ALGEBRAIC MULTIGRID PRECONDITIONERS FOR MULTIPHASE FLOW IN POROUS MEDIA

QUAN BUI*, HOWARD ELMAN†, AND DAVID MOULTON‡

Abstract. Multiphase flow is a critical process in a wide range of applications, including carbon sequestration, contaminant remediation, and groundwater management. It is modeled by a nonlinear system of partial differential equations derived by considering the mass conservation of each phase (e.g., oil, water), along with constituitive laws for the relationship of phase velocity to phase pressure. The constraint that the phase saturations sum to one, along with initial and boundary conditions, closes this system. The nonlinearity of the constuitive laws, in conjuction with the coupling of the phases, often requires the use of implicit discretization in time for both stability and accuracy. In this work we study a model of immiscible two-phase flow in which the primary variables are the pressure of the wetting phase and the saturation of the nonwetting phase. We use a finite volume method for spatial discretization and the backward Euler method for time discretization of the coupled system, leading to a fully implicit solution method. In this setting the capillary pressure is the difference between the non-wetting phase and wetting phase pressures. If the capillary pressure changes quickly with respect to saturation, then the operators associated with each phase are diffusion-dominated, whereas if capillary pressures vary slowly, then the saturation depends strongly on an advection-dominated, nearly hyperbolic, operator. These variations in character affect the performance of AMG-based solvers. Here, we present our experience with the GMRES solution of the linear systems resulting from the linearization of the coupled equations with algebraic multigrid (AMG) as a preconditioner, and with two constrained pressure residual multigrid (CPR-AMG) preconditioners. AMG is implemented using the HYPRE software package from Lawrence Livermore National Laboratory. Numerical expermients demonstrate that GMRES with AMG preconditioning for the coupled systems works best in the diffusion-dominated case, but it suffers slow convergence and sometimes diverges in the advection-dominated case. Similarly, the previously established CPR-AMG method that uses AMG to solve the pressure block, does not scale optimally with problem size. The proposed CPR-AMG method that uses AMG to solve both the pressure block and the saturation block generally performs well for advection-dominated problems and scales optimally with problem size.

1. Introduction. Multiphase flow is a feature of many physical systems and models of it are used in many settings, including reservoir simulation, carbon sequestration, ground water management and contaminant transport. Modelling multiphase flow in highly heterogeneous media with complex geometries is difficult, especially when realistic processes such as capillary pressure are included. The system describing multiphase flow consists of nonlinear partial differential equations, constitutive laws and constraints. In this paper, we focus on the iterative solution of linear systems arising in a fully implicit cell-centered finite volume discretization of single component isothermal two-phase flow model with capillary pressure. This fully implicit time-stepping scheme is among the most robust for simulation of subsurface flow. Moreover, it can serve as a basis for modeling more complex processes in which the physical quantities are tightly coupled. These could be, for example, inclusion of additional components, miscibility between components, thermal effects, and phase transitions.

The fully implicit discretization gives rise to a nonlinear system of equations at each time step. We employ the standard Newton's method with an exact Jacobian of the discretized equations to solve this system. For the linear system, we use a preconditioned GMRES method with three different preconditioning strategies: (1) a direct AMG preconditioner for the global system, (2) two-stage CPR-AMG with correction for the pressure block, also known as the combinative two-stage approach, and (3) CPR-AMG with correction for both the pressure and saturation blocks, known as the two-stage additive approach. In this work, we report our experience with the performance and scalability of these strategies.

The simulator for two-phase flow is implemented within Amanzi, the computational engine of the Advanced Simulation Capability for Environmental Management (ASCEM) project [1]. Amanzi is interfaced with the Hypre package, a software library for high performance preconditioners and solvers for large, sparse linear systems developed by Lawrence Livermore National Laboratory. In particular, we use BoomerAMG [16] to define our algebraic multigrid preconditioners.

2. Problem Statement. We consider isothermal, immiscible two-phase flow through a porous medium. For example, often in reservoir simulation, one phase is oil (the nonwetting phase) and the other is pure water (wetting phase); alternatively, in groundwater management, one may consider a system of contaminated water that infiltrates a domain saturated with air.

*Applied Math, Stats, and Scientific Computation, University of Maryland, College Park, MD (mquanbui@math.umd.edu).
†Department of Computer Science, University of Maryland, College Park, MD (elman@cs.umd.edu).
‡Applied Math and Plasma Physics Group, T-5, Los Alamos National Laboratory, Los Alamos, NM (moulton@lanl.gov).

