Comput Geosci (2011) 15:647–659
DOI 10.1007/s10596-011-9233-7

ORIGINAL PAPER

Block-preconditioned Newton–Krylov solvers for fully coupled flow and geomechanics

Joshua A. White · Ronaldo I. Borja

Received: 18 March 2010 / Accepted: 14 March 2011 / Published online: 12 April 2011

Abstract The focus of this work is efficient solution methods for mixed finite element models of variably saturated fluid flow through deformable porous media. In particular, we examine preconditioning techniques to accelerate the convergence of implicit Newton–Krylov solvers. We highlight an approach in which preconditioners are built from block-factorizations of the coupled system. The key result of the work is the identification of effective preconditioners for the various sub-problems that appear within the block decomposition. We use numerical examples drawn from both linear and nonlinear hydromechanical models to test the robustness and scalability of the proposed methods. Results demonstrate that an algebraic multigrid variant of the block preconditioner leads to mesh-independent convergence, good parallel efficiency, and insensitivity to the material parameters of the medium.

Keywords Newton–Krylov · Coupled geomechanics · Algebraic multigrid · Mixed finite elements

J. A. White (✉)
Computational Geosciences Group,
Lawrence Livermore National Laboratory,
P.O. Box 808, L-286, Livermore,
CA, 94551, USA
e-mail: jawhite@llnl.gov

R. I. Borja
Civil and Environmental Engineering,
Stanford University, Stanford,
CA, 94305, USA
e-mail: borja@stanford.edu

1 Introduction

In many applications involving porous media, it is necessary to model hydromechanical processes in a tightly coupled way in order to make meaningful predictions. The focus of this work is efficient numerical methods for modeling these coupled systems. In particular, we focus on efficient preconditioning techniques to accelerate the convergence of implicit Newton–Krylov solvers.

Coupled fluid–structure interaction problems appear in many geotechnical and geoscientific applications, e.g., earthen dams, levees, and other geotechnical structures [1–5]; oil and gas formations [6–8]; natural and engineered geothermal reservoirs [9]; and carbon sequestration sites [10–12]. Of course, each class of problem presents its own modeling issues, and the choice of methods that may be applied vary widely. Nevertheless, there is widespread recognition that uncoupled or loosely coupled models are often insufficient.

In this broader context, this work focuses on a formulation suitable for modeling variably saturated flow in soils—a useful formulation for many geotechnical and hydrologic applications. Many ideas contained herein can also be extended to other subsurface systems. The resulting discretization leads to a nonlinear system of residual equations that must be solved to advance the solution in each timestep of the simulation. The roots of these residual equations are found using Newton's method, and it is this portion of the code that dominates the computational expense of a typical simulation. The key focus of this work is how to design a scalable nonlinear solver that can readily handle increasingly large problems by taking efficient advantage of today's parallel computing platforms.

Because the coupled formulation is a multi-field problem, the Jacobian system that must be solved in each Newton iteration inherits a useful block structure. Unfortunately, this system is typically ill-conditioned, making its solution using iterative solvers inefficient without good preconditioning. Efficient solvers can be designed, however, by using the underlying block structure to break the global problem into easier-to-solve sub-problems. In particular, here we explore the use of triangular block preconditioners for this class of problem. Preconditioners of this type were first explored by Bramble and Pasciak [13] for Stokes flow, and this and similar approaches have become increasingly popular in the fluid dynamics community for dealing with Stokes and Navier–Stokes systems [14]. This is particularly true in the geoscience community focused on magma migration and crustal dynamics [15, 16]. A key advantage of the block-focused approach is that the global preconditioner can be built from smaller algebraic or physics-based preconditioners, allowing for significant code reuse. The resulting approach naturally lends itself to an object-oriented design and an elegant algorithmic framework for incorporating single-physics expertise into multiphysics codes. Unfortunately, while the block preconditioning approach is quite flexible and can be applied to a wide range of problems, the success of the method depends on identifying effective preconditioners for the various sub-problems that appear in the block decomposition. Effective sub-preconditioners are problem specific, and the key focus of this work is identifying good strategies for the coupled hydromechanical model at hand.

The application of the block approach to coupled hydromechanical systems is quite natural, but unfortunately this avenue has not been sufficiently explored to date. One exception is the work of Toh et al. [17], who considered the application of a block-preconditioning approach to Biot's linear consolidation model [18, 19]. In this work, we consider a more general, nonlinear formulation and also examine the parallel implementation of the solver. We also advocate a different approach for preconditioning the various sub-problems that arise. We identify a particular sub-preconditioning strategy that leads to mesh-independent convergence, good parallel efficiency, and insensitivity to the material parameters of the problem.

As a general outline, Section 2 presents the governing formulation, derived from continuum theory of mixtures. Section 3 discusses the implementation of the model via a stabilized mixed finite element method. The heart of the work is Section 4, which considers the design of a scalable Newton–Krylov solver framework. Finally, Section 5 examines the scalability and parallel efficiency of the proposed methodology on several test problems and highlights important features of the variably saturated problem to consider when implementing preconditioning strategies.

2 Coupled formulation

2.1 Governing equations

We are interested in solution techniques for the following nonlinear model of a quasi-static, variably saturated soil,

$$\begin{aligned} &\theta\dot{\psi} + \psi\nabla\cdot\dot{\boldsymbol{u}} + \nabla\cdot\boldsymbol{w} = 0 \\ &\nabla\cdot(\boldsymbol{\sigma}' - \psi p\mathbf{1}) + \rho\boldsymbol{g} = \mathbf{0}. \end{aligned} \tag{1}$$

Here, θ is the soil porosity, ψ is the water-phase saturation, $\boldsymbol{u}$ is the solid displacement, $\boldsymbol{w}$ is the seepage (Darcy) velocity, $\boldsymbol{\sigma}'$ is the effective stress, $\mathbf{1}$ is a second-order unit (Kronecker) tensor, and $\rho\boldsymbol{g}$ is a body force due to the self-weight of the three-phase mixture. We solve the problem via a two-field mixed finite element formulation in which the primary unknowns are $\boldsymbol{u}$ and p.

Several assumptions about the soil behavior are necessary to arrive at this particular model. For near surface formations, a common assumption is that the air pressure remains in equilibrium with atmospheric pressure, or $p_a = 0$. This passive-gas assumption leads to a pseudo-three-phase formulation in which it is unnecessary to track the air-phase pressures. At typical pressures in applications of interest here, the solid and fluid phases are approximately incompressible, so that the bulk moduli $K_s = K_w \approx \infty$. Several compressibility terms have therefore been dropped. This also implies that the Biot coefficient,

$$b = 1 - \frac{K_{sk}}{K_s} \approx 1, \tag{2}$$

where K_{sk} is the bulk modulus for the solid skeleton. We have also adopted a particular form for the effective stress decomposition,

$$\boldsymbol{\sigma}' = \boldsymbol{\sigma} + \psi p\mathbf{1}. \tag{3}$$

All of these assumptions can be relaxed if they are inappropriate for a particular geologic system. For example, see Borja [20] for a formulation containing full phase compressibility, Ehlers et al. [2] for the active air-pressure case, and Gawin et al. [1] for hydrothermal–mechanical models. Additional discussion concerning

the appropriate choice of the effective stress decomposition can also be found in [20–25]. For our purposes here, the above model contains all the salient features necessary for a discussion of solution methodologies.

2.2 Constitutive models

To complete the formulation, several constitutive relationships are required. A generalized Darcy's law is used to relate the seepage velocity to the pressure and elevation potential,

$$\boldsymbol{w} = \frac{-k_{\mathrm{r}} k}{\eta} \left(\nabla p - \rho_{\mathrm{w}} \boldsymbol{g}\right). \tag{4}$$

Here, $k_{\mathrm{r}} \in [0, 1]$ is a relative permeability factor (dimensionless), k is the saturated intrinsic permeability (units of square meters), and η is the fluid viscosity (units of pascal second). For notational convenience, we also define a lumped coefficient $\kappa = (k_{\mathrm{r}} k)/\eta$. In this work, we use the van Genuchten [26] pressure/saturation and pressure/relative permeability relationships $\psi(p)$ and $k_{\mathrm{r}}(p)$.

