

PRECONDITIONERS FOR COMPUTING MULTIPLE SOLUTIONS IN THREE-DIMENSIONAL FLUID TOPOLOGY OPTIMIZATION*

IOANNIS P. A. PAPADOPOULOS[†] AND PATRICK E. FARRELL[‡]

Abstract. Topology optimization problems generally support multiple local minima, and real-world applications are typically three-dimensional. In previous work [I. P. A. Papadopoulos, P. E. Farrell, and T. M. Surowiec, *SIAM J. Sci. Comput.*, 43 (2021), pp. A1555–A1582], the authors developed the deflated barrier method, an algorithm that can systematically compute multiple solutions of topology optimization problems. In this work, we develop preconditioners for the linear systems arising in the application of this method to Stokes flow, making it practical for use in three dimensions. In particular, we develop a nested block preconditioning approach which reduces the linear systems to solving two symmetric positive-definite matrices and an augmented momentum block. An augmented Lagrangian term is used to control the innermost Schur complement, and we apply a geometric multigrid method with a kernel-capturing relaxation method for the augmented momentum block. We present multiple solutions in three-dimensional examples computed using the proposed iterative solver.

Key words. topology optimization, multiple solutions, deflated barrier method, preconditioning, multigrid

MSC codes. 65F08, 65K10, 65N55, 35Q35, 49M15, 90C26

DOI. 10.1137/22M1478598

1. Introduction. Topology optimization has become a prominent tool in engineering design. The objective is to find the optimal distribution of a continuum within a given domain that minimizes a problem-specific cost functional, without prior knowledge of the shape or topology of the solution [10].

Topology optimization problems are generally nonconvex and can support multiple local minima. A major challenge in topology optimization is to identify multiple local minimizers, so that the best one (in performance, manufacturability, or aesthetics) may be chosen. Continuation in model parameters is often used to promote convergence to “better” local minima. However, Stolpe and Svanberg [65] show that such methods can be ineffective even on the simplest of topology optimization models.

In recent work, Papadopoulos, Farrell, and Surowiec [51] developed an algorithm, called the deflated barrier method, that can systematically discover *multiple* stationary points of topology optimization problems formulated using a density approach. The deflated barrier method combines the techniques of barrier methods

*Submitted to the journal’s Computational Methods in Science and Engineering section February 16, 2022; accepted for publication (in revised form) June 9, 2023; published electronically November 29, 2023.

<https://doi.org/10.1137/22M1478598>

Funding: The work of the first author was supported by the EPSRC Centre for Doctoral Training in Partial Differential Equations: Analysis and Applications grant EP/L015811/1 and The MathWorks, Inc. Revisions were supported by the EPSRC grant “Spectral element methods for fractional differential equations, with applications in applied analysis and medical imaging” grant EP/T022132/1 and the Leverhulme Trust Research Project “Constructive approximation theory on and inside algebraic curves and surfaces” grant RPG-2019-144. The work of the second author was supported by the Engineering and Physical Sciences Research Council grants EP/R029423/1 and EP/W026163/1.

[†]Department of Mathematics, Imperial College London, London SW7 2AZ, UK (ioannis.papadopoulos13@imperial.ac.uk).

[‡]Mathematical Institute, University of Oxford, Oxford OX2 6GG, UK (patrick.farrell@maths.ox.ac.uk).

[32, 33, 35, 60, 61, 67, 70], primal-dual active set solvers [11], and deflation [25, 26]. In one example [51, Fig. 1], the authors found 42 stationary points of a fluid topology optimization problem constrained by the Navier–Stokes equations in a rectangular domain with five small holes removed. However, an application of the deflated barrier method requires the solution of many linear systems similar to those solved in an all-at-once (simultaneous analysis and design, SAND) approach. In [51], a sparse LU factorization was used to solve these linear systems, which is not generally affordable for three-dimensional problems.

In this work, we develop preconditioners for the linear systems arising in the deflated barrier method when applied to the model proposed by Borrvall and Petersson [14]. The goal of this model is to minimize the power dissipation of a fluid in Stokes flow, subject to a volume constraint restricting the proportion of the domain that the fluid can occupy. The preconditioners we develop make it feasible to identify multiple solutions to the Borrvall–Petersson problem in three dimensions with the deflated barrier method. A number of strategies have been proposed for the solution of three-dimensional topology optimization of fluid flow in other contexts [1, 2, 3, 9, 19, 20, 24, 41, 53, 58, 74, 76].

We choose a discontinuous Galerkin $H(\text{div})$ -conforming finite element discretization for the velocity and pressure pair [15, 16] due to its pressure-robustness [42] and its crisp characterization of the kernel of the divergence operator. We will show that block preconditioning can reduce the linear systems arising in the deflated barrier method to ones that resemble the systems arising in the discretization of the Stokes–Brinkman equations [24]. Then, we apply modern block preconditioning (pioneered by Wathen and coworkers [21, 64, 69]) to solve the inner linear systems. The inner linear systems are themselves solved with block preconditioning, using an augmented Lagrangian term to control the innermost Schur complement [34]. Finally, we develop a geometric multigrid method for the augmented momentum block with a vertex-star patch relaxation that captures the kernel of the augmented Lagrangian term [28, 30, 39, 62]. The multigrid scheme also requires a characterization of the active set on all levels of the mesh hierarchy, which we discuss.

We introduce the Borrvall–Petersson model in section 2 and the deflated barrier method in section 3. In section 4 we detail the discontinuous Galerkin finite element discretization and derive the linear systems that arise in the subproblems of the deflated barrier method. In section 5 we develop the preconditioning strategies for these linear systems. In section 6 we investigate a number of two- and three-dimensional examples. All the examples support multiple solutions which are discovered by the deflated barrier method, and all are solved with our proposed iterative methods.

2. Topology optimization of Stokes flow. Given a volume constraint on a fluid in a fixed Lipschitz domain $\Omega \subset \mathbb{R}^d$, $d \in \{2, 3\}$, the Borrvall–Petersson model attempts to minimize the energy lost by the flow due to viscous dissipation, while maximizing the flow velocities at the applied body force. More precisely, the objective is to find $(\mathbf{u}, \rho) \in H_{\mathbf{g}, \text{div}}^1(\Omega)^d \times C_\gamma$ that minimizes

$$(BP) \quad J(\mathbf{u}, \rho) := \frac{1}{2} \int_{\Omega} (\alpha(\rho)|\mathbf{u}|^2 + \nu|\nabla \mathbf{u}|^2 - 2\mathbf{f} \cdot \mathbf{u}) \, dx,$$

where $\mathbf{u} : \Omega \rightarrow \mathbb{R}^d$ denotes the velocity of the fluid, $\rho : \Omega \rightarrow \mathbb{R}$ is the material distribution of the fluid, and

$$H_{\mathbf{g}}^1(\Omega)^d := \{\mathbf{v} \in H^1(\Omega)^d : \mathbf{v}|_{\partial\Omega} = \mathbf{g}\},$$

$$H_{\mathbf{g},\text{div}}^1(\Omega)^d := \{\mathbf{v} \in H_{\mathbf{g}}^1(\Omega)^d : \text{div}(\mathbf{v}) = 0 \text{ a.e. in } \Omega\},$$

$$C_\gamma := \left\{ \eta \in L^\infty(\Omega) : 0 \leq \eta \leq 1 \text{ a.e.}, \int_\Omega \eta \, dx \leq \gamma |\Omega| \right\}.$$

Here, $H^1(\Omega)$ denotes the Sobolev space $W^{1,2}(\Omega)$, equipped with the inner product

$$(2.1) \quad (u, v)_{H^1(\Omega)} = \int_\Omega uv + \nabla u \cdot \nabla v \, dx,$$

which induces the norm $\|v\|_{H^1(\Omega)} := (v, v)_{H^1(\Omega)}^{1/2}$. $L^\infty(\Omega)$ denotes the vector space of essentially bounded measurable functions equipped with the essential supremum norm, i.e., $\|\eta\|_{L^\infty(\Omega)} := \inf\{c \geq 0 : |\eta(x)| \leq c \text{ for a.e. } x \in \Omega\}$. Furthermore, $\mathbf{f} \in L^2(\Omega)^d$ is a body force, $\nu > 0$ is the (constant) viscosity, and $\gamma \in (0, 1)$ is the volume fraction. The restriction, $|\partial\Omega|$, is to be understood in the boundary trace sense [23, Chap. 5.5]. Moreover, the boundary data $\mathbf{g} \in H^{1/2}(\partial\Omega)^d$ and $\mathbf{g} = \mathbf{0}$ on $\Gamma \subset \partial\Omega$, with $\mathcal{H}^{d-1}(\Gamma) > 0$, i.e., Γ has nonzero Hausdorff measure on the boundary. Borrvall and Petersson introduced the inverse permeability term, α , which models the influence of the material distribution on the flow. For values of ρ close to one, $\alpha(\rho)$ is small, permitting fluid flow; for values of ρ close to zero, $\alpha(\rho)$ is very large, restricting fluid flow. The function α satisfies the following properties:

- (A1) $\alpha : [0, 1] \rightarrow [\underline{\alpha}, \bar{\alpha}]$ with $0 \leq \underline{\alpha} < \bar{\alpha} < \infty$;
- (A2) α is monotonically decreasing;
- (A3) $\alpha(0) = \bar{\alpha}$ and $\alpha(1) = \underline{\alpha}$;
- (A4) α is twice continuously differentiable;
- (A5) α is strongly convex.

It generates an operator also denoted $\alpha : C_\gamma \rightarrow L^\infty(\Omega; [\underline{\alpha}, \bar{\alpha}])$. Typically, in the literature, α takes the form [14, 24]

$$(2.2) \quad \alpha(\rho) = \bar{\alpha} \left(1 - \frac{\rho(q+1)}{\rho+q} \right),$$

where $q > 0$ is a penalty parameter, so that $\lim_{q \rightarrow \infty} \alpha(\rho) = \bar{\alpha}(1 - \rho)$. Borrvall and Petersson [14, sect. 3.2] remark that as $q \rightarrow \infty$, the material distribution tends to a 0-1 solution.

The following existence theorem is due to Borrvall and Petersson [14, Thm. 3.1].

THEOREM 2.1. *Suppose that $\Omega \subset \mathbb{R}^d$ is a Lipschitz domain, with $d = 2$ or $d = 3$, and α satisfies properties (A1)–(A5). Then, there exists a pair $(\mathbf{u}, \rho) \in H_{\mathbf{g},\text{div}}^1(\Omega) \times C_\gamma$ that minimizes J (as defined in (BP)).*

3. The deflated barrier method. In this section, we summarize how the deflated barrier method finds multiple solutions to fluid topology optimization problems.

3.1. Forming the barrier functional. The first step in the algorithm is to form a barrier-like objective from the original objective functional. This is to both aid convergence and provide an opportunity for deflation to discover multiple local minima (described in section 3.2). Consider the following minimization problem: find the minimizers, $(\mathbf{u}, \rho) \in H_{\mathbf{g},\text{div}}^1(\Omega)^d \times C_\gamma$, of

$$(BF) \quad J_\mu^{\epsilon_{\log}}(\mathbf{u}, \rho) := J(\mathbf{u}, \rho) - \mu \int_\Omega \log(\rho + \epsilon_{\log}) + \log(1 + \epsilon_{\log} - \rho) \, dx,$$

where $\mu \geq 0$ is the barrier parameter and $0 < \epsilon_{\log} \ll 1$. We note that the box constraints imposed by the barrier-like terms are never active as $0 \leq \rho \leq 1$ a.e. in Ω .

A result similar to one shown by Evgrafov [24, sect. 4] guarantees the existence of a minimizer to (BF) [49, Prop. 4.2]. As with the original optimization problem, (BP), the minimizer is not necessarily unique.

The deflated barrier method targets the first-order optimality conditions of this problem, which we derive next. Define the following forms to be used in stating the first-order optimality conditions:

$$(3.1) \quad c(\rho, \eta; \mathbf{u}, \lambda) := \int_{\Omega} \left[\frac{1}{2} \alpha'(\rho) |\mathbf{u}|^2 - \frac{\mu}{\rho + \epsilon_{\log}} + \frac{\mu}{1 + \epsilon_{\log} - \rho} + \lambda \right] \eta \, dx,$$

$$(3.2) \quad a(\mathbf{u}, \mathbf{v}; \rho) := \int_{\Omega} \alpha(\rho) \mathbf{u} \cdot \mathbf{v} + \nu \nabla \mathbf{u} : \nabla \mathbf{v} \, dx, \quad l(\mathbf{v}) := \int_{\Omega} \mathbf{f} \cdot \mathbf{v} \, dx,$$

$$(3.3) \quad b(\mathbf{u}, p) := - \int_{\Omega} p \nabla \cdot \mathbf{u} \, dx, \quad d(\lambda, \zeta; \rho) := - \zeta \int_{\Omega} (\gamma - \rho) \, dx.$$

The proof of the following proposition, concerning the first-order optimality conditions of (BF), follows from the result in [49, Prop. 2.5].