Conservation of mass of each of the phases leads to the following coupled PDEs:

$$\phi \frac{\partial \rho_w s_w}{\partial t} + \nabla \cdot (\rho_w \mathbf{v}_w) = q_w \tag{1}$$

$$\phi \frac{\partial \rho_n S_n}{\partial t} + \nabla \cdot (\rho_n \mathbf{v}_n) = q_n \tag{2}$$

in which S_w, S_n are the saturation, ρ_w, ρ_n are the densities, q_w, q_n are the source terms of the wetting and non-wetting phases respectively, and ϕ is the porosity of the medium. We assume an extension of Darcy's law to multiphase flow and express the velocities $\mathbf{v}_w, \mathbf{v}_n$ as

$$\mathbf{v}_\alpha = -\frac{k_{r\alpha}\mathbf{K}}{\mu_\alpha}(\nabla P_\alpha - \rho_\alpha g \nabla D), \qquad \alpha = w, n. \tag{3}$$

Here, $\mathbf{K}$ is the absolute permeability tensor. The terms $k_{r\alpha}, \mu_\alpha, P_\alpha$ are the relative permeability, viscosity, and pressure of phase α respectively, g is gravitational constant, and D is the depth. We also define the phase mobility $\lambda_\alpha = k_{r\alpha}/\mu_\alpha$. To close the system, we also have the following constitutive law and constraint

$$Pc(S_w) = P_n - P_w \tag{4}$$

$$S_w + S_n = 1 \tag{5}$$

From equations (1) and (2), one can derive separate equations for pressure and saturation. The pressure equation is elliptic or parabolic (diffusion); the saturation equation is hyperbolic or convection-dominated. The pressure equation is solved implicitly, and depending on the time discretization strategies applied to the saturation equation, several methods have been developed. In the case where the saturation equation is discretized using an explicit method (e.g., forward Euler), it is referred as IMPES (implicit pressure explicit saturation)[3]; for an implicit time discretization of the saturation equation, the method is known as the sequential approach, which was first applied to the black-oil model by Watts in 1985 [29].

The appeal of these methods lies in the complete decoupling between pressure and saturation variables. Each equation can be solved separately. In addition, knowing the characteristics of each equation facilitates the design of efficient preconditioners, which is critical to achieving high performance. Both of these methods have been successfully applied to many problems where the fully implicit method is difficult to implement or shown to be too costly. However, the solution obtained from these approaches may lose accuracy if pressure and saturation are strongly dependent, or if capillary pressure changes very quickly. The lack of accuracy of these methods can be even more pronounced if more complex processes such as miscibility, thermal, and phase transitions are included in the model. For a more complete summary of the advantages and disadvantages of these approaches, we refer to [19].

Substitution of (3) and (4) into (1) and (2) and using the constraint (5) yields a system of two equations and two unknowns. Using one popular choice of primary variables, the pressure in the wetting phase and saturation in the nonwetting phase $\mathbf{u} = (P_w, S_n)$, we obtain

$$-\frac{\partial(\phi\rho_w S_n)}{\partial t} - \nabla \cdot \left(\rho_w \frac{k_{rw}(S_n)}{\mu_w}\mathbf{K}(\nabla P_w - \rho_w g \nabla D)\right) = q_w \tag{6}$$

$$\frac{\partial(\phi\rho_n S_n)}{\partial t} - \nabla \cdot \left(\rho_n \frac{k_{rn}(S_n)}{\mu_n}\mathbf{K}(\nabla(P_w + P_c(S_n)) - \rho_n g \nabla D)\right) = q_n \tag{7}$$

In this paper, we consider solving the system consisting of (6) and (7) together fully implicitly. We use a cell-centered finite volume method for spatial discretization and the backward Euler method for time discretization, similar to an approach defined in [13]. This will serve as a base model for adding more complexity in the future. The finite volume method described below is known for its mass conservation property. In addition, it can deal with the case of discontinuous permeability coefficients, and it is relatively straightforward to implement. Under appropriate assumptions, this method also falls into the mixed finite element framework [22]. For simplicity, we use a uniform grid. For each cell i, integration of the mass conservation equations and the divergence theorem gives

$$\frac{\partial}{\partial t}\int_{C_i} \xi_\alpha + \sum_{j \in \eta_i} \int_{\gamma_{ij}} \psi_\alpha \cdot \mathbf{n} = \int_{C_i} q_\alpha \tag{8}$$

where the *storage* $\xi_\alpha = \phi\rho_\alpha S_\alpha$ and the *flux* $\psi_\alpha = \rho_\alpha \mathbf{v}_\alpha$ terms are approximated using the mid-point rule which is second-order accurate:

$$\bar{\xi}_\alpha = \frac{1}{V_{C_i}} \int_{C_i} \xi_\alpha, \quad Q_\alpha = \frac{1}{V_{C_i}} \int_{C_i} q_\alpha. \tag{9}$$

The surface integrals are discretized using two-point flux approximation (TPFA); dropping the phase subscript, this gives

$$\int_{\gamma_{ij}} \psi \cdot \mathbf{n} = -\gamma_{ij} \Big(\rho \frac{k_r}{\mu} \mathbf{K}\Big)_{ij+1/2} \big(\omega_i - \omega_j\big)\,, \tag{10}$$

$$\omega_i = \frac{P_i - \rho_{ij+1/2}\; g D_i}{\Delta x_{ij+1/2}}\,, \tag{11}$$

in which γ_{ij} is the area of the face adjacent to cells i, j. The index $ij+1/2$ signifies an appropriate averaging of properties at the interface between cell i and j. The coefficients $(\rho k_r/\mu)_{ij+1/2}$ are approximated by upwinding based on the direction of the velocity field, i.e.,

$$\Big(\rho \frac{k_r}{\mu}\Big)_{ij+1/2} = \begin{cases} (\rho \frac{k_r}{\mu})_i, & \text{if } \mathbf{v} \cdot \mathbf{n} > 0 \\ (\rho \frac{k_r}{\mu})_j, & \text{otherwise} \end{cases} \tag{12}$$

and the absolute permeability tensor on the faces are computed using harmonic averaging,

$$\mathbf{K}_{ij} = (\Delta x_i + \Delta x_j)\Big(\frac{\mathbb{K}_i \mathbb{K}_j}{\Delta x_i \mathbb{K}_j + \Delta x_j \mathbb{K}_i}\Big). \tag{13}$$