An elastoplastic model is used to represent the soil behavior. For simplicity of presentation, we make a small-strain assumption and define the strain measure $\boldsymbol{\epsilon}$ as the symmetric gradient of the displacement,

$$\boldsymbol{\epsilon} = \nabla^s \boldsymbol{u} = \frac{1}{2}\left(\nabla \boldsymbol{u} + \nabla^{\mathsf{T}} \boldsymbol{u}\right). \tag{5}$$

The effective stress–strain relationship can be written in a general incremental form as

$$\Delta \boldsymbol{\sigma}' = \mathbb{C} : \Delta \boldsymbol{\epsilon}, \tag{6}$$

where $\mathbb{C}$ is a fourth-order tensor of tangential moduli. In the numerical examples to follow, we have adopted a relatively simple plasticity model—non-associative Drucker–Prager plasticity.

The solver techniques presented below can be readily extended to the finite deformation regime, but for the present, we wish to avoid clouding the presentation with the proliferation of geometric terms that would appear in the linearization of the variational form. The key difference between the small-strain and finite-strain regimes is that certain symmetries are lost in the operators when including geometric terms. The solution strategies proposed below do not rely on any particular symmetry properties for the system anyway, so the finite deformation case can be handled quite naturally.

2.3 Simplified models

Before concluding this section, we make some observations concerning certain limit states of the variably saturated model. The studied model contains both nonlinear solid and fluid behavior. If, however, the soil is fully saturated and the stress–strain response is linear, we recover a linear model of coupled consolidation [18, 19, 27],

$$\begin{aligned} \nabla \cdot \dot{\boldsymbol{u}} + \nabla \cdot \boldsymbol{w} &= 0 \\ \nabla \cdot (\boldsymbol{\sigma}' - p\mathbf{1}) + \rho \boldsymbol{g} &= \mathbf{0}. \end{aligned} \tag{7}$$

Also note that in the undrained limit, as the permeability of the medium $k \to 0$, the mass balance equation further reduces to the constraint equation,

$$\nabla \cdot \dot{\boldsymbol{u}} = 0. \tag{8}$$

That is, the fluid trapped in the pore space enforces volumetric incompressibility on the solid skeleton, and the fluid pressures can be reinterpreted as Lagrange multipliers enforcing this constraint. In the discrete setting, the presence of this constraint introduces an inf–sup stability restriction on the choice of finite element spaces that can be used without spurious pressure oscillations—in exactly the same manner as for mixed finite element implementations of Stokes flow or incompressible elasticity [28–31]. As a result, the mixed finite element implementation must either adopt a Ladyzhenskaya–Babuska–Brezzi (LBB)-stable finite element pair or employ a stabilized formulation that allows for a broader range of stable combinations [6, 32–34]. In this work, we adopt the stabilized formulation proposed in [32], which allows for equal-order linear interpolation of the displacement and pressure fields while maintaining optimal convergence. We emphasize though that the solution strategies proposed in this work can be used interchangeably with LBB-stable and stabilized formulations.

3 Implementation

3.1 Variational form

The mixture occupies a domain Ω with boundary Γ. Initial conditions at $t = 0$ are given as $\{\boldsymbol{u}_0, p_0\}$. The boundary is suitably decomposed into regions where essential and natural conditions are specified for both

the solid and fluid. For the two-field discretization, two spaces of trial functions are defined as

$$\mathcal{U} = \{\boldsymbol{u} : \Omega \to \mathbb{R}^3 \mid \boldsymbol{u} \in \boldsymbol{H}^1,\ \boldsymbol{u} = \overline{\boldsymbol{u}} \text{ on } \Gamma_u\} \tag{9}$$

$$\mathcal{P} = \{p : \Omega \to \mathbb{R} \mid p \in H^1,\ p = \overline{p} \text{ on } \Gamma_p\} \tag{10}$$

where H^1 is a Sobolev space of degree one. Corresponding spaces of weighting functions are similarly defined, with homogeneous conditions on the essential boundaries,

$$\mathcal{U}_0 = \{\boldsymbol{\eta} : \Omega \to \mathbb{R}^3 \mid \boldsymbol{\eta} \in \boldsymbol{H}^1,\ \boldsymbol{\eta} = \boldsymbol{0} \text{ on } \Gamma_u\} \tag{11}$$

$$\mathcal{P}_0 = \{\phi : \Omega \to \mathbb{R} \mid \phi \in H^1,\ \phi = 0 \text{ on } \Gamma_p\} \tag{12}$$

The resulting weak problem is to find $\{\boldsymbol{u}, p\} \in \mathcal{U} \times \mathcal{P}$ such that for all $\{\boldsymbol{\eta}, \phi\} \in \mathcal{U}_0 \times \mathcal{P}_0$,

$$\begin{aligned}\mathcal{R}_{\text{mom.}} = &-\int_\Omega \nabla^s \boldsymbol{\eta} : \boldsymbol{\sigma}' d\Omega + \int_\Omega \psi p \, \nabla \cdot \boldsymbol{\eta} \, d\Omega \\ &+ \int_\Omega \boldsymbol{\eta} \cdot \rho \boldsymbol{g} \, d\Omega + \int_{\Gamma_t} \boldsymbol{\eta} \cdot \overline{\boldsymbol{t}} \, d\Gamma = 0\end{aligned} \tag{13}$$

$$\begin{aligned}\mathcal{R}_{\text{mass}} = &\int_\Omega \phi \, \theta \dot{\psi} \, d\Omega + \int_\Omega \phi \psi \, \nabla \cdot \dot{\boldsymbol{u}} \, d\Omega \\ &+ \int_\Omega \nabla\phi \cdot \kappa \nabla p \, d\Omega - \int_\Omega \nabla\phi \cdot \kappa \, \rho_w \boldsymbol{g} \, d\Omega \\ &- \int_{\Gamma_q} \phi \overline{q} \, d\Gamma = 0\end{aligned} \tag{14}$$

Here, $\overline{q}$ and $\overline{\boldsymbol{t}}$ are specified flux and traction values at the boundaries. Note that in many hydrologic applications, more complicated mixed boundary conditions are also common.

3.2 Discrete form

These two nonlinear residual equations govern the behavior of the mixture and can be suitably discretized in space and time. For brevity, we will not review these details, as they are more or less standard. The spatial discretization uses mixed hexahedral elements for the displacement and pressure field, and the temporal discretization uses a trapezoidal integration scheme, i.e., either implicit Euler or Crank–Nicolson.

