

MONOLITHIC MULTIGRID FOR A REDUCED-QUADRATURE DISCRETIZATION OF POROELASTICITY*

JAMES H. ADLER[†], YUNHUI HE[‡], XIAOZHE HU[†], SCOTT MACLACHLAN[§],
AND PETER OHM[¶]

Abstract. Advanced finite-element discretizations and preconditioners for models of poroelasticity have attracted significant attention in recent years. The equations of poroelasticity offer significant challenges in both areas, due to the potentially strong coupling between unknowns in the system, saddle-point structure, and the need to account for wide ranges of parameter values, including limiting behavior such as incompressible elasticity. This paper was motivated by an attempt to develop monolithic multigrid preconditioners for the discretization developed in [C. Rodrigo et al., *Comput. Methods App. Mech. Engrg*, 341 (2018), pp. 467–484]; we show here why this is a difficult task and, as a result, we modify the discretization in [Rodrigo et al.] through the use of a reduced-quadrature approximation, yielding a more “solver-friendly” discretization. Local Fourier analysis is used to optimize parameters in the resulting monolithic multigrid method, allowing a fair comparison between the performance and costs of methods based on Vanka and Braess–Sarazin relaxation. Numerical results are presented to validate the local Fourier analysis predictions and demonstrate efficiency of the algorithms. Finally, a comparison to existing block-factorization preconditioners is also given.

Key words. Biot poroelasticity, reduced quadrature, finite elements, monolithic multigrid, local Fourier analysis

AMS subject classifications. 65F08, 65F10, 65M55, 65M60, 65N22, 76S05

DOI. 10.1137/21M1429072

1. Introduction. Poroelasticity describes a number of processes modeled by flows in deformable porous media, which are of interest in geoscience, biomedical science, and engineering. In this paper, we consider Biot’s model for linear poroelasticity [6, 7], a coupled, multiphysics system of partial differential equations (PDEs). There are many challenges in developing both discretizations and fast and robust solvers for these equations. For the discretization, using mixed finite elements, it is necessary to carefully choose approximation spaces in order to avoid spurious oscillations in the pressure field as well as to achieve robustness to variations in the PDE parameters, particularly in extreme limits, such as incompressibility. After discretization, the resulting linear system is of saddle-point type, requiring special solvers to deal with the indefiniteness and with the usual ill-conditioning of the

*Received by the editors June 28, 2021; accepted for publication (in revised form) March 14, 2022; published electronically June 24, 2022.

<https://doi.org/10.1137/21M1429072>

Funding: The work of the first, third, and fifth authors was partially supported by National Science Foundation grant DMS-1620063. The work of the fourth author was partially supported by an NSERC Discovery Grant. The work of the fifth author was supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Applied Mathematics program.

[†]Department of Mathematics, Tufts University, Medford, MA 02155 USA (James.Adler@tufts.edu, Xiaozhe.Hu@tufts.edu).

[‡]Department of Computer Science, University of British Columbia, Vancouver V6T 1Z4, BC, Canada, and Department of Mathematics and Statistics, Memorial University of Newfoundland, St. John’s A1C 5S7, NL, Canada (yunhui.he@ubc.ca).

[§]Department of Mathematics and Statistics, Memorial University of Newfoundland, St. John’s A1C 5S7, NL, Canada (smaclachlan@mun.ca).

[¶]Sandia National Laboratories, Albuquerque, NM 87185 USA (pohm@sandia.gov).

discretized system, and to achieve similar robustness with respect to the physical parameters.

Many different types of discretizations exist for the various formulations of Biot's model. For instance, a finite-volume method on a staggered grid is provided for the three-dimensional Biot poroelastic system in [48]. For the two-field formulation, where displacement and pressure are the unknowns, stable Taylor–Hood elements are used in [45, 46, 47]. In [51], a MINI element and a stabilized P1-P1 finite-element discretization are presented, and a stabilization term is introduced to remove non-physical oscillations, leading to monotone behavior of the stabilized schemes. More recently, a weak Galerkin finite-element method is proposed on general shape-regular polytopal meshes, which demonstrates the robustness of the proposed weak Galerkin discretization [35]. For three-field formulations, where displacement, pressure, and the Darcy velocity are the unknowns, a nonconforming finite-element approach for the three-field formulation, using Crouzeix–Raviart finite elements for the displacements, lowest-order Raviart–Thomas–Nédélec elements for the Darcy velocity, and a piecewise constant approximation for the pressure, is considered in [63] for the two-dimensional case on rectangular grids. It is extended to general cases in [36], where a mass-lumping technique is introduced for the Raviart–Thomas–Nédélec elements to eliminate the Darcy velocity, reducing the computational cost. A family of parameter-robust schemes is found in [34] and a general theory for the error analysis is introduced. More recently, hybridization schemes are developed in [26, 49]. For a four-field formulation, with the stress tensor, fluid flux, displacement, and pore pressure as unknowns, stable discretizations are developed in [40, 64].

In this work, we consider a stabilized finite-element method based on the popular P1-RT0-P0 discretization of the three-field formulation developed in [53], where face bubble functions are used to enrich the P1 space for the displacements. A perturbation of the bilinear form allows for local elimination of the bubble functions, leading to the same number of degrees of freedom (DoFs) as the P1-RT0-P0 discretization. This type of discretization is appealing, as it leads to a minimally sized system of equations, yet lends itself to robust linear solvers, independent of discretization and physical parameters. While we do not consider the perturbation in this paper, the main goal here is to extend this bubble-enriched discretization to make it amenable to efficient solvers, such as the monolithic multigrid solvers described below.

After discretization, large linear systems of equations must be solved to compute the finite-element approximation to the solution of the poroelasticity equations. This requires development of specialized preconditioners, and both block preconditioning and monolithic multigrid methods have been successfully applied, especially for Biot's model. For instance, robust block preconditioners are studied for the two-field formulation in [3, 15, 16, 21, 61] and for the three-field formulation in [2, 15, 22, 34]. A multigrid method using alternating line Gauss–Seidel relaxation for the three-dimensional Biot poroelasticity system is presented in [48], which focuses on the study of the grid-transfer operators in the multigrid method. For the quasi-static Biot model, pointwise and linewise box Gauss–Seidel relaxation are investigated in [14], where local Fourier analysis (LFA) is used to help analyze and predict performance of the algorithms. In [41], an Uzawa relaxation is employed and analyzed using LFA. The fixed-stress split method is used as a relaxation scheme for the two-field formulation of Biot's consolidation model in [29], where again LFA is applied to study the convergence of the multigrid method. Similarly, a new version of the fixed-stress splitting method [9] is proposed for solving coupled flow and geomechanics in porous media, modeled by

a two-field formulation of Biot's equations. Finally, multigrid waveform relaxation based on a pointwise Vanka relaxation method is proposed for solving a collocated finite-difference discretization of the linear Biot model in [25].

Despite the work mentioned above, applications of monolithic multigrid for the discretized systems of Biot's model are rare. In particular, given the scalable preconditioning results shown in [2] for the discretization from [53], a natural question to ask is whether monolithic multigrid can compete with efficient block preconditioners. As discussed below, initial work for this paper focused on the extension of typical monolithic multigrid relaxation schemes, known as Braess–Sarazin [10] and Vanka [60] relaxation, to the three-field discretization from [53]. While direct extensions of these methods lead to efficient preconditioners for some physical parameters, we found that they did not extend effectively to the limit of an incompressible material. Following [54, 55, 56], we recognize this as an inherent consequence of the fact that the bubble-enriched P1 space does not admit a local basis for the space of divergence-free functions and, as such, standard multigrid approaches for the elasticity block are not parameter-robust. To overcome this difficulty, we modify the discretization from [53] to make use of the reduced-quadrature approach [8, 17, 43, 54, 55, 56], which replaces exact integration of the divergence terms with that of an L^2 projection. Such a modification has been adopted for poroelasticity problems in [65] in order to handle locking issues when $\lambda \rightarrow \infty$. Here, we find that it also provides a “solver-friendly” discretization. One of our contributions in this work is to show that using the reduced-quadrature approach still results in a well-posed discretization, which is parameter-robust, and does not lose accuracy in comparison to the discretization of [53].

Having constructed the reduced-quadrature discretization, the remainder of this paper focuses on the development and analysis of optimal monolithic multigrid preconditioners for it. In particular, we apply LFA [59, 62] to the components of the multigrid method in order to optimize parameters within the commonly used Braess–Sarazin and Vanka relaxation schemes. In recent years, LFA has been widely used for this purpose in many contexts; for systems of PDEs, such as we consider here, it has been applied to discretizations of the Stokes equations [28, 31, 32, 42, 52] and, in a more limited manner, to discretizations of poroelasticity [41]. Numerical results confirm the accuracy of the LFA predictions.

In what follows, we address how the incompressibility constraint associated with the elasticity block of the coupled system affects the convergence of our proposed multigrid algorithm. In particular, we show that the ideas of reduced-quadrature discretization and divergence-free interpolation, originally proposed and analyzed for the incompressible elasticity subproblem, can be extended to the fully coupled Biot model. We show that the modified discretization remains well-posed, and that we are able to develop a robust monolithic multigrid approach for the resulting three-field formulation. Specifically, this paper is organized as follows. In section 2, we introduce the stabilized finite-element discretization provided in [53] for the three-field formulation of Biot's model, as well as the reduced-quadrature discretization, for which proofs of well-posedness and error estimates are given. In section 3, we review monolithic multigrid, with focus on both the choice of relaxation scheme for solving the discretized system and the use of divergence-preserving interpolation operators to achieve robustness in the nearly incompressible case. LFA for this discretization is considered in section 4. In section 5, numerical results are presented to show the efficiency of the proposed solvers, and comparisons are given between existing block preconditioning approaches and the monolithic multigrid methods proposed here. Finally, conclusions and remarks are drawn in section 6.

2. Biot’s three-field formulation and its discretization. The mathematical model of the *three-field formulation* of the consolidation process is described by the following system of PDEs in a domain $\Omega \subset \mathbb{R}^d, d = 2, 3$, with sufficiently smooth boundary, $\Gamma = \partial\Omega$:

$$(2.1) \quad -\operatorname{div}(2\mu\varepsilon(\mathbf{u})) - \lambda\nabla(\operatorname{div}\mathbf{u}) + \alpha\nabla p = \rho\mathbf{g},$$

$$(2.2) \quad \mathbf{K}^{-1}\mu_f\mathbf{w} + \nabla p = \rho_f\mathbf{g},$$

$$(2.3) \quad \frac{\partial}{\partial t} \left(\frac{1}{M}p + \alpha\operatorname{div}\mathbf{u} \right) + \operatorname{div}\mathbf{w} = f.$$

Here, μ_f is the viscosity of the fluid, M is the Biot modulus, ρ and ρ_f are the bulk density and fluid density, respectively, and $\alpha = 1 - \frac{K_b}{K_s}$ is the Biot–Willis constant, with K_b and K_s denoting the drained and the solid-phase bulk moduli, respectively. The absolute permeability tensor is given by \mathbf{K} which is symmetric positive definite. The strain tensor is denoted by $\varepsilon(\mathbf{u}) = \frac{1}{2}(\nabla\mathbf{u} + \nabla\mathbf{u}^\top)$. The unknown functions are the displacement vector \mathbf{u} , the pore pressure p , and the percolation velocity of the fluid, or Darcy velocity, relative to the soil, \mathbf{w} . The vector-valued function \mathbf{g} represents the gravitational force. Finally, $\mu = \frac{E}{2+2\nu}$ and $\lambda = \frac{E\nu}{(1-2\nu)(1+\nu)}$ are the Lamé coefficients where ν is the Poisson ratio and E is Young’s modulus. As $\nu \rightarrow 0.5$, we have $\lambda \rightarrow \infty$, the incompressible limit that causes difficulties in numerical simulations. Other limits that cause numerical difficulties are when the permeability, $\mathbf{K} \rightarrow \mathbf{0}$, and (2.2) is dominated by its first term or, when discretized, the time step goes to zero and (2.3) is dominated by the term from time stepping. Finally, this system is subject to boundary conditions of various forms. One typical example is

$$\begin{aligned} p = 0 \quad \text{for } x \in \bar{\Gamma}_t, \quad 2\mu\varepsilon(\mathbf{u})\mathbf{n} + \lambda\operatorname{div}(\mathbf{u})\mathbf{n} = \mathbf{0} \quad \text{for } x \in \Gamma_t, \\ \mathbf{u} = \mathbf{0} \quad \text{for } x \in \bar{\Gamma}_c, \quad \frac{\partial p}{\partial \mathbf{n}} = 0 \quad \text{for } x \in \Gamma_c, \end{aligned}$$

where \mathbf{n} is the outward unit normal to the boundary, $\bar{\Gamma} = \bar{\Gamma}_t \cup \bar{\Gamma}_c$, with Γ_t and Γ_c being open (with respect to Γ) subsets of Γ with nonzero measure. Appropriate initial conditions for the pressure and displacement (more precisely, for $\operatorname{div}\mathbf{u}$) are also needed.

2.1. Finite-element discretization. Following [53], we consider a variational problem such that for each $t \in (0, T]$, $(\mathbf{u}(t), p(t), \mathbf{w}(t)) \in \mathbf{V} \times Q \times \mathbf{W}$, with

$$\begin{aligned} \mathbf{V} &= \{\mathbf{u} \in \mathbf{H}^1(\Omega) \mid \mathbf{u}|_{\bar{\Gamma}_c} = \mathbf{0}\}, & Q &= L^2(\Omega), \\ \mathbf{W} &= \{\mathbf{w} \in \mathbf{H}(\operatorname{div}, \Omega) \mid (\mathbf{w} \cdot \mathbf{n})|_{\Gamma_c} = 0\}, \end{aligned}$$

where $\mathbf{H}^1(\Omega)$ is the space of square integrable vector-valued functions whose first derivatives are also square integrable, and $\mathbf{H}(\operatorname{div}, \Omega)$ contains the square integrable vector-valued functions with square integrable divergence.