Discretization in time using the backward Euler method gives a fully discrete system of nonlinear equations,

$$(\bar{\xi})_i^{n+1} - (\bar{\xi})_i^n = -\frac{\triangle t}{V_{C_i}} \sum_{j \in \eta_i} \gamma_{ij} (\rho \frac{k_r}{\mu} \mathbb{K})_{ij+1/2}^{n+1} \Big(\omega_i^{n+1} - \omega_j^{n+1}\Big) - Q^{n+1}. \tag{14}$$

3. Solution Algorithm. The system of nonlinear equations (14) can be written generically as $F(u) = 0$ where $F : \mathbb{R}^n \to \mathbb{R}^n$. We solve the system using Newton's method, which requires solution of a linear system at each iteration k:

$$\frac{\partial F}{\partial u}\Big|_{u=u_k} (u_{k+1} - u_k) = -F(u_k). \tag{15}$$

In our case, the solution vector u consists of all the pressure and saturation unknowns at all the cell centers. The Jacobian system resulting from the derivative $\partial F/\partial u$ is often very difficult to solve using iterative methods, and preconditioning is critical for rapid convergence of Krylov subspace methods such as GMRES [25]. In the next section, we will discuss the linear system arising from the Newton's method and give a detailed description of the solution algorithm we will use to solve this system.

3.1. Linear System. For the set of primary variables $u = (P_w, S_n)$, assuming unknowns corresponding to physical variables are grouped together and unknowns associated with nodal points in the domain are ordered lexicographically, each nonlinear Newton iteration entails the solution of a linear system of the form

$$\begin{pmatrix} -\nabla \cdot (\lambda_w K \nabla) & -\dfrac{\phi}{\partial t} - \nabla \cdot (\mathbf{v}_w) \\ -\nabla \cdot (\lambda_n K \nabla) & \dfrac{\phi}{\partial t} + \nabla \cdot (\mathbf{v}_n) + \nabla \cdot (\lambda_n P_c' K \nabla) \end{pmatrix} \begin{pmatrix} \delta P_w \\ \delta S_n \end{pmatrix} = - \begin{pmatrix} q_w \\ q_n \end{pmatrix} \tag{16}$$

in which

$$\mathbf{v}_w = -\lambda_w' K \nabla \tilde{P}_w \tag{17}$$

$$\mathbf{v}_n = -\lambda_n' K \nabla \tilde{P}_n + \lambda_n K \nabla (P_c') \tag{18}$$

All the coefficients in equation (16) are evaluated at the linearization point $\tilde{P}_w, \tilde{S}_n$. In a more concise form, the Jacobian matrix of the system has 2×2 block structure

$$J = \begin{pmatrix} A_{pp} & A_{ps} \\ A_{sp} & A_{ss} \end{pmatrix} \tag{19}$$

The characteristics of this matrix have been discussed in numerous papers [27, 6, 13, 17]. We summarize some main points here:

- J is nonsymmetric and indefinite
- The block A_{pp} has the structure of a purely elliptic problem for pressure.
- The coupling block A_{ps} has the structure of a first-order hyperbolic problem in the non-wetting phase saturation.
- The coupling block A_{sp} has the structure of a convection-free parabolic problem in the wetting phase pressure.
- The block A_{ss} has the structure of a parabolic (convection-diffusion) problem for saturation when capillary pressure is present. Without capillary pressure, it has the form of a hyperbolic problem.
- Under mild conditions, i.e. modest time-step size, all A_{pp}, A_{ps}, A_{ss} blocks are diagonally dominant.

In this paper, we want to present some numerical results that show how different models of capillary pressure affect the algebraic properties of the (2,2)-block A_{ss} in particular and the global system in general, which consequently determines the success of AMG algorithms.

3.2. Decoupling Operators. Designing effective preconditioners for the Jacobian system (19) is challenging because of the strong coupling between pressure and saturation. To weaken the coupling between pressure and saturation unknowns, one can often construct a decoupling operator D, which acts as a matrix scaling to the original Jacobian

$$\tilde{J} \equiv D^{-1} J \equiv \begin{pmatrix} \tilde{A_{pp}} & \tilde{A_{ps}} \\ \tilde{A_{sp}} & \tilde{A_{ss}} \end{pmatrix} \tag{20}$$

A natural choice for D that promotes code reuse is an IMPES type operator [13]

$$D = \begin{pmatrix} D_{pp} & D_{ps} \\ D_{sp} & D_{ss} \end{pmatrix} = \begin{pmatrix} \mathrm{diag}(A_{pp}) & \mathrm{diag}(A_{ps}) \\ \mathrm{diag}(A_{sp}) & \mathrm{diag}(A_{ss}) \end{pmatrix} \tag{21}$$

Other approaches include quasi-IMPES [18] and alternate block factorization [4] strategies. In general, the decoupling operator applied to the Jacobian system results in clustering of the eigenvalues around one, and this operator can act an effective preconditioner. For a more detailed discussion of the effects of this decoupling step, we refer to [17]. Here however, we do not employ any decoupling strategies, in order to examine the robustness of BoomerAMG for fully coupled systems. Moreover, when more complex processes are added to the model, there is no guarantee that a good decoupling strategy exists.

