SIAM J. SCI. COMPUT.
Vol. 39, No. 5, pp. S662–S680

ALGEBRAIC MULTIGRID PRECONDITIONERS FOR MULTIPHASE FLOW IN POROUS MEDIA*

QUAN M. BUI†, HOWARD C. ELMAN‡, AND DAVID J. MOULTON§

Abstract. Multiphase flow is a critical process in a wide range of applications, including carbon sequestration, contaminant remediation, and groundwater management. Typically, this process 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 constitutive laws for the relationship of phase velocity to phase pressure. In this study, we develop and study efficient solution algorithms for solving the algebraic systems of equations derived from a fully coupled and time-implicit treatment of models of multiphase flow. We explore the performance of several preconditioners based on algebraic multigrid (AMG) for solving the linearized problem, including "black-box" AMG applied directly to the system, a new version of constrained pressure residual multigrid (CPR-AMG) preconditioning, and a new preconditioner derived using an approximate Schur complement arising from the block factorization of the Jacobian. We show that the new methods are the most robust with respect to problem character, as determined by varying effects of capillary pressures, and we show that the block factorization preconditioner both is efficient and scales optimally with problem size.

Key words. algebraic multigrid, preconditioning, multiphase flow, porous media, Krylov methods

AMS subject classifications. 65F08, 35Q35

DOI. 10.1137/16M1082652

1. Introduction.

Multiphase flow is a feature of many physical systems, and models of it are used in applications such as reservoir simulation, carbon sequestration, ground water management, and contaminant transport. Modeling 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 (PDEs), 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 a 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. This additional complexity could include adding more components, miscibility between components, thermal effects, and phase transitions.

*Received by the editors July 5, 2016; accepted for publication (in revised form) May 3, 2017; published electronically October 26, 2017.

http://www.siam.org/journals/sisc/39-5/M108265.html

Funding: The work of the first and second authors was supported by the U.S. Department of Energy under grant DE-SC0009301 and by the U.S. National Science Foundation under grant DMS1418754. The work of the first author was also supported by the Department of Energy at Los Alamos National Laboratory under contract DE-AC52-06NA25396 through the DOE Advanced Simulation Capability for Environmental Management (ASCEM) program and the LANL Summer Student Program at the Center for Nonlinear Studies (CNLS).

†Applied Math, Stats, and Scientific Computation, University of Maryland, College Park, MD 20742 (mquanbui@math.umd.edu).

‡Department of Computer Science, University of Maryland, College Park, MD 20742 (elman@cs.umd.edu).

§Applied Math and Plasma Physics Group, T-5, Los Alamos National Laboratory, Los Alamos, NM 87544 (moulton@lanl.gov).

The fully implicit discretization gives rise to a nonlinear system of equations at each time step. We employ a variant of Newton's method with an exact Jacobian of the discretized equations to solve this system. For the linear system, we use a preconditioned generalized minimal residual (GMRES) method [30]. There is a vast literature on different approaches to preconditioning the Jacobian system. A very popular approach is to use incomplete LU factorization (ILU) for constructing the preconditioner. Though popular for its general applicability, ILU-based preconditioners are neither effective nor scalable in many cases. Another approach is to consider decoupled preconditioners for the coupled system [6]. This methodology is based on a direct solution of the decoupled pressure system, followed by an iterative solution using ILU for the global system. This formulation was refined in [34], where it was proposed to solve the pressure system iteratively, giving rise to the decoupled implicit pressure explicit saturation (IMPES) preconditioner. The effect of the decoupling is to weaken the coupling between pressure and saturation. Thus, it is often used as a preprocessing step to produce a modified Jacobian system, for which new preconditioners can be developed [10, 32]. Another approach to breaking up the coupled problem into a sequence of simpler problems includes operator splitting techniques, developed in [14, 15, 26]. With recent development of algebraic multigrid (AMG) algorithms, the pressure block can be solved efficiently, resulting in the constrained pressure residual multigrid (CPR-AMG) approach. Recently, AMG has also been applied to solving the coupled system with some success [10, 32], although developing a general AMG algorithm for these types of problems remains a topic of ongoing research [35]. Since the Jacobian matrix has a block structure, one can also consider a block LU decomposition with an approximate Schur complement, which has been successfully applied to other models of fluid dynamics [24, 37]. Besides AMG-based methods, geometric multigrid has also been applied successfully to solve these types of problems [4, 5]. Our focus in this study is on methods based on AMG because of its general applicability.

In this paper, we develop a new block preconditioner designed to respect the coupling inherent in models of multiphase flow, and we report our experience with the performance and scalability of four different preconditioning strategies: (1) a direct AMG preconditioner for the global system; (2) a two-stage CPR-AMG method with correction for the pressure block, also known as the combinative two-stage approach; (3) CPR-AMG with corrections for both the pressure and saturation blocks, known as the two-stage additive approach; and (4) the block factorization (BF) preconditioner. An outline of the paper is as follows. In section 2, we present the mathematical formulation for two-phase flow in porous media and discretization schemes. In section 3, we describe the solution algorithms for the linearized system. Numerical results for the algorithms are presented in section 4. We conclude with some remarks and discussion of future research directions in section 5.

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.

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 nonwetting phases, respectively, and ϕ is the porosity of the medium. We assume a common extension of Darcy's law to multiphase flow and express the phase 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 $k_{r\alpha}$, μ_α, P_α are the relative permeability, viscosity, and pressure, respectively, of phase α; g is the gravitational constant; and D is the depth. Note that the relative permeability terms $k_{r\alpha}(s_\alpha)$ are nonlinear functions of the saturation. 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 (1) and (2), one can derive separate equations for pressure and saturation. The pressure equation is elliptic due to incompressibility; the saturation equation is of convection-diffusion type. Depending on the applications and capillary pressure models, the saturation equation can be diffusion-dominated, convection-dominated, or even purely hyperbolic (in the absence of capillary pressure). 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 [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 [36].

The appeal of these methods lies in the sequential 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 or thermal and phase transitions are included in the model. For a more complete summary of the advantages and disadvantages of these approaches, we refer the reader to [23].