Using backward Euler as a time discretization on a time interval $(0, T]$ with constant time step size τ , the discrete variational form for Biot’s three-field consolidation model, (2.1)–(2.3), is written as, Find $(\mathbf{u}_h^m, p_h^m, \mathbf{w}_h^m) \in \mathbf{V}_h \times Q_h \times \mathbf{W}_h$ such that

$$(2.4) \quad a(\mathbf{u}_h^m, \mathbf{v}_h) - (\alpha p_h^m, \operatorname{div}\mathbf{v}_h) = (\rho\mathbf{g}, \mathbf{v}_h) \quad \forall \mathbf{v}_h \in \mathbf{V}_h,$$

$$(2.5) \quad \tau(\mathbf{K}^{-1}\mu_f\mathbf{w}_h^m, \mathbf{r}_h) - \tau(p_h^m, \operatorname{div}\mathbf{r}_h) = \tau(\rho_f\mathbf{g}, \mathbf{r}_h) \quad \forall \mathbf{r}_h \in \mathbf{W}_h,$$

$$(2.6) \quad -\left(\frac{1}{M}p_h^m, q_h\right) - (\alpha\operatorname{div}\mathbf{u}_h^m, q_h) - \tau(\operatorname{div}\mathbf{w}_h^m, q_h) = -(\hat{f}, q_h) \quad \forall q_h \in Q_h,$$

where (\cdot, \cdot) denotes the standard $L^2(\Omega)$ inner product. Here, $(\mathbf{u}_h^m, p_h^m, \mathbf{w}_h^m)$ is an approximation to $(\mathbf{u}(\cdot, t_m), p(\cdot, t_m), \mathbf{w}(\cdot, t_m))$, at time $t_m = m\tau$, $m = 1, 2, \dots$, $(\hat{f}, q_h) = \tau(f, q_h) + (\frac{1}{M}p_h^{m-1}, q_h) + (\alpha \operatorname{div} \mathbf{u}_h^{m-1}, q_h)$ and $a(\mathbf{u}, \mathbf{v}) = 2\mu(\varepsilon(\mathbf{u}), \varepsilon(\mathbf{v})) + \lambda(\operatorname{div} \mathbf{u}, \operatorname{div} \mathbf{v})$ is the usual weak form for linear elasticity. Note that (2.5) has been scaled by τ and (2.6) has been scaled by -1 to make the system symmetric.

For finite-element spaces, we consider linear elements (P1), enriched with bubble functions on faces for $\mathbf{V}_h \subset \mathbf{V}$. These face-normal bubble functions are quadratic in two dimensions and cubic in three dimensions. Their DoFs are defined as the integrated normal displacement across the associated faces. This space is covered in depth in Chapter 2.1 of [30]. We choose $Q_h \subset Q$ as the piecewise constant space (P0) for the pressure, and $\mathbf{W}_h \subset \mathbf{W}$ as the standard lowest-order Raviart–Thomas space (RT0) for the Darcy velocity. It has been shown that this discretization is a stable finite-element approximation; see [53].

Finally, this discrete variational form can be represented in block matrix form as

$$(2.7) \quad \mathcal{A} \begin{pmatrix} \mathbf{u} \\ \mathbf{w} \\ p \end{pmatrix} = \mathbf{b} \quad \text{with} \quad \mathcal{A} = \begin{pmatrix} A_{\mathbf{u}} & 0 & \alpha B_{\mathbf{u}}^\top \\ 0 & \tau M_{\mathbf{w}} & \tau B_{\mathbf{w}}^\top \\ \alpha B_{\mathbf{u}} & \tau B_{\mathbf{w}} & -\frac{1}{M} M_p \end{pmatrix}.$$

The blocks in the matrix \mathcal{A} correspond to the following bilinear forms:

$$\begin{aligned} a(\mathbf{u}_h, \mathbf{v}_h) &\rightarrow A_{\mathbf{u}}, & -(\operatorname{div} \mathbf{u}_h, q_h) &\rightarrow B_{\mathbf{u}}, & -(\operatorname{div} \mathbf{w}_h, q_h) &\rightarrow B_{\mathbf{w}}, \\ (\mathbf{K}^{-1} \mu_f \mathbf{w}_h, \mathbf{r}_h) &\rightarrow M_{\mathbf{w}}, & (p_h, q_h) &\rightarrow M_p. \end{aligned}$$

2.2. Solver incompatibility. While the above discretization is well-posed and, as shown in [53], is robust to variations in the physical and discretization parameters, solving the resulting linear system in a similarly parameter-robust manner is not straightforward. A block preconditioning framework was proposed in [2] for the solution of the linear system and the proposed approaches were proven to be parameter-robust under the assumption that each diagonal block of the preconditioner can be solved in a parameter-robust manner. While [2] contains a detailed parameter study, the primary measure of convergence there was in outer iterations of FGMRES, where the inner iterations (to approximate solves with the diagonal blocks of the block preconditioners) were done to fixed tolerances with AMG-preconditioned GMRES. As numerical results presented below in section 5 will show, while the outer iterations reported in [2] are robust to the physical parameters (in particular, the incompressible limit), the inner iterations are not.

In preliminary investigations for this paper, similar behavior was seen for the monolithic multigrid methods detailed below. There are several common relaxation schemes considered when applying monolithic multigrid to block-structured saddle-point problems, such as the system in (2.7), which will be described in more detail below. Braess–Sarazin approaches use approximations to the block factorization of \mathcal{A} as relaxation schemes. *Exact* Braess–Sarazin relaxation (BSR) is based on exact solution of the approximate Schur complement(s) in such a factorization, while *inexact* Braess–Sarazin methods also introduce an approximation to the Schur complement(s). An alternative approach is to use Vanka relaxation schemes (see subsection 3.2.1), which are block overlapping Schwarz methods, with small blocks chosen to reflect the saddle-point structure of the system. Table 2.1 shows that while exact BSR is effective in a parameter-independent manner, convergence suffers for both inexact Braess–Sarazin and Vanka relaxation schemes.

TABLE 2.1

Measured convergence factors for monolithic multigrid applied to (2.7), with $\mathbf{K} = k\mathbf{I}$, $k = 10^{-6}$, and varying ν , for a uniform mesh $h = 1/64$ of the unit square and a time step size of $\tau = 1.0$.

ν	0.0	0.2	0.4	0.45	0.49	0.499
Exact BSR	0.067	0.067	0.067	0.067	0.067	0.067
Inexact BSR	0.440	0.471	0.586	0.659	0.790	0.968
Vanka	0.515	0.513	0.589	0.659	0.794	0.970

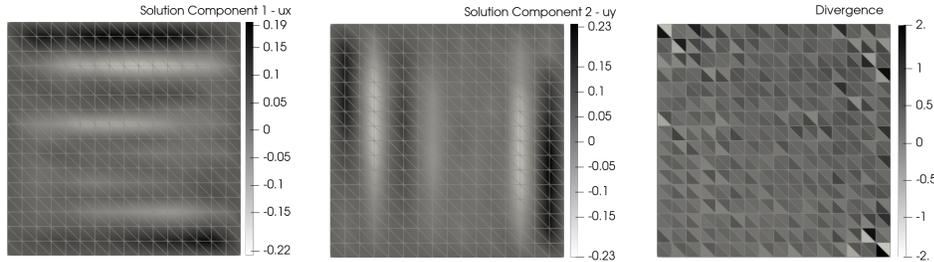


FIG. 2.1. Error and divergence of the error for displacement after 40 cycles of two-level multigrid with Vanka relaxation, for $\nu = 0.49$, applied to a problem with zero right-hand side and random initial guess. The divergence of the error illustrates neighboring element pairs with divergence of similar magnitude but opposite sign, indicating a globally supported divergence-free null-space.

The degradation in performance from exact to inexact BSR as $\nu \rightarrow 0.5$ in Table 2.1 was carefully studied. For both Braess–Sarazin variants considered, we took a Schur complement onto the displacement DoFs, and invested significant effort into constructing relaxation schemes for that Schur complement that would lead to a robust inexact Braess–Sarazin variant. The primary source of the problem became clear when looking at the dominant errors in the displacements after running two-grid cycles with either the inexact Braess–Sarazin or Vanka relaxation, visualized for the Vanka case in Figure 2.1. In essence, this error reflects a *globally supported* divergence-free null-space that is difficult to eliminate using local relaxation schemes. As we next show, this arises from the exact evaluation of the $(\operatorname{div} \mathbf{u}, \operatorname{div} \mathbf{v})$ term within the discretization, resulting in a discretization that is inherently not “solver-friendly,” due to the lack of a local basis for the space of (nearly) divergence-free functions. To address this, we modify the discretization using a reduced-quadrature approach [8, 43], as suggested in [65] for poroelasticity problems.

2.3. Reduced quadrature. As recognized in [54, 55, 56], the nonlocal nature of the basis for the divergence-free spaces arises from the direct evaluation of the $(\operatorname{div} \mathbf{u}, \operatorname{div} \mathbf{v})$ term in the weak form, since the discrete divergence of the displacement space is not a subset of the piecewise constant pressure space. To avoid this, we implement a reduced integration approach [8, 17, 43, 65] and replace $(\operatorname{div} \mathbf{u}, \operatorname{div} \mathbf{v})$ with $(P_{Q_h} \operatorname{div} \mathbf{u}, P_{Q_h} \operatorname{div} \mathbf{v})$, where P_{Q_h} is the L^2 projection from Q onto Q_h , the space of piecewise constant functions. With this reduced integration approach, a basis for the space of divergence-free functions is readily constructed with local support, allowing local relaxation schemes to be effective for divergence-free components.

To illustrate this further, consider that the discretization for displacements has a total of $2N_v + N_e$ DoFs, where N_v is the number of vertices in the mesh, and N_e is the number of edges. By direct computation, around each vertex in the mesh, we can introduce a local basis of three divergence-free functions, shown in Figure 2.2,

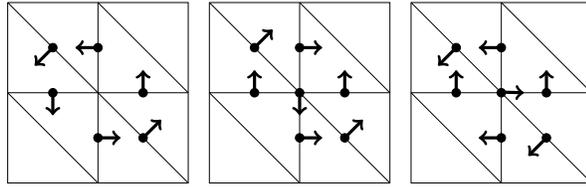


FIG. 2.2. The local divergence-free bases supported by the discretization.

resulting in $3N_v$ divergence-free basis functions. The reduced-quadrature approach constrains $\operatorname{div} \mathbf{u}$ to be in the piecewise constant pressure space, thus there are $N_T - 1$ divergence-free constraints, where N_T is the number of triangular elements. Now, subtracting the number of divergence-free constraints from the total DoFs, $(2N_v + N_e) - (N_T - 1) = 3N_v$, we get the number of divergence-free basis functions. Thus, the reduced-quadrature approach fully supports the divergence-free functions through the local basis functions in Figure 2.2.

Therefore, we define the bilinear form for the reduced-quadrature discretization as

$$a^{\text{RQ}}(\mathbf{u}, \mathbf{v}) := 2\mu (\varepsilon(\mathbf{u}), \varepsilon(\mathbf{v})) + \lambda(P_{Q_h} \operatorname{div} \mathbf{u}, P_{Q_h} \operatorname{div} \mathbf{v}).$$

Using this, the poroelastic system is then written as

$$(2.8) \quad \mathcal{A}^{\text{RQ}} = \begin{pmatrix} A_{\mathbf{u}}^{\text{RQ}} & 0 & \alpha B_{\mathbf{u}}^{\top} \\ 0 & \tau M_{\mathbf{w}} & \tau B_{\mathbf{w}}^{\top} \\ \alpha B_{\mathbf{u}} & \tau B_{\mathbf{w}} & -\frac{1}{M} M_p \end{pmatrix},$$

where $a^{\text{RQ}}(\mathbf{u}_h, \mathbf{v}_h) \rightarrow A_{\mathbf{u}}^{\text{RQ}}$. We next show that this reduced-quadrature approach remains well-posed independent of the physical and discretization parameters. To do this, we first introduce the following lemma concerning the Stokes inf-sup condition.

LEMMA 2.1. *Let the pair of finite-element spaces $\mathbf{V}_h \times Q_h$ be Stokes stable, i.e., satisfy the inf-sup condition [30],*

$$\sup_{\mathbf{v} \in \mathbf{V}_h} \frac{(\operatorname{div} \mathbf{v}, p)}{\|\mathbf{v}\|_1} \geq \gamma_B^0 \|p\| \quad \forall p \in Q_h,$$

where $\gamma_B^0 > 0$ is a constant that does not depend on mesh size. Then, for any $p \in Q_h$

$$(2.9) \quad \sup_{\mathbf{v} \in \mathbf{V}_h} \frac{(\operatorname{div} \mathbf{v}, p)}{\|\mathbf{v}\|_{A_{\mathbf{u}}^{\text{RQ}}}} \geq \frac{\gamma_B^0}{\sqrt{d\zeta}} \|p\| =: \frac{\gamma_B}{\zeta} \|p\|,$$

where $\|\mathbf{v}\|_{A_{\mathbf{u}}^{\text{RQ}}}^2 := a^{\text{RQ}}(\mathbf{v}, \mathbf{v})$, d is the dimension, and $\zeta := \sqrt{\lambda + 2\mu/d}$.

Proof. Using the properties of projection operators, we have that $\|P_{Q_h} \operatorname{div} \mathbf{v}\| \leq \|\operatorname{div} \mathbf{v}\|$ for all $\mathbf{v} \in \mathbf{V}_h$. This, along with the definitions of $A_{\mathbf{u}}^{\text{RQ}}$ and $A_{\mathbf{u}}$, yields

$$(2.10) \quad \|\mathbf{v}\|_{A_{\mathbf{u}}^{\text{RQ}}} \leq \|\mathbf{v}\|_{A_{\mathbf{u}}} \quad \forall \mathbf{v} \in \mathbf{V}_h.$$

Next, by direct computation and applying Young's inequality, we have that $(\operatorname{div} \mathbf{v}, \operatorname{div} \mathbf{v}) \leq d(\varepsilon(\mathbf{v}), \varepsilon(\mathbf{v}))$. This implies that $a(\mathbf{v}, \mathbf{v}) \leq (2\mu + d\lambda)(\varepsilon(\mathbf{v}), \varepsilon(\mathbf{v}))$, and, through another application of Young's inequality, we have $\|\mathbf{v}\|_{A_{\mathbf{u}}} \leq \sqrt{d\zeta} \|\mathbf{v}\|_1$, with $a(\mathbf{v}, \mathbf{v}) =: \|\mathbf{v}\|_{A_{\mathbf{u}}}^2$. Then, for any $p \in Q_h$,

$$(2.11) \quad \sup_{\mathbf{v} \in \mathbf{V}_h} \frac{(B_{\mathbf{u}} \mathbf{v}, p)}{\|\mathbf{v}\|_{A_{\mathbf{u}}}} \geq \frac{\gamma_B^0}{\sqrt{d}\zeta} \|p\| =: \frac{\gamma_B}{\zeta} \|p\|.$$

Thus, (2.11) and (2.10) give (2.9). \square

Note that since the norm $\|\cdot\|_{A_{\mathbf{u}}^{\text{RQ}}}$ is parameter-dependent, in the large λ limit, both sides of (2.9) behave as $1/\sqrt{\lambda}$. We now show that the reduced-quadrature discretization is well-posed, using the fact that the bubble-enriched P1-RT0-P0 discretization is *Stokes–Biot stable* (see Definition 3.1 in [53]).

THEOREM 2.2. *Let $\mathbf{X}_h = (\mathbf{V}_h, \mathbf{W}_h, Q_h)$ be Stokes–Biot stable, that is,*

- $\exists C_V > 0$ such that $a(\mathbf{u}, \mathbf{v}) \leq C_V \|\mathbf{u}\|_1 \|\mathbf{v}\|_1$ for all $\mathbf{u}, \mathbf{v} \in \mathbf{V}_h$;
- $\exists \alpha_V > 0$ such that $a(\mathbf{u}, \mathbf{u}) \geq \alpha_V \|\mathbf{u}\|_1^2$ for all $\mathbf{u} \in \mathbf{V}_h$;
- (\mathbf{W}_h, Q_h) is Poisson stable, satisfying the necessary stability and continuity conditions for the mixed formulation of Poisson’s equation; and
- the pair of spaces (\mathbf{V}_h, Q_h) is Stokes stable.