Let the vector $X_n = \{U_n, P_n\}$ represent the discrete solution at time t_n. The fully discrete system of equations for the solution at each timestep is given by

$$R(X_n, X_{n-1}) = \begin{bmatrix} R_{\text{mom.}} \\ R_{\text{mass}} \end{bmatrix} = 0. \tag{15}$$

This system is nonlinear due to nonlinear the constitutive behavior of the $\boldsymbol{\sigma}'(\boldsymbol{u})$, $\psi(p)$, and $\kappa(p)$ relationships and can exhibit very stiff behavior. A full Newton iteration is used to drive this residual system to zero. Because the residual vector R consists of two blocks, the Jacobian system that must be solve in each Newton update has a 2×2 block structure,

$$\begin{bmatrix} A & B_1 \\ B_2 & C \end{bmatrix} \begin{bmatrix} \Delta U \\ \Delta P \end{bmatrix} = - \begin{bmatrix} R_{\text{mom.}} \\ R_{\text{mass}} \end{bmatrix}. \tag{16}$$

The various matrices are assembled in the standard way from element contributions. For an implicit Euler time integration scheme, the respective sub-matrices are given by

$$[A]_e^{a,b} = -\int_{\Omega^e} \nabla^s \boldsymbol{\eta}^a : \mathbb{C} : \nabla^s \boldsymbol{\eta}^b \, d\Omega \tag{17}$$

$$\begin{aligned}[B_1]_e^{a,b} = &\int_{\Omega^e} \nabla \cdot \boldsymbol{\eta}^a \left(\psi_n + \frac{\partial \psi}{\partial p}\Big|_n p_n \right) \phi^b \, d\Omega \\ &+ \int_{\Omega^e} \boldsymbol{\eta}^a \cdot \left(\frac{\partial \rho}{\partial p}\Big|_n \boldsymbol{g} \right) \phi^b \, d\Omega\end{aligned} \tag{18}$$

$$[B_2]_e^{a,b} = \int_{\Omega^e} \phi^a \psi_n \nabla \cdot \boldsymbol{\eta}^b \, d\Omega \tag{19}$$

$$\begin{aligned}[C]_e^{a,b} = &\int_{\Omega^e} \phi^a \left(\theta \frac{\partial \psi}{\partial p}\Big|_n \right) \phi^b \, d\Omega \\ &+ \int_{\Omega^e} \phi^a \frac{\partial \psi}{\partial p}\Big|_n \nabla \cdot (\boldsymbol{u}_n - \boldsymbol{u}_{n-1}) \, \phi^b \, d\Omega \\ &+ \int_{\Omega^e} \nabla\phi^a \cdot (\Delta t \kappa_n) \nabla \phi^b \, d\Omega \\ &+ \int_{\Omega^e} \nabla\phi^a \cdot \Delta t \frac{\partial \kappa}{\partial p}\Big|_n (\nabla p_n - \rho_w \boldsymbol{g}) \, \phi^b \, d\Omega\end{aligned} \tag{20}$$

where a, b are degree-of-freedom indices and e is an element domain index. Note that in the saturated regime, the symmetry $B_2 = B_1^T$ exists, but this symmetry is lost when the saturation $\psi < 1$. Furthermore, for general models of solid and fluid constitutive behavior, the A and C blocks may lose symmetry as well. As a result, a general solver methodology must avoid relying on any symmetry assumptions for the system.

3.3 Stability

As noted earlier, the coupled formulation is subject to an inf–sup stability restriction in certain regimes. In particular, if we consider a saturated mixture and look

at the limit as the permeability goes to zero, we find that the pore fluid exerts an incompressibility condition on the deformation of the solid matrix. In discrete form, the Jacobian system becomes

$$\begin{bmatrix} A & B_1 \\ B_1^T & 0 \end{bmatrix} \begin{bmatrix} \Delta U \\ \Delta P \end{bmatrix} = - \begin{bmatrix} R_{\text{mom.}} \\ R_{\text{mass}} \end{bmatrix} \tag{21}$$

where the C-block vanishes. The saddle-point structure of this matrix is familiar from other applications involving constraints, including Stokes flow and Darcy flow [28–31].

We adopt a stabilizing modification to the variational form that allows us to circumvent the inf–sup condition and use equal-order elements even in the vanishing permeability regime. There are a variety of such schemes available, but here we use a variant of the polynomial pressure projection (PPP) technique. This scheme was developed by Dohrmann, Bochev, and Gunzburger for Stokes flow and Darcy flow [35, 36] and was further analyzed by Burman in [37]. In [32], the authors applied a PPP-type stabilization for mixed finite element models of coupled fluid flow and geomechanics, with good results. It is not our purpose here to describe the use of a stabilized formulation, as this has already be described extensively in [32, 38]. Other stabilization schemes for coupled consolidation problems have also been tested in [6, 33, 34]. Note that the solver methodology proposed here does not depend on the introduction of stabilizing terms and can just as well be applied to intrinsically stable discretizations.

4 Newton–Krylov solver

The efficiency of the solution process depends on two factors: the convergence rate of the outer Newton search and the efficiency of the linear solver chosen to handle the Newton update equation. While variably saturated soil models often exhibit very stiff behavior, a full Newton iteration typically shows good convergence properties—in particular, quadratic behavior when the current guess moves into a neighborhood of the solution. The time-dependent nature of the problem helps in this respect, as the solution to the previous timestep often provides a good initial guess for the subsequent step.

The key computational challenge is therefore to solve the large, ill-conditioned Jacobian systems in an efficient and scalable manner. As direct solvers quickly become memory-limited, Krylov methods are the strategy of choice. Unfortunately, given the sensitivity of Krylov solvers to the conditioning properties of J, quality preconditioning is essential. For simplicity, we focus on left preconditioning the linear systems, i.e.,

$$P^{-1} J x = P^{-1} b \tag{22}$$

The key question to be addressed in this work is how to choose P^{-1} for the coupled hydromechanical model such that a Krylov solver applied to (22) exhibits mesh-independent convergence.

4.1 Block preconditioner

The Jacobian admits a variety of block factorizations [39]. In particular, a block LU factorization is

$$J = \begin{bmatrix} A & 0 \\ B_2 & S \end{bmatrix} \begin{bmatrix} I & A^{-1} B_1 \\ 0 & I \end{bmatrix}, \tag{23}$$

where $S = C - B_2 A^{-1} B_1$ is the Schur complement—with respect to A—for the system. The key observation for our purposes here is that while the system J is ill-conditioned, its upper-triangular factor U only has a single distinct eigenvalue $\lambda = 1.0$. A Krylov-based iteration on U would therefore converge in at most two iterations. This implies that an effective preconditioning strategy is to choose $P \approx L$ such that $P^{-1} J \approx U$. Depending on the quality of this approximation, one would expect the Krylov solver to converge in only a few iterations. Note that the inverse of the "exact" preconditioner can be formally computed as

$$P^{-1} = \begin{bmatrix} A^{-1} & 0 \\ -S^{-1} B_2 A^{-1} & S^{-1} \end{bmatrix}. \tag{24}$$

Block-triangular preconditioners were first explored by Bramble and Pasciak [13] for the Stokes problem and have become increasingly popular for dealing with a variety of two-field problems in fluid and solid mechanics [14–17]. In practice, Eq. 24 only provides a template for forming the preconditioner, as the exact inverses A^{-1} and S^{-1} are too expensive to compute. Instead, these operators are replaced with their own preconditioners P_A^{-1} and P_S^{-1}, which by design are good approximations to the exact inverses,

$$P^{-1} = \begin{bmatrix} P_A^{-1} & 0 \\ -P_S^{-1} B_2 P_A^{-1} & P_S^{-1} \end{bmatrix}. \tag{25}$$

Because the block preconditioning strategy is based on a purely algebraic argument, it can be applied to a wide range of model problems. Unfortunately, the strategy only provides a framework for constructing a good global preconditioner when one has readily available methods for approximating the inverses A^{-1}

and S^{-1}. The behavior of these sub-operators is obviously problem specific, and so the ultimate success of this approach lies is finding good preconditioners P_A^{-1} and P_S^{-1} for the specific problem under consideration. The heart of this work is focused on good choices for these operators in the case of coupled hydromechanical models.

Note that the block-triangular preconditioner is asymmetric, and therefore, an asymmetric Krylov method must be employed, e.g., GMRES or BICGSTAB . In [13], the authors recovered a symmetric method for the Stokes system by using a non-standard inner product inside a CG iterations. In practice, if the preconditioner works well and the number of iterations is kept small, the disadvantage to using an asymmetric method will not be great. Further, except under limiting constitutive assumptions, the symmetry of the Jacobian for the variably saturated hydromechanical model is lost anyway, so an asymmetric formulation is often unavoidable.