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, $\mathbf{u} = (P_w, S_n)$ [38], we obtain

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

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

This formulation has the advantage that extending it to the case of compressible flow and multicomponent flow is quite straightforward. (See also [22, 18] for use of this

model.) In this paper, we consider solving the coupled system consisting of (6) and (7) 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 [25, 28]. For simplicity, we consider a uniform partitioning of the domain $\mathbf{\Omega}$ into equal sized cells C_i, i.e., $\mathbf{\Omega} = \bigcup_{i=1} C_i$. Let γ_{ij} denote the area of the face between cells C_i and C_j. For each cell C_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 midpoint rule, which is second-order accurate:

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

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

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

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

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.,

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

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

$$\mathbf{K}_{ij+1/2} = (\Delta x_i + \Delta x_j)\left(\frac{\mathbf{K}_i\mathbf{K}_j}{\Delta x_i\mathbf{K}_j + \Delta x_j\mathbf{K}_i}\right). \tag{13}$$

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

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

3. Solution algorithms. 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:

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

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. Next, we discuss the linear system arising from the Newton's method and give a detailed description of the solution algorithms we will use to solve this system.

3.1. Linear system. For the set of primary variables $u = (P_w, S_n)$, assuming that each physical variable is ordered lexicographically, then each nonlinear Newton iteration entails the solution of a discrete version of a block 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 (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,

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

and the linear system is $Jc = q$. The characteristics of the matrix have been discussed in numerous papers [5, 13, 20, 32]. We summarize important characteristics of the operators here:

- J is nonsymmetric and indefinite.
- The block A_{pp} has the structure of a discrete purely elliptic problem for pressure.
- The coupling block A_{ps} has the structure of a discrete first-order hyperbolic problem in the nonwetting phase saturation.
- The coupling block A_{sp} has the structure of a discrete convection-free parabolic problem in the wetting phase pressure.
- The block A_{ss} has the structure of a discrete parabolic (convection-diffusion) problem for saturation when capillary pressure is a nonconstant function of the saturation. When capillary pressure is zero or a constant, there is no diffusion term, and the block has the form of a hyperbolic problem.
- Under mild conditions, i.e., modest time step size, the blocks A_{pp}, A_{ps}, A_{ss} are diagonally dominant.

In this paper, we 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 solution algorithms. Our emphasis is on the development and use of preconditioning operators denoted $M \approx J$, for the purpose of solving preconditioned systems

$$JM^{-1}\hat{c} = q, \quad c = M^{-1}\hat{c}. \tag{20}$$

3.2. Algebraic multigrid. Multigrid is a highly efficient and scalable method available for solving large sparse linear systems [33, 39]. 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 [31] 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 wide range of problems on complicated domains and unstructured grids. Despite its successful application to scalar problems, use of 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, 32]. In this work, we use BoomerAMG [19], part of the Hypre package [16, 17], as a black-box AMG solver. We note that in our implementation of the coupled system, the Jacobian matrix passed to BoomerAMG is ordered by grid points, i.e.,

$$J = \begin{pmatrix} A_{11} & \dots & A_{1N} \\ \vdots & \ddots & \dots \\ A_{N1} & \dots & A_{NN} \end{pmatrix}, \tag{21}$$

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 [32].

3.3. Two-stage preconditioning with AMG. Unlike AMG, which has not been popular in reservoir simulation until recently, two-stage preconditioners are widely used [20]. This idea was developed and first appeared in the context of multiphase flow modeling in the work of Wallis, Kendall, and Little [34]. Following [13], 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 $\delta u_{k+1/2} = P_1^{-1} r_k$, compute intermediate solution $u_{k+1/2} = u_k + \delta u_{k+1/2}$.
3. Update the residual $r_{k+1/2} = r_k - A\delta u_{k+1/2}$.
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 $\delta u_{k+1/2} = P_1^{-1} r_k$, compute intermediate solution $u_{k+1/2} = u_k + \delta u_{k+1/2}$.
3. Update the residual $r_{k+1/2} = r_k - A\delta u_{k+1/2}$.
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 the spaces associated with pressure and saturation, respectively. That is, $R_p \in \mathbb{R}^{n \times 2n}$,

and for $u = \binom{p}{s}$

$$R_p u = p, \quad R_p^T u = \begin{pmatrix} p \\ 0 \end{pmatrix}, \qquad R_s u = s, \quad R_s^T u = \begin{pmatrix} 0 \\ s \end{pmatrix}. \tag{22}$$

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

$$\delta u = M_{comb}^{-1} r = (I - R_p^T A_{pp}^{-1} R_p (A - P_1)) P_1^{-1} r, \tag{23}$$

$$\delta u = M_{add}^{-1} r = (I - (R_p^T A_{pp}^{-1} R_p + R_s^T A_{ss}^{-1} R_s)(A - P_1)) P_1^{-1} r. \tag{24}$$

The preconditioner P_1 in step 2 of both algorithms is taken to be the ILU(0) factorization of A, i.e., the incomplete factorization with no fill applied to the global matrix. For the correction solve, we use AMG with one V-cycle iteration. The combinative approach with AMG was presented in [21]. 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 more 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 handle it efficiently.

3.4. Block factorization preconditioners. Consider the following decomposition of the Jacobian:

$$J = \begin{pmatrix} A_{pp} & A_{ps} \\ A_{sp} & A_{ss} \end{pmatrix} = \begin{pmatrix} I & A_{ps} A_{ss}^{-1} \\ 0 & I \end{pmatrix} \begin{pmatrix} S & 0 \\ 0 & A_{ss} \end{pmatrix} \begin{pmatrix} I & 0 \\ A_{ss}^{-1} A_{sp} & I \end{pmatrix},$$

where S is the Schur complement,

$$S = A_{pp} - A_{ps} A_{ss}^{-1} A_{sp}. \tag{25}$$

We could choose

$$M = \begin{pmatrix} I & A_{ps} A_{ss}^{-1} \\ 0 & I \end{pmatrix} \begin{pmatrix} S & 0 \\ 0 & A_{ss} \end{pmatrix} = \begin{pmatrix} S & A_{ps} \\ 0 & A_{ss} \end{pmatrix} \tag{26}$$

as an upper-triangular block preconditioner; this incorporates the effects of the coupling block A_{ps}. This block is important as it contains the time derivative and gravity terms (16). If the time step is small, then the coefficients on the diagonal of this block become large, and it is important that this term be included in the preconditioner. We use an approximation of the Schur complement in which A_{ss} is replaced by its diagonal values:

$$\tilde{S} = A_{pp} - A_{ps} \operatorname{diag}(A_{ss})^{-1} A_{sp}. \tag{27}$$

The purpose of this is to keep the Schur complement sparse so that the action of its inverse can be applied efficiently. This idea is the basis of the SIMPLE method used in other models of fluid dynamics [24]. A similar approach has also been applied to problems in single-phase flow coupled with geomechanics in [37].