For $\mathbf{x} = (\mathbf{u}, \mathbf{w}, p) \in \mathbf{X}_h$ and $\mathbf{y} = (\mathbf{v}, \mathbf{r}, q) \in \mathbf{X}_h$, define

$$(2.12) \quad \mathcal{B}(\mathbf{x}, \mathbf{y}) = a^{\text{RQ}}(\mathbf{u}, \mathbf{v}) - (\alpha p, \text{div } \mathbf{v}) + \tau(\mathbf{K}^{-1} \mu_f \mathbf{w}, \mathbf{r}) - \tau(p, \text{div } \mathbf{r}) \\ - \tau(\text{div } \mathbf{w}, q) - \left(\frac{1}{M} p, q\right) - (\alpha \text{div } \mathbf{u}, q),$$

$$(2.13) \quad \|\mathbf{x}\|_{\mathcal{D}^{\text{RQ}}}^2 = \|\mathbf{u}\|_{A_{\mathbf{u}}^{\text{RQ}}}^2 + c_p^{-1} \|p\|^2 + \tau \|\mathbf{w}\|_{M_{\mathbf{w}}}^2 + \tau^2 c_p \|\text{div } \mathbf{w}\|^2,$$

where $\|\mathbf{w}\|_{M_{\mathbf{w}}}^2 := (\mathbf{K}^{-1} \mu_f \mathbf{w}, \mathbf{w})$, and $c_p = (\frac{\alpha^2}{\zeta^2} + \frac{1}{M})^{-1}$. Then

$$(2.14) \quad \sup_{\mathbf{0} \neq \mathbf{x} \in \mathbf{X}_h} \sup_{\mathbf{0} \neq \mathbf{y} \in \mathbf{X}_h} \frac{\mathcal{B}(\mathbf{x}, \mathbf{y})}{\|\mathbf{x}\|_{\mathcal{D}^{\text{RQ}}} \|\mathbf{y}\|_{\mathcal{D}^{\text{RQ}}}} \leq \tilde{\zeta},$$

$$(2.15) \quad \inf_{\mathbf{0} \neq \mathbf{y} \in \mathbf{X}_h} \sup_{\mathbf{0} \neq \mathbf{x} \in \mathbf{X}_h} \frac{\mathcal{B}(\mathbf{x}, \mathbf{y})}{\|\mathbf{x}\|_{\mathcal{D}^{\text{RQ}}} \|\mathbf{y}\|_{\mathcal{D}^{\text{RQ}}}} \geq \tilde{\gamma},$$

where the constants $\tilde{\zeta}$ and $\tilde{\gamma}$ are independent of the physical and discretization parameters.

Proof. Using Lemma 2.1, we know that for a given $p \in Q_h$, there exists $\mathbf{z} \in \mathbf{V}_h$ such that $(p, \text{div } \mathbf{z}) \geq \frac{\gamma_B}{\zeta} \|p\|^2$ and $\|\mathbf{z}\|_{A_{\mathbf{u}}^{\text{RQ}}} = \|p\|$. Let $\mathbf{v} = \mathbf{u} - \psi_1 \mathbf{z}$, $\mathbf{r} = \mathbf{w}$, and $q = -p - \psi_2 \tau \text{div } \mathbf{w}$ for constants ψ_1 and ψ_2 that will be specified later. Then, by the Cauchy–Schwarz and Young’s inequalities,

$$\mathcal{B}(\mathbf{x}, \mathbf{y}) = \|\mathbf{u}\|_{A_{\mathbf{u}}^{\text{RQ}}}^2 - \psi_1 a^{\text{RQ}}(\mathbf{u}, \mathbf{z}) + \psi_1 \alpha (p, \text{div } \mathbf{z}) + \tau \|\mathbf{w}\|_{M_{\mathbf{w}}}^2 + \frac{1}{M} \|p\|^2 \\ + \psi_2 \tau \frac{1}{M} (p, \text{div } \mathbf{w}) + \psi_2 \alpha \tau (P_{Q_h} \text{div } \mathbf{u}, \text{div } \mathbf{w}) + \psi_2 \tau^2 \|\text{div } \mathbf{w}\|^2 \\ \geq \|\mathbf{u}\|_{A_{\mathbf{u}}^{\text{RQ}}}^2 - \frac{1}{2} \|\mathbf{u}\|_{A_{\mathbf{u}}^{\text{RQ}}}^2 - \frac{\psi_1^2}{2} \|\mathbf{z}\|_{A_{\mathbf{u}}^{\text{RQ}}}^2 + \psi_1 \frac{\alpha \gamma_B}{\zeta} \|p\|^2 + \tau \|\mathbf{w}\|_{M_{\mathbf{w}}}^2 + \frac{1}{M} \|p\|^2 \\ - \frac{3\psi_2}{2} \frac{1}{M^2} \|p\|^2 - \frac{\psi_2}{6} \tau^2 \|\text{div } \mathbf{w}\|^2 - \frac{\psi_2}{2} \alpha^2 \|P_{Q_h} \text{div } \mathbf{u}\|^2 - \frac{\psi_2}{2} \tau^2 \|\text{div } \mathbf{w}\|^2 \\ + \psi_2 \tau^2 \|\text{div } \mathbf{w}\|^2.$$

As in the proof of Lemma 2.1,

$$\frac{1}{d} (P_{Q_h} \text{div } \mathbf{u}, P_{Q_h} \text{div } \mathbf{u}) \leq \frac{1}{d} (\text{div } \mathbf{u}, \text{div } \mathbf{u}) \leq (\epsilon(\mathbf{u}), \epsilon(\mathbf{u})).$$

Then, by direct calculation and the definition of $A_{\mathbf{u}}^{\text{RQ}}$, we have

$$(2.16) \quad \|P_{Q_h} \operatorname{div} \mathbf{u}\| \leq \frac{1}{\zeta} \|\mathbf{u}\|_{A_{\mathbf{u}}^{\text{RQ}}}.$$

Combining terms and applying (2.16) gives

$$\begin{aligned} \mathcal{B}(\mathbf{x}, \mathbf{y}) &\geq \left(\frac{1}{2} - \frac{\psi_2 \alpha^2}{2 \zeta^2} \right) \|\mathbf{u}\|_{A_{\mathbf{u}}^{\text{RQ}}}^2 + \tau \|\mathbf{w}\|_{M_w}^2 + \frac{1}{3} \psi_2 \tau^2 \|\operatorname{div} \mathbf{w}\|^2 \\ &\quad + \left(\psi_1 \frac{\alpha \gamma_B}{\zeta} - \frac{\psi_1^2}{2} \right) \|p\|^2 + \left(1 - \frac{3}{4} \frac{2\psi_2}{M} \right) \frac{1}{M} \|p\|^2. \end{aligned}$$

Choosing $\psi_1 = \frac{\alpha \gamma_B}{2\zeta}$ and $\psi_2 = \frac{1}{2} \left(\frac{\alpha^2}{\zeta^2} + \frac{1}{M} \right)^{-1}$ then gives

$$\begin{aligned} \mathcal{B}(\mathbf{x}, \mathbf{y}) &\geq \left(\frac{1}{2} - \frac{1}{4} \right) \|\mathbf{u}\|_{A_{\mathbf{u}}^{\text{RQ}}}^2 + \tau \|\mathbf{w}\|_{M_w}^2 + \frac{1}{6} \tau^2 \left(\frac{\alpha^2}{\zeta^2} + \frac{1}{M} \right)^{-1} \|\operatorname{div} \mathbf{w}\|^2 \\ &\quad + \left(\frac{3\alpha^2 \gamma_B^2}{8\zeta^2} \right) \|p\|^2 + \left(1 - \frac{3}{4} \right) \frac{1}{M} \|p\|^2 \\ &\geq \bar{\gamma} \|(\mathbf{u}, \mathbf{w}, p)\|_{\mathcal{D}^{\text{RQ}}}^2, \end{aligned}$$

where $\bar{\gamma} = \min\{\frac{1}{6}, \frac{3\gamma_B^2}{8}\}$. Then, by the triangle inequality,

$$\|\mathbf{y}\|_{\mathcal{D}^{\text{RQ}}}^2 = \|\mathbf{v}\|_{A_{\mathbf{v}}^{\text{RQ}}}^2 + \left(\frac{\alpha^2}{\zeta^2} + \frac{1}{M} \right) \|q\|^2 + \tau \|\mathbf{r}\|_{M_w}^2 + \tau^2 c_p \|\operatorname{div} \mathbf{r}\|^2 \leq (\gamma^*)^2 \|\mathbf{x}\|_{\mathcal{D}^{\text{RQ}}}^2,$$

where $(\gamma^*)^2 = \max\{2, \frac{2\beta}{4}\}$. Thus, the bilinear form $\mathcal{B}(\cdot, \cdot)$ defined in (2.12) satisfies (2.15) with $\tilde{\gamma} = \gamma^*/\bar{\gamma}$. For the upper bound, (2.14), using Cauchy–Schwarz and (2.16), we have $\mathcal{B}(\mathbf{x}, \mathbf{y}) \leq 8\|\mathbf{x}\|_{\mathcal{D}^{\text{RQ}}}\|\mathbf{y}\|_{\mathcal{D}^{\text{RQ}}}$, which completes the proof. \square

Remark 2.3. To better understand the choice of the weighted norm (2.13), consider two limiting cases. When $\lambda \rightarrow \infty$, $\mathcal{B}(\mathbf{x}, \mathbf{y})$ is dominated by $\lambda(P_{Q_h} \operatorname{div} \mathbf{u}, P_{Q_h} \operatorname{div} \mathbf{v})$, which corresponds to the dominating term $\lambda\|P_{Q_h} \operatorname{div} \mathbf{u}\|^2$ in the weighted norm. When $\tau \rightarrow 0$, $\mathcal{B}(\mathbf{x}, \mathbf{y})$ reduces to $a^{\text{RQ}}(\mathbf{u}, \mathbf{v}) - (\alpha p, \operatorname{div} \mathbf{v}) - (\alpha \operatorname{div} \mathbf{u}, q) - \frac{1}{M}(p, q)$, which is a Stokes-like problem. The weighted norm (2.13), in this case, reduces to $\|\mathbf{u}\|_{A_{\mathbf{u}}^{\text{RQ}}}^2 + c_p^{-1} \|p\|^2$, which is a proper choice for Stokes-type problems. Thus, the weighted norm (2.13) is a proper choice in those limiting cases.

Remark 2.4. In [44], the *minimal Stokes–Biot stability condition* was proposed, under which a wider class of discretizations can be shown to be parameter-robust for solving the three-field formulation (2.1)–(2.3). That result also applies to the reduced-quadrature discretization presented here, and the conclusions of Theorem 2.2 still hold if we assume $\mathbf{X}_h = (\mathbf{V}_h, \mathbf{W}_h, Q_h)$ to be *minimal Stokes–Biot stable*, i.e., replacing the condition that (\mathbf{W}_h, Q_h) is Poisson stable by $\operatorname{div} \mathbf{W}_h \subset Q_h$. In fact, the proof of Theorem 2.2 uses only the minimal Stokes–Biot stability condition. This means that the reduced-quadrature technique can be applied to other discretizations that are minimal Stokes–Biot stable but not Stokes–Biot stable, e.g., the bubble-enriched P1-P1-P0 and P2-P1-P0 discretizations. We refer to [44] for further discussion of spaces that satisfy the minimal Stokes–Biot stability condition.

Remark 2.5. In [65], it has been shown that the reduced-quadrature discretization is well-posed independent of the discretization parameters by using the traditional

Brezzi theory for saddle-point systems [11]. Here, with the help of Stokes–Biot stability and properly chosen weighted norm, we show that the reduced-quadrature discretization is well-posed independent of the physical parameters as well. This implies that the reduced-quadrature approach is parameter-robust and also does not destroy the approximation properties of the bubble-enriched P1-RT0-P0 discretization [53].

3. Monolithic multigrid. Preconditioners for coupled systems, such as the reduced-quadrature discretization in (2.8), generally fall into two classes, those based on block-factorization approaches and those based on monolithic multigrid. The block-factorization approach was considered for the discretization from [53] in [2]; here, we focus on monolithic multigrid, extending recent studies in [31, 32, 41]. The defining feature of monolithic multigrid is the use of coupled relaxation schemes that are crafted to address the block structure of the system, along with a coarse-grid correction procedure that, again, couples the blocks within the system. Here, we consider geometric multigrid [59], combining coarse-grid correction based on geometric interpolation operators (modified, as discussed below, to account for divergence-free functions) with relaxation that aims to damp oscillatory error components on each grid level. We write the two-grid error-propagation operator as

$$(3.1) \quad E_{TG} = E_s^{\nu_2} E_{CGC} E_s^{\nu_1},$$

where ν_1 and ν_2 are the number of pre- and postrelaxation iterations, respectively. The error-propagation operator for relaxation is $E_s = I - \omega \mathcal{M}^{-1} \mathcal{A}$, where ω is a damping parameter, and $E_{CGC} = I - P \mathcal{A}_H^{-1} R \mathcal{A}$ for the coarse-grid correction (CGC) where P is the multigrid interpolation operator and R is the restriction operator. The coarse-grid operator, \mathcal{A}_H , is constructed either by rediscrretization or as the Galerkin operator, $R \mathcal{A} P$. As is typical for monolithic multigrid, the interpolation operator is determined blockwise, given as

$$(3.2) \quad P = \begin{pmatrix} P_{\mathbf{u}} & 0 & 0 \\ 0 & P_{\mathbf{w}} & 0 \\ 0 & 0 & P_p \end{pmatrix},$$

where $P_{\mathbf{u}}$ is the interpolation operator for displacements, $P_{\mathbf{w}}$ is that for the Darcy velocity, and P_p is the interpolation operator for pressure. We discuss the construction of $P_{\mathbf{u}}$ below; for $P_{\mathbf{w}}$ and P_p , we use the canonical finite-element interpolation operators for RT0 and P0. We fix $R = P^T$. While the Galerkin and rediscrretization coarse-grid operators coincide when the canonical finite-element operators are used for all fields, they will not do so here, due to the use of the divergence-preserving interpolation for $P_{\mathbf{u}}$ discussed below. Following the geometric multigrid structure, we use the rediscrretization operators instead of Galerkin, primarily because this allows easy extension from effective two-level solvers to the multilevel case.