Also, we observe that an alternative decomposition could take the Schur complement with respect to C instead, i.e., $S = A - B_1 C^{-1} B_2$. We have chosen the Schur complement with respect to A because the term $B_2 A^{-1} B_1$ is definite and relatively easy to precondition. In contrast, the term $B_1 C^{-1} B_2$ is singular and can make conditioning $S = A - B_1 C^{-1} B_2$ a more challenging problem. There is also the additional difficulty that for a LBB-stable discretization and the undrained limit, $C = 0$, and therefore C^{-1} does not exist. It is quite common for an undrained analysis to be performed to assess certain features of the geotechnical response, and so a special case preconditioner would need to be developed to handle this condition.

4.2 Sub-preconditioners

We now consider various strategies for approximating the necessary inverse operators to build the global block preconditioner. The more straightforward problem is finding a good candidate for P_A^{-1}. The A block of the Jacobian system is a stiffness matrix associated with the mechanical response of the porous medium. Over the years, many quality algebraic preconditioners have been developed for elastic and elastoplastic problems, and these can be used directly within the block-preconditioning framework. In the numerical experiments to follow, we test two basic strategies: incomplete LU (ILU) and algebraic multigrid (AMG) preconditioning.

The trickier problem is how to precondition the Schur complement $S = C - B_2 A^{-1} B_1$. While C is sparse, the presence of A^{-1} in the second term leads to a dense S which is awkward to deal with. Again, a variety of strategies are available. The first approach is to develop an explicit but sparse approximation to S and then compute the preconditioner based on this sparse approximation. Consider the simple approximation

$$S_D = C - B_2 \,\mathrm{diag}(A)^{-1} B_1 \tag{26}$$

This operator is easy to compute and has a good sparsity pattern. An algebraic preconditioner (such as ILU or AMG) is then applied to S_D. An appealing feature is that S_D accounts for possible asymmetries in B_2 and B_1. The quality of this approximation, however, will depend on the diagonal dominance of A.

A second, more interesting strategy also involves a sparse approximation to S. Consider the approximation

$$S_M = C - \alpha M_\mathrm{p}, \qquad [M_\mathrm{p}]_e^{a,b} = \int_{\Omega^e} \phi^a \phi^b \, d\Omega \tag{27}$$

Here, M_p is the pressure mass matrix. The use of the pressure mass matrix to approximate $B_2 A^{-1} B_1$ is motivated by observations in [40, 41] on discretizations of incompressible elasticity and Stokes flow. For a saturated, linear elastic model in the undrained limit ($k \to 0$), note that the Schur complement operator $B_2 A^{-1} B_1$ for the hydromechanical model is formally equivalent to that encountered in a Stokes flow discretization. One can show from an inf–sup argument [41] that the pressure mass matrix is spectrally equivalent to $B_2 A^{-1} B_1$, and therefore, it can serve as the basis for a sparse preconditioning operator.

The current problem is slightly different in that the C block is generally non-zero. The relative magnitudes of the C and $B_2 A^{-1} B_1$ terms depend on the given permeability and timestep. The scalar α is a weighting factor which is used to ensure that the two components M_p and C of the sparse approximation have appropriate relative magnitudes. One choice is

$$\alpha = \frac{-1}{\|\mathbb{C}\|} \tag{28}$$

where $\|\mathbb{C}\|$ is the norm of the current elastoplastic tangent operator.

Also note that the symmetry $B_2 = B_1^T$ is lost in the unsaturated regime. The question is then whether the symmetric mass matrix approximation continues to work well in the unsaturated regime. Note that when the soil is unsaturated, however, it is significantly more compressible than in the saturated state—due to the additional fluid storage available through slight changes in saturation—and thus the magnitude of C is large in comparison to $B_2 A^{-1} B_1$. As a result, for unsaturated

models or highly permeable media, the particular approximation to $B_2 A^{-1} B_1$ becomes less crucial in determining convergence behavior. The numerical results in the next section illustrate this particular behavior.

A final option to consider in preconditioning the sub-problems is to approximate the action of A^{-1} and/or S^{-1} through a few iterations of their own Krylov solvers. In this approach, note that the preconditioners become variable operators, and therefore, a flexible outer iteration such as GMRESR or FGMRES must be employed. The advantage of this approach is that the quality of the approximation of P_A^{-1} and P_S^{-1} can be controlled by using higher or lower tolerances, allowing one to reduce the number of global iterations in favor of a larger number of sub-iterations and vice versa. The disadvantage of this approach is that the complexity of the nested iterations can create significant computational overhead.

5 Numerical experiments

We now consider a series of numerical experiments to test the performance of the proposed block-preconditioning strategies. As indicated earlier, the block-preconditioning approach can lead to a host of alternative strategies depending on how the particular sub-operators are preconditioned. Here, six variants are studied. The first four are labeled BP-ILU-SD, BP-AMG-SD, BP-ILU-SM, and BP-AMG-SM. The "SD" preconditioners are based on the sparse Schur-complement approximation S_D, while "SM" variants use the mass matrix approximation S_M. "ILU" and "AMG" denote the algebraic strategy that is used to compute the sub-preconditioners P_A^{-1} and P_S^{-1}. We also test two additional strategies, labeled BP-EXACT and BP-INEXACT. Both of these strategies instead use inner solvers to approximate the action of A^{-1} and S^{-1} whenever they are required. Note that this leads to several levels of nested iterations. The difference between the two is that BP-EXACT uses a high tolerance for convergence (equal to the outer Krylov convergence tolerance), so that the approximation of the inverse operators is precise. This is not meant as a practical computational strategy, but is used to illustrate the convergence of the outer Krylov solver given the "exact" preconditioner (Eq. 24). In the inexact strategy, a much lower tolerance for convergence is specified—a more practical strategy. In both cases, these inner iterations can be left unpreconditioned, or take their own inner preconditioners. Here, we precondition the sub-iterations with AMG preconditioners for A and S_M.

The examples test both serial and parallel performance. The parallel tests were run on a distributed memory platform using 2.3-GHz quad-core processors. These multi-core CPUs are arranged into 16-processor nodes sharing 32 GB of memory, with InfiniBand switches between nodes. The serial computations were run on individual processors of the same platform. In general, all performance metrics have been averaged over many timesteps to provide a representative sample.

The finite element code used for these examples employs a number of widely available, open-source packages: the DEAL.II finite element library [42], the P4EST mesh-handling library [43, 44], and the TRILINOS suite of algorithms [45]. The ILU decompositions use the ILU(0) implementation in the TRILINOS-IFPACK package [46]. The AMG preconditioners use the smoothed-aggregation multigrid strategy implemented in the TRILINOS-ML package [47]. We use a Chebyshev smoother and a full W-cycle. The number of grid levels in the hierarchy is allowed to grow with the problem size to achieve a desired coarsening rate.

Constitutive and numerical parameters for all the examples are summarized in Table 1. In all examples, a relative reduction convergence tolerance of 10^{-8} is used for the linear solves. The particular Krylov solver employed is indicated in the discussion of each test problem. For the nonlinear model, an exact Newton algorithm is used, with updates in each timestep continued until a reduction criterion

$$\frac{\|R^k\|}{\|R^0\| + 1} < 10^{-8} \tag{29}$$

is satisfied. In general, a more efficient inexact Newton algorithm can be employed [48], but for simplicity we do not focus on this additional detail.