ALGORITHM 3. BF preconditioner.

1. At each iteration k let the residual be r_k.
2. Solve for the saturation $A_{ss} s_{k+1} = R_s r_k$ using AMG.
3. Compute the residual for pressure $r = R_p r_k - A_{ps} s_{k+1}$.

4. Solve for the pressure $\tilde{S}p_{k+1} = r$ using AMG.

An important advantage of this algorithm is that it does not rely on an ILU factorization. In matrix form,

$$M_{bf}^{-1} = \begin{pmatrix} \tilde{S}^{-1} & -\tilde{S}^{-1}A_{ps}A_{ss}^{-1} \\ 0 & A_{ss}^{-1} \end{pmatrix}. \tag{28}$$

4. Numerical results. In this section, we perform numerical experiments for the four aforementioned preconditioners. All of them are implemented in Amanzi, a parallel open-source multiphysics 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. ILU(0) is used sequentially for the two-dimensional examples, and parallel ILU(0) (also from Euclid) is used for the three-dimensional cases. GMRES is provided within Amanzi. For simplicity, we employ structured Cartesian grids for the test cases, but we can also use unstructured K-orthogonal grids. The test cases are run on Edison at the National Energy Research Scientific Computing Center (NERSC).[1] We run the two-dimensional test cases in serial, and the three-dimensional SPE10 problem using 256 cores. Amanzi and other libraries are compiled with OpenMPI 1.6.5 and gcc-4.9.2. The total time is measured in seconds. This section has three parts. In the first part, we show the results for a two-dimensional oil-water model problem. Although the problem is small, it is difficult to solve due to the heterogeneity of the permeability field. In the second part, we report the results for 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 and porosity.

4.1. Two-dimensional oil-water problem. The domain is a rectangle with dimensions 762×15.24 meters. The mesh is 100×20, which means that the problem is truly two-dimensional in the xz plane. The absolute permeability field is shown in Figure 1. 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:

$$\text{linear model: } P_c(S_w) = P_0(1 - \bar{S}_w), \qquad \text{Brooks–Corey: } P_c(S_w) = P_d \bar{S}_w^{-1/\lambda}, \tag{29}$$

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 $[0.2, 3.0]$ (see [4, 12]). In general, λ is greater than 2 for narrow distributions of

[1] A DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under contract DE-AC02-05CH11231.

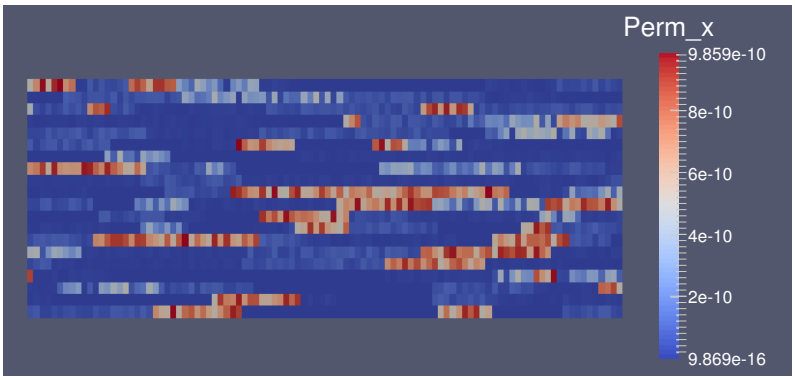

FIG. 1. *Permeability field obtained from SPE*10 *model* 1 *data. The x-direction is scaled down by* 1/20 *for visualization.*

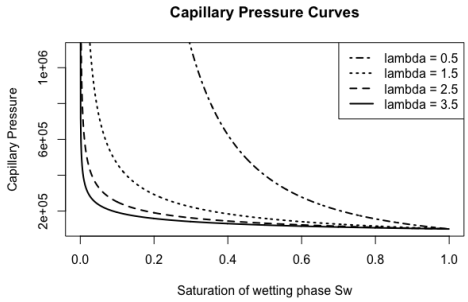

FIG. 2. *Capillary pressure curves for Brooks–Corey model with entry pressure* $P_d = 10^5\ Pa$.

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

pore sizes, and λ is less than 2 for wide distributions. For example, sandpacks with broader distributions of particle sizes have λ ranging from 1.8 to 3.7 [7]. The Brooks–Corey capillary pressure curves for various values of λ are plotted in Figure 2. Other parameters are listed in Table 1 and Example 1 (Ex 1) of Table 2.

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 $||J\delta u_k - F(u_k)|| \leq 10^{-12}||F(u_k)||$, which is the default in Amanzi. 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 [27], and the smoother is the forward hybrid

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

TABLE 3
Performance of three preconditioning strategies for the set of parameters in the diffusion-dominated Example 1.

Methods/Models	Linear				Brooks–Corey			
	NI	LI	LI/NI	Time	NI	LI	LI/NI	Time
AMG	32	368	11.5	27.2	36	470	13.1	37.24
CPR-AMG(1)	32	3695	115.5	324.15	36	5831	162	567.7
CPR-AMG(2)	32	899	28.1	103.94	36	1102	30.6	134.6
BF	32	524	16.4	33.17	36	599	16.6	46.2

Gauss–Seidel/successive over-relaxation scheme.