To simplify the notation, we rewrite

$$(3.3) \quad \mathcal{A}^{\text{RQ}} \mathbf{x} = \begin{pmatrix} A & B^\top \\ B & -C \end{pmatrix} \begin{pmatrix} \mathbf{y} \\ p \end{pmatrix},$$

where

$$A = \begin{pmatrix} A_{\mathbf{u}}^{\text{RQ}} & 0 \\ 0 & \tau M_{\mathbf{w}} \end{pmatrix}, \quad B = (\alpha B_{\mathbf{u}} \quad \tau B_{\mathbf{w}}), \quad C = \frac{1}{M} M_p, \quad \text{and } \mathbf{y} = \begin{pmatrix} \mathbf{u} \\ \mathbf{w} \end{pmatrix}.$$

Next, we detail the nonstandard aspects of our multigrid method, namely the divergence-preserving interpolation operator and the coupled relaxation schemes.

3.1. Divergence-preserving interpolation. As recognized in [54, 55, 56] (see also [20]), a key to achieving solvers for elasticity that are robust in the incompressible (large λ) limit is the interpolation of divergence-free functions on the coarse mesh to divergence-free functions on the fine mesh. If \mathbf{u}_H is a coarse-grid divergence-free function, then, by the divergence theorem,

$$\int_{\partial T} \mathbf{n}^\top \mathbf{u}_H \, ds = 0 \quad \forall T \in \mathcal{T}_H,$$

where the subscript H denotes the coarse grid whose elements form the set \mathcal{T}_H . Asking that the prolongation of \mathbf{u}_H to the fine grid also be divergence-free yields

$$(3.4) \quad \int_{\partial T} \mathbf{n}^\top (P_{\mathbf{u}} \mathbf{u}_H) \, ds = 0 \quad \forall T \in \mathcal{T}_h,$$

where we now impose the condition on the fine-mesh elements in \mathcal{T}_h .

The standard finite-element interpolation operator on the displacement space does not satisfy this condition. To build an operator that does, we consider the interpolation locally from each coarse-grid element, as pictured in Figure 3.1. The key step in the construction is to use the finite-element interpolation operator to fix all fine-mesh DoFs on the edges of the coarse-mesh triangle, and use the three edge DoFs on the “interior” fine-mesh triangle to enforce (3.4). A columnwise construction of the interpolation operator is then given by first computing $\mathbf{c}_i = \hat{P}_{\mathbf{u}} \mathbf{e}_i$, where \mathbf{e}_i is the i th canonical unit vector on the coarse mesh, and $\hat{P}_{\mathbf{u}}$ is the standard finite-element interpolation operator. Then, the entries in \mathbf{c}_i that correspond to the interior bubble DoFs depicted in Figure 3.1 are replaced by values that ensure satisfaction of (3.4). Consider the triangle, $t^{1,2,3}$, in Figure 3.1 with vertices labeled 1, 2, 3. Let $c_b^{v_1, v_2}$ denote the entry in \mathbf{c}_i associated with the bubble DoFs on the edge between vertices v_1 and v_2 , and let c_v denote the entries in \mathbf{c}_i associated with the x and y DoFs on vertex v . To make the function represented by \mathbf{c}_i divergence-free on $t^{1,2,3}$, we set the coefficients of the interior bubble DoFs, $c_b^{1,3}$, to cancel that from the remaining DoFs,

$$c_b^{1,3} = - \left(c_b^{1,2} + c_b^{2,3} + \frac{1}{|\partial t^{1,2,3}|} \sum_{v=1}^3 \int_{\partial t^{1,2,3}} \mathbf{n}^\top \mathbf{c}_v \lambda_v \, ds \right),$$

where \mathbf{n} is the outward normal, and λ_v is the linear basis function associated with vertex v . Note that this calculation is simplified by choosing the bubble DoFs to be defined directly as integrals over the associated edges.

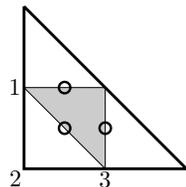


FIG. 3.1. A coarse-mesh element, $T \in \mathcal{T}_H$, and the four fine-mesh triangles that interpolate from it. The circles represent the bubble DoFs that are used to satisfy the divergence-free interpolation condition. The gray fine-grid triangle is referred to as the interior triangle, while the other fine-grid triangles are the “corner” triangles.

3.2. Monolithic multigrid relaxation. It is widely recognized that standard relaxation schemes, such as Jacobi or Gauss–Seidel, are not effective components of a multigrid algorithm for many saddle-point problems [59]. Instead, several families of relaxation schemes tailored to this setting have been proposed and studied. Here, we focus on two classes of such methods, Vanka and Braess–Sarazin relaxation.

3.2.1. Vanka relaxation scheme. Vanka relaxation, originally proposed in [60], has been adapted for a wide variety of discretizations and saddle-point problems [1, 37, 39, 42]. At its root, Vanka methods are overlapping block relaxation schemes, that can be considered in either additive (block-Jacobi) or multiplicative (block-Gauss–Seidel) form. While multiplicative variants have long been considered, the additive form has attracted recent interest, due to its natural parallelization [18, 19].

Given a decomposition of the set of DoFs into L (overlapping) blocks, a standard Schwarz method is most easily defined by defining the restriction operator, V_ℓ , from global vectors to local vectors on block ℓ . Then, given a current residual, $\mathbf{r}^{(j)} = \mathbf{b} - \mathcal{A}^{\text{RQ}}\mathbf{x}^{(j)}$, we can solve the projected system

$$V_\ell \mathcal{A}^{\text{RQ}} V_\ell^\top \hat{\mathbf{x}}_\ell = V_\ell \mathbf{r}^{(j)}$$

on each block. The weighted additive form of the relaxation is then

$$\mathbf{x}^{(j+1)} = \mathbf{x}^{(j)} + \omega \sum_\ell V_\ell^\top D_\ell \hat{\mathbf{x}}_\ell,$$

where ω is a damping parameter and D_ℓ is a diagonal weight matrix that is chosen to compensate for the fact that different (global) DoFs appear in different numbers of patches. Here, we consider D_ℓ to be given by the “natural weights” of the overlapping block decomposition, where each diagonal entry is equal to the reciprocal of the number of patches that the corresponding DoFs appears in.

The construction of the Vanka blocks is critically important to the success of the resulting multigrid method, with general principles being well-understood for their construction in several contexts [5, 19, 42]. Following the construction of the reduced-quadrature discretization above, our primary concern is in ensuring relaxation suitably handles the locally supported basis functions for the divergence-free space [54, 55, 56]. Since those basis functions are supported around the nodes of the mesh, as shown in Figure 2.2, we also use nodal patches for the Vanka blocks; see Figure 3.2. For the full poroelasticity system, we use the patches shown at right; those at left will be used within the BSR scheme discussed next.

3.2.2. Braess–Sarazin relaxation schemes. Braess–Sarazin-type algorithms were originally proposed as relaxation schemes for the Stokes’ equations [10], using an approximate block factorization as an approximation to the original system. Like Vanka relaxation, they have also been extended to many discretizations and systems [1, 31, 32, 37, 39], and are closely related to Uzawa schemes [41]. Using the 2×2 block structure in (3.3), given a residual $\mathbf{r}^{(j)}$, *exact* BSR updates the approximation as

$$(3.5) \quad \begin{pmatrix} \mathbf{y}^{(j+1)} \\ p^{(j+1)} \end{pmatrix} = \begin{pmatrix} \mathbf{y}^{(j)} \\ p^{(j)} \end{pmatrix} + \omega \begin{pmatrix} F & B^\top \\ B & -C \end{pmatrix}^{-1} \mathbf{r}^{(j)},$$

where F is an approximation of A , often taken to be $\omega_b I$ or $\omega_b \text{diag}(A)$, with weight ω_b chosen to separately damp the correction to the variables in \mathbf{y} from that given by the global parameter, ω .

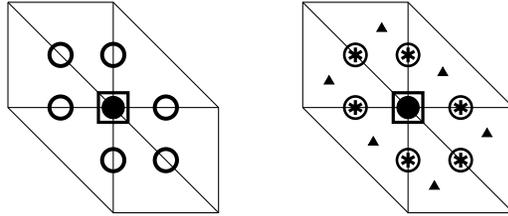


FIG. 3.2. Choices of DoFs for blocks within Vanka relaxation on the displacement subsystem (left) and full poroelasticity system (right). In both figures, filled circles and squares denote the DoFs associated with the linear component of the displacement, while empty circles show the bubble DoFs. At right, asterisks are used to denote the RT0 DoFs for the Darcy velocity space, and triangles denote the P0 DoFs for the pressure space.

The matrix inversion in (3.5) can be carried out in two stages as solving

$$(3.6) \quad \begin{aligned} S\delta p &= BF^{-1}\mathbf{r}_y^{(j)} - \mathbf{r}_p^{(j)}, \\ F\delta \mathbf{y} &= \mathbf{r}_y^{(j)} - B^\top \delta p, \end{aligned}$$

where $S = C + BF^{-1}B^\top$, and $\mathbf{r}_y^{(j)}$ and $\mathbf{r}_p^{(j)}$ are the first and second block components of $\mathbf{r}^{(j)}$ in this decomposition. In *exact* BSR, there is a significant cost associated with the inversion of the Schur complement, S , in (3.6). For this reason, inexact BSR methods were proposed, where the exact solution of the Schur complement equation is replaced by a suitable iterative method applied to (3.6), typically given by a few steps of a relaxation scheme or of a multigrid cycle for that subsystem.

Here, we make use of the block structure of A , to note that

$$BA^{-1}B^\top = \alpha^2 B_u (A_u^{\text{RQ}})^{-1} B_u^\top + \tau B_w M_w^{-1} B_w^\top,$$

and that, particularly in the large λ limit, $B_u (A_u^{\text{RQ}})^{-1} B_u^\top$ is well-approximated by a scaled mass matrix on the pressure space. This idea is motivated by the inf-sup condition (2.9) and is, essentially, the well-known “fixed-stress” approximation [38]. Thus, we first approximate

$$S \approx \frac{1}{M} M_p + \frac{\alpha^2}{\lambda + 2\mu/d} M_p + \tau B_w D_w^{-1} B_w^\top,$$

where D_w is the diagonal of M_w , and refer to the method with exact inversion of this system in (3.6) as exact BSR. This is in combination with a single sweep of a Jacobi iteration on M_w to approximate the \mathbf{w} component of \mathbf{y} , and a single iteration of the Vanka relaxation with patches chosen as shown at left of Figure 3.2 to approximate the inversion of A_u^{RQ} to approximate the \mathbf{u} component of \mathbf{y} . For inexact BSR, we replace the exact solve with the approximation to S by a single sweep of weighted Jacobi (with relaxation weight ω_J) on (3.6).

A downside of these relaxation schemes is their dependence on multiple relaxation parameters in their component parts. While some general principles exist to help us choose those parameters, often they are fixed by expensive brute-force testing. Here, we will make use of LFA to make these choices.

4. Local Fourier analysis. LFA is a common and useful tool to predict and analyze actual performance of algorithms for the solution of discretized PDEs [59, 62].

In particular, it has been used to analyze the construction and optimization of the components of a multigrid algorithm, such as relaxation schemes and grid-transfer operators [29, 31, 32, 33, 41, 42]. In this paper, we apply the LFA framework developed in [29, 31, 32, 42] to monolithic multigrid methods for the discretized Biot model in (2.8), in order to optimize the relaxation parameters described above.

4.1. Two-grid LFA. Following [59, 62], we first consider two-dimensional infinite uniform grids, $\mathbf{G}_h = \{\mathbf{x} := (x_1, x_2) = (k_1, k_2)h, (k_1, k_2) \in \mathbb{Z}^2\}$. Let L_h be a scalar Toeplitz operator defined as $L_h w_h(\mathbf{x}) = \sum_{\boldsymbol{\kappa} \in \mathbf{S}} s_{\boldsymbol{\kappa}} w_h(\mathbf{x} + \boldsymbol{\kappa}h)$, $\boldsymbol{\kappa} = (\kappa_1, \kappa_2) \in \mathbf{S}$, with constant coefficients $s_{\boldsymbol{\kappa}} \in \mathbb{R}$ (or \mathbb{C}), and where $w_h(\mathbf{x})$ is a function in $l^2(\mathbf{G}_h)$. Here, $\mathbf{S} \subset \mathbb{Z}^2$ is a finite index set over which the stencil is nonzero. Because L_h is formally diagonalized by the Fourier modes $\varphi(\boldsymbol{\theta}, \mathbf{x}) = e^{i\boldsymbol{\theta} \cdot \mathbf{x}/h} = e^{i\theta_1 x_1/h} e^{i\theta_2 x_2/h}$, where $\boldsymbol{\theta} = (\theta_1, \theta_2)$, we use $\varphi(\boldsymbol{\theta}, \mathbf{x})$ as a Fourier basis with $\boldsymbol{\theta} \in [-\frac{\pi}{2}, \frac{3\pi}{2})^2$ (or any pair of intervals with length 2π). High and low frequencies for standard coarsening (as considered here) are given by

$$\boldsymbol{\theta} \in T^{\text{low}} = \left[-\frac{\pi}{2}, \frac{\pi}{2}\right)^2, \boldsymbol{\theta} \in T^{\text{high}} = \left[-\frac{\pi}{2}, \frac{3\pi}{2}\right)^2 \setminus \left[-\frac{\pi}{2}, \frac{\pi}{2}\right)^2.$$

DEFINITION 4.1. *If for all grid functions $\varphi(\boldsymbol{\theta}, \mathbf{x})$, $L_h \varphi(\boldsymbol{\theta}, \mathbf{x}) = \tilde{L}_h(\boldsymbol{\theta}) \varphi(\boldsymbol{\theta}, \mathbf{x})$, we call $\tilde{L}_h(\boldsymbol{\theta}) = \sum_{\boldsymbol{\kappa} \in \mathbf{S}} s_{\boldsymbol{\kappa}} e^{i\boldsymbol{\theta} \cdot \boldsymbol{\kappa}}$ the symbol of L_h .*

For simple scalar operators (such as second-order finite-difference or finite-element discretizations of constant-coefficient diffusion equations), the performance of a standard relaxation method, such as the weighted Jacobi or Gauss–Seidel iterations, is easily analyzed by considering the symbol of the relaxation scheme [59, 62]. From the heuristic argument that coarse-grid correction effectively reduces error in T^{low} , the LFA smoothing factor for a relaxation scheme with error-propagation operator given by $I - \omega M_h^{-1} L_h$ is introduced as $\mu = \sup_{\boldsymbol{\theta} \in T^{\text{high}}} |1 - \omega \tilde{M}_h(\boldsymbol{\theta})^{-1} \tilde{L}_h(\boldsymbol{\theta})|$, where ω is a damping parameter.

While the LFA smoothing factor provides excellent predictions of true multigrid performance for simple discretizations of simple operators, it is known to provide poor predictions when used on complicated or higher-order operators [33]. In such settings, it is more reliable to use the two-grid LFA convergence factor, which takes into account the coarse-grid correction process. To do this, we define the harmonic modes by taking $\boldsymbol{\theta}^\alpha = (\theta_1^{\alpha_1}, \theta_2^{\alpha_2}) = \boldsymbol{\theta}^{00} + \pi \cdot \boldsymbol{\alpha}$, $\boldsymbol{\alpha} = (\alpha_1, \alpha_2) \in \{(0, 0), (1, 0), (0, 1), (1, 1)\}$, and $\boldsymbol{\theta}^{00} \in T^{\text{low}}$. That is, for each low-frequency mode $\boldsymbol{\theta} \in T^{\text{low}}$, we define a four-dimensional harmonic space, $\mathcal{F}(\boldsymbol{\theta}) = \text{span}\{\varphi(\boldsymbol{\theta}^\alpha, \cdot) : \boldsymbol{\alpha} \in \{(0, 0), (1, 0), (0, 1), (1, 1)\}\}$, which is invariant for standard full-coarsening two-grid algorithms.