3.3. Algebraic Multigrid. Multigrid is one of the most most efficient and scalable methods available to solve large sparse linear systems [30]. Geometric multigrid uses a hierarchy of nested grids, whose construction depends on the geometry of the problem and *a priori* knowledge of the grids. AMG methods such as those developed in [26] have the advantage of not requiring an explicit hierarchy of nested grids. AMG constructs coarse grids based on the matrix values only, which makes it suitable for solving a wider range of problems on more complicated domains and unstructured grids. Despite its successful application to scalar problems, using AMG for coupled systems is still relatively limited. Some attempts to use AMG to solve fully coupled systems encountered in modeling multiphase flow for reservoir simulation include [10, 27]. In this work, we use BoomerAMG [16], part of the Hypre package [14, 15] as a black-box AMG solver. We note that in order to use BoomerAMG for the coupled system in our case, the Jacobian matrix needs to be ordered by grid points, i.e.

$$J = \begin{pmatrix} A_{11} & \dots & A_{1N} \\ \vdots & \ddots & \dots \\ A_{N1} & \dots & A_{NN} \end{pmatrix} \tag{22}$$

in which N is the number of grid points, and A_{ij} are 2×2 matrices representing the couplings between pressure and saturation at points i and j. This is called the "point" method in [27].

3.4. Two-stage Preconditioning with AMG. Unlike AMG, which has not been popular in reservoir simulation until recently, two-stage preconditioners are widely used [17]. However, the idea of two-stage iterative methods is not new [21], and first appeared in the context of multiphase flow modeling in the work of Wallis [28]. We refer to this method as the constrained pressure residual (CPR) approach. There are many variants of two-stage preconditioners. We discuss two algorithms here: the two-stage combinative preconditioner - CPR-AMG(1), and the two-stage additive preconditioner - CPR-AMG(2) [2].

Algorithm 1. Two-stage Combinative - CPR-AMG(1)

1. At each iteration k let the residual be r_k.
2. Solve $u_{k+1/2} = P_1^{-1} r_k$.
3. Update the residual $r_{k+1/2} = r_k - A u_{k+1/2}$.
4. Solve for the pressure correction $A_{pp}\delta_p = R_p r_{k+1/2}$
5. Update the solution $u_{k+1} = u_{k+1/2} + R_p^T \delta_p$

Algorithm 2. Two-stage Additive - CPR-AMG(2)

1. At each iteration k let the residual be r_k.
2. Solve $u_{k+1/2} = P_1^{-1} r_k$.
3. Update the residual $r_{k+1/2} = r_k - A u_{k+1/2}$.
4. Solve for the pressure correction $A_{pp}\delta_p = R_p r_{k+1/2}$
5. Solve for the saturation correction $A_{ss}\delta_s = R_s r_{k+1/2}$
6. Update the solution $u_{k+1} = u_{k+1/2} + R_p^T \delta_p + R_s^T \delta_s$

The matrices R_p, R_s denote the restriction of the global unknown vector to those associated with pressure and saturation respectively. That is, $R_p \in \mathbb{R}^{n \times 2n}$ and for $u = \begin{pmatrix} p \\ s \end{pmatrix}$

$$R_p u = p, \quad R_p^T u = \begin{pmatrix} p \\ 0 \end{pmatrix}; \qquad R_s u = s \quad R_s^T u = \begin{pmatrix} 0 \\ s \end{pmatrix} \tag{23}$$

Then, in matrix form, the action of the two-stage preconditioners can be expressed as

$$u = M_{comb}^{-1} r = (I - R_p^T A_{pp}^{-1} R_p (A - P_1)) P_1^{-1} r \tag{24}$$

$$u = M_{add}^{-1} r = (I - (R_p^T A_{pp}^{-1} R_p + R_s^T A_{ss}^{-1} R_s)(A - P_1)) P_1^{-1} r \tag{25}$$

For the preconditioner P_1 in step 2 of both algorithms, we use the incomplete factorization with no fill ILU(0) method. For the correction solve, we apply AMG with one V-cycle iteration. The combinative approach with AMG was presented in [18, 20]. However, this method does not work well in the presence of fast changing capillary pressure. We confirm this observation in the next section. To deal with fast changing capillary pressure, we employ an additive CPR-AMG approach, which involves one extra AMG solve for the correction of the saturation block. The intuition is that when the absolute value of the derivative of capillary pressure $|dP_c/dS_w|$ is large, the block A_{ss} becomes diffusion dominated, and AMG can deal with it efficiently.