PROPOSITION 3.1. *Fix $\mu \geq 0$, $0 < \epsilon_{\log} \ll 1$ and suppose that $(\mathbf{u}, \rho) \in H_{\mathbf{g}, \text{div}}^1(\Omega)^d \times C_{\gamma}$ is a strict local or global minimizer of (BF). Then, there exist unique Lagrange multipliers $p \in L_0^2(\Omega)$ and $\lambda \in \mathbb{R}$, such that, for all $(\eta, \mathbf{v}, q, \zeta) \in C_{[0,1]} \times H_0^1(\Omega)^d \times L_0^2(\Omega) \times \mathbb{R}$, the following necessary first-order optimality conditions are satisfied:*

$$(3.4) \quad c(\rho, \eta - \rho; \mathbf{u}, \lambda) \geq 0,$$

$$(3.5) \quad a(\mathbf{u}, \mathbf{v}; \rho) + b(\mathbf{v}, p) = l(\mathbf{v}),$$

$$(3.6) \quad b(\mathbf{u}, q) = 0,$$

$$(3.7) \quad d(\lambda, \zeta; \rho) = 0.$$

Here $C_{[0,1]} := \cup_{\gamma \in [0,1]} C_{\gamma}$, i.e., we relax the volume constraint on the variation in the material distribution.

Under suitable conditions, (3.4)–(3.7) can be reformulated as a complementarity problem. After a suitable discretization, throughout the deflated barrier method, the nonlinear complementarity problems are solved using the Benson–Munson (BM) algorithm [11], a Newton-like algorithm that can enforce the true box constraints on ρ . A thorough description of this strategy is given in [51, sect. 3.1] and subsection 4.2. We emphasize that the barrier-like terms are *not* used to enforce box constraints, which are handled by BM, and the enlarged box constraints $[-\epsilon_{\log}, 1 + \epsilon_{\log}]$ are never active for any $\mu \geq 0$. The barrier-like terms are included solely to aid nonlinear convergence via continuation, and to find multiple solutions via deflation.

3.2. Deflation. Deflation is a mechanism for systemically discovering multiple solutions of a nonlinear system with a Newton-like algorithm [25, 26]. Let Z and Y be Banach spaces. Consider the nonlinear system $F(z) = 0$, $F : Z \rightarrow Y$ that has the solutions z_1, \dots, z_l . Given a suitable initial guess, a Newton-like algorithm might converge to the solution z_1 . Now, deflation modifies the nonlinear system in order to remove the known solution while retaining undiscovered solutions. This is done via a deflation operator $\mathcal{M}(z; z_1) : Y \rightarrow Y$ that ensures that $\mathcal{M}(z; z_1)F(z)$ has a root if the original problem $F(z) = 0$ has an unknown solution, and ensures that a Newton-like method applied to the newly deflated system does not converge to $z = z_1$. In this work, the following shifted deflation operator [25] is used,

$$(3.8) \quad \mathcal{M}(z; z_1) = \left(\frac{1}{\|\rho - \rho_1\|_{L^2(\Omega)}^2} + 1 \right) \mathcal{I},$$

where $\mathbf{z} = (\rho, \mathbf{u}, p, \lambda)$, and $\mathcal{I} : Y \rightarrow Y$ is the identity operator. Deflation can be very efficiently implemented; the BM update of the deflated system can be expressed as a (nonlinear) scaling of the BM update of the *undeflated* system evaluated at the same iterate [25, sect. 3]. This implementation detail is crucial for this work as it implies that preconditioning strategies for the undeflated BM systems can be immediately applied to the deflated systems. A discussion of how the nonlinear scaling is cheaply computed is given in [51, sect. 3.2].

3.3. Prediction, continuation, and deflation. The deflated barrier method [51, Alg. 3.1] is an iterative procedure consisting of three steps. The algorithm is initialized with barrier parameter $\mu = \mu_0$, and the idea is to follow branches (a generalization of the central path) of solutions as $\mu \rightarrow 0$. At iteration k of the deflated barrier method, the barrier functional subproblem at $\mu = \mu_k$ is constructed. In the *prediction step*, a cheap initial guess is computed for the barrier functional subproblem from the solution on the same branch at $\mu = \mu_{k-1}$. In the *continuation step*, the BM solver is used to find a solution to the first-order optimality conditions (3.4)–(3.7) at $\mu = \mu_k$, using the initial guess computed in the prediction step, while deflating away all other known solutions at $\mu = \mu_k$. In the *deflation step*, we search for new branches of solutions at $\mu = \mu_k$, using solutions discovered at $\mu = \mu_{k-1}$ as initial guesses. The deflation step is terminated if a pre-determined number of branches have already been found or the BM solver fails to converge (by reaching a specified number of iterations without converging). The algorithm terminates when the first-order optimality conditions have been solved (for multiple solutions) at a barrier parameter value of zero.

4. Discretization and linearization.

4.1. Finite element discretization. The linear systems that arise in the deflated barrier method will be tied to our choice of finite element discretization. In this work, we choose a discontinuous Galerkin (DG) k th-order $\text{BDM}_k \times \text{DG}_{k-1}$ Brezzi–Douglas–Marini finite element discretization for the velocity and the pressure [15, 16]. Let \mathcal{T}_h denote a shape-regular triangulation of the domain Ω . We define the spaces $\mathbf{X}_{\text{BDM}_k}$ and $\mathbf{X}_{\text{BDM}_k^g}$ as

$$(4.1) \quad \mathbf{X}_{\text{BDM}_k} := \{\mathbf{v} \in L^2(\Omega)^d : \nabla \cdot \mathbf{v} \in L^2(\Omega), \mathbf{v}|_T \in \mathcal{P}_k(T)^d \forall T \in \mathcal{T}_h\},$$

$$(4.2) \quad \mathbf{X}_{\text{BDM}_k^g} := \{\mathbf{v} \in \mathbf{X}_{\text{BDM}_k} : (\mathbf{v} - \mathbf{g}) \cdot \mathbf{n} = 0 \text{ on } \partial\Omega\},$$

where \mathcal{P}_k denotes the set of polynomials of order k . The degrees of freedom for two dimensions are given in [16, sect. 2] and for three dimensions in [15, sect. 2]. Similarly, $X_{\text{DG}_{k-1}}$ denotes the set of discontinuous piecewise $(k-1)$ th order polynomials,

$$(4.3) \quad X_{\text{DG}_{k-1}} := \{v \in L^2(\Omega) : v|_T \in \mathcal{P}_{k-1}(T) \forall T \in \mathcal{T}_h\}.$$

Some finite element methods for fluid flow, such as the Taylor–Hood finite element pair, do not satisfy the incompressibility constraint $\nabla \cdot \mathbf{u} = 0$ pointwise. Failure to satisfy the incompressibility constraint pointwise has been observed to support instabilities that result in nonphysical solutions [42, 45]. In the BDM finite element pair, $\nabla \cdot \mathbf{X}_{\text{BDM}_k} \subset X_{\text{DG}_{k-1}}$. Therefore, for any solution \mathbf{u}_h satisfying the incompressibility constraint (3.6), we have that $\|\nabla \cdot \mathbf{u}_h\|_{L^2(\Omega)} = 0$. Hence, the solution is pointwise divergence-free.

The finite element space for the material distribution is denoted by $C_{[0,1],h}$ and is defined by

$$(4.4) \quad C_{[0,1],h} := \{\eta_h \in X_{\text{DG}_0} : 0 \leq \eta_h \leq 1\},$$

where X_{DG_0} is the set of piecewise constant finite element functions. As the test functions of DG_0 take the value of one in their respective element and are zero elsewhere, in practice we discretize the material distribution with DG_0 functions and allow the Benson–Munson strategy described below to handle the box constraints.

Since the discontinuous velocity space is not H^1 -conforming, we use an interior penalty to penalize jumps across edges and faces. We now describe the DG discretization of (3.4)–(3.7) as found in [36, sect. 7.1] and [48, sect. 3]. The interior penalty only arises in the discretization of (3.5).

Let the set \mathcal{F}_h denote the set of all facets of the triangulation \mathcal{T}_h , and let h_F represent the diameter of each facet $F \in \mathcal{F}_h$. We split the set of facets into the union $\mathcal{F}_h = \mathcal{F}_h^i \cup \mathcal{F}_h^\partial$ where \mathcal{F}_h^i is the subset of interior facets and \mathcal{F}_h^∂ collects all Dirichlet boundary facets $F \subset \partial\Omega$. For every facet $F \in \mathcal{F}_h$, we assign a unit norm vector \mathbf{n}_F , where if $F \in \mathcal{F}_h^\partial$, then \mathbf{n}_F is the outer unit normal vector \mathbf{n} . If $F \in \mathcal{F}_h^i$, then $F = \partial K^+ \cap \partial K^-$ for two elements $K^-, K^+ \in \mathcal{T}_h$, and \mathbf{n}_F points in an arbitrary but fixed direction. Let $\phi \in (X_{\text{DG}_k})^d$ and $\Phi \in (X_{\text{DG}_k})^{d \times d}$ be any piecewise vector- or matrix-valued function, with traces from within the interior of K^\pm denoted by ϕ^\pm and Φ^\pm , respectively. We define the jump $[\![\cdot]\!]_F$ and the average $\{\!\{\cdot\}\!\}_F$ operators across interior facets $F \in \mathcal{F}_h^i$ by

$$(4.5) \quad [\![\phi]\!]_F = \phi^+ \otimes \mathbf{n}_F^+ + \phi^- \otimes \mathbf{n}_F^- \quad \text{and} \quad \{\!\{\Phi\}\!\}_F = \frac{1}{2}(\Phi^+ + \Phi^-).$$

If $F \in \mathcal{F}_h^\partial$, we set $[\![\phi]\!]_F = \phi \otimes \mathbf{n}_F$ and $\{\!\{\Phi\}\!\}_F = \Phi$. We note that, for any $F \in \mathcal{F}_h^\partial$, $\int_F |[\![\phi]\!]_F|^2 ds = \int_F |\phi|^2 ds$. Finally, for a sufficiently large penalization parameter $\sigma > 0$, we define the broken form $a_h(\mathbf{u}, \mathbf{v}; \rho)$ by

$$(4.6) \quad \begin{aligned} a_h(\mathbf{u}, \mathbf{v}; \rho) := & \sum_{K \in \mathcal{T}_h} \int_K \alpha(\rho) \mathbf{u} \cdot \mathbf{v} + \nu \nabla \mathbf{u} : \nabla \mathbf{v} \, dx + \nu \sum_{F \in \mathcal{F}_h} \sigma h_F^{-1} \int_F [\![\mathbf{u}]\!]_F : [\![\mathbf{v}]\!]_F \, ds \\ & - \nu \sum_{F \in \mathcal{F}_h} \int_F \{\!\{\nabla \mathbf{u}\}\!\}_F : [\![\mathbf{v}]\!]_F \, ds - \nu \sum_{F \in \mathcal{F}_h} \int_F [\![\mathbf{u}]\!]_F : \{\!\{\nabla \mathbf{v}\}\!\}_F \, ds, \end{aligned}$$

and the linear functional l_h as

$$(4.7) \quad l_h(\mathbf{v}) := \int_\Omega \mathbf{f} \cdot \mathbf{v} \, dx + \nu \sum_{F \in \mathcal{F}_h^\partial} \sigma h_F^{-1} \int_F [\![\mathbf{g}]\!]_F : [\![\mathbf{v}]\!]_F \, ds - \nu \sum_{F \in \mathcal{F}_h^\partial} \int_F [\![\mathbf{g}]\!]_F : \{\!\{\nabla \mathbf{v}\}\!\}_F \, ds.$$