To compute the LFA two-grid convergence factor, we must obtain an LFA representation of all components of the multigrid cycle. This requires finding symbols for not just the fine-grid operator and relaxation scheme, but also for the interpolation and restriction operators, and for the coarse-grid operator. The *symbol* of the two-grid algorithm is a 4×4 matrix that describes the action of the two-grid algorithm, and comes from noting that structured constant-coefficient interpolation and restriction operators map naturally between the four fine-grid harmonic modes in $\mathcal{F}(\boldsymbol{\theta})$ and the coarse-grid mode $2\boldsymbol{\theta}$. Writing \tilde{L}_{2h} for the symbol of the coarse-grid operator and \tilde{P}_h and \tilde{R}_h for the symbols of the interpolation and restriction operators, the Fourier representation of the two-grid error-propagation operator is defined as

$$\tilde{\mathbf{E}}_{TG}(\boldsymbol{\theta}) = \tilde{\mathbf{E}}_s^{\nu_2}(\boldsymbol{\theta})(I - \tilde{\mathbf{P}}_h(\boldsymbol{\theta})(\tilde{\mathbf{L}}_{2h}(2\boldsymbol{\theta}))^{-1}\tilde{\mathbf{R}}_h(\boldsymbol{\theta})\tilde{\mathbf{L}}_h(\boldsymbol{\theta}))\tilde{\mathbf{E}}_s^{\nu_1}(\boldsymbol{\theta}),$$

where

$$\begin{aligned}\tilde{\mathbf{L}}_h(\boldsymbol{\theta}) &= \text{diag} \left\{ \tilde{\mathbf{L}}_h(\boldsymbol{\theta}^{00}), \tilde{\mathbf{L}}_h(\boldsymbol{\theta}^{10}), \tilde{\mathbf{L}}_h(\boldsymbol{\theta}^{01}), \tilde{\mathbf{L}}_h(\boldsymbol{\theta}^{11}) \right\}, \\ \tilde{\mathbf{E}}_s(\boldsymbol{\theta}) &= \text{diag} \left\{ \tilde{\mathbf{E}}_s(\boldsymbol{\theta}^{00}), \tilde{\mathbf{E}}_s(\boldsymbol{\theta}^{10}), \tilde{\mathbf{E}}_s(\boldsymbol{\theta}^{01}), \tilde{\mathbf{E}}_s(\boldsymbol{\theta}^{11}) \right\}, \\ \tilde{\mathbf{R}}_h(\boldsymbol{\theta}) &= \left(\tilde{\mathbf{R}}_h(\boldsymbol{\theta}^{00}), \tilde{\mathbf{R}}_h(\boldsymbol{\theta}^{10}), \tilde{\mathbf{R}}_h(\boldsymbol{\theta}^{01}), \tilde{\mathbf{R}}_h(\boldsymbol{\theta}^{11}) \right), \\ \tilde{\mathbf{P}}_h(\boldsymbol{\theta}) &= \left(\tilde{\mathbf{P}}_h(\boldsymbol{\theta}^{00}); \tilde{\mathbf{P}}_h(\boldsymbol{\theta}^{10}); \tilde{\mathbf{P}}_h(\boldsymbol{\theta}^{01}); \tilde{\mathbf{P}}_h(\boldsymbol{\theta}^{11}) \right).\end{aligned}$$

Here, $\text{diag}\{T_1, T_2, T_3, T_4\}$ denotes the block-diagonal matrix with diagonal blocks, T_1, T_2, T_3 , and T_4 [59, 62]. With this, we define the two-grid LFA convergence factor.

DEFINITION 4.2. *The two-grid LFA convergence factor, ρ_{LFA} , is defined as*

$$(4.1) \quad \rho_{LFA} = \sup \left\{ \rho(\tilde{\mathbf{E}}_{TG}(\boldsymbol{\theta})) : \boldsymbol{\theta} \in T^{\text{low}} \right\},$$

where $\rho(\tilde{\mathbf{E}}_{TG}(\boldsymbol{\theta}))$ denotes the spectral radius of matrix $\tilde{\mathbf{E}}_{TG}(\boldsymbol{\theta})$.

As described above, it is natural to introduce algorithmic parameters when designing multigrid methods for complicated problems. It is for this purpose that we introduce LFA here. While it is often possible to optimize the LFA smoothing factor for simple problems through analytical means (see, for example, [31]), optimizing the two-grid LFA convergence factor for more complicated problems and algorithms is a challenging task [12]. Here, we will develop LFA representations of the monolithic multigrid algorithms above and optimize the two-grid convergence factor in (4.1) using brute-force sampling. In particular, while the true two-grid LFA convergence factor is most naturally defined as a supremum over a continuous range of values of $\boldsymbol{\theta}$, we will use a discrete sampling at a finite number of evenly spaced frequencies in the domain $(-\frac{\pi}{2}, \frac{\pi}{2})^2$, but without any change of notation.

4.2. LFA representation of discretized system. To extend Fourier analysis to the full discretized system in (2.8), we must account for the fact that the system is not readily extended to a Toeplitz operator on an infinite grid, unlike in the scalar case. This occurs in two ways. First, as is clear, the discretization of a coupled system of PDEs leads, at best, to a block operator with Toeplitz blocks. Second, even within a single block, such as $A_{\mathbf{u}}^{\text{RQ}}$, there are different “types” of DoFs, leading to nested block-Toeplitz structure.

The key concept in enabling LFA is in expressing the block-Toeplitz structure of the multigrid hierarchy and relaxation operator relative to the infinite grid, \mathbf{G}_h . With triangular cells and face- and cell-based DoFs, this is slightly nonintuitive. Figure 4.1 shows the DoFs in a typical pair of elements on the mesh, constructed by “cutting” a quadrilateral cell into two triangles. With this arrangement of DoFs, we have natural periodic structure for the P1 components of the displacement (2 DoFs, 1 for each component of the two-dimensional displacement vector, \mathbf{u}), but also for the 6 face-based DoFs, coming in two pairs of 3 DoFs, corresponding to the normal displacement bubble component along each face and the face-based Raviart–Thomas DoFs for the Darcy velocity. Note that we do not “collapse” the Fourier representation of the face-based DoFs to a single component within the symbol; this is not possible, since the matrix connections between face-based DoFs along (for example) horizontal edges

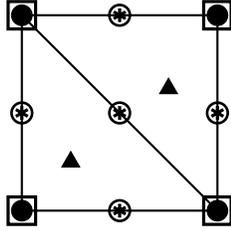


FIG. 4.1. *Cut quadrilateral mesh cell showing DoFs. Filled circles and squares denote the DoFs associated with the linear component of the displacement, while empty circles show the bubble DoFs. Asterisks are used to denote the RT0 DoFs for the Darcy velocity space, and triangles denote the P0 DoFs for the pressure space.*

will be different from those along diagonal edges. Instead, we will maintain an entry in the Fourier symbol for each “type” of face-based DoF. Similarly, the connections between the P0 DoFs in the lower-left triangles and the other variables in the cell may be different from those with the P0 DoFs in the upper-right triangles. Thus, we introduce Fourier representations of both of these DoFs. In total, this yields a 10×10 block Fourier symbol for the operator, \mathcal{A}^{RQ} . With this structure, it is a straightforward (but tedious) task to compute the Fourier symbol of \mathcal{A}^{RQ} . We outline the main ideas here, but leave the technical details to the Supplementary Material (to_supp.pdf [local/web 341KB]) for the interested reader.

First, we “expand” its block structure from the canonical 3×3 form to that of a 10×10 block-structured linear system, with 1 block for each DoF identified above and in Figure 4.1. In this ordering, each diagonal block is a Toeplitz matrix, whose symbol can be calculated according to Definition 4.1. Off-diagonal blocks in this structure are also Toeplitz matrices, although we also account for the offsets between DoF locations in the mesh in the Fourier symbols, in a similar manner to what was done in [31, 32]. Details of these calculations are presented in section SM1.

Similarly, Fourier representations of the grid-transfer operators can also be computed in block form. Taking the block-diagonal interpolation operator from (3.2), we separately compute Fourier representations of each interpolation operator, accounting for block structure of the DoFs and the details of the interpolation schemes. Since we have a 10-dimensional space associated with each Fourier frequency, and interpolation and restriction map between four harmonic frequencies on the fine mesh and a single frequency on the coarse mesh, this results in a 40×10 symbol for interpolation and a 10×40 symbol for restriction, which can be broken into 10×10 blocks giving the part of the symbol associated with each individual frequency in the harmonic set. These 10×10 blocks can be broken down further, based on the block-diagonal form in (3.2), to a 5×5 block associated with displacements, a 3×3 block for Darcy velocities, and a 2×2 block for pressures. It is somewhat more natural to compute Fourier representations of the restriction operators and use (scaled) transposes of these symbols for interpolation, which is the approach followed in section SM2.

Finally, Fourier representations of the relaxation schemes can be computed. For Vanka relaxation, this follows the approach presented in [18], where the Fourier representation of a residual at given frequency is restricted, via V_ℓ , to a Vanka patch, and the action of the local solve is computed exactly on this basis, with accounting for the overlap between patches. Details are given in section SM3. For Braess-Sarazin relaxation, the symbols of F , S and the approximation to S are readily computed in the same manner as the symbols above, and the incorporation of a relaxation scheme in place of an exact inversion of S is done similarly. See section SM4 for the details.

4.3. Validation and optimization. While we are primarily interested in the use of monolithic multigrid as a preconditioner for GMRES, we begin by studying its use as a stationary iteration, for the purposes of optimizing parameters in the methods. We use LFA to predict convergence factors associated with given choices of parameters, and compare to measured performance of a stationary iteration, approximating the asymptotic convergence factor of the iteration as $\rho_N = \frac{\|\mathbf{r}^{(j)}\|}{\|\mathbf{r}^{(j-1)}\|}$, where $\mathbf{r}^{(j)}$ is the residual at the j th iteration. To ensure a good approximation of the asymptotic convergence factor, iterations are run until the change in the measured convergence factor between iterations is less than 10^{-3} . While LFA can be made exact in the case of periodic boundary conditions, the numerical tests were performed using Dirichlet boundary conditions as is more common. We consider $\Omega = [0, 1]^2$ covered with a uniform triangular grid with mesh spacing $h = 1/64$. As a test problem, we consider a zero right-hand side, with a random initial guess for a single time step with $\tau = 1$. To demonstrate the impact of the physical parameters, the permeability, \mathbf{K} , and the Poisson ratio ν are varied. In all test cases, we consider a diagonal permeability tensor $\mathbf{K} = k\mathbf{I}$. Additionally, $\alpha = 1$, $\mu_f = 1$, $M = 10^6$, and $E = 3 \times 10^4$. LFA is performed using 32 evenly spaced sample points in each coordinate direction, offset so that no sample is taken at the origin in Fourier space. Note that the two-grid LFA convergence factor, (4.1), is a function of the damping parameter, ω . In order to obtain an efficient algorithm, we use brute-force sampling to optimize the LFA-predicted two-grid convergence factors over choices of ω , with steps of size 0.02.

In Tables 4.1 and 4.2, we present LFA-optimized parameters and both LFA-predicted and numerically measured two-grid convergence factors for monolithic multigrid using Vanka (with $\nu_1 = \nu_2 = 2$) and inexact BSR schemes (with $\nu_1 = \nu_2 = 1$), respectively. To validate the parameters for inexact BSR, we first perform LFA for

TABLE 4.1

Optimized relaxation parameter (ω_{opt}), observed convergence factor (ρ_N) with Dirichlet boundary conditions, and optimal two-grid LFA predictions (ρ_{LFA}) for additive Vanka relaxation ($\nu_1 = \nu_2 = 2$) on the full system with the 20-DoF vertex-based patch (Figure 3.2, right), varying k and ν .

$\nu \backslash k$		1	10^{-2}	10^{-4}	10^{-6}	10^{-8}	10^{-10}
		$\nu = 0$	ω_{opt}	0.92	0.92	0.92	0.92
	ρ_{LFA}	0.705	0.705	0.705	0.702	0.490	0.552
	ρ_N	0.722	0.722	0.722	0.722	0.475	0.547
$\nu = 0.2$	ω_{opt}	0.90	0.90	0.90	0.90	0.86	0.76
	ρ_{LFA}	0.624	0.624	0.624	0.622	0.474	0.557
	ρ_N	0.610	0.610	0.610	0.611	0.468	0.552
$\nu = 0.4$	ω_{opt}	0.80	0.80	0.80	0.80	0.78	0.76
	ρ_{LFA}	0.410	0.410	0.410	0.410	0.436	0.562
	ρ_N	0.403	0.403	0.403	0.404	0.432	0.557
$\nu = 0.45$	ω_{opt}	0.76	0.76	0.76	0.76	0.76	0.76
	ρ_{LFA}	0.492	0.492	0.492	0.492	0.498	0.564
	ρ_N	0.489	0.489	0.489	0.489	0.495	0.560
$\nu = 0.49$	ω_{opt}	0.74	0.74	0.74	0.74	0.74	0.76
	ρ_{LFA}	0.572	0.572	0.572	0.572	0.572	0.573
	ρ_N	0.569	0.569	0.569	0.570	0.570	0.569
$\nu = 0.499$	ω_{opt}	0.72	0.72	0.72	0.72	0.72	0.74
	ρ_{LFA}	0.600	0.600	0.600	0.600	0.600	0.599
	ρ_N	0.596	0.596	0.596	0.596	0.596	0.596

TABLE 4.2

Optimized relaxation parameters ($\omega_{J,\text{opt}}, \omega_{\text{opt}}$), observed convergence factor (ρ_N) with Dirichlet boundary conditions, and optimal two-grid LFA predictions (ρ_{LFA}) for inexact BSR (using one sweep of damped Jacobi for the approximate solve of the Schur complement and additive Vanka for the displacement block), varying k and ν .