4. Numerical Results. In this section, we perform numerical experiments for the three aforementioned preconditioners. All of them are implemented in Amanzi, a parallel open-source multi-physics C++ code developed as a part of the ASCEM project [1]. Although Amanzi was first designed for simulation of subsurface flow and reactive transport, its modular framework and concept of process kernels [11] allow new physics to be added relatively easily for other applications. The two-phase flow simulator employed in this work is one such example. Amanzi works on a variety of platforms, from laptops to supercomputers. It also leverages several popular packages for mesh infrastructure and solvers through a unified input file. Here, all of our experiments use a classical AMG solver through BoomerAMG in Hypre. The ILU(0) method is from Euclid, also a part of Hypre. GMRES is provided within Amanzi. We use structured Cartesian grids. This section has three parts. The first case is an oil-water model on a two-dimensional domain that is small but hard to solve. The second part studies a three-dimensional example. In the last part, we examine the scalability of the three preconditioning strategies. Unless specified otherwise, we use the benchmark problem SPE10 [8] for permeability data.

4.1. Two-dimensional oil-water problem. The domain is a rectangle of dimensions 762×15.24 meters. The mesh is 100×20, which means that the problem is truly two-dimensional in the xz plane. We

inject pure water into the domain through the boundary at the lower left corner, and oil and water exit the domain through the top right corner. These correspond to the $S_w = 1.0,\ \lambda_w \nabla P_w \cdot \mathbf{n} = -50\ \ m^3/day$ at the inlet, and $S_w = 0.2,\ P_w = 0$ at the outlet. The simulation is run for 200 days with time step $\Delta t = 20$ days.

For capillary pressure models, we employ a simple linear model and the Brooks-Corey [7] model:

$$\text{Linear model: } P_c(S_w) = P_0(1 - \bar{S}_w), \qquad \text{Brooks-Corey: } P_c(S_w) = P_d \bar{S}_w^{-1/\lambda} \tag{26}$$

in which $\bar{S}_w$ is the effective saturation, P_d is the entry pressure, and λ is related to the pore-size distribution. For the Brooks-Corey model, the typical range of λ is $0.2, 3.0$ [5, 12]. where $\lambda > 2$ indicates narrow distributions of pore sizes and $\lambda < 2$ wide distributions. For example, sandpacks with broader distributions of particle sizes have λ ranging from 1.8 to 3.7 [23]. The Brooks-Corey capillary pressure curves for various values of λ are plotted in Figure 1. Other parameters are listed in Table 1 and example 1 of Table 2.

Fig. 1: Capillary pressure curves for Brooks-Corey model with entry pressure $P_d = 10^5\ \ Pa$.

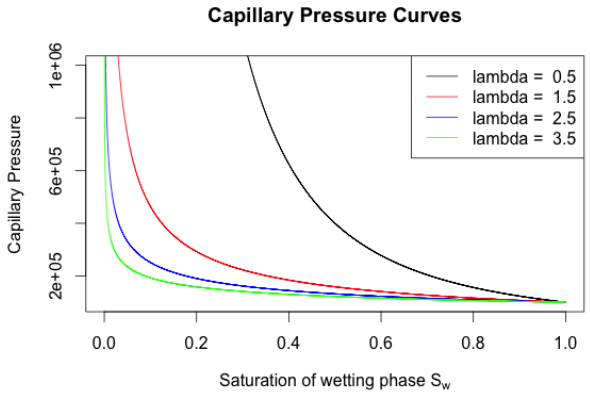

For all of the simulations presented here, the convergence tolerance for Newton's method is $||F(x)|| \leq 10^{-6}$, and the linear tolerance for GMRES is $||J\delta u_k - F(u_k)|| \leq 10^{-12}||F(u_k)||$. BoomerAMG is used as a preconditioner. The number of V-cycle steps is set to 1. The coarsening strategy is the parallel Cleary-Luby-Jones-Plassman (CLJP) coarsening [9]. The interpolation method is the classical interpolation defined in [24], and the smoother is the forward hybrid Gauss-Seidel / SOR scheme.

Table 1: Input data for the quarter-five spot problem.

Initial wetting phase pressure	$10^5\ Pa$
Initial nonwetting phase saturation	0.8
Residual wetting phase saturation	0.0
Nonwetting phase density	700 kg/m^3
Wetting phase density	1000 kg/m^3
Nonwetting phase viscosity	10.0 cP
Wetting phase viscosity	1 cP
Porosity	0.2

Table 2: Parameters for capillary pressure models

Parameters	Ex 1	Ex 2	Ex 3	Ex 4
Linear entry pressure P_0	10^5	10^4	10^3	10^6
Brooks-Corey entry pressure P_d	10^6	10^5	2×10^4	10^6
Brooks-Corey λ	2.5	0.8	2.5	0.8

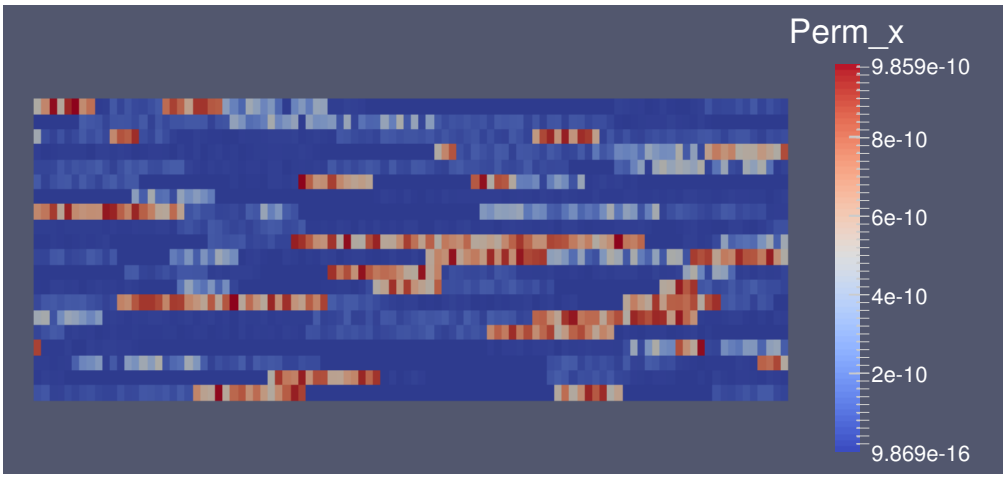

Fig. 2: Permeability field obtained from SPE10 model 1 data.
The x-direction is scaled down by 1/20 for visualization.