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 (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. Example 4 represents another diffusion-dominated case, and it is only used in the scaling test in section 4.5. 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 and 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 is summarized 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 iteration, and Time the total time in seconds of the whole simulation. For the diffusion-dominated problem for which the results are shown in Table 3, AMG is the most efficient method, about 25% more efficient than the block preconditioner in terms of both iteration counts (linear iterations per Newton step) and total run time. Note that, in this example, the diffusion term in the (2,2)-block (A_{ss}) is large, and the block is close to a scalar elliptic problem. Hence, it is not surprising that AMG performs well in this case. For the linear model, the BF approach still takes about 8 times fewer linear iterations, and it is about 10 times faster in total run time than CPR-AMG(1). The reason for this discrepancy is that CPR-AMG(1) is a two-stage preconditioner, and it requires an extra global solve using ILU. The BF preconditioner does not rely on ILU, which helps improve the run time significantly. CPR-AMG(2) also performs well in this case; although it requires one more AMG solve per Newton iteration than CPR-AMG(1),

TABLE 4
Performance of three preconditioning strategies for the set of parameters in the advection-dominated Example 2.

Methods/Models	Linear				Brooks–Corey			
	NI	LI	LI/NI	Time	NI	LI	LI/NI	Time
AMG	37	2575	69.6	138.8	-	-	-	-
CPR-AMG(1)	37	1919	51.9	175.5	55	4851	88.2	605.7
CPR-AMG(2)	37	1222	33.0	157.1	55	3701	67.3	506.8
BF	37	684	18.5	51.7	55	1633	29.7	131.1

TABLE 5
Performance of three preconditioning strategies for the set of parameters in the strongly advection-dominated Example 3.

Methods/Models	Linear				Brooks–Corey			
	NI	LI	LI/NI	Time	NI	LI	LI/NI	Time
AMG	-	-	-	-	-	-	-	-
CPR-AMG(1)	43	1079	25.1	122.8	48	2173	45.3	247.6
CPR-AMG(2)	43	1442	35.5	169.8	48	4805	100.1	560.5
BF	43	1002	23.3	69.8	48	1829	38.1	121.8

it still outperforms CPR-AMG(1) 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.

The results reported in Table 4 reveal the lack of robustness of AMG when applied to the coupled system. In contrast to the diffusion-dominated case, for the linear model of capillary pressure, AMG requires the greatest number of linear iterations per Newton step for the advection-dominated case, and it even diverges for the Brooks–Corey model. The BF preconditioner still shows good performance, taking about half the number of iterations and running four times faster than the next best method, which is CPR-AMG(2). CPR-AMG(1) is still the least effective method in this case for both capillary pressure models.

For the strongly advection-dominated problem with parameters in Example 3, AMG diverges for both the linear and Brooks–Corey capillary pressure models, as shown in Table 5. The performance of CPR-AMG(2) is also affected in this case, trailing that of CPR-AMG(1). CPR-AMG(2) is still more robust than direct application of AMG, however, since unlike AMG, this method still converges. The BF preconditioner is again the most effective method, requiring fewer iterations and about half the run time of CPR-AMG(1). This suggests that when the diffusion term in the A_{ss} block gets small, the coupling block A_{ps}, which has the structure of a discrete first-order hyperbolic problem for the saturation, becomes important and needs to be taken into account. The BF method does exactly this. Moreover, it takes advantage of the effectiveness of AMG for scalar problems. Recall that, in the BF approach, AMG is applied to the approximate Schur block $\tilde{S}$ and to the block A_{ss}. $\tilde{S}$ has the form of a perturbed elliptic problem, and therefore we believe it is similar in character to the original pressure block A_{pp}. Thus, AMG is a natural choice for approximating the action of the inverse of $\tilde{S}$. Similarly, the block A_{ss} is a discrete version of a convection-diffusion problem, for which AMG should work well.

4.2. Two-dimensional problem with gravity. In this example, we compare the performance of the different strategies for a problem in which gravity plays a

TABLE 6
Iteration counts for the diffusion-dominated case with gravity, time step $dt = 10$ days.

Methods/Mesh sizes	20^2	40^2	80^2	160^2
AMG	7	7	7	7
CPR-AMG(1)	15.1	25.9	49.4	95.2
CPR-AMG(2)	22.0	30.9	38.1	40.7
BF	19.9	21.0	21.1	21.1

TABLE 7
Iteration counts for the advection-dominated case with gravity, time step $dt = 4$ days.

Methods/Mesh sizes	20^2	40^2	80^2	160^2
AMG	-	-	-	17.3
CPR-AMG(1)	17.9	17.8	16.1	25.2
CPR-AMG(2)	30.3	29.8	22.7	30.2
BF	13.0	17.0	21.1	23.9

dominant role. The domain is a square box of size 20×20 meters. The absolute permeability is a homogeneous field of 100 mD. Water is injected into the domain through the boundary at the top left corner, and the outlet is at the top right corner. The rate of injection is 5 m^3/day. For spatial discretization, we use uniform grids of size 20×20, 40×40, 80×80, and 160×160. The initial conditions are the same as the heterogeneous two-dimensional example above. The time steps are 10, 4, and 1 days, and the final times are 20, 8, and 2 days, respectively.

The diffusion-dominated case, shown in Table 6, exhibits the same pattern as in the previous example: the AMG preconditioner is the most efficient method, followed by the BF method, CPR-AMG(2), and CPR-AMG(1). AMG and the BF method exhibit optimal performance with respect to problem size. The number of iterations for CPR-AMG(2) also seems to reach a plateau as the mesh size is refined. In contrast, the performance of CPR-AMG(1) does not scale well with respect to mesh size for this case, taking about twice the number of iterations for each level of mesh refinement.

The results for the advection-dominated case are shown in Table 7. The AMG method is not robust and converges only for the largest mesh size (for which it takes the fewest iterations). The BF preconditioner is highly robust and also appears to require iteration counts tending to a constant as the mesh is refined. The performance of CPR-AMG(2) is consistent except for the 80×80 mesh. Although it requires more iterations than CPR-AMG(1), this method shows promising scaling properties, similar to the previous example, since the number of iterations does not grow as the mesh is refined. CPR-AMG(1) performs quite well for this case, but it still exhibits poor scalability, as the number of iterations grow quickly between 80×80 and 160×160.