Then, given a barrier parameter μ and a penalization parameter $\sigma > 0$, the discretized deflated barrier subproblem is to find $(\rho_h, \mathbf{u}_h, p_h, \lambda_h) \in C_{\gamma,h} \times \mathbf{X}_{\text{BDM}_k^g} \times X_{\text{DG}_{k-1}} \times \mathbb{R}$ such that, for all $(\eta_h, \mathbf{v}_h, q_h, \zeta_h) \in C_{[0,1],h} \times \mathbf{X}_{\text{BDM}_k^o} \times (X_{\text{DG}_{k-1}} \setminus \mathbb{R}) \times \mathbb{R}$, we have

$$(4.8) \quad c(\rho_h, \eta_h - \rho_h; \mathbf{u}_h, \lambda_h) \geq 0,$$

$$(4.9) \quad a_h(\mathbf{u}_h, \mathbf{v}_h; \rho_h) + b(\mathbf{v}_h, p_h) = l_h(\mathbf{v}_h),$$

$$(4.10) \quad b(\mathbf{u}_h, q_h) = 0,$$

$$(4.11) \quad d(\lambda_h, \zeta_h; \rho_h) = 0.$$

Let the broken H^1 -norm $\|\cdot\|_{H_g^1(\mathcal{T}_h)}^2$ be defined as

$$(4.12) \quad \begin{aligned} \|\mathbf{v}\|_{H_g^1(\mathcal{T}_h)}^2 &:= \|\mathbf{v}\|_{L^2(\Omega)}^2 + \sum_{K \in \mathcal{T}_h} \|\nabla \mathbf{v}\|_{L^2(K)}^2 \\ &+ \sum_{F \in \mathcal{F}_h^i} \int_F h_F^{-1} |[\![\mathbf{v}]\!]_F|^2 ds + \sum_{F \in \mathcal{F}_h^\partial} \int_F h_F^{-1} |[\![\mathbf{v} - \mathbf{g}]\!]_F|^2 ds. \end{aligned}$$

Building on previous work [52], it was shown by Papadopoulos [48] that, for every isolated minimizer (\mathbf{u}, ρ, p) of (BP), there exists a sequence of discretized solutions $(\mathbf{u}_h, \rho_h, p_h, \lambda_h)$ to (4.8)–(4.11), such that, as $h \rightarrow 0$, $\|\mathbf{u} - \mathbf{u}_h\|_{H_g^1(\mathcal{T}_h)} \rightarrow 0$, $\|\rho - \rho_h\|_{L^s(\Omega)} \rightarrow 0$, $s \in [1, \infty)$, and $\|p - p_h\|_{L^2(\Omega)} \rightarrow 0$. By extrapolating known results of a BDM discretization for the Stokes and Stokes–Brinkman equations [15, 16, 43], we expect a first-order BDM₁ discretization to converge at a rate of $\mathcal{O}(h)$ in the broken H^1 -norm for the velocity, the L^2 -norm for the pressure, and the L^2 -norm for the material distribution. Numerical evidence for these rates of convergence is given in [48, Fig. 2].

Remark 4.1. The value of $\alpha(\rho)$ in the momentum equation (3.5) can range between 0 and $\bar{\alpha}$ where in practice $\bar{\alpha} \sim \mathcal{O}(10^4)$. Hence, an element that is robust to the transition between Stokes and Darcy flows may have better accuracy, e.g., the Mardal–Tai–Winther (MTW) finite element [46]. The MTW finite element also produces pointwise divergence-free solutions, and the characterization of the kernel of the divergence operator is known [46, sect. 4.2]. In particular, it has been shown that vertex-star patch relaxation is also effective for multigrid cycles involving MTW finite elements [7, sect. 7.2]. Hence, we believe the preconditioner described below would still be effective. However, unlike BDM finite elements with an interior penalty, there are currently no convergence results for an MTW finite element discretization to minimizers of the Borrvall–Petersson problem.

4.2. The Benson–Munson linear system. The BM solver [11] attempts to find a solution of a complementarity problem via linearizations of the residual constrained to the inactive set. First, the discrete Newton system is formed and the active set is defined. The active set contains the degrees of freedom that satisfy a strict complementarity condition in the primal and residual vectors. Next, the rows and columns of the Jacobian in the Newton system associated with the active set degrees of freedom are set to those of the identity. Finally, the rows on the right-hand side vector associated with the active set degrees of freedom are fixed to zero. Once the update, $\delta \mathbf{z}$, of this modified system is computed, the new iterate $\mathbf{z}^{k+1} = \mathbf{z}^k + \delta \mathbf{z}$ is componentwise projected onto the box constraints.

Let $\mathbf{f}: \mathbb{R}^n \rightarrow \mathbb{R}^n$, and consider the box constraints $\mathbf{a}, \mathbf{b} \in \mathbb{R}^n$ where $\mathbf{a}_i < \mathbf{b}_i$ for all $i = 1, \dots, n$. Consider the mixed complementarity problem given by

$$(4.13) \quad \text{either } \mathbf{a}_i < \mathbf{z}_i < \mathbf{b}_i \quad \text{and } \mathbf{f}(\mathbf{z})_i = 0,$$

$$(4.14) \quad \text{or } \mathbf{a}_i = \mathbf{z}_i \quad \text{and } \mathbf{f}(\mathbf{z})_i \geq 0,$$

$$(4.15) \quad \text{or } \mathbf{z}_i = \mathbf{b}_i \quad \text{and } \mathbf{f}(\mathbf{z})_i \leq 0.$$

We define $\hat{\mathbf{f}}: \{\mathbf{z} \in \mathbb{R}^n : \mathbf{a}_i \leq \mathbf{z}_i \leq \mathbf{b}_i, i = 1, \dots, n\} \rightarrow \mathbb{R}^n$ as follows, for $i = 1, \dots, n$:

$$(4.16) \quad [\hat{\mathbf{f}}(\mathbf{z})]_i := \begin{cases} \mathbf{f}(\mathbf{z})_i & \text{if } \mathbf{a}_i < \mathbf{z}_i < \mathbf{b}_i, \\ \min\{\mathbf{f}(\mathbf{z})_i, 0\} & \text{if } \mathbf{z}_i = \mathbf{a}_i, \\ \max\{\mathbf{f}(\mathbf{z})_i, 0\} & \text{if } \mathbf{z}_i = \mathbf{b}_i. \end{cases}$$

Algorithm 4.1 Benson–Munson solver [11, sect. 3.2].

1: Initialize:

$k \leftarrow 0$	▷ Initial iteration number
\mathbf{a}, \mathbf{b}	▷ Box constraints
\mathbf{z}^0	▷ Feasible initial guess
tol	▷ Approximate solve tolerance

2: while $\|\hat{\mathbf{f}}(\mathbf{z}^k)\|_2 > \text{tol}$ **do**
3: Define the active set \mathcal{A}^k and the inactive set \mathcal{I}^k as:

$$\mathcal{A}^k := \{i : \mathbf{z}_i^k = \mathbf{a}_i \text{ and } \mathbf{f}(\mathbf{z}^k)_i > 0\} \cup \{i : \mathbf{z}_i^k = \mathbf{b}_i \text{ and } \mathbf{f}(\mathbf{z}^k)_i < 0\},$$

$$\mathcal{I}^k := \{i\}_{i=1}^n \setminus \mathcal{A}^k.$$

4: Solve the BM linear system: $\mathbf{f}'(\mathbf{z}^k)_{\mathcal{I}^k, \mathcal{I}^k} \delta \mathbf{z} = -\mathbf{f}(\mathbf{z}^k)_{\mathcal{I}^k}$.

5: $\mathbf{z}^{(k+1)} \leftarrow \pi(\mathbf{z}^k + \beta^k \delta \mathbf{z})$, where β^k is a (possibly adaptive) linesearch.

6: $k \leftarrow k + 1$
7: end while

We denote the componentwise projection onto the true box constraints by π , i.e.,

$$(4.17) \quad \pi(\mathbf{z})_i = \begin{cases} \mathbf{a}_i & \text{if } \mathbf{z}_i < \mathbf{a}_i, \\ \mathbf{z}_i & \text{if } \mathbf{a}_i \leq \mathbf{z}_i \leq \mathbf{b}_i, \\ \mathbf{b}_i & \text{if } \mathbf{z}_i > \mathbf{b}_i. \end{cases}$$

Consider a matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ and the subset of indices $S \subset \{1, \dots, n\}$. Then, the matrix $\mathbf{A}_{S,S} \in \mathbb{R}^{n \times n}$ is defined by zeroing the rows and columns of \mathbf{A} in $\{1, \dots, n\} \setminus S$ and placing a 1 on the diagonal of the zeroed rows and columns. Similarly, for any column vector $\mathbf{x} \in \mathbb{R}^n$, the vector $\mathbf{x}_S \in \mathbb{R}^n$ is constructed by zeroing the rows in the set $\{1, \dots, n\} \setminus S$ of \mathbf{x} .

Remark 4.2. In the original BM solver found in [11, sect. 3.2], the rows and columns are eliminated rather than zeroed, resulting in a reduced matrix. However, the BM solvers as described in Algorithm 4.1 and [11, sect. 3.2] are equivalent.

When applied to a linear elliptic problem, and if the active and inactive sets of the two algorithms coincide, the updates computed by BM and the primal-dual active set strategy of Hintermüller, Ito, and Kunisch [38] are identical [51, Thm. A.1]. Under suitable conditions, the primal-dual active set strategy can be shown to be a semismooth Newton method [55, 56, 66]. Although the problems we consider here do not satisfy the assumptions made in Theorem A.1 of [51], we do observe local superlinear convergence in practice.

We now derive the linear systems that arise when using the BM active set strategy to solve the nonlinear system (4.8)–(4.11). Denote the basis functions of the finite element spaces of the material distribution, the velocity, the pressure, and \mathbb{R} by η_i, ϕ_i, ψ_i , and r , respectively and the number of degrees of freedom by $n_\rho, n_{\mathbf{u}}, n_p$, and 1, respectively, so that the total number of degrees of freedom of the system is given by $n = n_\rho + n_{\mathbf{u}} + n_p + 1$. Consider the finite element BM iterate $\mathbf{z}_h^k = (\rho_h^k, \mathbf{u}_h^k, p_h^k, \lambda_h^k)$. Let $\mathbf{f}^k : \mathbb{R}^n \rightarrow \mathbb{R}^n$ denote the nonlinear residual induced by the complementarity reformulation of (4.8)–(4.11). In the following, we drop the superscript iteration number k for clarity. Let \mathbf{z} denote the discrete coefficient vector of \mathbf{z}_h . In the context of Algorithm 4.1, the box constraints take the values $\mathbf{a}_i = 0$ and $\mathbf{b}_i = 1$ for all degrees

of freedom associated with ρ_h and $\mathbf{a}_i = -\infty$, $\mathbf{b}_i = +\infty$, otherwise. The BM linear system, as solved on line 4 of Algorithm 4.1, is the following:

$$(4.18) \quad \mathbf{H}_{\rho, \mathbf{u}, p, \lambda} \delta \mathbf{z} = \begin{pmatrix} \mathbf{C}_\mu & \mathbf{D}^\top & \mathbf{0} & \mathbf{E}^\top \\ \mathbf{D} & \mathbf{A} & \mathbf{B}^\top & \mathbf{0} \\ \mathbf{0} & \mathbf{B} & \mathbf{0} & \mathbf{0} \\ \mathbf{E} & \mathbf{0} & \mathbf{0} & \mathbf{0} \end{pmatrix} \begin{pmatrix} \delta \rho \\ \delta \mathbf{u} \\ \delta \mathbf{p} \\ \delta \lambda \end{pmatrix} = - \begin{pmatrix} \mathbf{f}_\rho \\ \mathbf{f}_\mathbf{u} \\ \mathbf{f}_p \\ \mathbf{f}_\lambda \end{pmatrix} = -\mathbf{f},$$

where $\delta \rho, \delta \mathbf{u}, \delta \mathbf{p}$, and $\delta \lambda$ denote the discrete coefficient vector BM updates for ρ, \mathbf{u}, p , and λ , and $\mathbf{f}_\rho, \mathbf{f}_\mathbf{u}, \mathbf{f}_p$, and \mathbf{f}_λ are the corresponding blocks of the nonlinear residual with the active set rows, $i \in \mathcal{A}$, in \mathbf{f}_ρ zeroed. The entries of \mathbf{A} and \mathbf{B} are given by

$$(4.19) \quad [\mathbf{A}]_{ij} = a_h(\phi_j, \phi_i; \rho_h) \text{ and } [\mathbf{B}]_{ij} = b(\phi_j, \psi_i).$$

Furthermore, if $j \in \mathcal{I}$, then

$$(4.20) \quad [\mathbf{D}]_{ij} = \int_{\Omega} (\alpha'(\rho_h) \mathbf{u}_h \cdot \phi_i) \eta_j \, dx, \quad [\mathbf{E}]_{ij} = -r \int_{\Omega} \eta_j \, dx.$$

Otherwise if $j \in \mathcal{A}$, then $[\mathbf{D}]_{ij} = 0$ and $[\mathbf{E}]_{ij} = 0$ for all i . Finally if $i, j \in \mathcal{I}$, then

$$(4.21) \quad [\mathbf{C}_\mu]_{ij} = \int_{\Omega} \left[\frac{1}{2} \alpha''(\rho_h) |\mathbf{u}_h|^2 + \frac{\mu}{(\rho_h + \epsilon_{\log})^2} + \frac{\mu}{(1 + \epsilon_{\log} - \rho_h)^2} \right] \eta_i \eta_j \, dx.$$

Otherwise, if $i \in \mathcal{A}$ or $j \in \mathcal{A}$, then $[\mathbf{C}_\mu]_{ij} = \delta_{ij}$, where δ_{ij} is the Kronecker delta.