$\nu \backslash k$		1	10^{-2}	10^{-4}	10^{-6}	10^{-8}	10^{-10}
$\nu = 0$	$\omega_{J,\text{opt}}$	1.10	1.15	1.06	1.26	0.98	0.96
	ω_{opt}	0.72	0.68	0.76	0.62	0.88	0.98
ρ_{LFA}		0.648	0.649	0.656	0.645	0.556	0.417
	ρ_N	0.647	0.646	0.650	0.636	0.535	0.552
$\nu = 0.2$	$\omega_{J,\text{opt}}$	1.30	1.23	1.00	0.97	1.14	0.95
	ω_{opt}	0.60	0.62	0.71	0.82	0.73	1.07
ρ_{LFA}		0.660	0.652	0.683	0.647	0.568	0.435
	ρ_N	0.652	0.620	0.682	0.616	0.578	0.418
$\nu = 0.4$	$\omega_{J,\text{opt}}$	1.27	1.16	1.17	0.94	1.18	0.79
	ω_{opt}	0.69	0.74	0.74	0.74	0.77	1.14
ρ_{LFA}		0.684	0.663	0.670	0.690	0.659	0.509
	ρ_N	0.680	0.655	0.659	0.654	0.658	0.507
$\nu = 0.45$	$\omega_{J,\text{opt}}$	1.16	1.07	1.10	0.91	1.26	0.76
	ω_{opt}	0.72	0.72	0.72	0.72	0.74	1.17
ρ_{LFA}		0.732	0.732	0.732	0.732	0.732	0.570
	ρ_N	0.723	0.723	0.723	0.722	0.731	0.567
$\nu = 0.49$	$\omega_{J,\text{opt}}$	1.21	1.27	1.07	1.18	1.29	1.06
	ω_{opt}	0.69	0.69	0.69	0.69	0.69	1.00
ρ_{LFA}		0.772	0.772	0.772	0.772	0.772	0.688
	ρ_N	0.757	0.757	0.757	0.757	0.741	0.681
$\nu = 0.499$	$\omega_{J,\text{opt}}$	0.85	1.29	1.00	0.86	0.85	1.42
	ω_{opt}	0.68	0.68	0.68	0.68	0.68	0.73
ρ_{LFA}		0.786	0.786	0.786	0.786	0.786	0.772
	ρ_N	0.783	0.783	0.783	0.783	0.783	0.751

the exact BSR scheme discussed above (not shown here). For values of ν larger than 0.4, we find identical performance between exact and inexact BSR, except for the case of $k = 10^{-10}$, where inexact BSR slightly outperforms exact BSR for $\nu = 0.45$. Exact BSR performance notably improves as ν decreases, achieving convergence factors around 0.48 for $\nu = 0$ and larger values of k . While this is a slight improvement in convergence over the inexact BSR case, it relies on the prohibitively expensive exact inversion of the approximate Schur complement. Note that we also optimize for the Jacobi weight, ω_J , for approximately solving the Schur complement.

In general, we see good agreement between the LFA predictions and the measured factors, and that the two-grid schemes are robust to both the incompressible limit, $\nu \rightarrow 0.5$, and extremely small values of k . We note some irregularity in both the convergence factors themselves and the match between prediction and measurement in the small k limit, which appears to be due to ill-conditioning of the Fourier symbols when k is so small. This also leads to some irregularity in the optimal relaxation parameters also in this limit.

In these tests, we focus on the optimization of only the outer relaxation parameter, ω , using LFA. While it is possible to introduce more relaxation parameters (e.g., in the inner Vanka relaxation for inexact BSR, or the weighting matrix, D_ℓ), preliminary experiments showed that these did not greatly improve convergence. It is also important to note both that the optimal relaxation parameter varies with ν and that good choices for one value of ν do not lead to good performance across all values

considered here. With Vanka relaxation, for fixed $\omega = 0.9$ (close to the optimal value for $\nu = 0$), we see divergence for all tested values of $\nu > 0.2$. For fixed $\omega = 2/3$ (close to the optimal value for $\nu \rightarrow 1/2$), we see strong degradation in convergence as ν gets small, with divergence for all tested values of $\nu < 0.49$. We also note that, because these relaxation weights are used in multiplicative combination with coarse-grid correction, the performance of multigrid-preconditioned FGMRES, as is considered in section 5, is also sensitive to these choices.

5. Numerical results. We now consider performance of the reduced-quadrature discretization and the monolithic multigrid preconditioners, extending the two-level results shown above to the multilevel case. To allow fair comparison between the relaxation schemes, we have implemented both Vanka and inexact BSR in a single codebase, namely the HAZmath package [4]: a simple finite element, graph, and solver library. All timed numerical results are done using a workstation with an 8-core 3-GHz Intel Xeon Sandy Bridge CPU and 256 GB of RAM. This also allows direct comparison to timings for the block preconditioners from [2].

5.1. Steady-state model. Here, we use a single four-level V-cycle of the monolithic multigrid method as a preconditioner for FGMRES using a relative residual stopping tolerance of 10^{-6} and compare the performance with the block upper-triangular preconditioner previously used in [2], with form

$$(5.1) \quad B_U = \begin{pmatrix} A_{\mathbf{u}}^{\text{RQ}} & \alpha B_{\mathbf{u}}^{\top} & 0 \\ 0 & \left(\frac{\alpha^2}{\zeta^2} + \frac{1}{M}\right) M_p & -\tau B_{\mathbf{w}} \\ 0 & 0 & \tau M_{\mathbf{w}} + \tau^2 \left(\frac{\alpha^2}{\zeta^2} + \frac{1}{M}\right)^{-1} A_{\mathbf{w}} \end{pmatrix}.$$

Notice that (5.1) is applied to a permuted form of the discretization, as was considered in [2]. Similar to [2], each diagonal block in the preconditioner is solved to a relative residual tolerance of 10^{-3} using preconditioned FGMRES preconditioned with algebraic multigrid for the pressure and Darcy blocks and FGMRES preconditioned with geometric multigrid using the Vanka relaxation presented in subsection 3.2.1 for the displacement block.

In this example, the right-hand-side functions \mathbf{g} and f are chosen so that the exact solution is given by

$$\begin{aligned} \mathbf{u}(x, y, t) &= \text{curl } \varphi = \begin{pmatrix} \partial_y \varphi \\ -\partial_x \varphi \end{pmatrix}, \quad \varphi(x, y) = [xy(1-x)(1-y)]^2, \\ p(x, y, t) &= 1, \quad \mathbf{w}(x, y, t) = \mathbf{0}. \end{aligned}$$

The material parameters are the same as those used in the LFA validation above. Finally, starting with a zero initial guess, we set $\tau = 1$ and $t_{\max} = 1$, so that we only perform one time step, and fix the mesh spacing to be $h = 1/64$ (the four-level V-cycle has a direct solve on the coarse mesh with spacing $h = 1/8$). Table 5.1 presents results for monolithic multigrid with both Vanka and inexact Braess–Sarazin relaxation, and for the block preconditioner.

There are several takeaways from these results. First, the monolithic multigrid with Vanka relaxation method is robust with respect to the physical parameters, though we do see a slight degradation for small ν (i.e., the compressible case). This is not surprising, as the methods developed here were developed specifically for the limit as ν approaches $1/2$. Second, the multilevel monolithic multigrid with inexact BSR relaxation struggles when the permeability constant k is small, in contrast to the

TABLE 5.1

CPU time in seconds (and iterations to convergence) for FGMRES preconditioned by monolithic multigrid with additive Vanka relaxation on the full system with the 20-DoF vertex-based patch (Figure 3.2, right) and inexact BSR, and preconditioned by a block upper-triangular system from [2] on steady-state problem.

Scheme	k		ν					
	ν		1	10^{-2}	10^{-4}	10^{-6}	10^{-8}	10^{-10}
Vanka	0.0		2.417 (18)	1.966 (18)	2.037 (18)	2.049 (18)	1.029 (10)	1.002 (10)
BSR	0.0		0.457 (9)	0.458 (9)	0.451 (9)	0.613 (12)	0.502 (10)	0.459 (9)
Block	0.0		0.733 (16)	0.738 (16)	0.845 (16)	0.611 (13)	0.499 (12)	0.415 (8)
Vanka	0.2		1.863 (15)	1.605 (15)	1.544 (15)	1.550 (15)	1.025 (10)	1.103 (10)
BSR	0.2		0.615 (12)	0.561 (11)	0.667 (10)	0.556 (11)	0.454 (9)	0.507 (10)
Block	0.2		0.695 (15)	0.698 (15)	0.718 (15)	0.637 (12)	0.484 (11)	0.448 (8)
Vanka	0.4		1.095 (9)	1.027 (9)	1.000 (9)	0.923 (9)	1.006 (9)	1.089 (10)
BSR	0.4		0.554 (11)	0.608 (12)	0.758 (15)	0.664 (13)	1.124 (22)	0.658 (13)
Block	0.4		0.784 (15)	0.785 (15)	0.834 (15)	0.850 (13)	0.596 (11)	0.516 (9)
Vanka	0.45		1.091 (9)	0.921 (9)	1.003 (9)	0.922 (9)	1.018 (10)	0.944 (10)
BSR	0.45		0.658 (13)	0.844 (13)	0.921 (14)	0.763 (15)	28.5 (452)	0.819 (16)
Block	0.45		1.047 (15)	1.047 (15)	1.521 (16)	1.443 (15)	0.876 (12)	0.742 (11)
Vanka	0.49		1.347 (11)	1.142 (11)	1.121 (11)	1.128 (11)	1.125 (11)	0.957 (11)
BSR	0.49		0.917 (18)	0.911 (18)	0.971 (19)	0.966 (19)	17.4 (299)	1.126 (22)
Block	0.49		1.081 (15)	1.080 (15)	1.073 (15)	1.833 (16)	1.184 (12)	1.094 (11)
Vanka	0.499		1.565 (13)	1.354 (13)	1.318 (13)	1.318 (13)	1.437 (13)	1.187 (14)
BSR	0.499		1.221 (24)	1.163 (23)	1.220 (24)	1.656 (26)	1.341 (24)	2.808 (55)
Block	0.499		2.142 (15)	2.140 (15)	2.146 (15)	2.144 (15)	2.194 (16)	2.099 (15)

robust two-level results in section 4. It may be that W-cycles, or other approaches, are needed to achieve robustness in this case, but we do not investigate this here. Note, however, that for larger k , the total computational time when using inexact BSR is slightly faster than that for Vanka relaxation. Comparing the monolithic multigrid performance with that of the block preconditioner, we see that the block preconditioner performance is similar to that of the monolithic multigrid with inexact BSR for small ν . However, there is a clear degradation in performance of the block preconditioner in the incompressible limit, where monolithic multigrid is more robust. Since the degradation in CPU time is much worse than that in iteration count, we infer that the required iterations of the inner (block) solvers must be increasing in this limit.

To verify that the reduced-quadrature formulation does accurately approximate the problem, we perform a convergence study of the finite-element discretization with respect to the mesh size, given by $h = 1/(N - 1)$, where N is the number of vertices in each dimension. Here, we set $\tau = 1.0$ and $k = 10^{-6}$ as an example, with results shown in Figure 5.1 for $\nu = 0.4$ and $\nu = 0.499$. The displacement displays a first-order convergence with respect to the H^1 -seminorm, as expected, with no difference in error values for the different values of ν . The pressure displays second-order convergence despite only using P0 elements, with slight improvement as $\nu \rightarrow 0.5$. This superconvergence is due to having a very smooth solution (p is a constant) and using a uniform mesh. Additionally, Figure 5.2 shows that the monolithic multigrid approach (with exact solve on a coarsest mesh with $h = 1/8$) with either the Vanka or inexact BSR relaxation methods follows the expected $O(N^2 \log(N^2))$ scaling in CPU time with respect to problem size, even as $\nu \rightarrow 0.5$. We note that while iterations to convergence are independent of problem size for both values of ν when using Vanka relaxation, degradation to $O(\log(N^2))$ iterations is seen for BSR relaxation as $\nu \rightarrow 0.5$. These results further indicate the better performance of Vanka as $\nu \rightarrow 0.5$, and the loss of robustness in ν for multilevel BSR.

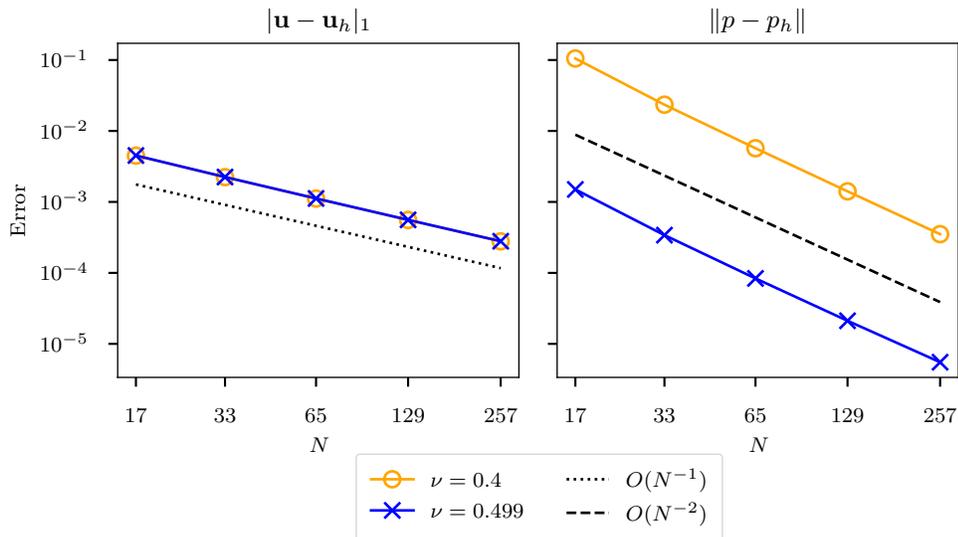


FIG. 5.1. Convergence study for steady-state problem, using FGMRES preconditioned by monolithic multigrid with the additive Vanka relaxation scheme using $\nu = 0.4$ and $\nu = 0.499$. Left: H^1 -seminorm error for displacement versus mesh size. Right: L^2 -error for pressure versus mesh size.

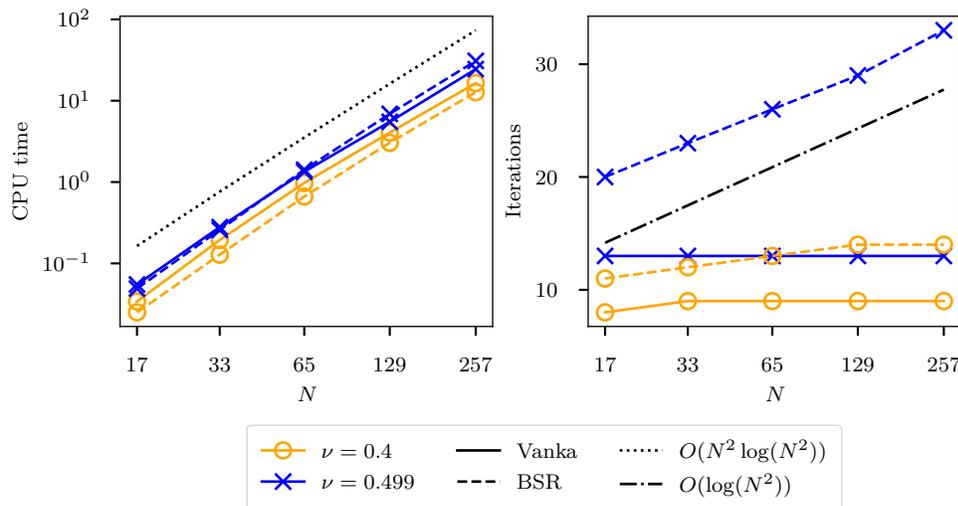


FIG. 5.2. CPU time in seconds (at left) and iterations to convergence (at right) for FGMRES preconditioned by a monolithic multigrid V-cycle with additive Vanka and inexact BSR relaxation schemes versus mesh size on steady-state problem for $\nu = 0.4$ and $\nu = 0.499$.

