drops below 2.5%.

To highlight further that these weak tolerance updates enable the multiscale solution to better approximate the fine-scale solution, we plot the flux ratios for four macro-edges, E, F, G, and H of Fig. 11. in Fig. 15. First, it is apparent that significant deviations of approximately 10% from the initial flux-ratio occur as the saturation front passes through each macro-edge. In the case of 10 updates, only macro-edge E, which has the slowest and smoothest evolution of α , is well approximated by these updates. For the other macro-edges, the updates generate a lagged profile that may completely step over features, such as the drop in α around 0.4PVI for macro-edge F. As noted in the previous section, it is the error in the saturation that dominates the error in the flux ratios, and hence, decreasing ε_r will not improve the accuracy of the flux ratios or the multiscale solution. Although, we note that the long time behavior is captured well even with 10 updates. Thus, it is not surprising that we observed a marked improvement in the error in saturation at later times (Fig. 14), even with only 10 updates. However, to improve the accuracy of the multiscale solution shortly after breakthrough, it is apparent that more frequent updates are required. In particular, the faster features are not better resolved with 20 updates (approximately every 0.1 PVI). Consequently, we did not observe a significant improvement in the accuracy for any of the producers shortly after breakthrough (Fig. 14) using only 20 updates. In fact, 100 updates were necessary (approximately every 0.02 PVI) to drive the error in producer A below 5% for all times after breakthrough, which is comparable to the error present in the fine-scale solution.

Thus, developing update strategies that adequately resolve important fine-scale temporal features is critical to achieving a specified level of accuracy in the quantities of interest. Moreover, the overall computational efficiency of the M^3 method is maintained because a weak convergence tolerance on the PCG(AMG) algorithm delivers sufficient accuracy in the coarsening parameters.

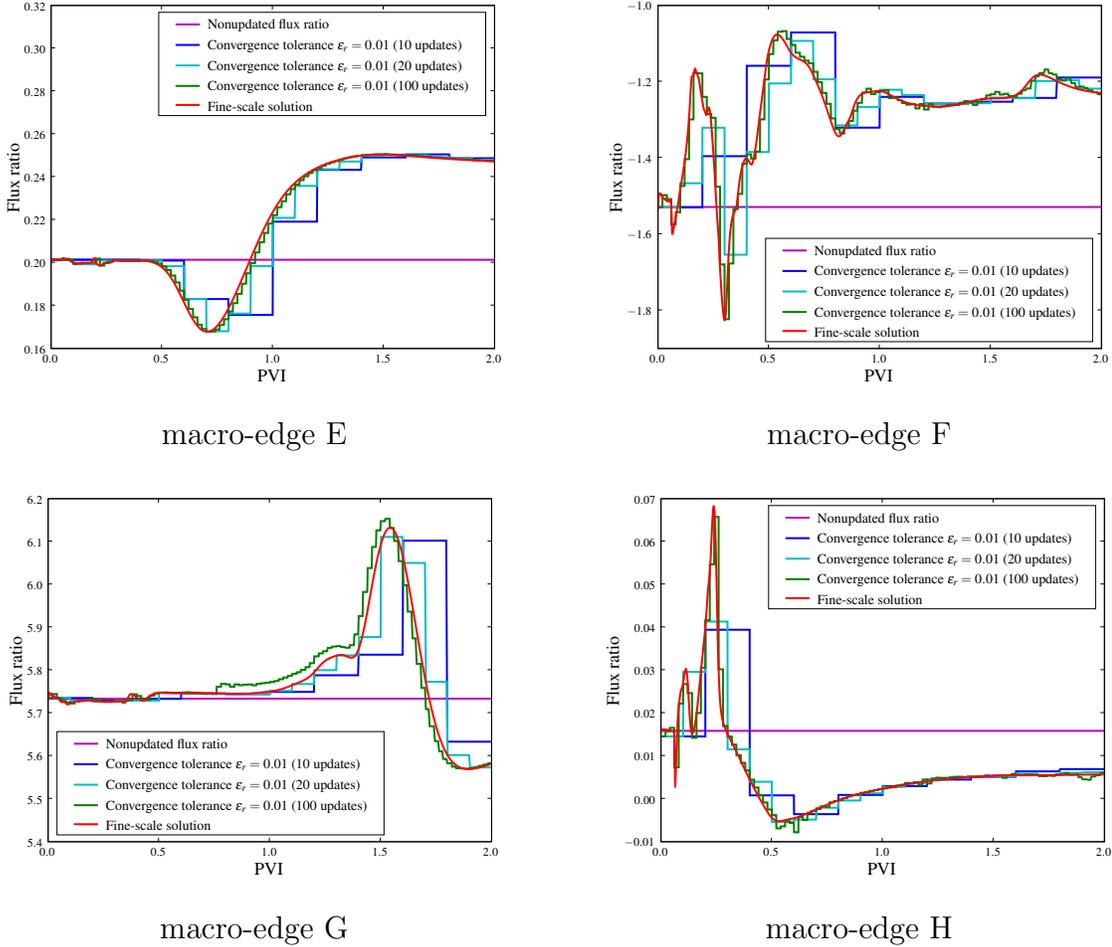


Figure 15: The comparison of the flux ratios obtained with different update frequencies.