Remark 4.3. \mathbf{E} is a row vector of size $1 \times n_\rho$.

In the remainder of this subsection we discuss the invertibility of the matrix in (4.18) and its subblocks.

PROPOSITION 4.4. *The matrix $\mathbf{A} \in \mathbb{R}^{n_\mathbf{u} \times n_\mathbf{u}}$ is symmetric, and provided the penalization parameter $\sigma > 0$ is sufficiently large, then it is also positive-definite.*

Proof. Symmetry is realized by swapping the indices i and j of the basis functions in their respective definitions and noting that the resulting integrals are equal. Positive-definiteness of \mathbf{A} , for sufficiently large $\sigma > 0$, follows from $\alpha(\rho) \geq 0$ and [39, sect. 3.3]. \square

PROPOSITION 4.5. *The matrix \mathbf{C}_μ is symmetric positive semidefinite. Moreover, if either $\mu > 0$ or $|\mathbf{u}_h| > 0$ a.e., then \mathbf{C}_μ is symmetric positive-definite.*

Proof. The symmetry of \mathbf{C}_μ is realized by swapping the indices i and j of the basis functions in its definition and noting that the resulting integrals are equal. Consider the unmodified matrix $\hat{\mathbf{C}}_\mu$, defined by (4.21) for all i and j . Pick an arbitrary function $\eta_h \in X_{\text{DG}_0}$ with discrete coefficient vector $\boldsymbol{\eta} \in \mathbb{R}^{n_\rho}$. We note that

$$(4.22) \quad \boldsymbol{\eta}^\top \hat{\mathbf{C}}_\mu \boldsymbol{\eta} = \int_{\Omega} \left[\frac{1}{2} \alpha''(\rho_h) |\mathbf{u}_h|^2 + \frac{\mu}{(\rho_h + \epsilon_{\log})^2} + \frac{\mu}{(1 + \epsilon_{\log} - \rho_h)^2} \right] |\eta_h|^2 \, dx.$$

All the terms in the integral (4.22) are nonnegative. Hence, $\hat{\mathbf{C}}_\mu$ is positive semidefinite.

Property (A5) implies that $\alpha''(\rho_h) > 0$. If $\mu > 0$, then the rational expressions are strictly greater than zero as $0 \leq \rho_h \leq 1$. Otherwise, if $|\mathbf{u}_h| > 0$ a.e., then the term $\alpha''(\rho_h) |\mathbf{u}_h|^2 > 0$ a.e. Hence, if either $\mu > 0$ or $|\mathbf{u}_h| > 0$ a.e., the right-hand side of (4.22) is equal to zero if and only if $\eta_h = 0$, which is true if and only if $\boldsymbol{\eta} = \mathbf{0}$. Therefore, if either $\mu > 0$ or $|\mathbf{u}_h| > 0$ a.e., $\boldsymbol{\eta}^\top \hat{\mathbf{C}}_\mu \boldsymbol{\eta} \geq 0$ with equality if and only if $\boldsymbol{\eta} = \mathbf{0}$. Hence, $\hat{\mathbf{C}}_\mu$ is symmetric positive-definite.

Since the discretization for ρ is piecewise constant, $\hat{\mathbf{C}}_\mu$ is a diagonal matrix, and therefore all diagonal entries must be positive. The procedure of zeroing rows and columns associated with the BM active set and replacing the diagonal entry with a 1 will result in a diagonal matrix with nonnegative (positive if either $\mu > 0$ or $|\mathbf{u}_h| > 0$ a.e.) diagonal entries. We conclude that \mathbf{C}_μ must be symmetric positive semidefinite, and if either $\mu > 0$ or $\mathbf{u}_h > 0$ a.e., then it is symmetric positive-definite. \square

Remark 4.6. For problems with a pure Dirichlet boundary condition, $\mathbf{H}_{\rho,\mathbf{u},p,\lambda}$ will have a nullspace of at least one dimension associated with the fact that the pressure is only unique up to a constant [21, Chap. 4]. Since we will use an FGMRES Krylov method on the outermost level of the linear solve, this particular nullspace can be handled by FGMRES without the method breaking down [21, Chap. 9.3.5]. To show that the nullspace is indeed the space of constants with respect to the pressure, we assume that \mathbf{B} has been modified so that it is of full row rank, e.g., by picking a pressure degree of freedom and fixing it to zero.

PROPOSITION 4.7. *Suppose that either $\mu > 0$ or $|\mathbf{u}_h| > 0$ a.e. Moreover, assume that $\mathbf{A} - \mathbf{D}\mathbf{C}_\mu^{-1}\mathbf{D}^\top$ is symmetric positive-definite and \mathbf{B} has been modified so that it has full row rank. Then, the matrix $\mathbf{H}_{\rho,\mathbf{u},p,\lambda}$, as defined in (4.18), is invertible.*

Proof. If either $\mu > 0$ or $|\mathbf{u}_h| > 0$ a.e., then \mathbf{C}_μ is symmetric positive-definite by Proposition 4.5 and, therefore, invertible. Thus the matrix $\mathbf{A} - \mathbf{D}\mathbf{C}_\mu^{-1}\mathbf{D}^\top$ is well-defined. By assumption, $\mathbf{A} - \mathbf{D}\mathbf{C}_\mu^{-1}\mathbf{D}^\top$ is symmetric positive-definite. Therefore, since \mathbf{A} is symmetric positive-definite by Proposition 4.4, we have that the following matrix is also symmetric positive-definite [75, Thm. 1.12]:

$$(4.23) \quad \mathbf{G} := \begin{pmatrix} \mathbf{C}_\mu & \mathbf{D}^\top \\ \mathbf{D} & \mathbf{A} \end{pmatrix}.$$

We define $\mathbf{B}_0 := (\mathbf{0} \quad \mathbf{B}) \in \mathbb{R}^{n_p \times (n_\rho + n_u)}$ and $\mathbf{E}_0 := (\mathbf{E} \quad \mathbf{0}) \in \mathbb{R}^{1 \times (n_\rho + n_u)}$ and reblock $\mathbf{H}_{\rho,\mathbf{u},p,\lambda}$ as follows:

$$(4.24) \quad \mathbf{H}_{\rho,\mathbf{u},p,\lambda} = \begin{pmatrix} \mathbf{G} & \mathbf{B}_0^\top & \mathbf{E}_0^\top \\ \mathbf{B}_0 & \mathbf{0} & \mathbf{0} \\ \mathbf{E}_0 & \mathbf{0} & \mathbf{0} \end{pmatrix}.$$

The matrix on the right-hand side of (4.24) is of double-saddle point type. \mathbf{E}_0 is a nonzero row vector and thus must have full row rank. Moreover, by assumption, \mathbf{B} has been modified so that it has full row rank. Hence, \mathbf{B}_0 also has full row rank. Consider any vector $\mathbf{x} = (\mathbf{x}_1^\top \quad \mathbf{x}_2^\top)^\top \in \mathbb{R}^{n_\rho + n_u}$. We note that $\mathbf{E}_0^\top \mathbf{x}^\top = \mathbf{E}^\top \mathbf{x}_1^\top$ and $\mathbf{B}_0^\top \mathbf{x}^\top = \mathbf{B}^\top \mathbf{x}_2^\top$. Hence, $\text{range}(\mathbf{B}_0^\top) \cap \text{range}(\mathbf{E}_0^\top) = \{\mathbf{0}\}$. Thus $\mathbf{H}_{\rho,\mathbf{u},p,\lambda}$ is invertible [4, Prop. 2.3]. \square

Remark 4.8. The proof of the invertibility of $\mathbf{H}_{\rho,\mathbf{u},p,\lambda}$ (provided the pressure null space has been eliminated) relied heavily on the assumption that $\mathbf{A} - \mathbf{D}\mathbf{C}_\mu^{-1}\mathbf{D}^\top$ is symmetric positive-definite. Although the symmetry is guaranteed, positive-definiteness is not. Numerically, we found that $\mathbf{A} - \mathbf{D}\mathbf{C}_\mu^{-1}\mathbf{D}^\top$ is often positive-definite for large ranges of μ , choices of α , and BM iterates $\rho_{h,k}$ and $\mathbf{u}_{h,k}$. Despite this, there are choices where $\mathbf{A} - \mathbf{D}\mathbf{C}_\mu^{-1}\mathbf{D}^\top$ has negative eigenvalues. In practice we have not encountered a case where $\mathbf{A} - \mathbf{D}\mathbf{C}_\mu^{-1}\mathbf{D}^\top$ is singular.

Remark 4.9. A full proof of the invertibility of $\mathbf{H}_{\rho,\mathbf{u},p,\lambda}$ in the case where $\mathbf{A} - \mathbf{D}\mathbf{C}_\mu^{-1}\mathbf{D}^\top$ is indefinite or singular is beyond the scope of this work. However, we

note the following. The first-order optimality conditions (3.4)–(3.7) can be rewritten as a semismooth system of equations. Suppose, after discretization, we denote this system as $F: \mathbb{R}^{n_u+3n_\rho+n_p+1} \rightarrow \mathbb{R}^{n_u+3n_\rho+n_p+1}$ where the extra $2n_\rho$ degrees of freedom are associated with the Lagrange multipliers that enforce the box constraints on ρ . Then, provided one can show that F is locally Lipschitz continuous, F is semismooth, and $F'(\mathbf{z}_*)$ is invertible at \mathbf{z}_* where \mathbf{z}_* satisfies $F(\mathbf{z}_*) = 0$, then one may be able to invoke a semismooth analogue of the Rall–Rheinboldt theory [26]. Hence, provided the iterate \mathbf{z}_k is sufficiently close to \mathbf{z}_* , $F'(\mathbf{z}_k)$ is then invertible. However, we note that although the BM solver has been shown to be almost equivalent to a semismooth Newton method for linear elliptic problems in [51, App. A], that theory does not extend to the systems we are solving here. Hence the invertibility of $F'(\mathbf{z}_k)$ would not necessarily imply the invertibility of $\mathbf{H}_{\rho, \mathbf{u}, p, \lambda}$.

5. Preconditioning. In this section, we develop a preconditioner for solving (4.18). As discussed in subsection 3.2, preconditioning strategies that are robust for the undeflated system can also be used to compute solutions of the deflated systems. On the outermost level of the deflated barrier method, we perform continuation in the barrier parameter μ . Next, at a given μ , we use the BM solver to find a solution to (4.8)–(4.11). A direct sparse LU factorization of the matrix in (4.18) is infeasible on fine meshes of three-dimensional problems. Thus we turn to preconditioning techniques to reduce the cost of each inner linear solve. The preconditioning is made difficult by the saddle point nature of the matrix in (4.18) and the barrier-like terms in \mathbf{C}_μ . In the following subsections we introduce a nested block preconditioning method for solving (4.18), where the Schur complements are controlled with an augmented Lagrangian term. As outermost Krylov solver, we use a preconditioned FGMRES method [59].