5.2. Smooth test problem. We next consider a slightly more realistic test problem, now with a time-dependent smooth solution, taken from [27]. The manufactured solution is defined on $\Omega = [0, 1]^2$ as

$$\mathbf{u}(x, y, t) = e^{-t} \begin{pmatrix} \sin(\pi y) \left(-\cos(\pi x) + \frac{1}{\mu+\lambda} \sin(\pi x) \right) \\ \sin(\pi x) \left(\cos(\pi y) + \frac{1}{\mu+\lambda} \sin(\pi y) \right) \end{pmatrix},$$

$$p(x, y, t) = e^{-t} \sin(\pi x) \sin(\pi y),$$

$$\mathbf{w}(x, y, t) = -k \nabla p,$$

with right-hand sides chosen appropriately. We consider Dirichlet boundary conditions on all sides for displacement and pressure. The physical parameters are $\alpha = 1$, $\mu_f = 1$, $M = 10^6$, and $E = 3 \times 10^4$. We perform all simulations from time $t = 0$ to $t = 0.5$. Here, we use a relative residual stopping tolerance for FGMRES of 10^{-10} , as preliminary experiments showed that this was needed to accurately resolve the pressure solution. Moreover, we only consider the additive Vanka method, as it proved more robust in the multilevel setting. Additionally, as we are mostly concerned with the incompressible limit, we focus on values of the Poisson ratio above 0.4.

Parameter robustness for the solver is demonstrated in Table 5.2, showing the average solver iteration count and average CPU time over 64 time steps with time step size $\tau = 1/128$. The mesh spacing is fixed to $h = 1/64$, and the values of k and ν are varied. Robustness with respect to discretization parameters, h and τ , is shown in Figure 5.3, for $k = 10^{-6}$ and both $\nu = 0.4$ and $\nu = 0.499$. We test on meshes with $N = 2^\ell + 1$, for $\ell = 4$ to 8, with $\tau = 2^{-m}$ for $m = 4$ to 8. Here, we see nearly identical CPU times with expected $O(N^2 \log(N^2))$ scaling for all values of τ . The corresponding LFA parameters from Tables 4.1 and 4.2 are used. The averaged iteration counts (not shown here) remain consistently in the range of 13 to 16 across all parameter values.

TABLE 5.2

Average CPU time in seconds (iterations) over 64 time steps for FGMRES preconditioned by monolithic multigrid with an additive Vanka relaxation scheme for the smooth solution problem with varying physical parameters k and ν .

$\nu \backslash k$	1	10^{-2}	10^{-4}	10^{-6}	10^{-8}	10^{-10}
0.4	1.536 (14.0)	1.521 (14.0)	1.543 (14.0)	1.520 (14.0)	1.679 (16.0)	1.506 (14.0)
0.45	1.431 (13.0)	1.402 (13.0)	1.404 (13.0)	1.525 (14.0)	1.441 (14.0)	1.467 (14.0)
0.49	1.536 (14.0)	1.518 (14.0)	1.513 (14.0)	1.513 (14.0)	1.344 (14.0)	1.294 (14.0)
0.499	1.663 (15.0)	1.628 (15.0)	1.622 (15.0)	1.620 (15.0)	1.312 (14.0)	1.331 (15.0)

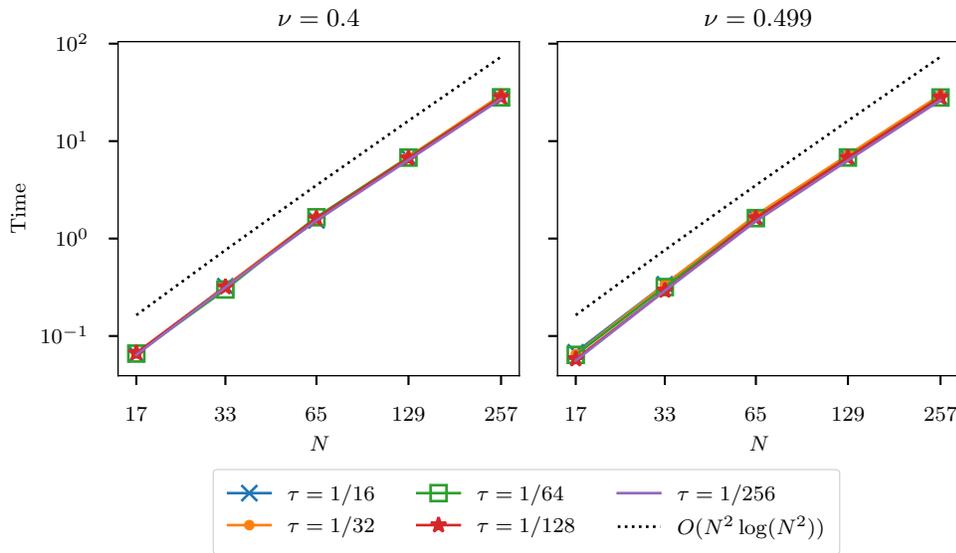


FIG. 5.3. Average CPU time in seconds for FGMRES preconditioned by monolithic multigrid for the smooth solution problem with $k = 10^{-6}$ and $\nu = 0.4$ (at left) and 0.499 (at right) with varying discretization parameters, N and τ .

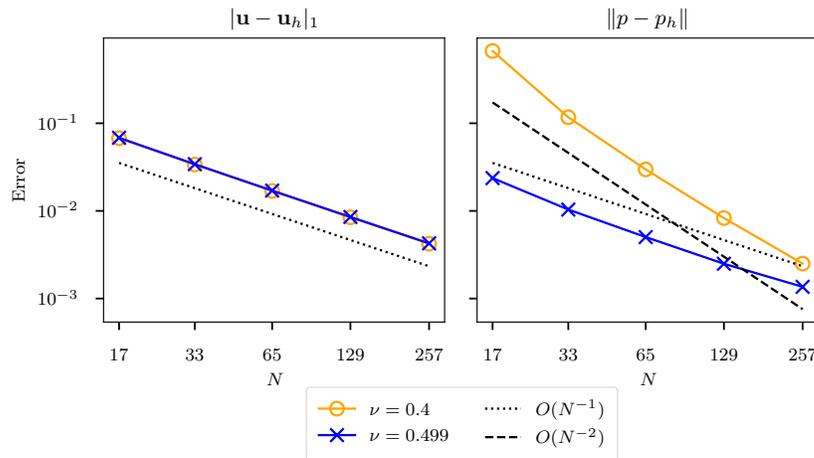


FIG. 5.4. Convergence study for smooth solution problem, using FGMRES preconditioned by monolithic multigrid with the additive Vanka relaxation scheme for $\nu = 0.4$ and $\nu = 0.499$ and $\tau = h = \frac{1}{N-1}$. Left: H^1 -seminorm error for displacement versus mesh size. Right: L^2 -error for pressure versus mesh size.

Again, to validate the discretization, we show finite-element convergence with respect to mesh size and time step size in Figure 5.4, fixing $h = \tau$, with $k = 10^{-6}$ and both $\nu = 0.4$ and $\nu = 0.499$. Expected $O(h + \tau)$ convergence is seen for both the H^1 -seminorm of \mathbf{u} and the L^2 norm of p .

5.3. Terzaghi's problem. Finally, we consider a standard benchmark in poroelasticity. The Terzaghi consolidation problem models a fluid-saturated column of a poroelastic material subject to a loading force on the top [57, 58]; the cylinder height and width are 1.0, so, once again we take $\Omega = [0, 1]^2$. The physical parameters are $\alpha = 1$, $\mu_f = 1$, and $E = 3 \times 10^4$, but we take $M = \infty$ as the Biot modulus. This means that the diagonal block of \mathcal{A}^{RQ} corresponding to the pressure is zero, resulting in vertex-based Vanka blocks that are difficult to invert. To resolve this, a small positive weight of 10^{-8} is added to the diagonal of the Vanka blocks. This test problem has an analytical solution defined by an infinite series,

$$\mathbf{u}(x, y, t) = \frac{p_0}{\lambda + 2\mu} \begin{pmatrix} 1 - x - \sum_{i=0}^{\infty} \frac{8}{\pi^2} \frac{1}{(2i+1)^2} e^{-(2i+1)^2 \pi^2 k(\lambda+2\mu)t/4} \cos\left(\frac{(2i+1)\pi x}{2}\right) \\ 0 \end{pmatrix},$$

$$p(x, y, t) = \frac{4p_0}{\pi} \sum_{i=0}^{\infty} \frac{1}{(2i+1)} e^{-(2i+1)^2 \pi^2 k(\lambda+2\mu)t/4} \cos\left(\frac{(2i+1)\pi x}{2}\right),$$

$$\mathbf{w}(x, y, t) = -k\nabla p,$$

with initial conditions $\mathbf{u}(x, y, 0) = \mathbf{0}$ and $p(x, y, 0) = p_0 = 1.0$. The problem is designed to have $\mathbf{0}$ as the right-hand side.

Parameter robustness for the monolithic multigrid solver is demonstrated in Table 5.3, showing the average solver FMGRES iteration count and average CPU time over 10 time steps using the additive Vanka relaxation applied to the whole system with the 20-DoF vertex-based patch (Figure 3.2, right). A relative residual stopping tolerance of 10^{-6} is used, with mesh spacing fixed to $h = 1/64$, and the values of k and ν are varied. Due to the wide range of physical parameters considered, there is no reasonable single time step size for use with all parameter combinations. Thus,

TABLE 5.3

CPU time in seconds (iterations) per time step, averaged over 10 time steps with $\tau = \hat{\tau}/100$ and $h = 1/64$ for FGMRES preconditioned by monolithic multigrid with additive Vanka relaxation for the Terzaghi problem with varying physical parameters k and ν .

$\nu \backslash k$	10^{-4}	10^{-6}	10^{-8}	10^{-10}
0.4	0.796 (7.3)	0.803 (7.3)	0.796 (7.3)	0.802 (7.3)
0.45	0.827 (7.6)	0.828 (7.6)	0.828 (7.6)	0.822 (7.6)
0.49	0.903 (8.7)	0.900 (8.7)	0.898 (8.7)	0.903 (8.7)
0.499	1.378 (14.5)	1.367 (14.5)	1.376 (14.5)	1.372 (14.5)

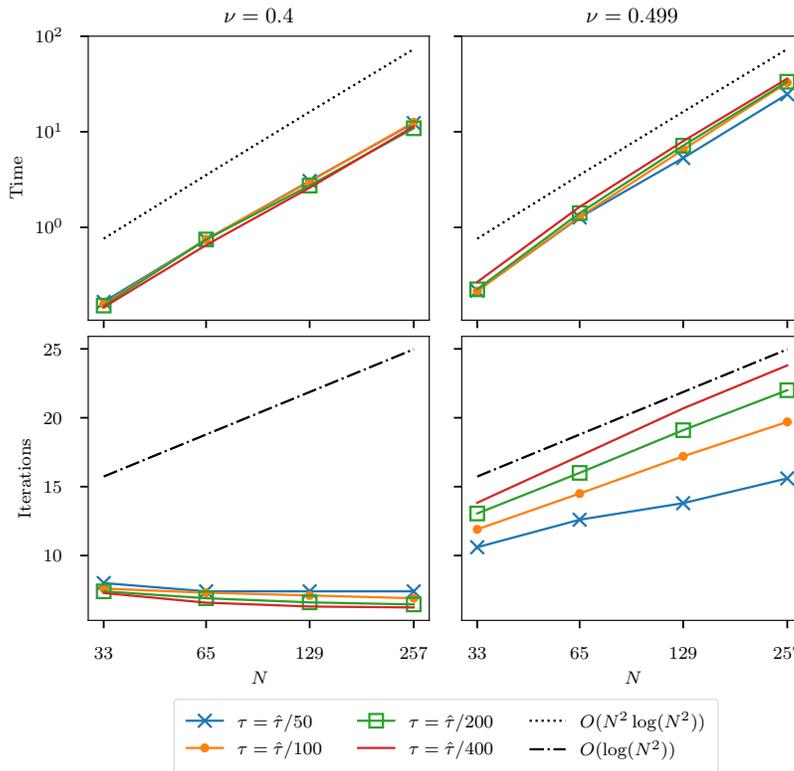


FIG. 5.5. Average CPU time in seconds (top) and iterations to convergence (bottom) per time step over time interval $[0, \hat{\tau}/10]$ for FGMRES preconditioned by monolithic multigrid with additive Vanka relaxation for the Terzaghi problem. Here, $k = 10^{-6}$ and $\nu = 0.4$ (left) and 0.499 (right).

we determine a parameter-dependent time scale, $\hat{\tau} = \frac{1}{0.25\pi^2 k(\lambda+2\mu)}$, derived from the form of the time-dependence in the analytical solution. All tests below simulate from time $t = 0$ to $t = \hat{\tau}/10$. Note that this physical time step size can vary over several orders of magnitude as we vary k and ν , thus the optimal parameters for the steady-state model problem may not be suitable here, so these parameters were recomputed for the Terzaghi problem. Additionally, large values of permeability are not realistic for this type of test problem, so we only consider values of k in the range 10^{-4} to 10^{-10} . The results in Table 5.3 highlight the robustness of the monolithic multigrid method as well as the utility of the LFA relaxation parameter predictions.

In Figure 5.5, we explore the robustness with respect to the time step size, τ , and mesh size (number of points in one direction), N , with $h = 1/(N - 1)$, for $k = 10^{-6}$ and both $\nu = 0.4$ and $\nu = 0.499$, for monolithic multigrid using Vanka relaxation as described above. Note that with a smaller time step size, more time steps are needed to get to the same final time. In all cases, the iteration counts remain stable (no worse than $O(\ln(N^2))$), and the computational time scales as $O(N^2 \log(N^2))$.

6. Conclusions. In this paper, we investigate the construction of parameter-robust preconditioners for three-field models of Biot poroelasticity. Following [2, 53], we consider a bubble-enriched P1-RT0-P0 finite-element discretization; however, in order to allow for robust solvers, we introduce a reduced-quadrature approximation and show that the discretization quality does not suffer from this change. With this, and suitable treatment of divergence-free displacements in both the relaxation and interpolation operators, we derive robust monolithic multigrid methods to solve this problem, with both Vanka and inexact Braess–Sarazin relaxation schemes. In numerical tests, we see that the additive form of Vanka relaxation is more robust than inexact Braess–Sarazin. Both approaches outperform the block-triangular preconditioner of [2], particularly in the incompressible limit. Improving robustness of inexact BSR in the small permeability and nearly incompressible limits is an interesting question for future work.