In order to explore the effects of different models for capillary pressure on solver performance, we use the four sets of parameters listed in Table 2. In Example 1, the parameters are chosen such that the L_∞ norm of the derivative of capillary pressure P_c' is large, leading to a diffusion-dominated case (see equation (16)). In Example 2, the parameters are tuned to reduce the L_∞ norm of P_c', leading to an advection-dominated case. Example 3 is a more extreme case of example 2, in which P_c' is further decreased, leading to a strongly advection-dominated case. We also note the difference between the linear model and the Brooks-Corey model for capillary pressure. The derivative P_c' for the linear model is a constant value, which means that the character of the problem, i.e. diffusion-dominated or advection-dominated, is the same everywhere for the whole domain. In the Brooks-Corey model, P_c' depends on the saturation of the wetting phase, and the problem can be diffusion-dominated in one part of the domain, while advection-dominated in another part. This can cause further difficulties for AMG based solvers, whose optimal performance is sensitive to the characteristics of the problem.

The performance of the three strategies are given in Tables 3, 4, and 5. NI denotes the number of nonlinear iterations, LI the number of linear iterations, LI/NI the average number of linear iterations per nonlinear iterations, and Time the total time of the whole simulation in seconds.

Table 3: Performance of three preconditioning strategies for set of parameters in example 1 of Table 2

Methods/Models	Linear				Brooks Corey			
	NI	LI	LI/NI	Time	NI	LI	LI/NI	Time
AMG	33	771	23.3	67.6	51	1534	30.1	200.5
CPR-AMG(1)	33	3370	102.1	428.8	51	8786	172.3	1713.6
CPR-AMG(2)	33	1175	35.6	171.9	51	2195	43.0	467.5

From Table 3, it is clear that AMG is the best method for both capillary pressure models in terms of both the number of linear iterations per Newton step and the total run time in the diffusion-dominated case. For the linear model, AMG takes about 4 times fewer linear iterations than CPR-AMG(1), and it is more than 6 times faster in total run time. The reason for this discrepancy is that CPR-AMG(1) is a two-stage preconditioner, and it requires an extra global solve using ILU. Similarly, CPR-AMG(2) needs one more AMG solve than CPR-AMG(1), and thus the run time per linear iteration of CPR-AMG(2) is higher than that of CPR-AMG(1). However, CPR-AMG(2) still outperforms CPR-AMG(1) for both capillary pressure models in terms of both the number of linear iterations per Newton step and the total run time. The same conclusion can be made for the Brooks-Corey model.

Table 4: Performance of three preconditioning strategies for set of parameters in example 2 of Table 2

Methods/Models	Linear				Brooks Corey			
	NI	LI	LI/NI	Time	NI	LI	LI/NI	Time
AMG	37	4665	126.1	332.7	65	4082	62.5	523.4
CPR-AMG(1)	37	1916	51.8	293.1	65	5462	84.0	3788.5
CPR-AMG(2)	37	1212	32.8	207.4	65	3462	53.3	877.6

Table 5: Performance of three preconditioning strategies for set of parameters in example 3 of Table 2

Methods/Models	Linear				Brooks Corey			
	NI	LI	LI/NI	Time	NI	LI	LI/NI	Time
AMG	-	-	-	-	-	-	-	-
CPR-AMG(1)	49	1474	30.1	312.1	61	2683	44.0	633.9
CPR-AMG(2)	49	3091	63.1	602.5	61	8492	139.2	1929.7

The results for the advection-dominated case reported in Table 4 show that AMG applied to the coupled system does not perform well, especially for the linear model of capillary pressure. In this case, AMG is the worst method in terms of both the number of linear iterations per Newton step and total run time. CPR-AMG(2) is the best method for the same performance measures. For the Brooks-Corey model, although CPR-AMG(2) takes the smallest number of linear iterations per Newton step, its total run time is slower than that of AMG. Again, this is due to the fact that CPR-AMG(2) needs one global ILU solve and two AMG solves for the A_{pp} and A_{ss} blocks.

For the strongly advection-dominated problem with parameters in example 3, AMG diverges for both linear and Brooks-Corey capillary pressure models. The performance of CPR-AMG(2) is also affected in this case, and CPR-AMG(1) is the most efficient method. This seems to suggest that using AMG to solve for the correction for the saturation block A_{ss} does not work well when the operator associated with the A_{ss} block is strongly advection-dominated or near hyperbolic. However, the two-stage preconditioner CPR-AMG(2) is still more robust than direct AMG, since unlike AMG, this method still converges.