5.1 Cryer's sphere

The first example examines the behavior of a linear elastic, fully saturated soil. In this case, the governing equations are linear, and only a single Krylov solve is required per timestep. As a test problem, we consider Cryer's sphere, a classic consolidation problem for which an analytical solution exists [27, 49]. A three-dimensional, poro-elastic sphere of radius r is loaded on its outer boundary by a confining pressure p_0, while its outer boundary is allowed to drain freely. Figure 1 illustrates the geometry of the test problem as well as a comparison of the analytical and numerical solutions for the pore pressure at the center of the sphere, for several values of the Poisson ratios ν. The problem is non-dimensionalized in terms of the consolidation

Table 1 Constitutive and numerical parameters used for the numerical examples

Parameter	Symbol	Cryer's sphere	Unsaturated levee	Units
Porosity	θ	0.5	0.5	
Solid density	ρ_s	1.0	2.0	Mg/m^3
Fluid density	ρ_w	1.0	1.0	Mg/m^3
Bulk modulus	K_{sk}	10.0	2.0	MPa
Poisson ratio	ν	0.25	0.30	
Cohesion	c	–	5.0	kPa
Friction angle	φ_f	–	25.0	deg
Dilatancy angle	φ_d	–	0.0	deg
Intrinsic permeability	k	5.56×10^{-8}	10^{-12}	m^2
Dynamic viscosity	η	10^{-6}	10^{-6}	kPa s
Residual saturation	ψ_1	–	0	
Maximum saturation	ψ_2	–	1.0	
Scaling suction	s_a	–	10.0	kPa
Shape parameter	n	–	2.0	
Shape parameter	m	–	0.5	
Timestep	Δt	1.0	1,800.0	s
Integration parameter	θ	0.5	1.0	
Krylov tolerance	ϵ_k	10^{-8}	10^{-8}	
Newton tolerance	ϵ_n	–	10^{-8}	

coefficient c and demonstrates a sharp rise in pore pressure before subsequent dissipation—see [27, 49] for further details. For comparison, we have included numerical results using both a stabilized equal-order $Q1Q1$ interpolation and an intrinsically stable $Q2Q1$ (Taylor–Hood) discretization. For the preconditioner tests to follow, we use the less expensive $Q1Q1$ discretization and a Poisson ratio of $\nu = 0.25$. The size of the test problem at various levels of refinement is indicated in Table 2. Levels 1 to 3 are examined in single-processor studies, while levels 4 to 7 are used for a weak-scaling study.

We first consider the serial performance of the six suggested preconditioning strategies upon mesh refinement. In this study, FGMRES is used, as a flexible iteration is required for BP-INEXACT. All performance metrics are summarized in Table 3. The setup time refers to the time taken to initialize the various sub-preconditioners, e.g., computing the multigrid hierarchy for the AMG preconditioners, or performing the

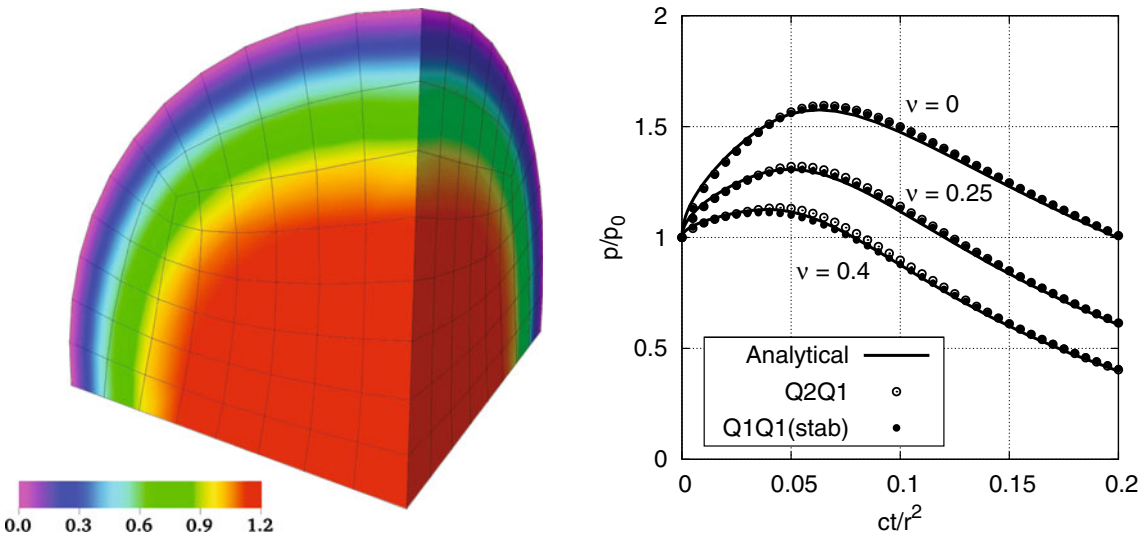

(a) Pressure (kPa), $\nu = 0.25$, $t = 0.02$ s (b) Nondimensional pressure vs. time

Fig. 1 Comparison of numerical and analytical pressure solutions to Cryer's sphere problem, using both a LBB-stable $Q2Q1$ discretization and a stabilized $Q1Q1$ formulation

Table 2 Cryer's sphere problem dimensions at various refinement levels

Ref.	Total unknowns	Displacement	Pressure
1	2,388	1,791	597
2	16,292	12,219	4,073
3	120,132	90,099	30,033
4	922,244	691,683	230,561
5	7.2 million	5.4 million	1.8 million
6	57.2 million	42.9 million	14.3 million
7	455.3 million	341.5 million	113.8 million

Levels 1 to 3 are addressed with a serial scaling study, while levels 4 to 7 are addressed with a parallel scaling study

decomposition for ILU. While these are relatively expensive operations, for the linear test problem, they can be performed once at the beginning of the simulation and remain fixed over subsequent timesteps. The table also records the average number of iterations to convergence in each timestep and the average solve time.

As expected, BP-EXACT converges in only two iterations regardless of the refinement level. While this provides a convenient check that the block preconditioner is properly implemented, it is by far the most expensive strategy. To improve the performance, BP-INEXACT uses a lower convergence tolerance, allowing additional outer iterations in favor of cheaper sub-preconditioner operations. In this particular case, choosing a sub-tolerance of 10^{-3} led to the best performance, but the strategy is still very expensive. It appears that the complexity associated with the various levels of iteration strongly affects performance. The more direct, single-iteration strategies are significantly cheaper.

Table 3 Serial performance of various preconditioning strategies applied to Cryer's sphere problem

Strategy	Ref.	Setup (s)	Average iterations	Average solve (s)
BP-EXACT	1	0.1	2.0	2.0
	2	0.5	2.0	34.9
	3	3.4	2.0	324.5
BP-INEXACT	1	0.1	8.5	1.2
	2	0.5	8.0	16.6
	3	3.4	8.0	165.0
BP-ILU-SD	1	0.1	26.9	0.1
	2	0.7	42.7	1.1
	3	5.3	87.4	18.5
BP-ILU-SM	1	0.1	28.7	0.1
	2	0.5	46.3	1.1
	3	3.9	94.6	21.8
BP-AMG-SD	1	0.1	26.4	0.1
	2	0.6	29.9	1.3
	3	4.3	30.0	11.0
BP-AMG-SM	1	0.1	21.0	0.1
	2	0.5	19.4	0.7
	3	4.0	19.9	6.4

Solver is FGMRES

While both of the ILU preconditioner variants are significantly cheaper than the inner-iteration strategies, both of them exhibit a significant growth in iterations as the mesh is refined. While these may therefore serve as the basis for useful strategies on small problems, they do not exhibit the desired scaling properties. The algebraic multigrid variants, in particular BP-AMG-SM, are the strongest performers. BP-AMG-SM exhibits mesh-independent behavior and the quickest solve times. From this example alone, however, both AMG strategies appear quite satisfactory.