5.1. Block preconditioning. Consider the well-posed linear system

$$(5.1) \quad \begin{pmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{C} & \mathbf{D} \end{pmatrix} \begin{pmatrix} \mathbf{x} \\ \mathbf{y} \end{pmatrix} = \begin{pmatrix} \mathbf{c} \\ \mathbf{d} \end{pmatrix},$$

where $\mathbf{A} \in \mathbb{R}^{n_1 \times n_1}$ is invertible, $\mathbf{B} \in \mathbb{R}^{n_1 \times n_2}$, $\mathbf{C} \in \mathbb{R}^{n_2 \times n_1}$, and $\mathbf{D} \in \mathbb{R}^{n_2 \times n_2}$. Then, under suitable conditions [12, sect. 3.2] the inverse of the matrix in (5.1) admits a full block factorization of the form

$$(5.2) \quad \begin{pmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{C} & \mathbf{D} \end{pmatrix}^{-1} = \begin{pmatrix} \mathbf{I} & -\mathbf{A}^{-1}\mathbf{B} \\ 0 & \mathbf{I} \end{pmatrix} \begin{pmatrix} \mathbf{A}^{-1} & 0 \\ 0 & \mathbf{S}^{-1} \end{pmatrix} \begin{pmatrix} \mathbf{I} & 0 \\ -\mathbf{C}\mathbf{A}^{-1} & \mathbf{I} \end{pmatrix},$$

where $\mathbf{S} = \mathbf{D} - \mathbf{C}\mathbf{A}^{-1}\mathbf{B}$. Preconditioners for (5.1) can be found by developing cheap approximations to \mathbf{A}^{-1} and \mathbf{S}^{-1} and substituting them into (5.2) [47, 69].

The subspace spanned by the volume constraint Lagrange multiplier λ is one-dimensional and can be handled by at most one iteration of a Krylov subspace solver or via block preconditioning. Experimentally, we found that a full block preconditioner of the real block performed best. Writing the density-momentum-pressure block,

$$(5.3) \quad \mathbf{H}_{\rho, \mathbf{u}, p} := \begin{pmatrix} \mathbf{C}_\mu & \mathbf{D}^\top & \mathbf{0} \\ \mathbf{D} & \mathbf{A} & \mathbf{B}^\top \\ \mathbf{0} & \mathbf{B} & \mathbf{0} \end{pmatrix},$$

we choose $\mathbf{A} = \mathbf{H}_{\rho, \mathbf{u}, p}$, $\mathbf{B} = \mathbf{E}^\top$, $\mathbf{C} = \mathbf{E}$, and $\mathbf{D} = \mathbf{0}$. The Schur complement,

$$(5.4) \quad \mathbf{S} = \mathbf{S}_0 := -\mathbf{E}\mathbf{H}_{\rho, \mathbf{u}, p}^{-1}\mathbf{E}^\top,$$




FIG. 1. Main components of the deflated barrier method solver. The remainder of this section focuses on developing preconditioners for the item in blue.

is a 1×1 matrix and can be inverted by taking its reciprocal. Hence, the difficulty now lies in solving linear systems involving (5.3). Since we have only decreased the size of the linear system by one dimension, an LU factorization is still infeasible, and we consider block preconditioners for (5.3). We summarize the initial components of the solver in Figure 1.

We precondition (5.3) by taking the Schur complement with respect to the momentum-pressure block. This approach was utilized by Evgrafov for preconditioning the linear systems arising in a similar solver [24, sect. 5]. In the notation of (5.1), $\mathbf{A} = \mathbf{C}_\mu$, $\mathbf{B} = (\mathbf{D}^\top \ \mathbf{0})$, $\mathbf{C} = (\mathbf{D} \ \mathbf{0})^\top$, and,

$$(5.5) \quad \mathbb{D} = \mathbf{Q} := \begin{pmatrix} \mathbf{A} & \mathbf{B}^\top \\ \mathbf{B} & \mathbf{0} \end{pmatrix}.$$

We know that \mathbf{C}_μ is invertible by Proposition 4.5. Hence, we write $\mathbf{A}^{-1} = \mathbf{C}_\mu^{-1}$. The matrix \mathbf{Q} resembles the linear system that arises in the discretization of the Stokes–Brinkman equations. The Schur complement is given by

$$(5.6) \quad \mathbb{S} = \mathbf{S}_1 := \mathbf{Q} - \begin{pmatrix} \mathbf{D} \\ \mathbf{0} \end{pmatrix} \mathbf{C}_\mu^{-1} (\mathbf{D}^\top \ \mathbf{0}) = \begin{pmatrix} \mathbf{A} - \mathbf{D}\mathbf{C}_\mu^{-1}\mathbf{D}^\top & \mathbf{B}^\top \\ \mathbf{B} & \mathbf{0} \end{pmatrix}.$$

The reason we use a DG_0 piecewise constant discretization for the material distribution is to ensure that \mathbf{S}_1 is sparse. Since the material distribution is discretized with DG_0 finite elements, \mathbf{C}_μ is a diagonal matrix, $\mathbf{D}\mathbf{C}_\mu^{-1}\mathbf{D}^\top$ is still sparse, and hence \mathbf{S}_1 is also sparse.

Thus far, the application of block preconditioning has reduced the solution of the full matrix (4.18) to the following:

1. Apply an outer FGMRES solver;
2. apply the reciprocal of $\mathbf{S}_0 \in \mathbb{R}$;
3. invert the diagonal matrix \mathbf{C}_μ ;
4. apply the inverse of the 2×2 block matrix \mathbf{S}_1 , which is the same size as the matrix that arises in a discretized pure Stokes problem.

For now we assume that \mathbf{S}_1 is invertible and defer discussions of invertibility to later sections.

We must now develop solvers for \mathbf{S}_1 as given in (5.6). One option is to use a direct solver. However, we can further reduce the computational work with another application of block preconditioning. Consider taking the inner Schur complement in \mathbf{S}_1 with respect to the pressure block. Using the notation of (5.1), $\mathbb{B} = \mathbf{B}^\top$, $\mathbb{C} = \mathbf{B}$, and $\mathbb{A} = \mathbf{A} - \mathbf{D}\mathbf{C}_\mu^{-1}\mathbf{D}^\top$. The innermost Schur complement takes the form

$$(5.7) \quad \mathbb{S} = \mathbf{S}_2 := -\mathbf{B}(\mathbf{A} - \mathbf{D}\mathbf{C}_\mu^{-1}\mathbf{D}^\top)^{-1}\mathbf{B}^\top.$$

This time, \mathbf{S}_2 is dense, and we employ an augmented Lagrangian approach.

We first recall some known results for the Stokes and Stokes–Brinkman equations. Let $-\Delta_h$ denote the negative discretized Laplacian matrix. In the context of the incompressible Stokes equations, $\mathbf{B}(-\Delta_h)^{-1}\mathbf{B}^\top$ is spectrally equivalent to the viscosity-scaled pressure mass matrix $\nu^{-1}\mathbf{M}_p$ [63, 68]; see also [21, Thm. 5.22]. \mathbf{M}_p is a sparse mass matrix and can be cheaply factorized or solved with a multigrid method. Therefore, a good approximation to the Schur complement of the pure Newtonian Stokes problem is given by $\nu^{-1}\mathbf{M}_p$. The idea is that the momentum block can then be solved with a direct solver, multigrid methods, or other alternative solvers. Unfortunately, in the context of the Stokes–Brinkman equations, Popov [54] noted that the presence of the Brinkman term $\alpha(\rho_h)\mathbf{u}_h$ in the momentum block \mathbf{A} renders the approximation given by $\nu^{-1}\mathbf{M}_p$ ineffective. Popov proposed a Schur complement preconditioning technique based on incomplete LU factorization, but such factorizations do not generally yield mesh-independent preconditioners. An alternative is a preconditioning scheme utilized by Borrvall and Petersson in their original paper [14, sect. 2.6] based on the work of Cahouet and Chabard [18]. However, during numerical experiments, we found that an augmented Lagrangian approach performed better for the problems in this work; see [49, Chap. 5.3.1].

We now propose an augmented Lagrangian strategy to control the second Schur complement \mathbf{S}_2 . The essential idea is to add a term to the momentum equation that does not change the exact solution but does change the Schur complement \mathbf{S}_2 ; in particular, by scaling the augmentation appropriately, \mathbf{S}_2 can be almost perfectly approximated with a scaled pressure mass matrix. The cost of this approach is that it makes the augmented momentum block more difficult to solve. The augmented Lagrangian approach has been shown to be robust for a variety of difficult saddle point systems, such as the stationary Navier–Stokes equations at high Reynolds number [30], implicitly constituted anisothermal non-Newtonian flow [31], and magnetohydrodynamics [44]. Hence, this approach has potential for extension to different fluid topology optimization problems.

There are two possible augmented Lagrangian approaches: continuous and discrete. These approaches are mathematically equivalent for exactly divergence-free elements such as the one employed here. We choose to introduce the method in the discrete setting. Postdiscretization, the matrix \mathbf{A} in (4.18) is modified by adding an augmented Lagrangian term

$$(5.8) \quad \mathbf{A}_{\gamma_d} := \mathbf{A} + \gamma_d \mathbf{B}^\top \mathbf{M}_p^{-1} \mathbf{B},$$

where $\gamma_d \gg 0$, and the right-hand side of (4.18) is modified so that the solution of the linear system remains unchanged (since $\mathbf{B}\mathbf{u}$ is known). In particular, if the current velocity iterate is divergence-free, then $\gamma_d \mathbf{B}^\top \mathbf{M}_p^{-1} \mathbf{B} \delta \mathbf{u} = \mathbf{0}$, and no modification to the right-hand side is required. While it does not change the solution, the addition of the augmented Lagrangian term influences the nature of the inner Schur complement. In particular, \mathbf{S}_1 becomes

$$(5.9) \quad \mathbf{S}_{1,\gamma_d} = \begin{pmatrix} \mathbf{A}_{\gamma_d} - \mathbf{D}\mathbf{C}_\mu^{-1}\mathbf{D}^\top & \mathbf{B}^\top \\ \mathbf{B} & \mathbf{0} \end{pmatrix},$$

and \mathbf{S}_2 becomes

$$(5.10) \quad \mathbf{S}_{2,\gamma_d} = -\mathbf{B}(\mathbf{A}_{\gamma_d} - \mathbf{D}\mathbf{C}_\mu^{-1}\mathbf{D}^\top)^{-1}\mathbf{B}^\top.$$

Remark 5.1. As noted in Remark 4.9, it is possible that $\mathbf{A}_{\gamma_d} - \mathbf{D}\mathbf{C}_\mu^{-1}\mathbf{D}^\top$ is indefinite, and therefore it is unclear if \mathbf{S}_{1,γ_d} is invertible. If we were to encounter such

a case, then the BM linear system solve would fail. Here, the deflated barrier method would repeatedly halve the step size in the barrier parameter μ (keeping the initial guess the same). We did not observe this behavior in our numerical experiments. Henceforth, we assume that \mathbf{S}_{1,γ_d} is invertible. We note that if \mathbf{S}_{1,γ_d} is invertible, then \mathbf{S}_{2,γ_d} is invertible. This is seen by applying Sylvester's law of inertia to the Schur complement decomposition (5.2) [21, Chap. 4].

The action of $\mathbf{S}_{2,\gamma_d}^{-1}$ is required during the solve. As already mentioned, \mathbf{S}_{2,γ_d} is dense, and hence it is expensive to assemble and apply its inverse. In the next proposition we show that the action of $\mathbf{S}_{2,\gamma_d}^{-1}$ can be approximated with $-\gamma_d \mathbf{M}_p^{-1}$ with increasing accuracy as $\gamma_d \rightarrow \infty$. In practice, values of $\gamma_d \sim \mathcal{O}(10^5)$ were sufficient to approximate the action of $\mathbf{S}_{2,\gamma_d}^{-1}$ with $-\gamma_d \mathbf{M}_p^{-1}$ to tolerances of $\mathcal{O}(10^{-10})$.

PROPOSITION 5.2. *Suppose that $\mathbf{A} - \mathbf{D}\mathbf{C}_\mu^{-1}\mathbf{D}^\top$, $\mathbf{A}_{\gamma_d} - \mathbf{D}\mathbf{C}_\mu^{-1}\mathbf{D}^\top$, \mathbf{S}_2 and \mathbf{S}_{2,γ_d} are invertible, and \mathbf{B} has been modified to ensure it has full rank (see Remark 4.6). Consider the eigenvalue problem*

$$(5.11) \quad \gamma_d \mathbf{M}_p^{-1} \mathbf{S}_{2,\gamma_d} \mathbf{q}_i = \lambda_i \mathbf{q}_i, \quad i = 1, \dots, n_p.$$

Then, $\lambda_i \rightarrow -1$, $i = 1, \dots, n_p$, as $\gamma_d \rightarrow \infty$.