4.2. 3D Case. We use a homogeneous permeability field of 100 millidarcy, but the grid is stretched to induce anisotropy. The model dimensions are $250 \times 1000 \times 60$ meters and the cell size is $5 \times 10 \times 0.5$ meter. Thus, the mesh is $50 \times 100 \times 120$, and the problem has 1.2 million unknowns in total. Water is injected into the domain at one bottom corner and the outlet is at the opposite corner. The injection rate is 5 m^3/day. The parameters for the capillary pressure model is from example 1 of Table 2. The simulation is run for 100 days with time step $\Delta t = 20$ days.

Table 6: Performance in the 3D case for the set of parameters in example 1 of Table 2

Methods/Models	Linear				Brooks Corey			
	NI	LI	LI/NI	Time	NI	LI	LI/NI	Time
AMG	16	282	17.63	103.1	20	452	22.6	144.7
CPR-AMG(1)	16	2698	168.63	803.2	20	6069	303.45	1940.8
CPR-AMG(2)	16	712	35.6	299.5	20	1900	95.0	741.1

Table 6 shows the performance results of the diffusion-dominated case for this 3D example, which are consistent with those of the previous 2D example. AMG is still the best method in terms of both the linear iteration counts per Newton step and the time it takes to complete the simulation for both capillary pressure models. CPR-AMG(2) does not perform quite as well as AMG, but it is much more efficient than CPR-AMG(1) for both performance measures and capillary pressure models.

4.3. Scaling Results. To perform a scalability study, we run a test problem on a box of dimensions $20 \times 20 \times 20$ meters. The initial mesh is $20 \times 20 \times 20$ and is repeatedly refined in the z-direction. The domain

has constant material properties. The parameters for the water retention models are listed in example 4 of Table 2. Note that this set of parameters corresponds to a diffusion-dominated problem.

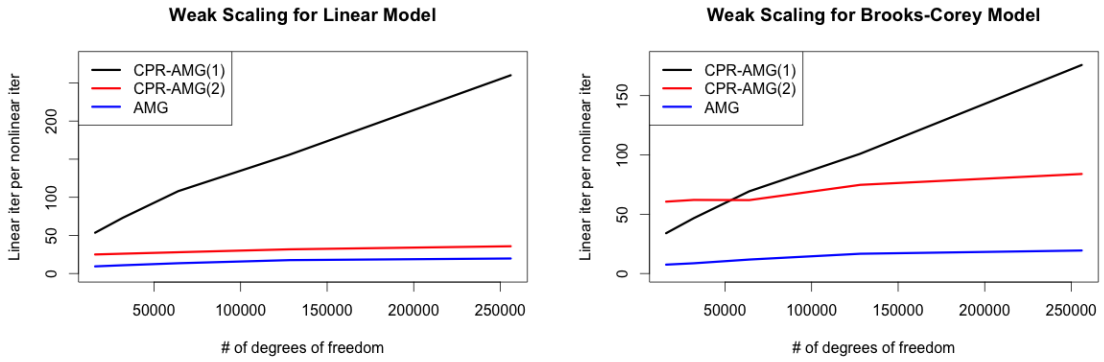

Fig. 3: Weak scaling for different strategies.

The results shown in Figure 3. indicate that the performance of CPR-AMG(2) and AMG methods is independent of the mesh size. The number of linear iterations per Newton step does not grow as the mesh is refined which is optimal multigrid performance. The average number of linear iterations per Newton step for CPR-AMG(2) is quite high compared to AMG in the Brooks-Corey case. CPR-AMG(1), however, does not scale as well as the other two methods. The linear iteration counts for CPR-AMG(1) grows linearly as the mesh is refined.

5. Conclusions. In this work, we have implemented a fully implicit parallel two-phase flow simulator along with three different preconditioning strategies to solve the linear systems resulting from linearization of the coupled equations. Numerical experiments demonstrate that GMRES with AMG preconditioning for the coupled systems works best in the diffusion-dominated case in both two-dimensional and three-dimensional examples. However, this method exhibits slow convergence and sometimes diverges for the advection-dominated or near hyperbolic case. The combinative CPR-AMG(1) method is the most robust, but it does not scale optimally with problem size. It is also slower than other methods in most cases, with the exception of the near hyperbolic case when AMG diverges. The additive CPR-AMG(2) method performs well in most cases except the near hyperbolic case. It is the best method for our advection-dominated two-dimensional example. It also scales optimally with problem size for both advection-dominated and diffusion-dominated case.