A second study reveals the sharp difference between BP-AMG-SD and BP-AMG-SM. We again consider the Cryer sphere problem but examine the sensitivity of the preconditioners to the permeability of the medium. BICGSTAB is used for the simulations, as a flexible iteration is no longer required. Figure 2 illustrates the computational results. In the high permeability regime, both the preconditioners exhibit very similar performance properties and little growth in iteration count when moving from one refinement level to the next. When the permeability is large, the C term in the Schur-complement operator $S = C - B_2 A^{-1} B_1$ is dominant, and the particular approximation to the dense component is unimportant. In the low-permeability regime, however, the reverse situation holds and there is a marked difference in preconditioner behavior. While there is some mild growth in iteration count when moving toward the undrained regime, the mass matrix

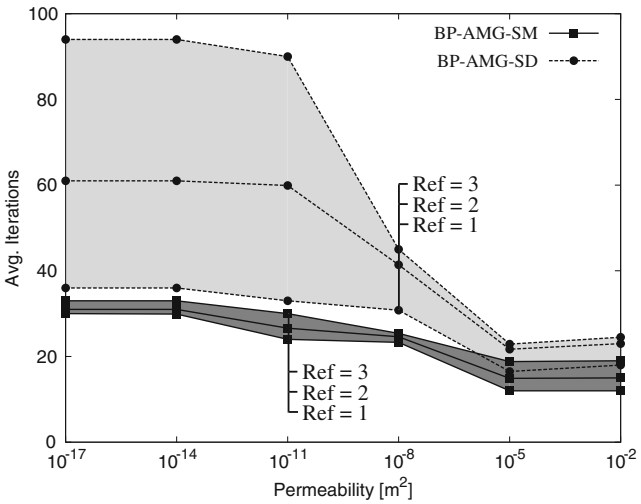

Fig. 2 Sensitivity of preconditioner performance to permeability. Solver is BICGSTAB

Table 4 Parallel performance of various preconditioning strategies applied to the linear consolidation problem

Strategy	Ref.	Processors	Setup (s)	Average iterations	Average solve (s)
BP-AMG-SD	3	1	4.3	19.9	13.5
	4	8	5.8	16.7	14.6
	5	64	9.0	18.2	47.6
	6	512	17.7	20.1	66.9
	7	4,096	44.0	18.8	70.1
BP-AMG-SM	3	1	3.4	12.0	6.9
	4	8	4.8	11.8	9.5
	5	64	6.6	13.2	33.4
	6	512	15.3	11.9	42.5
	7	4,096	39.8	11.8	56.1

Solver is BICGSTAB

approximation S_M continues to exhibit approximately constant iteration counts. In contrast, the S_D approximation appears quite poor and shows significant growth upon refinement. From this example, there appears little to recommend BP-MG-SD over BP-AMG-SM.

As a final linear example, we consider a weak scaling study with the Cryer sphere problem to see if the block preconditioner provides a viable parallel solution strategy. We again focus only on the BP-AMG-SD and BP-AMG-SM variants as they are the most competitive, and use BICGSTAB. Table 4 presents the computational results. Beginning with a single processor, each time the mesh is refined, the number of processors is increased by a factor of 8, leading to a constant number of elements per processor. The size of the various problems is reported in Table 2. The permeability of the medium (5.56×10^{-8} m^2) leads to a regime where either the S_M or S_D approximations work well, and both exhibit mesh-independent behavior. The mass matrix approximation is still the best, leading to the quickest convergence behavior. While the iteration counts are mesh-independent, the current implementation clearly exhibits some growth in wall-clock time—a factor of 8.1 for a problem 4,096 times as large. Additional optimization of the sub-preconditioning strategy can certainly improve this behavior further. Nevertheless, we see that solution times for a problem with 455 million unknowns are still very reasonable.

5.2 Unsaturated levee problem

The previous example focused on a fully saturated, linear elastic medium, so that the problem is linear in each timestep. This example features a variably saturated, elastoplastic soil for which both the solid and fluid flow responses are nonlinear. The block-preconditioned solver is therefore embedded within an outer Newton iteration to drive the coupled residual equations to zero.

This example considers a typical geotechnical application: analyzing the behavior of an L-shaped levee (Fig. 3) during a flood event. The levee rests on a rigid, impermeable layer so that all flow and deformation takes place within the levee itself and the foundation layer may be ignored. The levee is 6 m tall and 24 m wide along the generating cross section. The levee walls have a 1:1.5 slope ratio. The protected side on the levee lies to the south and east, while a flood plain is located to the north and west. The relevant material parameters are summarized in Table 1. The geometry of the levee has a mirror symmetry, so it is only necessary to discretize half of the domain.

A flooding event is simulated, described by a time-dependent water level $h(t)$ on the upstream face.