Proof. We note that \mathbf{M}_p is symmetric positive-definite and therefore invertible. Our assumptions satisfy the requirements of [13, Lem. 4.1], and thus

$$(5.12) \quad -\mathbf{S}_{2,\gamma_d}^{-1} = -\mathbf{S}_2^{-1} + \gamma_d \mathbf{M}_p^{-1}.$$

By left multiplying (5.11) by $\gamma_d^{-1} \mathbf{M}_p$, then $\mathbf{S}_{2,\gamma_d}^{-1}$, and applying (5.12) we see that

$$(5.13) \quad \lambda_i [\mathbf{S}_2^{-1} \gamma_d^{-1} \mathbf{M}_p - \mathbf{I}] \mathbf{q}_i = \mathbf{q}_i.$$

As $\mathbf{M}_p^{-1} \mathbf{S}_{2,\gamma_d}$ is invertible, $\lambda_i \neq 0$ for all $i = 1, \dots, n_p$. Hence, by rearranging (5.13) we observe that

$$(5.14) \quad \mathbf{S}_2^{-1} \mathbf{M}_p \mathbf{q}_i = \gamma_d (\lambda_i^{-1} + 1) \mathbf{q}_i.$$

Since $\mathbf{S}_2^{-1} \mathbf{M}_p$ is invertible, its eigenvalues (denoted δ_i) are nonzero and satisfy

$$(5.15) \quad \delta_i = \gamma_d (\lambda_i^{-1} + 1), \quad i = 1, \dots, n_p.$$

Therefore, by taking the limit $\gamma_d \rightarrow \infty$ we conclude that $\lambda_i \rightarrow -1$ for all $i = 1, \dots, n_p$. \square

If assembled naïvely, the triple matrix product $\mathbf{B}^\top \mathbf{M}_p^{-1} \mathbf{B}$, as it occurs in the augmented Lagrangian term, is expensive to compute. However, it can be verified that the augmented Lagrangian term $\gamma_d \mathbf{B}^\top \mathbf{M}_p^{-1} \mathbf{B}$ corresponds to augmenting the weak form $a_h(\mathbf{u}_h, \mathbf{v}_h; \rho_h)$ in (4.9) by

$$(5.16) \quad \gamma_d \int_{\Omega} \Pi(\nabla \cdot \mathbf{u}_h) \Pi(\nabla \cdot \mathbf{v}_h) dx,$$

where Π is the projection onto the discretized pressure space. The projection is the identity for the BDM-DG pair. Therefore, assembling $\mathbf{B}^\top \mathbf{M}_p^{-1} \mathbf{B}$ is equivalent to assembling the matrix associated with the bilinear form $\int_{\Omega} \nabla \cdot \phi_i \nabla \cdot \phi_j dx$, where ϕ_i , $i = 1, \dots, n_u$, are the basis functions of the velocity finite element space.

With the proposed nested block preconditioning, we have reduced solving linear systems involving (4.18) to the following steps:




FIG. 2. The preconditioning strategy for solving the density-momentum-pressure block (5.3). We develop a geometric multigrid scheme for the item in pink in subsection 5.2.

1. Apply an outer FGMRES solver;
2. apply the reciprocal of $\mathbf{S}_0 \in \mathbb{R}$;
3. invert the diagonal matrix \mathbf{C}_μ ;
4. factorize and solve the block-diagonal pressure mass matrix \mathbf{M}_p ;
5. apply the action of the inverse of the augmented momentum block $\mathbf{A}_{\gamma_d} - \mathbf{D}\mathbf{C}_\mu^{-1}\mathbf{D}^\top$.

Factorizing \mathbf{M}_p , and $\mathbf{A}_{\gamma_d} - \mathbf{D}\mathbf{C}_\mu^{-1}\mathbf{D}^\top$ with a direct solver such as MUMPS [5] is faster than factorizing the full matrix in (4.18). We note that \mathbf{M}_p only needs to be factorized once at the start of the algorithm. With ideal inner solvers, most of the computational time during the run of the deflated barrier method is spent on factorizing $\mathbf{A}_{\gamma_d} - \mathbf{D}\mathbf{C}_\mu^{-1}\mathbf{D}^\top$ at each BM iteration.

In Figure 2, we summarize the block preconditioning strategy for solving linear systems involving the density-momentum-pressure block (5.3). In the next section we develop a specialized geometric multigrid scheme to efficiently solve linear systems involving $\mathbf{A}_{\gamma_d} - \mathbf{D}\mathbf{C}_\mu^{-1}\mathbf{D}^\top$ (highlighted in pink in Figure 2) in order to reduce the computational time further when the problem is discretized on a fine mesh.

5.2. A specialized multigrid scheme for $\mathbf{A}_{\gamma_d} - \mathbf{D}\mathbf{C}_\mu^{-1}\mathbf{D}^\top$. As already mentioned, the trade-off for using an augmented Lagrangian term to control the Schur complement \mathbf{S}_{2,γ_d} is that $\mathbf{A}_{\gamma_d} - \mathbf{D}\mathbf{C}_\mu^{-1}\mathbf{D}^\top$ becomes difficult to solve, due to the semi-definite term with a large coefficient $\gamma_d \gg 0$. In the past two decades, there has been progress on specialized multigrid schemes, based on the pioneering work of Schöberl [62], to handle the effects of the augmented Lagrangian term in \mathbf{A}_{γ_d} . Strategies based on Schöberl's work have been shown to be extremely effective in parameter-robust preconditioning of the three-dimensional incompressible Navier–Stokes equations [29, 30], Oseen–Frank models of cholesteric liquid crystals [71], implicitly constituted non-Newtonian incompressible flow [27, 31], and magnetohydrodynamics [44]. Related multigrid schemes have also been analyzed in the context of the $\mathbf{H}(\text{div}; \Omega)$ and $\mathbf{H}(\text{curl}; \Omega)$ Riesz maps [6] and, more relevant to our problem, an $\mathbf{H}(\text{div}; \Omega)$ -conforming discretization of the Stokes equations [39].

In this work, we construct a mesh hierarchy and construct our multilevel hierarchy geometrically. The coarse-level operators are induced from (4.8)–(4.11) via discretization. Moreover, in all our examples, we choose a direct solver for the coarse-level solver. Schöberl's analysis gives sufficient conditions on the relaxation method and transfer operators to achieve robustness in the context of multigrid cycles applied to symmetric positive-definite problems augmented with a parameter-dependent positive semidefinite term. The first condition is that the relaxation method must stably capture the kernel of the semidefinite term. The second requirement is that the prolongation operator must have a continuity constant that is independent of γ_d . As noted by Hong et al. [39, sect. 1], in a nested mesh hierarchy, an exactly divergence-

free function on the coarse-grid will be divergence-free on the fine-grid. Therefore, in our context, the natural prolongation operator suffices thanks to our choice of discretization, and we only discuss the relaxation method in this work. The kernel of the semidefinite term involving γ_d is

$$(5.17) \quad \mathcal{N}_h = \{\mathbf{w}_h \in \mathbf{X}_{\text{BDM}_k} : (\nabla \cdot \mathbf{w}_h, \nabla \cdot \mathbf{v}_h)_{L^2(\Omega)} = 0 \text{ for all } \mathbf{v}_h \in \mathbf{X}_{\text{BDM}_k}\},$$

i.e., all functions with divergence zero. For large γ_d , \mathbf{A}_{γ_d} becomes increasingly singular. Common relaxation methods like Jacobi and Gauss–Seidel do not offer γ_d -robust smoothing and yield ineffective multigrid cycles. To understand the degradation of Jacobi and Gauss–Seidel as $\gamma_d \rightarrow \infty$, it is fruitful to view the relaxation method as a subspace correction method [72, 73]. Consider the space decomposition

$$(5.18) \quad \mathbf{X}_{\text{BDM}_k} = \sum_i \mathbf{X}_i,$$

where the sum is not necessarily direct. A subspace correction method solves for an approximation of the error in each subspace and combines them (additively or multiplicatively) to form an updated guess for the solution. In the classical Jacobi and Gauss–Seidel iterations, the decomposition, (5.18), is given by $\{\mathbf{X}_i\} = \{\phi_i\}$ where ϕ_i , $i = 1, \dots, n_u$, are the velocity basis functions. The difference between Jacobi and Gauss–Seidel is whether the updates are applied additively (Jacobi) or multiplicatively (Gauss–Seidel).

A sufficient condition for the subspace correction method induced by the decomposition (5.18) to be robust in γ_d for a symmetric positive-definite matrix is that the decomposition captures the kernel \mathcal{N}_h in the following sense [13, 62, 30, 39]:

$$(5.19) \quad \mathcal{N}_h = \sum_i \mathbf{X}_i \cap \mathcal{N}_h.$$

In other words, the decomposition must be sufficiently rich so that any divergence-free velocity can be written as a combination of divergence-free functions from the subspaces \mathbf{X}_i .¹ Jacobi fails this criterion, as each ϕ_i is not divergence-free. A decomposition satisfying (5.19) for the BDM-DG discretization was developed by Hong et al. [39, sect. 4.5], where the decomposition is the so-called *star* patch around every vertex of the mesh. This decomposition is visualized in Figure 3 for a BDM₁ discretization in two dimensions. The same topological decomposition extends to higher orders and three dimensions. Since $\mathbf{A}_{\gamma_d} - \mathbf{D}\mathbf{C}_\mu^{-1}\mathbf{D}^\top$ is not guaranteed to be positive-definite, the theory does not guarantee robust convergence. Nevertheless, we find that a small number of FGMRES iterations preconditioned with the vertex-star patch iteration is very effective as a smoother, as reported in [30] and subsequent works.

Injecting the active set. A complication arises in the representation of $\mathbf{A}_{\gamma_d} - \mathbf{D}\mathbf{C}_\mu^{-1}\mathbf{D}^\top$ on the coarser levels. By first ignoring the BM active set, we note that $\mathbf{A}_{\gamma_d} - \mathbf{D}\mathbf{C}_\mu^{-1}\mathbf{D}^\top$ can be assembled by injecting the current finite element iterates, \mathbf{u}_h and ρ_h , to the relevant level on the mesh hierarchy, assembling the submatrices \mathbf{A} , \mathbf{D} , \mathbf{D}^\top and \mathbf{C}_μ , applying the Dirichlet boundary conditions of the injected velocity to the relevant rows and columns of \mathbf{A} , \mathbf{D} and \mathbf{D}^\top , and subtracting the triple matrix product $\mathbf{D}\mathbf{C}_\mu^{-1}\mathbf{D}^\top$ from \mathbf{A} . The triple matrix product is sparse and cheap to compute as \mathbf{C}_μ is diagonal on all levels. However, we found that an accurate representation

¹In addition, the decomposition must be stable, but we do not elaborate here.




FIG. 3. The patch of degrees of freedom (black dots inside the blue patch) around a vertex (red diamond) used in the multigrid relaxation for a BDM₁ discretization in two dimensions. Each vertex-star patch contains 12 degrees of freedom in two dimensions with this lowest-order element pair.

of the active set on the coarser levels is essential for the convergence of the multigrid scheme. This is similar to other experiences reported in the literature [22, 40]. Hence, the difficulty lies in defining the active set on the coarser levels. An obvious choice is to use the definitions of \mathcal{A} and \mathcal{I} in Algorithm 4.1 defined via the injected material distribution iterate. However, in numerical experiments we found that this choice resulted in poor iteration counts.

Consider a two-grid method with the fine-level triangulation \mathcal{T}_h , $h = H/2$, obtained by a uniform refinement of the simplices in coarse-level triangulation \mathcal{T}_H . As the material distribution is discretized with DG₀ elements, each degree of freedom i associated with the fine-level material distribution iterate can be associated with an element $K_h \in \mathcal{T}_h$ in the fine level and, analogously, with the degrees of freedom of the coarse-level material distribution iterate with elements in the coarse level. We say that a fine-level element $K_h \in \mathcal{T}_h$ is in the active set \mathcal{A}_h (written as $K_h \in \mathcal{A}_h$) if the degree of freedom associated with K_h is in the active set \mathcal{A}_h . This definition naturally extends to the coarse-level elements and active sets.