REFERENCES

[1] ASCEM. http://esd.lbl.gov/research/projects/ascem/thrusts/hpc/, 2009.
[2] O. Axelsson. *Iterative Solution Methods*. Cambridge University Press, 1994.
[3] K. Aziz and A. Settari. *Petroleum Reservoir Simulation*. Applied Science Publishers, 1979.
[4] R.E. Bank, T.F. Chan, W.M. Coughran Jr., and R.K. Smith. The alternate-block-factorization procedure for systems of partial differential equations. *BIT Numerical Mathematics*, 29(4):938–954, 1989.
[5] P. Bastian. Numerical computation of multiphase flow in porous media. Habilitationsschrift, 1999.
[6] P. Bastian and R. Helmig. Efficient fully-coupled solution techniques for two-phase flow in porous media: Parallel multigrid solution and large scale computations. *Advances in Water Resources*, 23(3):199 – 216, 1999.
[7] R. Brooks and A. Corey. *Hydraulic Properties of Porous Media*. Hydrology Papers, Colorado State Univeristy, 1964.
[8] M.A. Christie and M.J. Blunt. Tenth SPE comparative solution project: A comparison of upscaling techniques. In *SPE Reservoir Simulation Symposium*. Society of Petroleum Engineers (SPE), 2001.
[9] A.J. Cleary, R.D. Falgout, V.E. Henson, and J.E. Jones. Coarse-grid selection for parallel algebraic multigrid. In Alfonso Ferreira, Jos Rolim, Horst Simon, and Shang-Hua Teng, editors, *Solving Irregularly Structured Problems in Parallel*, volume 1457 of *Lecture Notes in Computer Science*, pages 104–115. Springer Berlin Heidelberg, 1998.
[10] T. Clees and L. Ganzer. An efficient algebraic multigrid solver strategy for adaptive implicit methods in oil-reservoir simulation. *SPE Journal*, 15(03):670–681, 2010.
[11] E. Coon, J. D. Moulton, and S. Painter. Managing complexity in simulations of land surface and near-surface processes. Technical Report LA-UR 14-25386, Applied Mathematics and Plasma Physics Group, Los Alamos National Laboratory, 2014.
[12] A. T. Corey. *Mechanics of Immiscible Fluids in Porous Media*. Water Resources Pubns, 1995.
[13] C.N. Dawson, H. Klie, M.F. Wheeler, and C.S. Woodward. A parallel, implicit, cell-centered method for two-phase flow with a preconditioned NewtonKrylov solver. *Computational Geosciences*, 1(3-4):215–249, 1997.
[14] R.D. Falgout, J.E. Jones, and U.M. Yang. The design and implementation of hypre, a library of parallel high performance preconditioners. In AreMagnus Bruaset and Aslak Tveito, editors, *Numerical Solution of Partial Differential Equations on Parallel Computers*, volume 51 of *Lecture Notes in Computational Science and Engineering*, pages 267–294. Springer Berlin Heidelberg, 2006.
[15] R.D. Falgout and U.M. Yang. HYPRE: a library of high performance preconditioners. In *Preconditioners, Lecture Notes in Computer Science*, pages 632–641, 2002.
[16] V.E. Henson and U.M. Yang. BoomerAMG: a parallel algebraic multigrid solver and preconditioner. *Applied Numerical Mathematics*, 41:155–177, 2000.
[17] H. Klie, M. Rame, and M.F. Wheeler. Two-stage preconditioners for inexact Newton methods in multiphase reservoir simulation. Technical report, Center for Research on Parallel Computation, Rice University, 1996.
[18] S. Lacroix, Yu. Vassilevski, J. Wheeler, and M.F. Wheeler. Iterative solution methods for modeling multiphase flow in porous media fully implicitly. *SIAM J. Sci. Comput.*, 25(3):905–926, jan 2003.
[19] B. Lu. *Iteratively Coupled Reservoir Simulation for Multiphase Flow in Porous Media*. PhD thesis, Center for Subsurface Modeling, ICES, The University of Texas - Austin, 2008.
[20] R. Masson, R. Scheichl, P. Quandalle, C. Guichard, R. Eymard, R. Herbin, and K. Brenner et al. CPR-AMG preconditioner for multiphase porous media flows: a few experiments, 2012.
[21] N.K. Nichols. On the convergence of two-stage iterative processes for solving linear equations. *SIAM J. Numer. Anal.*, 10(3):460–469, 1973.
[22] D.W. Peaceman. *Fundamentals of Numerical Reservoir Simulation*. Elsevier, Amsterdam, 1977.
[23] PetroWiki. Relative permeability models. http://petrowiki.org/Relative_permeability_models.
[24] J. W. Ruge and K. Stben. *4. Algebraic Multigrid*, chapter 4, pages 73–130. 1986.
[25] Y. Saad and M.H. Schultz. Gmres: A generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM J. Sci. Stat. Comput.*, 7(3):856–869, July 1986.
[26] K. Stueben. A review of algebraic multigrid. *Journal of Computational and Applied Mathematics*, 128(12):281 – 309, 2001. Numerical Analysis 2000. Vol. VII: Partial Differential Equations.
[27] K. Stueben, T. Clees, H. Klie, B. Lu, and M.F. Wheeler. Algebraic multigrid methods (AMG) for the efficient solution of fully implicit formulations in reservoir simulation. In *SPE Reservoir Simulation Symposium*. Society of Petroleum Engineers (SPE), 2007.
[28] J.R. Wallis, R.P. Kendall, and L.E. Little. Constrained residual acceleration of conjugate residual methods. 1985.
[29] J.W. Watts. A compositional formulation of the pressure and saturation equations. In *SPE Reservoir Simulation Symposium*. Society of Petroleum Engineers (SPE), SPE 12244, 1985.
[30] U.M. Yang. Parallel algebraic multigrid methods high performance preconditioners. In *Numerical Solution of Partial Differential Equations on Parallel Computers*, volume 51 of *Lecture Notes in Computational Science and Engineering*, pages 209–236. Springer Berlin Heidelberg, 2006.