In the strongly advection-dominated case, AMG diverges for all mesh sizes. The new BF is the most efficient method in this case, requiring the fewest iterations across all mesh sizes (see Table 8). Here, CPR-AMG(1) is more efficient than CPR-AMG(2), requiring about half the number of iterations. Both CPR-AMG(1) and CPR-AMG(2) show good scaling properties in this case. The scaling result for the BF method is not as clear as in the diffusion- and advection-dominated cases, but we suspect that the mesh is not fine enough for a consistent pattern to emerge.

Besides varying the mesh size, we also experimented with changing the time step size for a fixed mesh of 80×80 for the same problem. The final time for the simulation is eight days. The results are reported in Table 9.

TABLE 8
Iteration counts for highly advection-dominated case with gravity, time step $dt = 1$ day.

Methods/Mesh sizes	20^2	40^2	80^2	160^2
AMG	-	-	-	-
CPR-AMG(1)	18.6	19.6	19.7	18.8
CPR-AMG(2)	31.5	34.6	36.6	36.3
BF	13.1	9.3	12.0	16.4

TABLE 9
Results for the advection-dominated case with gravity $P_0 = 10^4$. NI/TS is the number of Newton iteration per time step.

Methods/Time steps	$dt = 1$ day		$dt = 2$ days		$dt = 4$ days		$dt = 8$ days	
	NI/TS	LI/NI	NI/TS	LI/NI	NI/TS	LI/NI	NI/TS	LI/NI
CPR-AMG(1)	12.4	16.9	17	15.9	23	16.1	28	19.4
CPR-AMG(2)	12.4	29.0	17	24.0	23	22.7	28	27.2
BF	12.4	16.7	17	19.0	23	21.1	28	23.0

Since AMG does not converge in this experiment, we exclude it from the results. From Table 9, it is clear that as the time step gets larger, Newton's method takes more iterations to converge. For $dt = 8$ days, there is only one time step, and it is the most difficult case. The number of iterations for CPR-AMG(1) is not significantly affected by the time step except for the largest time step size of eight days. Meanwhile, the number of iterations for CPR-AMG(2) decreases as the time step gets larger, but goes up again at $dt = 8$ days. The BF method shows a consistent increase in the number of iterations for larger time steps. Overall, there is not much of a difference in terms of iteration counts for these three methods, but it is worth noting that the BF method is much faster than the others in terms of run time, as it does not require a global ILU solve.

4.3. Behavior of eigenvalues. It is often possible to obtain insight into the properties of preconditioning operators from the eigenvalues of the preconditioned matrix JM^{-1}. In particular, recall a standard analysis [29] of the convergence behavior of GMRES for solving the preconditioned system (20). Assume that the preconditioned matrix is diagonalizable, $JM^{-1} = V\Lambda V^{-1}$, where Λ is a diagonal matrix containing the eigenvalues of the preconditioned matrix and the columns of V are the corresponding eigenvectors. If $c_k = M^{-1}\hat{c}_k$ are the iterates obtained at the kth step of GMRES iteration, with residual $r_k = q - Jc_k$, then

$$\frac{\|r_k\|_2}{\|r_0\|_2} \leq \|V\|_2 \, \|V^{-1}\|_2 \min_{p_k(0)=1} \max_{\lambda \in \sigma(JM^{-1})} |p_k(\lambda)|, \tag{30}$$

where the minimum in (30) is over all polynomials of degree at most k that have the value 1 at the origin, $\sigma(JM^{-1})$ is the set of eigenvalues of JM^{-1}, and the norm is the vector Euclidean norm. Thus, a good preconditioner tends to produce a preconditioned operator with a compressed spectrum whose entries are not near the origin. In this section, we explore the behavior of the eigenvalues of the preconditioned matrix with an eye toward understanding the effects of features of the discrete problem such as discretization mesh size and qualitative features of the model such as the relative weights of diffusion and advection and the degree of coupling between the components.