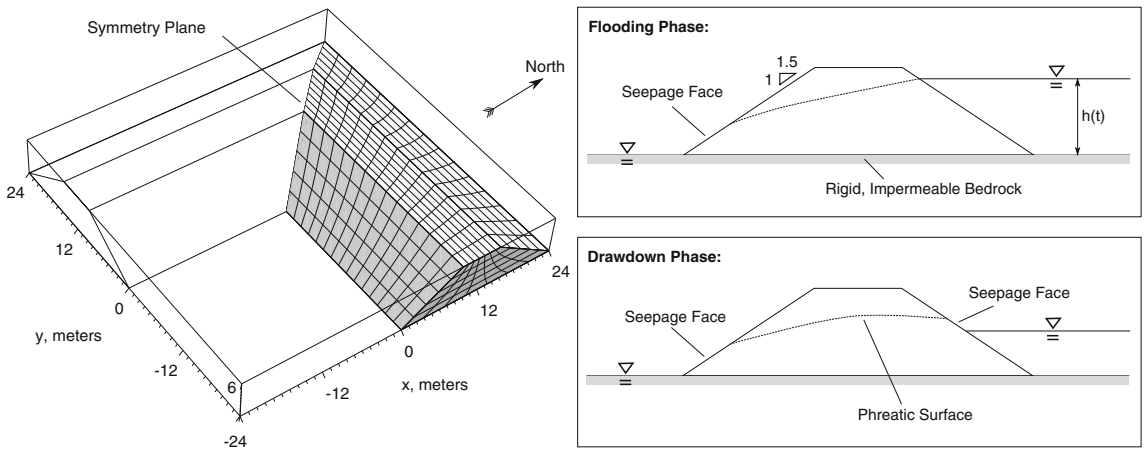

Fig. 3 Computational mesh and boundary conditions for the levee analysis

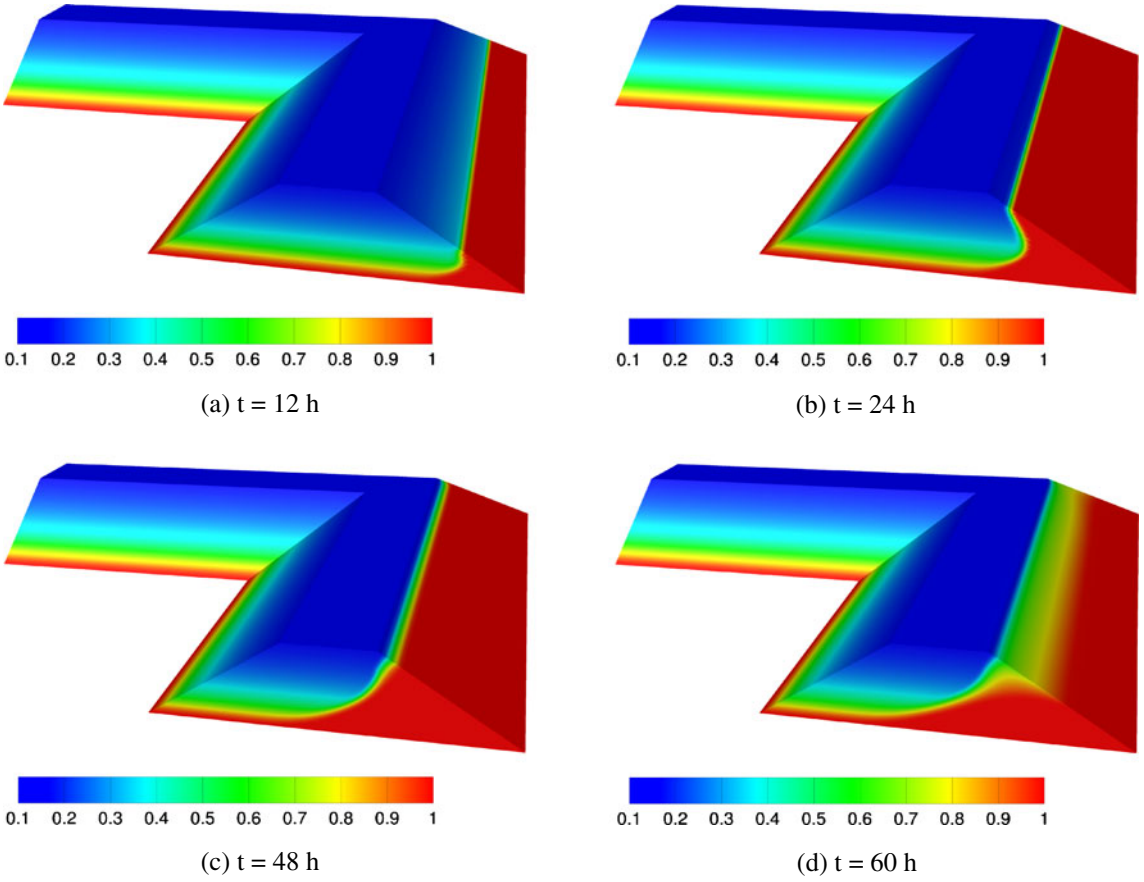

Fig. 4 Saturation within the levee at several time steps during flooding and drawdown

Figure 3 illustrates a typical cross-section through the levee during both the flooding and drawdown phases. At the start of flooding, the water level rises linearly over day 1 to a maximum height of 5 m. After this, the flood height remains constant for day 2, before linearly drawing down to zero again over day 3. On the protected side, it is assumed that the water table remains level with the ground surface for all time. A uniform timestep of 0.5 h is used throughout. Note that the evolution of the water level is rapid in comparison to the hydraulic conductivity of the medium. As a result, the phreatic surface is not necessarily in equilibrium with the external water level and seepage faces may develop at the boundaries. Contours of saturation within the levee at several timesteps are shown in Fig. 4. For this particular example, the material parameters of the soil are such that the levee remains stable during the entire loading and unloading process. In a weaker levee, however, the process of rapid drawdown may lead to shear localization and dramatic failure.

All simulations were run on 16 processors. The dimensions of the problem at various refinement levels are indicated in Table 5. We again focus only on the BP-AMG-SD and BP-AMG-SM preconditioner variants and use a BICGSTAB solver. Table 6 summarizes the various performance metrics as a function of refinement level. The average Newton iterations refer to the number of Newton iterations necessary to achieve convergence in a given timestep. This convergence behavior does not depend on the particular preconditioner chosen and so is the same for both strategies. We see that the Newton iteration is robust for this particular problem and is able to drive the nonlinear residual to zero with only a few updates. The average Krylov iterations refers to the average number of Krylov iterations per Newton

Table 5 Unsaturated levee problem dimensions at various refinement levels

Ref.	Total unknowns	Displacement	Pressure
1	2,236	1,677	559
2	14,900	11,175	3,725
3	108,388	81,291	27,097

Table 6 Performance of various preconditioning strategies applied to the unsaturated levee problem

Strategy	Ref.	Average Newton iterations	Average Krylov iterations	Setup (s)	Re-initialize (s)	Solve (s)
BP-AMG-SD	1	3.3	21.1	0.02	0.01	0.08
	2	3.6	15.1	0.10	0.08	0.30
	3	4.3	15.0	0.55	0.37	2.81
BP-AMG-SM	1	3.3	21.0	0.02	0.01	0.08
	2	3.6	15.1	0.10	0.08	0.28
	3	4.3	13.7	0.50	0.36	2.32

Solver is BICGSTAB

update, not per timestep. The actual number of Krylov iterations varies substantially with the timestep and Newton iteration, and so these metrics merely provide a representative mean value. The setup time again refers to the time taken to initialize the various sub-preconditioners. In the nonlinear problem, however, the Jacobian matrix entries change in each Newton update, and so the preconditioner is re-initialized at every Newton iteration. Much of the algebraic multigrid hierarchy can be re-used, however, so this re-initialization process is faster than the initial setup.

We see similar performance trends as in the previous examples. Both the strategies perform well, with BP-AMG-SM only slightly outperforming BP-AMG-SD. As mentioned earlier, for unsaturated problems, the medium becomes significantly more compressible due to the additional storage available through slight changes in saturation, and thus, the C component of the Schur-complement approximation tends to dominate. As a result, the performance here is less sensitive to the particular approximation chosen for $B_2 A^{-1} B_1$. Nevertheless, given that BP-AMG-SM outperforms BP-AMG-SD and continues to perform well in the saturated, low-permeability regime, the mass matrix approximation again appears to be the strategy of choice.

6 Conclusion

In this work, we have presented a block-preconditioned Newton–Krylov solver for two-field hydromechanical models. The key issue that has been addressed is how to efficiently precondition the sub-problems that appear in the block-decomposition—in particular, how to precondition the Schur-complement operator for the coupled hydromechanical model. The proposed BP-AMG-SM preconditioner demonstrates mesh-independent convergence, good parallel efficiency, and insensitivity to the material parameters of the medium. The method outperforms a number of alternative strategies.

Current work is focused on preconditioning alternative discretizations of hydromechanical models. In particular, ideas from this work can be extended to cover three-field, locally conservative formulations in which the displacement, pressure, and seepage velocity are chosen as primary variables.

Acknowledgements This work was performed under the auspices of the US Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. The first author is grateful for the support of the Lawrence Postdoctoral Fellowship Program. The second author was supported by the US National Science Foundation under Contract Numbers CMMI-0824440 and CMMI-0936421 to Stanford University.

References

1. Gawin, D., Baggio, P., Schrefler, B.A.: Coupled heat, water and gas flow in deformable porous media. Int. J. Numer. Methods Fluids **20**, 969–987 (1995)
2. Ehlers, W., Graf, T., Ammann, M.: Deformation and localization analysis of partially saturated soil. Comput. Methods Appl. Mech. Eng. **193**(27–29), 2885–2910 (2004)
3. Young, Y.L., White, J.A., Xiao, H., Borja, R.I.: Tsunami-induced liquefaction failure of coastal slopes. Acta Geotech. **4**, 17–34 (2009)
4. Borja, R.I., White, J.A.: Continuum deformation and stability analyses of a steep hillside slope under rainfall infiltration. Acta Geotech. 1–14 (2010)
5. Ferronato, M., Bergamaschi, L., Gambolati, G.: Performance and robustness of block constraint preconditioners in finite element coupled consolidation problems. Int. J. Numer. Methods Eng. **81**(3), 381–402 (2010)
6. Wan, J.: Stabilized Finite Element Methods for Coupled Geomechanics and Multiphase Flow. Ph.D. thesis, Stanford University (2002)
7. Minkoff, S.E., Stone, C.M., Bryant, S., Peszynska, M., Wheeler, M.F.: Coupled fluid flow and geomechanical deformation modeling. J. Pet. Sci. Eng. **38**(1–2), 37–56 (2003)
8. Jha, B., Juanes, R.: A locally conservative finite element framework for the simulation of coupled flow and reservoir geomechanics. Acta Geotech. **2**(3), 139–153 (2007)
9. Hayashi, K., Willis-Richards, J., Hopkirk, R.J., Niibori, Y.: Numerical models of HDR geothermal reservoirs—a review of current thinking and progress. Geothermics **28**(4–5), 507–518 (1999)
10. Johnson, J.W., Nitao, J.J., Morris, J.P.: Reactive transport modeling of cap rock integrity during natural and engineered CO_2 storage. In: Benson, S. (ed.) CO_2 Capture Project Summary, vol. 2. Elsevier, Amsterdam (2004)