Another natural topic for future work is extending the preconditioners developed here for the “bubble-eliminated” system described in [2, 53], where an approximate Schur complement is used to remove the face-based displacement DoFs. Additionally, more complicated models of poroelasticity will be considered, including their implementation for three-dimensional models, and for nonlinear models that describe porous materials with fractures, see [13, 23, 24, 50] and references therein. Developing robust multigrid solvers for the linearizations of these systems will aid in the development of fast simulations for real-world problems in the geosciences and biomedical research.

REFERENCES

- [1] J. H. ADLER, T. R. BENSON, AND S. P. MACLACHLAN, *Preconditioning a mass-conserving discontinuous Galerkin discretization of the Stokes equations*, Numer. Linear Algebra Appl., 24 (2017), e2047.
- [2] J. H. ADLER, F. J. GASPAR, X. HU, P. OHM, C. RODRIGO, AND L. T. ZIKATANOV, *Robust preconditioners for a new stabilized discretization of the poroelastic equations*, SIAM J. Sci. Comput., 42 (2020), pp. B761–B791.
- [3] J. H. ADLER, F. J. GASPAR, X. HU, C. RODRIGO, AND L. T. ZIKATANOV, *Robust block preconditioners for Biot’s model*, in Domain Decomposition Methods in Science and Engineering XXIV, Lect. Notes Comput. Sci. Eng. 125, 2018, pp. 3–16.
- [4] J. H. ADLER, X. HU, AND L. T. ZIKATANOV, *HAZmath: A Simple Finite Element, Graph, and Solver Library*, 2014–2020, <https://bitbucket.org/hazmath/hazmath/wiki/Home>.
- [5] D. N. ARNOLD, R. S. FALK, AND R. WINTHER, *Multigrid in $H(\text{div})$ and $H(\text{curl})$* , Numer. Math., 85 (2000), pp. 197–217.
- [6] M. A. BIOT, *General theory of three-dimensional consolidation*, J. Appl. Phys., 12 (1941), pp. 155–164.
- [7] M. A. BIOT, *Theory of elasticity and consolidation for a porous anisotropic solid*, J. Appl. Phys., 26 (1955), pp. 182–185.
- [8] D. BOFFI, F. BREZZI, AND M. FORTIN, *Mixed Finite Element Methods and Applications*, Springer Ser. Comput. Math., Springer, Heidelberg, 2013.
- [9] M. BORREGALES, K. KUMAR, F. A. RADU, C. RODRIGO, AND F. J. GASPAR, *A partially parallel-in-time fixed-stress splitting method for Biot’s consolidation model*, Comput. Math. Appl., 77 (2019), pp. 1466–1478.

- [10] D. BRAESS AND R. SARAZIN, *An efficient smoother for the Stokes problem*, Appl. Numer. Math., 23 (1997), pp. 3–19.
- [11] F. BREZZI, *On the existence, uniqueness and approximation of saddle-point problems arising from Lagrangian multipliers*, RAIRO Anal. Numero. 8 (1974), pp. 129–151.
- [12] J. BROWN, Y. HE, S. MACLACHLAN, M. MENICKELLY, AND S. M. WILD, *Tuning multigrid methods with robust optimization and local Fourier analysis*, SIAM J. Sci. Comput., 43 (2021), pp. A109–A138.
- [13] A. BUDIŠA AND X. HU, *Block preconditioners for mixed-dimensional discretization of flow in fractured porous media*, Comput. Geosci., 25 (2021), pp. 671–686.
- [14] C. RODRIGO CARDIEL, *Geometric Multigrid Methods on Semi-Structured Triangular Grids*, Ph.D. thesis, Universidad de Zaragoza, 2010.
- [15] N. CASTELLETTO, J. A. WHITE, AND M. FERRONATO, *Scalable algorithms for three-field mixed finite element coupled poromechanics*, J. Comput. Phys., 327 (2016), pp. 894–918.
- [16] N. CASTELLETTO, J. A. WHITE, AND H. A. TCHELEPI, *Accuracy and convergence properties of the fixed-stress iterative solution of two-way coupled poromechanics*, Internat. J. Numer. Anal. Methods Geomechanics, 39 (2015), pp. 1593–1618.
- [17] M. S. ENGELMAN, R. L. SANI, P. M. GRESHO, AND M. BERCOVIER, *Consistent vs. reduced integration penalty methods for incompressible media using several old and new elements*, Internat. J. Numer. Methods Fluids, 2 (1982), pp. 25–42.
- [18] P. E. FARRELL, Y. HE, AND S. P. MACLACHLAN, *A local Fourier analysis of additive Vanka relaxation for the Stokes equations*, Numer. Linear Algebra Appl., 28 (2021), e2306.
- [19] P. E. FARRELL, M. G. KNEPLEY, L. MITCHELL, AND F. WECHSUNG, *PCPATCH: Software for the topological construction of multigrid relaxation methods*, ACM Trans. Math. Software, 47 (2021).
- [20] P. E. FARRELL, L. MITCHELL, L. R. SCOTT, AND F. WECHSUNG, *Robust Multigrid Methods for Nearly Incompressible Elasticity Using Macro Elements*, arXiv:2002.02051, 2020.
- [21] M. FERRONATO, L. BERGAMASCHI, AND G. GAMBOLATI, *Performance and robustness of block constraint preconditioners in finite element coupled consolidation problems*, Internat. J. Numer. Methods Engrg., 81 (2010), pp. 381–402.
- [22] M. FERRONATO, A. FRANCESCHINI, C. JANNA, N. CASTELLETTO, AND H. A. TCHELEPI, *A general preconditioning framework for coupled multiphysics problems with application to contact- and poro-mechanics*, J. Comput. Phys., 398 (2019), 108887.
- [23] B. FLEMISCH, I. BERRE, W. BOON, A. FUMAGALLI, N. SCHWENCK, A. SCOTTI, I. STEFANSSON, AND A. TATOMIR, *Benchmarks for single-phase flow in fractured porous media*, Adv. Water Resources, 111 (2018), pp. 239–258.
- [24] B. FLEMISCH, A. FUMAGALLI, AND A. SCOTTI, *A review of the XFEM-based approximation of flow in fractured porous media*, in Advances in Discretization Methods: Discontinuities, Virtual Elements, Fictitious Domain Methods, Giulio Ventura and Elena Benvenuti, eds., SEMA SIMAI Springer Ser., Springer, Cham, 2016, pp. 47–76.
- [25] S. R. FRANCO, C. RODRIGO, F. J. GASPAR, AND M. A. V. PINTO, *A multigrid waveform relaxation method for solving the poroelasticity equations*, Comput. Appl. Math., 37 (2018), pp. 4805–4820.
- [26] M. FRIGO, N. CASTELLETTO, M. FERRONATO, AND J. A. WHITE, *Efficient solvers for hybridized three-field mixed finite element coupled poromechanics*, Comput. Math. Appl., 91 (2021), pp. 36–52.
- [27] G. FU, *A high-order HDG method for the Biot’s consolidation model*, Comput. Math. Appl., 77 (2019), pp. 237–252.
- [28] F. J. GASPAR, Y. NOTAY, C. W. OOSTERLEE, AND C. RODRIGO, *A simple and efficient segregated smoother for the discrete Stokes equations*, SIAM J. Sci. Comput., 36 (2014), pp. A1187–A1206.
- [29] F. J. GASPAR AND C. RODRIGO, *On the fixed-stress split scheme as smoother in multigrid methods for coupling flow and geomechanics*, Comput. Methods Appl. Mech. Engrg., 326 (2017), pp. 526–540.
- [30] V. GIRAULT AND P.-A. RAVIART, *Finite Element Methods for Navier-Stokes Equations*, Springer Ser. Comput. Math., Springer-Verlag, Berlin, 1986.
- [31] Y. HE AND S. P. MACLACHLAN, *Local Fourier analysis of block-structured multigrid relaxation schemes for the Stokes equations*, Numer. Linear Algebra Appl., 25 (2018), e2147.
- [32] Y. HE AND S. P. MACLACHLAN, *Local Fourier analysis for mixed finite-element methods for the Stokes equations*, J. Comput. Appl. Math., 357 (2019), pp. 161–183.
- [33] Y. HE AND S. P. MACLACHLAN, *Two-level Fourier analysis of multigrid for higher-order finite-element discretizations of the Laplacian*, Numer. Linear Algebra Appl., 27 (2020), e2285.

- [34] Q. HONG AND J. KRAUS, *Parameter-robust stability of classical three-field formulation of Biot's consolidation model*, Electron. Trans. Numer. Anal., 48 (2018), pp. 202–226.
- [35] X. HU, L. MU, AND X. YE, *Weak Galerkin method for the Biot's consolidation model*, Comput. Math. Appl., 75 (2018), pp. 2017–2030.
- [36] X. HU, C. RODRIGO, F. J. GASPAS, AND L. T. ZIKATANOV, *A nonconforming finite element method for the Biot's consolidation model in poroelasticity*, J. Comput. Appl. Math., 310 (2017), pp. 143–154.
- [37] V. JOHN AND L. TOBISKA, *Numerical performance of smoothers in coupled multigrid methods for the parallel solution of the incompressible Navier-Stokes equations*, Internat. J. Numer. Methods Fluids, 33 (2000), pp. 453–473.
- [38] J. KIM, H. A. TCHELEPI, AND R. JUANES, *Stability and convergence of sequential methods for coupled flow and geomechanics: Fixed-stress and fixed-strain splits*, Comput. Methods Appl. Mech. Engrg., 200 (2011), pp. 1591–1606.
- [39] M. LARIN AND A. REUSKEN, *A comparative study of efficient iterative solvers for generalized Stokes equations*, Numer. Linear Algebra Appl., 15 (2008), pp. 13–34.
- [40] J. J. LEE, *Robust error analysis of coupled mixed methods for Biot's consolidation model*, J. Sci. Comput., 69 (2016), pp. 610–632.
- [41] P. LUO, C. RODRIGO, F. J. GASPAS, AND C. W. OOSTERLEE, *On an Uzawa smoother in multigrid for poroelasticity equations*, Numer. Linear Algebra Appl., 24 (2017), e2074.
- [42] S. P. MACLACHLAN AND C. W. OOSTERLEE, *Local Fourier analysis for multigrid with overlapping smoothers applied to systems of PDEs*, Numer. Linear Algebra Appl., 18 (2011), pp. 751–774.
- [43] D. S. MALKUS AND T. J. R. HUGHES, *Mixed finite element methods—reduced and selective integration techniques: A unification of concepts*, Comput. Methods Appl. Mech. Engrg., 15 (1978), pp. 63–81.
- [44] K.-A. MARDAL, M. E. ROGNES, AND T. B. THOMPSON, *Accurate discretization of poroelasticity without darcy stability*, BIT, 61 (2021), pp. 941–976.
- [45] M. A. MURAD AND A. F. D. LOULA, *Improved accuracy in finite element analysis of Biot's consolidation problem*, Comput. Methods Appl. Mech. Engrg., 95 (1992), pp. 359–382.
- [46] M. A. MURAD AND A. F. D. LOULA, *On stability and convergence of finite element approximations of Biot's consolidation problem*, Internat. J. Numer. Methods Engrg., 37 (1994), pp. 645–667.
- [47] M. A. MURAD, V. THOMÉE, AND A. F. D. LOULA, *Asymptotic behavior of semidiscrete finite-element approximations of Biot's consolidation problem*, SIAM J. Numer. Anal., 33 (1996), pp. 1065–1083.
- [48] A. NAUMOVICH AND F. J. GASPAS, *On a multigrid solver for the three-dimensional Biot poroelasticity system in multilayered domains*, Comput. Vis. Sci., 11 (2008), pp. 77–87.
- [49] C. NIU, H. RUI, AND X. HU, *A stabilized hybrid mixed finite element method for poroelasticity*, Comput. Geosci., 25 (2021), pp. 757–774.
- [50] J. M. NORDBOTTEN, W. M. BOON, A. FUMAGALLI, AND E. KEILEGAVLEN, *Unified approach to discretization of flow in fractured porous media*, Comput. Geosci., 23 (2019), pp. 225–237.
- [51] C. RODRIGO, F. J. GASPAS, X. HU, AND L. T. ZIKATANOV, *Stability and monotonicity for some discretizations of the Biot's consolidation model*, Comput. Methods Appl. Mech. Engrg., 298 (2016), pp. 183–204.
- [52] C. RODRIGO, F. J. GASPAS, AND F. J. LISBONA, *On a local Fourier analysis for overlapping block smoothers on triangular grids*, Appl. Numer. Math., 105 (2016), pp. 96–111.
- [53] C. RODRIGO, X. HU, P. OHM, J. H. ADLER, F. J. GASPAS, AND L. T. ZIKATANOV, *New stabilized discretizations for poroelasticity and the Stokes' equations*, Comput. Methods Appl. Mech. Engrg., 341 (2018), pp. 467–484.
- [54] J. SCHÖBERL, *Robust multigrid preconditioning for parameter-dependent problems I: The Stokes-type case*, in Multigrid Methods V, W. Hackbusch and G. Wittum, eds., Springer, Berlin, 1998, pp. 260–275.
- [55] J. SCHÖBERL, *Multigrid methods for a parameter dependent problem in primal variables*, Numer. Math, 84 (1999), pp. 97–119.
- [56] J. SCHÖBERL, *Robust Multigrid Methods for Parameter Dependent Problems*, Ph.D. thesis, Johannes Kepler Universität Linz, Linz, Austria, 1999.
- [57] K. TERZAGHI, *Erdbaumechanik auf Bodenphysikalischer Grundlage*, F. Deuticke, Vienna, 1925.
- [58] K. TERZAGHI, *Theoretical Soil Mechanics*, Wiley, New York, 1943.
- [59] U. TROTTEMBERG, C. W. OOSTERLEE, AND A. SCHÜLLER, *Multigrid*, Academic Press, San Diego, CA, 2001.
- [60] S. P. VANKA, *Block-implicit multigrid calculation of two-dimensional recirculating flows*, Comput. Methods Appl. Mech. Engrg., 59 (1986), pp. 29–48.

- [61] J. A. WHITE, N. CASTELLETTO, AND H. A. TCHELEPI, *Block-partitioned solvers for coupled poromechanics: A unified framework*, *Comput. Methods Appl. Mech. Engrg.*, 303 (2016), pp. 55–74.
- [62] R. WIENANDS AND W. JOPPICH, *Practical Fourier Analysis for Multigrid Methods*, CRC Press, Boca Raton, FL, 2004.
- [63] S.-Y. YI, *A coupling of nonconforming and mixed finite element methods for Biot's consolidation model*, *Numer. Methods Partial Differential Equations*, 29 (2013), pp. 1749–1777.
- [64] S.-Y. YI, *Convergence analysis of a new mixed finite element method for Biot's consolidation model*, *Numer. Methods Partial Differential Equations*, 30 (2014), pp. 1189–1210.
- [65] S.-Y. YI, *A study of two modes of locking in poroelasticity*, *SIAM J. Numer. Anal.*, 55 (2017), pp. 1915–1936.