Figure 3 gives a representative depiction of the eigenvalues of preconditioned

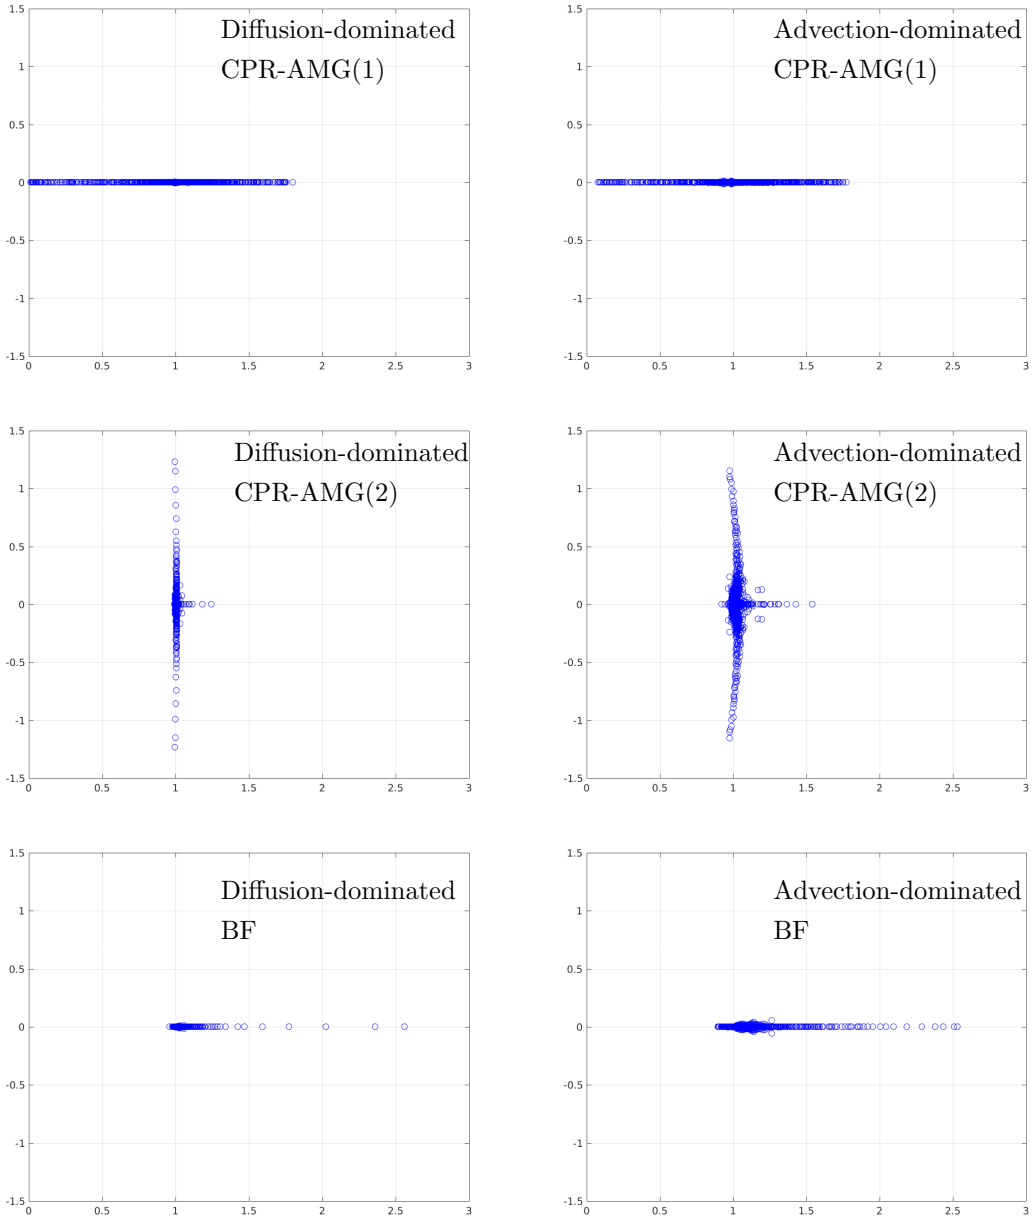

FIG. 3. *Eigenvalues of preconditioned systems for different strategies, applied to the diffusion-dominated Example* 1 *(left) and advection-dominated Example* 2 *(right).*

operators for three of the preconditioners considered. These results are for benchmark problems for which performance is considered in section 4.1, the two-dimensional linear oil-water model discretized on a 100×20 grid. The plots on the left side of the figure show eigenvalues for the diffusion-dominated case (for which solution performance is shown in Table 3), and those on the right show eigenvalues for the

TABLE 10
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	17.6	103.1	20	452	22.6	144.7
CPR-AMG(1)	16	2698	168.6	803.2	20	6069	303.45	1940.8
CPR-AMG(2)	16	712	44.5	299.5	20	1900	95.0	741.1
BF	16	355	22.2	133.6	20	752	37.6	231.1

advection-dominated case (performance in Table 4).[2]

These displays indicate that the spectra for the preconditioned systems for the CPR-AMG(2) and BF preconditioners are bounded away from the origin, whereas for the CPR-AMG(1) preconditioner there are many small eigenvalues. Performance of CPR-AMG(1) improves in the advection-dominated case, and the smallest associated eigenvalues are somewhat further from the origin. In contrast, the latter two preconditioners are largely unchanged in the advection-dominated case, where they are still effective, and the associated eigenvalues are also contained in similarly structured regions far from the origin. We believe the superior performance of the BF preconditioner comes from its greater emphasis on the coupling between pressure and saturation, derived from use of the approximate Schur complement (27).

4.4. Three-dimensional problem. We use a homogeneous permeability field of 100 mD, and the grid is stretched to induce anisotropy. The model dimensions are $25 \times 100 \times 6$ m, and the cell size is $0.5 \times 1 \times 0.05$ m. 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 0.75 m^3/day. The parameters for the capillary pressure model are from Example 1 of Table 2. The simulation is run for 100 days with time step $\Delta t = 20$ days. Table 10 shows the performance results of the diffusion-dominated case for this three-dimensional example, which are consistent with those of the previous two-dimensional example. The AMG preconditioner shows the best results for both the 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. As in the two-dimensional case, the new BF method performs well, requiring fewer than half the iterations of CPR-AMG(2) for both the linear and Brooks–Corey models, and running in about one third the CPU time.

We also tested the three-dimensional SPE10 problem with the linear model of capillary pressure for the different preconditioning strategies. Here, AMG diverges even for the diffusion-dominated case ($P_0 = 10^6$ Pa), even though it was the most efficient method for the two-dimensional example. The BF method is about four times faster than CPR-AMG(2) and five times faster than CPR-AMG(1) in the diffusion-dominated case. CPR-AMG(2) still outperforms CPR-AMG(1) both in terms of iteration counts and run time, but the margin is smaller than for the two-dimensional problem (Table 11). In the advection-dominated case ($P_0 = 10^5$ Pa), unlike in the two-

[2]These computations were done using the `eig` function in MATLAB, and they use MATLAB backslash to perform the actions of the inverses of A_{11}, A_{22} and the modified Schur complement. This contrasts with the solution algorithms tested, which approximate these operations using one AMG V-cycle.