6 Conclusions

We developed a new multilevel multiscale mimetic (M^3) method for the IMPES time discretization of two-phase flow models. In this formulation, the solution of the pressure equation on each time step is the dominant part of the simulation cost. Hence, a critical component of the M^3 method is the new multilevel coarsening algorithm for this equation. This algorithm brings together the recently proposed subgrid modeling algorithm with the algebraic multigrid for accurate calculation of the flux coarsening parameters. Through a sequence of algebraic manipulations, this algorithm generates a hierarchy of pressure equations that all have the same form. By design, the solution on all levels of the hierarchy is locally mass conservative independently of the accuracy of the coarsening parameters. Thus, on each time step, we first solve the pressure equation on the coarsest level and then interpolate the solution to obtain the fine-grid fluxes that are needed to advect the saturation.

Although, the basic elements of this approach are similar to existing multiscale

methods, there are important features that set it apart. To explore these new features and demonstrate the accuracy of the M^3 method we considered two models of highly heterogeneous permeability fields. In these numerical experiments we showed that M^3 is a truly multilevel method that is designed to use many levels and to enable large total coarsening factors. Specifically, we used up to seven levels and coarsening factors up to 64 in each direction. In contrast, typical studies consider a two-level approach with a coarsening factor of approximately 10 in each direction. Moreover, even with this large coarsening factor we obtained good accuracy in the water-cut and relative error in saturation at the producer wells.

The M^3 method uses a robust and efficient multigrid algorithm that naturally incorporates global flow information into estimates of the coarsening parameters α . When these parameters were computed very accurately only once, at the initial time step, the multiscale solution with many levels in the hierarchy still estimated the breakthrough time quite well. However, temporal features of multiscale solution tended to drift and be damped after breakthrough, producing relative errors in saturation of approximately 5 – 15%. The new method readily addresses this problem by updating the α 's infrequently throughout the simulation. Specifically, we use PCG(AMG) iterations with a convergence tolerance ε_r to update α . We demonstrated that with this multigrid approach, we achieved the required accuracy in α , even with a weak convergence tolerance. Moreover, we observed that using more frequent updates, even with this weak convergence tolerance, generated a more reliable approximation of the fine-scale solution. For example, in the SPE benchmark updating 100 times (every 0.02 PVI or approximately every 500 time steps) with $\varepsilon_r = 0.01$ reduced the error in producer A to below 5% throughout the simulation. This level of error, relative to the reference solution, is comparable to the error in the fine-scale solution. Yet, the cost of these updates is a negligible part of the overall simulation cost.

In the future, we are interested in treating more general distorted and unstructured polyhedral grids and well models for three dimensional reservoirs. Also, further gains in efficiency may result from incorporating the estimation of the coarsening parameters α directly into the hierarchy of first-order models, possibly by exploiting a connection to a smoothed aggregation based AMG method.

Acknowledgments

This work was carried out under the auspices of the National Nuclear Security Administration of the U.S. Department of Energy at Los Alamos National Laboratory under Contract No. DE-AC52-06NA25396 and the DOE Office of Science Advanced Scientific Computing Research (ASCR) Program in Applied Mathematics Research.

References

- [1] J. E. Aarnes. On the use of a mixed multiscale finite element method for greater flexibility and increased speed or improved accuracy in reservoir simulation. *Multiscale Model Sim.*, 2(3):421 – 439, 2004.
- [2] J. E. Aarnes, V. Kippe, and K.-A. Lie. Mixed multiscale finite elements and streamline methods for reservoir simulation of large geomodels. *Adv. Water Resour.*, 28(3):257 – 271, 2005.
- [3] I. Babuška and J. E. Osborn. Generalized finite element methods: Their performance and their relation to mixed methods. *SIAM J. Numer. Anal.*, 20(3):510–536, 1983.
- [4] A Bensoussan, J.-L. Lions, and G. Papanicolaou. *Asymptotic Analysis For Periodic Structures*, volume 5 of *Studies in Mathematics and its Applications*. North-Holland, New York, 1978.
- [5] M. Brezina, R. Falgout, S. MacLachlan, T. Manteuffel, S. McCormick, and J. Ruge. Adaptive smoothed aggregation (α SA) multigrid. *SIAM Rev.*, 47(2):317 – 346, 2005.
- [6] F. Brezzi and M. Fortin. *Mixed and Hybrid Finite Element Methods*. Number 15 in Springer Series in Computational Mathematics. Springer-Verlag, New York, 1991.
- [7] F. Brezzi, K. Lipnikov, M. Shashkov, and V. Simoncini. A new discretization methodology for diffusion problems on generalized polyhedral meshes. *Comput. Method Appl. M.*, 196(37-40):3682 – 3692, Aug 2007.
- [8] F. Brezzi, K. Lipnikov, and V. Simoncini. A family of mimetic finite difference methods on polygonal and polyhedral meshes. *Math. Mod. Meth. Appl. S.*, 15(10):1533–1551, October 2005.
- [9] Z. Chen and T. Y. Hou. A mixed multiscale finite element method for elliptic problems with oscillating coefficients. *Math. Comput.*, 72(242):541 – 76, 2003.
- [10] Z. Chen, G. Huan, and Y. Ma. *Computational Methods for Multiphase Flows in Porous Media*. Computational Science and Engineering Series. SIAM, Philadelphia, PA, 2006.
- [11] J. E. Dendy. Black box multigrid. *J. Comput. Phys.*, 48:366–386, 1982.
- [12] J. E. Dendy. Two multigrid methods for three-dimensional problems with discontinuous and anisotropic coefficients. *SIAM J. Sci. Stat. Comput.*, 8(2):673–685, 1987.