11. Rutqvist, J., Birkholzer, J.T., Tsang, C.F.: Coupled reservoir–geomechanical analysis of the potential for tensile and shear failure associated with CO2 injection in multilayered reservoir–caprock systems. Int. J. Rock Mech. Min. Sci. **45**(2), 132–143 (2008)
12. Morris, J.P., Detwiler, R.L., Friedmann, S.J., Vorobiev, O.Y., Hao, Y.: The large-scale effects of multiple CO2 injection sites on formation stability. Energy Procedia **1**(1), 1831–1837 (2009)
13. Bramble, J.H., Pasciak, J.E.: A preconditioning technique for indefinite systems resulting from mixed approximation of elliptic problems. Math. Comput. **50**(181), 1–17 (1988)
14. Elman, H., Howle, V.E., Shadid, J., Shuttleworth, R., Tuminaro, R.: A taxonomy and comparison of parallel block multilevel preconditioners for the incompressible Navier–Stokes equations. J. Comput. Phys. **227**(3), 1790–1808 (2008)
15. May, D.A., Moresi, L.: Preconditioned iterative methods for Stokes flow problems arising in computational geodynamics. Phys. Earth Planet. Inter. **171**(1–4), 33–47 (2008)
16. Burstedde, C., Ghattas, O., Stadler, G., Tu, T., Wilcox, L.C.: Parallel scalable adjoint-based adaptive solution of variable-viscosity Stokes flow problems. Comput. Methods Appl. Mech. Eng. **198**(21–26), 1691–1700 (2009)
17. Toh, K.C., Phoon, K.K., Chan, S.H.: Block preconditioners for symmetric indefinite linear systems. Int. J. Numer. Methods Eng. **60**, 1361–1381 (2004)
18. Biot, M.A.: General theory of three-dimensional consolidation. J. Appl. Phys. **12**(2), 155–164 (1941)
19. Biot, M.A.: Theory of elasticity and consolidation for a porous anisotropic solid. J. Appl. Phys. **26**(2), 182–185 (1955)
20. Borja, R.I.: On the mechanical energy and effective stress in saturated and unsaturated porous continua. Int. J. Solids Struct. **43**(6), 1764–1786 (2006)
21. Terzaghi, K.: Theoretical Soil Mechanics. Wiley, New York (1943)
22. Bishop, A.W.: The principle of effective stress. Tekn. Ukebl. **39**, 859–863 (1959)
23. Skempton, A.W.: Effective stress in soils, concrete and rocks. In: Pore Pressure and Suction in Soils, pp. 4–16. Butterworths, London (1961)
24. Nur, A., Byerlee, J.D.: An exact effective stress law for elastic deformation of rock with fluids. J. Geophys. Res. **76**, 6414–6419 (1971)
25. Borja, R.I., Koliji, A.: On the effective stress in unsaturated porous continua with double porosity. J. Mech. Phys. Solids (2009). doi:10.1016/j.jmps.2009.04.014
26. van Genuchten, M.T.: A closed-form equation for predicting the hydraulic conductivity of unsaturated soils. Soil Sci. Soc. Am. J. **44**(5), 892–898 (1980)
27. Cryer, C.W.: A comparison of the three-dimensional consolidation theories of Biot and Terzaghi. Q. J. Mech. Appl. Math. **16**(4), 401–412 (1963)
28. Brezzi, F.: On the existence, uniqueness and approximation of saddle-point problems arising from Lagrangian multipliers. RAIRO Anal. Numer. **8**, 129–151 (1974)
29. Brezzi, F.: A discourse on the stability conditions for mixed finite element formulations. Comput. Methods. Appl. Mech. Eng. **82**(1–3), 27–57 (1990)
30. Arnold, D.N.: Mixed finite element methods for elliptic problems. Comput. Methods Appl. Mech. Eng. **82**, 281–300 (1990)
31. Murad, M.A., Loula, A.F.D.: On stability and convergence of finite element approximations of Biot's consolidation problem. Int. J. Numer. Methods Eng. **37**, 645–667 (1994)
32. White, J.A., Borja, R.I.: Stabilized low-order finite elements for coupled solid-deformation/fluid-diffusion and their application to fault zone transients. Comput. Methods Appl. Mech. Eng. **197**(49–50), 4353–4366 (2008)
33. Pastor, M., Li, T., Liu, X., Zienkiewicz, O.C., Quecedo, M.: A fractional step algorithm allowing equal order of interpolation for coupled analysis of saturated soil problems. Mech. Cohes.-Frict. Mater. **5**(7), 511–534 (2000)
34. Truty, A., Zimmermann, T.: Stabilized mixed finite element formulations for materially nonlinear partially saturated two-phase media. Comput. Methods Appl. Mech. Eng. **195**, 1517–1546 (2006)
35. Dohrmann, C.R., Bochev, P.B.: A stabilized finite element method for the Stokes problem based on polynomial pressure projections. Int. J. Numer. Methods Fluids **46**, 183–201 (2004)
36. Bochev, P.B., Dohrmann, C.R.: A computational study of stabilized, low-order C^0 finite element approximations of Darcy equations. Comput. Mech. **38**, 323–333 (2006)
37. Burman, E.: Pressure projection stabilizations for Galerkin approximations of Stokes' and Darcy's problem. Numer. Methods Partial Differ. Equ. **24**(1), 127–143 (2007)
38. White, J.A.: Stabilized Finite Element Methods for Coupled Flow and Geomechanics. Ph.D. thesis, Stanford University, Stanford, CA (2009)
39. Benzi, M., Golub, G.H., Liesen, J: Numerical solution of saddle point problems. Acta Numer. **14**, 1–137 (2005)
40. Verfürth, R.: Error estimates for a mixed finite element approximation of the Stokes equations. RAIRO. Anal. Numér. **18**(2), 175–182 (1984)
41. Elman, H.C., Silvester, D.J., Wathen, A.J.: Iterative methods for problems in computational fluid dynamics. In: Iterative Methods in Scientific Computing, p. 271 (1997)
42. Bangerth, W., Hartmann, R., Kanschat, G.: Deal.II—a general purpose object oriented finite element library. ACM Trans. Math. Softw. **33**(4), 24 (2007)
43. Burstedde, C., Wilcox, L.C., Ghattas, O.: `p4est`: scalable algorithms for parallel adaptive mesh refinement on forests of octrees. SIAM J. Sci. Comput. (in press, 2011)
44. Bangerth, W., Burstedde, C., Heister, T., Kronbichler, M.: Algorithms and data structures for massively parallel generic adaptive finite element codes. ACM Trans. Math. Softw. (submitted, 2011)
45. Heroux, M.A., Bartlett, R.A., Howle, V.E., Hoekstra, R.J., Hu, J.J., Kolda, T.G., Lehoucq, R.B., Long, K.R., Pawlowski, R.P., Phipps, E.T., et al.: An overview of the Trilinos project. ACM Trans. Math. Softw. **31**(3), 397–423 (2005)
46. Sala, M., Heroux, M.: Robust algebraic preconditioners with IFPACK 3.0. Technical Report SAND-0662, Sandia National Laboratories (2005)
47. Gee, M.W., Siefert, C.M., Hu, J.J., Tuminaro, R.S., Sala, M.G.: ML 5.0 smoothed aggregation user's guide. Technical Report SAND2006-2649, Sandia National Laboratories (2006)
48. Eisenstat, S.C., Walker, H.F.: Choosing the forcing terms in an inexact Newton method. SIAM J. Sci. Comput. **17**, 16–32 (1996)
49. Verruijt, A.: Theory of Consolidation. In: An Introduction to Soil Dynamics, pp. 65–90 (2010)