TABLE 11
Performance for the three-dimensional SPE10 *model, diffusion-dominated case.*

Methods/Models	Linear			
	NI	LI	LI/NI	Time (s)
AMG	-	-	-	-
CPR-AMG(1)	17	2410	141.8	614.14
CPR-AMG(2)	17	1661	97.7	448.21
BF	17	490	28.8	121.71

TABLE 12
Performance for the three-dimensional SPE10 *model, advection-dominated case.*

Methods/Models	Linear			
	NI	LI	LI/NI	Time (s)
AMG	-	-	-	-
CPR-AMG(1)	18	1122	62.3	354.38
CPR-AMG(2)	18	1554	86.3	657.12
BF	18	474	26.3	157.24

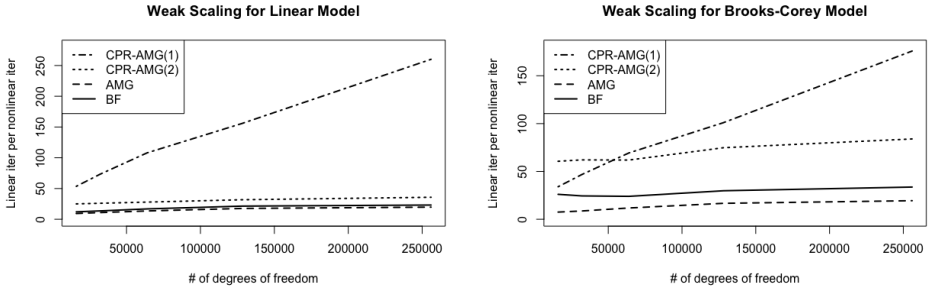

FIG. 4. *Weak scaling for different strategies.*

dimensional example, CPR-AMG(1) is more efficient than CPR-AMG(2), requiring about 23% fewer iterations and 45% the run time. The BF approach is still the most efficient method, taking fewer than half the number of iterations and less than half the run time of CPR-AMG(1) (Table 12). We also note that the number of iterations for the BF method is very consistent with respect to the characteristics of the problem; i.e., it does not change significantly whether the problem is diffusion-dominated, advection-dominated, or strongly advection-dominated.

4.5. 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 time step is fixed at $\Delta t = 20$ days. The domain has constant material properties. The parameters for the capillary pressure models are listed in Example 4 of Table 2. Note that this set of parameters corresponds to a diffusion-dominated problem. The results shown in Figure 4 indicate that the performance of the BF, 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 BF method's performance is nearly identical to that of AMG for the linear model, and still quite close for the

Brooks–Corey model, compared to CPR-AMG(2). CPR-AMG(1), however, does not scale as well as the other two methods. The linear iteration counts for CPR-AMG(1) grow linearly as the mesh is refined.

5. Conclusions. In this work, we have implemented a fully implicit parallel isothermal two-phase flow simulator along with four different preconditioning strategies to solve the linear systems resulting from linearization of the coupled equations, and we have tested the performance of these methods as preconditioners for GMRES. We have also developed a new BF preconditioner whose performance is robust and efficient across all benchmark problems studied. In contrast, although AMG preconditioning applied to the coupled systems is the most efficient choice in some cases (both two-dimensional and three-dimensional diffusion-dominated examples), it exhibits slow convergence and sometimes diverges for advection-dominated cases. The new BF preconditioner achieves consistently low iteration counts across all the tests and varying examples of capillary pressure, and it scales optimally with problem size. The combinative CPR-AMG(1), though robust across all the tests, is the least efficient method, with the exception of the near hyperbolic case, where it is faster than CPR-AMG(2). The additive CPR-AMG(2) method performs well in most cases except the strongly advection-dominated case. It also scales optimally with problem size for both advection-dominated and diffusion-dominated cases.

REFERENCES

[1] U.S. Department of Energy, *ASCEM*, software project, http://esd.lbl.gov/research/projects/ascem/thrusts/hpc/, 2009.

[2] O. Axelsson, *Iterative Solution Methods*, Cambridge University Press, Cambridge, UK, 1994.

[3] K. Aziz and A. Settari, *Petroleum Reservoir Simulation*, Applied Science Publishers, London, UK, 1979.

[4] P. Bastian, *Numerical Computation of Multiphase Flow in Porous Media*, Habilitationsschrift, Department of Mathematics, University of Kiel, 1999.

[5] P. Bastian and R. Helmig, *Efficient fully coupled solution techniques for two-phase flow in porous media: Parallel multigrid solution and large scale computations*, Adv. Water Res., 23 (1999), pp. 199–216, https://doi.org/10.1016/S0309-1708(99)00014-7.

[6] A. Behie and P. Vinsome, *Block iterative methods for fully implicit reservoir simulation*, SPE J., 22 (1982), pp. 658–668, https://doi.org/10.2118/9303-pa.

[7] R. Brooks and A. Corey, *Hydraulic Properties of Porous Media*, Hydrology Papers, Colorado State University, Fort Collins, CO, 1964.

[8] M. Christie and M. Blunt, *Tenth SPE comparative solution project: A comparison of upscaling techniques*, in Proceedings of the SPE Reservoir Simulation Symposium, Society of Petroleum Engineers (SPE), Richardson, TX, 2001, https://doi.org/10.2118/66599-ms.

[9] A. J. Cleary, R. D. Falgout, V. E. Henson, and J. E. Jones, *Coarse-grid selection for parallel algebraic multigrid*, in Solving Irregularly Structured Problems in Parallel, A. Ferreira, J. Rolim, H. Simon, and S.-H. Teng, eds., Lecture Notes in Comput. Sci. 1457, Springer, New York, Berlin, 1998, pp. 104–115, https://doi.org/10.1007/BFb0018531.

[10] T. Clees and L. Ganzer, *An efficient algebraic multigrid solver strategy for adaptive implicit methods in oil-reservoir simulation*, SPE J., 15 (2010), pp. 670–681, https://doi.org/10.2118/105789-pa.

[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, Los Alamos, NV, 2014.

[12] A. T. Corey, *Mechanics of Immiscible Fluids in Porous Media*, Water Resources Publications, Highlands Ranch, CO, 1995.

[13] C. Dawson, H. Klie, M. Wheeler, and C. Woodward, *A parallel, implicit, cell-centered method for two-phase flow with a preconditioned Newton–Krylov solver*, Comput. Geosci., 1 (1997), pp. 215–249, https://doi.org/10.1023/A:1011521413158.

[14] J. Douglas and D. W. Peaceman, *Numerical solution of two-dimensional heat-flow problems*, AIChE J., 1 (1955), pp. 505–512, https://doi.org/10.1002/aic.690010421.

[15] J. Douglas and H. H. Rachford, *On the numerical solution of heat conduction problems in two and three space variables*, Trans. Amer. Math. Soc., 82 (1956), pp. 421–421, https://doi.org/10.1090/s0002-9947-1956-0084194-4.