We now utilize an idea inspired by the work of Hoppe [40] and Engel and Griebel [22] to define the coarse-level active sets. A coarse-level element $K_H \in \mathcal{T}_H$ containing the parent fine-level elements $K_{h,1}, \dots, K_{h,s} \in \mathcal{T}_h$ is defined to be in the coarse-level active set \mathcal{A}_H if

$$(5.20) \quad |\{K_{h,j} \in \mathcal{A}_h : j = 1, \dots, s\}| \geq m,$$

where $m \in [1, s]$ and $s = 4$ in two dimensions and $s = 8$ in three dimensions. In other words, the coarse-level element is in the coarse-level active set if it contains m or more fine-level parent elements that are in the fine-level active set. By starting at the finest-level active set that is defined by Algorithm 4.1, we recursively define all the active sets in mesh hierarchy via (5.20). Experiments revealed that a good choice for m is $m = s/2$, i.e., a coarse-level element is active if at least half of its parent fine-level elements are active. A summary of the multigrid strategy is given in Figure 4.

Remark 5.3. The choice of (5.20) is more generous than utilizing the definition of the fine-level active set directly with the injected material distribution iterate. In particular, some coarse cells that are “borderline” between the active set and inactive set are placed in the active set by (5.20) but in the inactive set when defining the active set using the injected material distribution. Numerically, we see that the iteration counts suffer if the criteria for a coarse cell to be in the coarse-level active set are too strict.

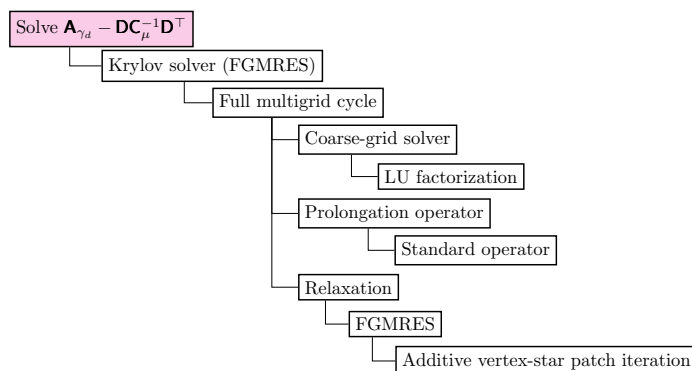


FIG. 4. The multigrid solver strategy of subsection 5.2 for solving $\mathbf{A}_{\gamma_d} - \mathbf{D}\mathbf{C}_{\mu}^{-1}\mathbf{D}^T$.




FIG. 5. The material distribution of the straight-channel (left) and double-ended wrench (right) solutions of the double-pipe optimization problem.

6. Numerical results. All examples in this paper were implemented with the finite element software Firedrake [57]. Block preconditioning and Krylov subspace methods were implemented using Firedrake [57] and PETSc [8], and sparse LU factorizations were performed with MUMPS [5]. Vertex-star patch relaxation is implemented via the PCPATCH functionality [28] recently introduced to PETSc. The meshes were created in Firedrake or Gmsh [37]. The uniqueness of the pressure was enforced by orthogonalizing against the nullspace of constants in the Krylov method. The coarsest-level correction in the multigrid scheme of subsection 5.2 is computed via an LU factorization. The BM updates are scaled with a (damped) l^2 -minimizing linesearch [17, Alg. 2], and we do not use a prediction step in any examples. Wherever (F)GMRES is used, it is not restarted.

6.1. Double-pipe. The first example is the two-dimensional double-pipe problem first introduced by Borrvall and Petersson [14, sect. 4.5]. The double-pipe problem is posed on a rectangular domain $\Omega = (0, 3/2) \times (0, 1)$ with two inlets and two outlets fixed by the following Dirichlet boundary condition:

$$(6.1) \quad \mathbf{g}(x, y) = \begin{cases} (1 - 144(y - 3/4)^2, 0)^\top & \text{if } 2/3 \leq y \leq 5/6, x = 0 \text{ or } 3/2, \\ (1 - 144(y - 1/4)^2, 0)^\top & \text{if } 1/6 \leq y \leq 1/3, x = 0 \text{ or } 3/2, \\ (0, 0)^\top & \text{elsewhere on } \partial\Omega. \end{cases}$$

We choose a volume fraction of $\gamma = 1/3$, and the inverse permeability α is given in (2.2), with $\bar{\alpha} = 2.5 \times 10^4$ and $q = 1/10$. The problem supports two minima: a local minimum of two straight channels from each inlet to its opposite outlet, and a global minimum in the shape of a double-ended wrench. These are depicted in Figure 5.

The two strategies we utilize for solving the linear systems are the following:

TABLE 1

The total cumulative number of iterations to compute both minimizers of the double-pipe problem over all the subproblems with the (aL1) preconditioner. BM stands for the number of Benson–Munson iterations, and OK stands for the number of outer Krylov FGMRES iterations. The numbers in brackets in the OK column are the number of average Krylov iterations per BM iteration.

h	Dofs	(aL1)	
		BM	OK
0.0361	25,201	274	461 (1.68)
0.0180	100,401	626	1180 (1.88)
0.0090	400,801	733	1283 (1.75)
0.0045	1,601,601	809	1607 (1.99)

TABLE 2

The total cumulative number of iterations to compute both minimizers of the double-pipe problem over all the subproblems with the (aL2) preconditioner. BM stands for the number of Benson–Munson iterations, OK stands for the number of outer Krylov FGMRES iterations, and IK is the number of inner Krylov FGMRES iterations preconditioned with the geometric multigrid method of subsection 5.2 to solve linear systems involving $\mathbf{A}_{\gamma_d} - \mathbf{D}\mathbf{C}_{\mu}^{-1}\mathbf{D}^{\top}$. The numbers in brackets in the OK and IK columns are the number of average Krylov iterations per BM iteration and per $\mathbf{A}_{\gamma_d} - \mathbf{D}\mathbf{C}_{\mu}^{-1}\mathbf{D}^{\top}$ solve, respectively.

h	Dofs	(aL2) 2-grid		
		BM	OK	IK
0.0180	100,401	627	1270 (2.03)	100,109 (13.14)
0.0090	400,801	735	1360 (1.85)	86,745 (10.63)

- (aL1) The nested block preconditioning approach of subsection 5.1 on (4.18) with $\gamma_d = 10^4$, and an LU factorization for the augmented momentum block $\mathbf{A}_{\gamma_d} - \mathbf{D}\mathbf{C}_{\mu}^{-1}\mathbf{D}^{\top}$;
- (aL2) the nested block preconditioning approach of subsection 5.1 on (4.18) with $\gamma_d = 10^4$, and the geometric multigrid method of subsection 5.2 for approximating the action of the inverse of $\mathbf{A}_{\gamma_d} - \mathbf{D}\mathbf{C}_{\mu}^{-1}\mathbf{D}^{\top}$. We fix the relaxation to five FGMRES iterations preconditioned with a vertex-star patch iteration, and a full multigrid cycle is used.

We opt for a first-order Brezzi–Douglas–Marini $\mathbf{BDM}_1 \times \mathbf{DG}_0$ mixed finite element discretization for the velocity-pressure pair, with interior penalty parameter $\sigma = 10$, and a \mathbf{DG}_0 discretization for the material distribution. This choice of discretization makes both \mathbf{M}_p and \mathbf{C}_{μ} diagonal. For all mesh sizes, we initialize the deflated barrier method at $\mu_0 = 105$ and perform deflation immediately to find the second branch. The first and second branches converge to the straight-channel and double-ended wrench solutions, respectively, as $\mu \rightarrow 0$. The nonlinear solves are terminated with an absolute tolerance of 10^{-5} . The outer FGMRES solver’s absolute and relative tolerances are both set to 10^{-7} . In Tables 1 and 2, we list the iteration counts for the strategies (aL1) and (aL2) on meshes with decreasing mesh sizes. In the (aL2) strategy, the augmented block solve is approximated to an absolute tolerance of 10^{-8} or a relative tolerance of 10^{-9} .

We see that the Krylov iterations per BM iteration are robust to the mesh size for both preconditioning strategies. The preconditioning strategy with an LU factorization for the augmented momentum block (aL1) averages to under two preconditioned FGMRES iterations per BM iteration. Similarly, with (aL2), where the augmented

block solve is approximated with FGMRES preconditioned with a two-grid multigrid cycle, the outer FGMRES iterations remain under two preconditioned FGMRES iterations per BM iteration on average. Moreover, the average inner FGMRES based on the kernel-preserving multigrid scheme iterations remain under 13 iterations per augmented momentum block solve over all mesh sizes. In particular, the average inner FGMRES iterations decrease on the fine mesh relative to the coarse mesh. We note that, unlike the conforming discretizations in [51], the number of BM iterations slowly increases with decreasing mesh size. This may be due to the discretization or could be related to the fact that the linear systems are not being solved exactly. We do not report the timings of the solves in this example. In two dimensions a direct solve, using MUMPS with $\gamma_d = 0$, of the full BM system (4.18) is faster than the block preconditioning and multigrid strategy proposed here. However, we note that in three dimensions, a direct solve often fails due to the ill-conditioning of the system. Moreover, assembling fine-grid matrices has large memory requirements. As the (aL2) strategy only requires an assembly of the coarsest-grid augmented block, this allows for finer mesh linear solves. This is further discussed in subsection 6.3.

A natural question is the sensitivity of the solver with respect to the augmented Lagrangian parameter γ_d . We wish to explore the following:

- (H1) The effect of roundoff error as γ_d increases;
- (H2) the improvement of the approximation of $-\gamma_d \mathbf{M}_p^{-1}$ for $\mathbf{S}_{2,\gamma_d}^{-1}$;
- (H3) the robustness of the vertex-star patch multigrid cycle for increasing γ_d .

An unfortunate side effect of increasing γ_d is the expectation of high roundoff error. This would manifest in the nonlinear solver stagnating at higher residual norm values than desired. However, as γ_d increases, we expect that the Schur complement approximation by the pressure mass matrix will improve as shown in Proposition 5.2. Hence, the number of outer Krylov FGMRES iterations per BM iteration should decrease as $\gamma_d \rightarrow \infty$. Finally, the multigrid cycle should be robust to the value of γ_d . Hence, we expect that the number of inner Krylov FGMRES iterations preconditioned by the multigrid cycle to approximate the inverse of the augmented momentum block will stay roughly constant with increasing γ_d . In Table 3 we test these hypotheses.

We investigate (H1) and (H2) via the (aL1) solver (with varying γ_d). We fix a mesh with mesh size $h = 0.0361$ and choose the discretization used previously. We consider the deflated barrier method subproblem at $\mu = 105$ and $\mu = 1$ for the first branch. The nonlinear solves are terminated when the decrease in the residual norm stagnates. The outer FGMRES solver's absolute and relative tolerances are set to 10^{-15} and 10^{-5} , respectively. The middle four columns confirm our first two hypotheses. At $\gamma_d = 10^7$ we achieve the optimal number of outer FGMRES iterations per BM iteration. The approximation of the Schur complement with the pressure mass matrix is near-perfect. However, we note that for each increase in the order of magnitude after $\gamma_d = 10^2$, the roundoff error causes the nonlinear solver to stagnate at an order of magnitude higher. We believe a good compromising choice is $\gamma_d = 10^4$.

To investigate (H3) we utilize the (aL2) solver (with varying γ_d). We fix a mesh with mesh size $h = 0.0361$ for the coarse level and uniformly refine the mesh once for the fine level. We consider the first outer Krylov FGMRES iteration at $\mu = 105$ and tabulate the average number of multigrid preconditioned inner Krylov FGMRES iterations per augmented momentum block solve required. The augmented momentum block solve is approximated to an absolute tolerance of 10^{-8} or a relative tolerance of 10^{-9} . We see that with increasing γ_d the iterations stay roughly constant. This confirms that the cycles are robust in γ_d .

TABLE 3

The sensitivity of the (aL1) and (aL2) solver strategies to the value of γ_d . For the (aL1) strategy, the columns labeled avg. OK show the average number of outer Krylov FGMRES iterations per BM iteration to solve the deflated barrier method subproblem at $\mu = 105$ and $\mu = 1$ for $h = 0.0361$. The columns labeled res. norm hold the values of the smallest residual norm before the nonlinear solver stagnates due to roundoff errors. For the (aL2) strategy the coarse level has a mesh size of $h = 0.0361$, and we perform one uniform refinement. The column labeled avg. IK shows the average number of multigrid preconditioned inner Krylov FGMRES iterations per augmented momentum block solve as seen in the first outer Krylov FGMRES iteration at subproblem $\mu = 105$.