- [13] C. V. Deutsch and A. G. Journel. *GSLIB: Geostatistical software library and user's guide*. Oxford University Press, New York, 2nd edition, 1998.
- [14] Y. Efendiev, V. Ginting, T. Hou, and R. Ewing. Accurate multiscale finite element methods for two-phase flow simulations. *J. Comput. Phys.*, 220(1):155–174, 2006.
- [15] P.A. Forsyth. Numerical simulation of gas venting for NAPL site remediation. *Soc Pet Eng AIME Pap SPE*, 18415:85 – 96, 1989.
- [16] P.A. Forsyth. A control volume finite element approach to NAPL groundwater contamination. *SIAM Journal on Scientific and Statistical Computing*, 12(5):1029–1057, 1991.
- [17] V. Gvozdev. discretization of the diffusion and Maxwell equations on polyhedral meshes. Technical Report Ph.D. Thesis, University of Houston, 2007.
- [18] T.Y. Hou and X.H. Wu. A multiscale finite element method for elliptic problems in composite materials and porous media. *J. Comput. Phys.*, 134:169–189, 1997.
- [19] P. Jenny, S. H. Lee, and H. A. Tchelepi. Adaptive multiscale finite-volume method for multiphase flow and transport in porous media. *Multiscale Modeling and Simulation*, 3(1):50 – 64, September 2004; 2004-2005.
- [20] V.V. Jikov, S.M. Kozlov, and O.A. Oleinik. *Homogenization of Differential Operators and Integral Functionals*. Springer-Verlag, 1994. Translated from Russian.
- [21] V. Kippe, J.E. Aarnes, and K.-A. Lie. A comparison of multiscale methods for elliptic problems in porous media flow. *Computat. Geosci.*, 2007. to appear in Special Issue on Multiscale Methods.
- [22] Yu. Kuznetsov. Mixed finite element methods on polyhedral meshes for diffusion equations. In *Computational modeling with PDEs in science and engineering*. Springer-Verlag, Berlin, 2008. (to appear).
- [23] Yu. A. Kuznetsov. Mixed finite element method for diffusion equations on polygonal meshes with mixed cells. *J. Numer. Math.*, 14(4):305–315, 2006.
- [24] S.P. MacLachlan and J.D. Moulton. Multilevel upscaling through variational coarsening. *Water Resour. Res.*, 42, 2006.
- [25] C. Prakash. Examination of the upwind (donor-cell) formulation in control volume finite-element methods for fluid flow and heat transfer. *Numerical Heat Transfer*, 11(4):401 – 16, 1987 1987.
- [26] J. W. Ruge and K. Stüben. Algebraic multigrid (AMG). In S. F. McCormick, editor, *Multigrid Methods*, volume 3 of *Frontiers in Applied Mathematics*, pages 73–130. SIAM, Philadelphia, PA, 1987.

- [27] T. Strouboulis, L. Zhang, and I. Babuška. p-version of the generalized FEM using mesh-based handbooks with applications to multiscale problems. *Int. J. Numer. Meth. Eng.*, 60(10):1639 – 1672, 2004.
- [28] K. Stüben. Algebraic multigrid (AMG): Experiences and comparisons. *Appl. Math. Comput.*, 13:419–451, 1983.
- [29] K. Stüben. A review of algebraic multigrid. *J. Comp. Appl. Math.*, 128(1/2):281 – 309, 2001.
- [30] P. Vanek, J. Mandel, and M. Brezina. Algebraic multigrid by smoothed aggregation for second and fourth order elliptic problems. *Computing*, 56(3):179 – 96, 1996.
- [31] Shi-tien Yang and A.F. Henry. A finite element synthesis method. *Nucl. Sci. Eng.*, 59(1):63 – 67, 1976.