[16] R. D. Falgout, J. E. Jones, and U. M. Yang, *The design and implementation of hypre, A library of parallel high performance preconditioners*, in Numerical Solution of Partial Differential Equations on Parallel Computers, A. Bruaset and A. Tveito, eds., Lect. Notes Comput. Sci. Eng. 51, Springer, New York, 2006, pp. 267–294, https://doi.org/10.1007/3-540-31619-1_8.

[17] R. Falgout and U. Yang, *HYPRE: A library of high performance preconditioners*, in Preconditioners, Lecture Notes in Comput. Sci. 2331, 2002, pp. 632–641.

[18] I. B. Gharbia and J. Jaffré, *Gas phase appearance and disappearance as a problem with complementarity constraints*, Math. Comput. Simulation, 99 (2014), pp. 28–36.

[19] V. Henson and U. Yang, *BoomerAMG: A parallel algebraic multigrid solver and preconditioner*, Appl. Numer. Math., 41 (2000), pp. 155–177.

[20] H. Klie, M. Rame, and M. Wheeler, *Two-stage preconditioners for inexact Newton methods in multiphase reservoir simulation*, technical report, Center for Research on Parallel Computation, Rice University, Houston, TX, 1996.

[21] S. Lacroix, Yu. Vassilevski, J. Wheeler, and M. Wheeler, *Iterative solution methods for modeling multiphase flow in porous media fully implicitly*, SIAM J. Sci. Comput., 25 (2003), pp. 905–926, https://doi.org/10.1137/S106482750240443X.

[22] A. Lauser, C. Hager, R. Helmig, and B. Wohlmuth, *A new approach for phase transitions in miscible multi-phase flow in porous media*, Adv. Water Res., 38 (2011), pp. 957–966.

[23] B. Lu, *Iteratively Coupled Reservoir Simulation for Multiphase Flow in Porous Media*, Ph.D. thesis, Center for Subsurface Modeling, ICES, The University of Texas - Austin, 2008.

[24] S. Patankar and D. Spalding, *A calculation procedure for heat, mass and momentum transfer in three-dimensional parabolic flows*, Int. J. Heat Mass Transfer, 15 (1972), pp. 1787–1806, https://doi.org/http://dx.doi.org/10.1016/0017-9310(72)90054-3.

[25] D. Peaceman, *Fundamentals of Numerical Reservoir Simulation*, Elsevier, Amsterdam, 1977.

[26] D. W. Peaceman and H. H. Rachford, Jr., *The numerical solution of parabolic and elliptic differential equations*, J. SIAM, 3 (1955), pp. 28–41, https://doi.org/10.1137/0103003.

[27] J. W. Ruge and K. Stueben, *Algebraic multigrid*, in Multigrid Methods, Frontiers in Appl. Math. 3, SIAM, Philadelphia, PA, 1987, Chapt. 4, pp. 73–130, https://doi.org/10.1137/1.9781611971057.ch4.

[28] T. F. Russell and M. F. Wheeler, *Finite element and finite difference methods for continuous flows in porous media*, in The Mathematics of Reservoir Simulation, Frontiers in Appl. Math. 1, SIAM, Philadelphia, 1983, Chapt. 2, pp. 35–106, https://doi.org/10.1137/1.9781611971071.ch2.

[29] Y. Saad, *Iterative Methods for Sparse Linear Systems*, 2nd ed., SIAM, Philadelphia, 2003, https://doi.org/10.1137/1.9780898718003.

[30] Y. Saad and M. H. Schultz, *GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems*, SIAM J. Sci. Stat. Comput., 7 (1986), pp. 856–869, https://doi.org/10.1137/0907058.

[31] K. Stueben, *A review of algebraic multigrid*, J. Comput. Appl. Math., 128 (2001), pp. 281–309, https://doi.org/http://dx.doi.org/10.1016/S0377-0427(00)00516-1.

[32] K. Stueben, T. Clees, H. Klie, B. Lu, and M. Wheeler, *Algebraic multigrid methods (AMG) for the efficient solution of fully implicit formulations in reservoir simulation*, in Proceedings of the SPE Reservoir Simulation Symposium, Society of Petroleum Engineers (SPE), Richardson, TX, 2007, https://doi.org/10.2118/105832-ms.

[33] U. Trottenberg, C. W. Oosterlee, and A. Schuller, *Multigrid*, Academic Press, Orlando, FL, 2001.

[34] J. Wallis, R. Kendall, and L. Little, *Constrained residual acceleration of conjugate residual methods*, in Proceedings of the SPE Reservoir Simulation Symposium, Society of Petroleum Engineers (SPE), Richardson, TX, 1985.

[35] L. Wang, D. Osei-Kuffuor, R. D. Falgout, I. D. Mishev, and J. Li, *Multigrid reduction for coupled flow problems with application to reservoir simulation*, in Proceedings of the SPE Reservoir Simulation Conference, Society of Petroleum Engineers (SPE), Richardson, TX, 2017, SPE-182723-MS.

[36] J. Watts, *A compositional formulation of the pressure and saturation equations*, in Proceedings of the SPE Reservoir Simulation Symposium, Society of Petroleum Engineers (SPE), Richardson, TX, 1985, SPE 12244.

[37] J. A. White and R. I. Borja, *Block-preconditioned Newton–Krylov solvers for fully coupled flow and geomechanics*, Comput. Geosci., 15 (2011), pp. 647–659, https://doi.org/10.1007/s10596-011-9233-7.

[38] Y.-S. Wu and P. Forsyth, *On the selection of primary variables in numerical formulation for modeling multiphase flow in porous media*, J. Contaminant Hydrol., 48 (2001), pp. 277–304.

[39] U. Yang, *Parallel algebraic multigrid methods—High-performance preconditioners*, in Numerical Solution of Partial Differential Equations on Parallel Computers, Lect. Notes Comput. Sci. Eng. 51, Springer, New York, 2006, pp. 209–236, https://doi.org/10.1007/3-540-31619-1_6.