γ_d	(aL1)				(aL2)
	$\mu = 105$		$\mu = 1$		$\mu = 105$
	avg. OK	res. norm	avg. OK	res. norm	avg. IK
10^{-5}	157.2	9.90×10^{-12}	66.0	6.40×10^{-12}	10.0
10^{-1}	133.5	9.93×10^{-12}	93.9	6.73×10^{-12}	11.3
10^0	77.8	1.09×10^{-11}	40.6	6.92×10^{-12}	12.8
10^1	30.0	1.15×10^{-11}	16.3	8.99×10^{-12}	10.0
10^2	8.23	3.92×10^{-11}	4.44	4.73×10^{-11}	9.50
10^3	3.31	3.75×10^{-10}	3.12	4.67×10^{-10}	9.17
10^4	1.85	3.71×10^{-9}	2.20	4.57×10^{-9}	9.50
10^5	1.33	3.70×10^{-8}	1.67	4.68×10^{-8}	9.50
10^6	1.25	3.77×10^{-7}	1.27	4.57×10^{-7}	9.50
10^7	1.14	3.73×10^{-6}	1.20	4.59×10^{-6}	9.50
10^8	1.18	3.63×10^{-5}	1.50	4.92×10^{-5}	9.50




FIG. 6. Setup of the 3D cross-channel problem. This problem features a unit cube domain with two inlets and two outlets arranged in a cross pattern.

6.2. 3D cross-channel. The first three-dimensional example we consider is the cross-channel problem as found in Sá et al. [58, sect. 7.5]. The domain is the unit cube, $\Omega = (0, 1)^3$, with two circular inlets and two circular outlets that are arranged in a cross pattern as visualized in Figure 6. The volume fraction is given by $\gamma = 1/10$, and we use (2.2) as our choice of α , with $\bar{\alpha} = 2.5 \times 10^4$ and $q = 1/10$. The Dirichlet boundary condition on \mathbf{u} is

$$(6.2) \quad \mathbf{g}(x, y, z) = (1 - 12\pi((y - a)^2 + (z - b)^2), 0, 0)^\top$$

if $12\pi((y - a)^2 + (z - b)^2) \leq 1$ and $x = 0$ with $a = 1/2$, $b \in \{1/4, 3/4\}$, or $x = 1$ with $a \in \{1/4, 3/4\}$, $b = 1/2$, and $\mathbf{g}(x, y, z) = (0, 0, 0)^\top$ elsewhere on $\partial\Omega$.

We apply the same first-order BDM discretization, with interior penalty penalization parameter $\sigma = 10$ and run the deflated barrier method twice. The first pass is on a $20 \times 20 \times 20$ mesh resulting in 391,201 degrees of freedom. The augmented Lagrangian parameter is chosen to be $\gamma_d = 10^6$. Due to the nested block preconditioning




FIG. 7. The material distribution of the solutions discovered by the deflated barrier method to the 3D cross-channel optimization problem discretized with 3,100,801 degrees of freedom. The power dissipation values are $J_h = 14.51, 14.62$, and 13.08 for branches 0, 1, and 2, respectively.

tioning, the action of the inverse of the augmented momentum block $\mathbf{A}_{\gamma_d} - \mathbf{D}\mathbf{C}_{\mu}^{-1}\mathbf{D}^{\top}$ must be applied six times per outer FGMRES iteration. On this relatively coarse mesh, it is cheaper to factorize $\mathbf{A}_{\gamma_d} - \mathbf{D}\mathbf{C}_{\mu}^{-1}\mathbf{D}^{\top}$ with MUMPS at the start of each BM iteration and reuse the factorization, rather than iteratively solve $\mathbf{A}_{\gamma_d} - \mathbf{D}\mathbf{C}_{\mu}^{-1}\mathbf{D}^{\top}$ with multigrid each time an inverse action is required, i.e., we use the augmented Lagrangian preconditioner (aL1) for the linear systems. We initialize the barrier parameter at $\mu_0 = 100$. The nonlinear solves are terminated with an absolute tolerance of 10^{-5} , and the outer FGMRES solver's absolute tolerance is set to 5×10^{-7} (all linear solves converged to the absolute tolerance). A second branch of solutions is found at $\mu = 38.74$ and a third branch at $\mu = 34.87$.

These coarse-mesh solutions at $\mu = 0$ are interpolated onto a finer $40 \times 40 \times 40$ mesh resulting in 3,100,801 degrees of freedom. The deflated barrier method is then reinitialized at $\mu_0 = 10^{-6}$; we found that the BM solver often diverges if initialized at $\mu_0 = 0$. On this finer mesh, we again apply the augmented Lagrangian preconditioner ($\gamma_d = 10^5$). The nonlinear solves are terminated with an absolute tolerance of 10^{-4} if $\mu > 0$ and 10^{-6} if $\mu = 0$. Moreover, the outer FGMRES solver's absolute tolerance is set to 10^{-7} if $\mu > 0$ and 10^{-9} if $\mu = 0$, and its relative tolerance is set to 10^{-7} . Now a direct solve of the augmented momentum block is prohibitive, and we switch to the (aL2) strategy where each approximate inverse of the augmented momentum block is solved to an absolute or relative tolerance of 10^{-8} or 10^{-9} , respectively, with the (two-grid) multigrid scheme of subsection 5.2. For the relaxation on the fine level, we use five FGMRES iterations preconditioned with the vertex-star patch relaxation.

The resulting three solutions are shown in Figure 7. Two of these are symmetric straight channel solutions where the inlets swap which outlet they exit from. Their symmetry results in similar costs. A third global minimizer comes in the form of a merged channel solution; the two channels briefly merge in the middle of the box domain before splitting to exit via the two outlets.

The iteration counts for the initial search on a $20 \times 20 \times 20$ mesh and the refinement on a $40 \times 40 \times 40$ mesh are given in Table 4. We see that our preconditioner is effective in both cases. When using a direct solve of the augmented momentum block, we average slightly more than one outer FGMRES iterations per BM iteration. Similarly, the tolerances for the augmented momentum block solve on the fine mesh are strict enough so that the outer FGMRES iterations average between 2.24 and 2.55.

TABLE 4

Cumulative number of BM iterations, outer FGMRES iterations (OK), and, for the fine mesh, of inner FGMRES iterations preconditioned with the multigrid scheme of subsection 5.2 (IK) for the 3D cross-channel problem. The bracketed numbers in the OK and IK columns show the average number of outer FGMRES iterations per BM iteration and average number of inner FGMRES iterations per augmented momentum block solve, respectively. The barrier parameter is initialized at $\mu_0 = 100$ on the coarse mesh and $\mu_0 = 10^{-6}$ on the fine mesh.

Branch	Coarse mesh, $\gamma_d = 10^6$		Fine mesh, $\gamma_d = 10^5$		
	BM	OK	BM	OK	IK
0	828	829 (1.00)	49	110 (2.24)	10,405 (15.77)
1	444	445 (1.00)	52	121 (2.33)	12,176 (16.77)
2	409	422 (1.03)	53	135 (2.55)	13,085 (16.15)




FIG. 8. Setup of the 3D five-holes quadruple-pipe problem. This problem features four inlets and four outlets. The domain is a box, $\Omega = (0, 3/2) \times (0, 1) \times (0, 1)$, with five internal holes in the shape of cubes with edge length $1/10$ that are centered at $(3/4, 1/4, 1/4)$, $(3/4, 1/4, 3/4)$, $(3/4, 3/4, 1/4)$, $(3/4, 3/4, 3/4)$, and $(3/4, 1/2, 1/2)$.

Moreover, each augmented momentum block solve requires an average in the range of 15.77–16.77 multigrid preconditioned FGMRES iterations to reach the prescribed tolerances.

6.3. 3D five-holes quadruple-pipe. In [51, sect. 4.4], it was observed that introducing holes in a rectangular domain caused a significant increase in the number of solutions. We now extend this idea to three dimensions and introduce the generalization of the five-holes double-pipe problem. This problem features a box domain $\Omega = (0, 3/2) \times (0, 1) \times (0, 1)$ with five internal holes in the shape of cubes, of edge length $1/10$, with centers at $(3/4, 1/4, 1/4)$, $(3/4, 1/4, 3/4)$, $(3/4, 3/4, 1/4)$, $(3/4, 3/4, 3/4)$, and $(3/4, 1/2, 1/2)$. There are four inlets and four outlets. The circular inlets of radius $1/\sqrt{12\pi}$ are positioned on the face $x = 0$ with the centers $(y, z) = (1/4, 1/4)$, $(1/4, 3/4)$, $(3/4, 1/4)$, and $(3/4, 3/4)$. The circular outlets of the same radius are positioned on the face $x = 3/2$ with the same centers. The domain setup is depicted in Figure 8. We impose a parabolic Dirichlet boundary condition on the inlets and outlets and a zero Dirichlet boundary condition elsewhere on the boundary (including the boundary of the five internal holes), i.e., the Dirichlet boundary condition is given by

$$(6.3) \quad \mathbf{g}(x, y, z) = (1 - 12\pi((y - a)^2 + (z - b)^2), 0, 0)^\top$$

if $12\pi((y - a)^2 + (z - b)^2) \leq 1$, where $a, b \in \{1/4, 3/4\}$, $x = 0$, or $3/2$ and

$$(6.4) \quad \mathbf{g}(x, y, z) = (0, 0, 0)^\top$$

TABLE 5

Cumulative number of BM iterations, outer FGMRES iterations (OK), and, for the fine mesh, of inner FGMRES iterations preconditioned with the multigrid scheme of subsection 5.2 (IK) for the 3D five-holes quadruple-pipe problem. The bracketed numbers in the OK and IK columns show the average number of outer FGMRES iterations per BM iteration and average number of inner FGMRES iterations per augmented momentum block solve, respectively. The barrier parameter is initialized at $\mu_0 = 200$ on the coarse mesh and $\mu_0 = 10^{-5}$ on the fine mesh.

Branch	Coarse mesh		Fine mesh		
	BM	OK	BM	OK	IK
0	438	438 (1)	41	109 (2.66)	15,828 (24.2)
1	338	338 (1)	36	97 (2.69)	13,736 (23.6)
2	396	396 (1)	49	130 (2.65)	18,999 (24.4)
3	254	254 (1)	46	114 (2.48)	16,464 (24.1)
7	221	221 (1)	53	132 (2.49)	20,061 (25.3)
11	183	183 (1)	64	180 (2.81)	27,738 (25.7)
13	187	187 (1)	60	169 (2.82)	28,542 (28.1)

elsewhere on $\partial\Omega$, including the boundaries of the five internal holes. We choose a volume fraction of $\gamma = 1/5$, and the inverse permeability term α is given by (2.2), with $\bar{\alpha} = 2.5 \times 10^4$ and $q = 1/10$.

The domain is nonconvex and has a complicated geometry due to the cuboidal internal holes, which we do not mesh. We build the coarse-level mesh with Gmsh; the finer levels are built by uniform refinement of the coarse mesh. We choose a first-order BDM discretization for the velocity-pressure pair, with interior penalty penalization parameter $\sigma = 10^3$, and run the deflated barrier method twice. The larger choice for σ is required to sufficiently enforce the boundary conditions in the tangential directions. The first run is on a (relatively) coarse mesh with 30,848 elements which results in 256,745 degrees of freedom. The barrier parameter is initialized at $\mu_0 = 200$, and we use the augmented Lagrangian preconditioner (aL1) for the linear systems, with an augmented Lagrangian parameter value of $\gamma_d = 10^5$. The nonlinear solves are terminated with an absolute tolerance of 10^{-4} . The outer FGMRES solver's absolute tolerance is set to 5×10^{-7} (all linear solves converged to the absolute tolerance).

In total we find 14 coarse-grid solutions. Branches 1 and 2 are found at $\mu = 53.81$. Branches 3, 4, 5, and 6 are found at $\mu = 11.39$. Branch 7 is found at $\mu = 10.25$. Branch 8 is found at $\mu = 9.23$. Branch 9 is found at $\mu = 6.73$, branch 10 at $\mu = 6.73$, branches 11 and 12 are found at $\mu = 6.05$, and, finally, branch 13 is found at $\mu = 4.41$.

We interpolate the coarse-level solutions onto the first refinement, which results in 2,014,113 degrees of freedom. The deflated barrier method is initialized at $\mu_0 = 10^{-5}$, using the coarse-level solutions as initial guesses. We damp the l^2 -minimizing line-search in the nonlinear solver by a factor of $1/2$. The nonlinear solves are terminated with an absolute tolerance of 10^{-4} if $\mu > 0$ and 10^{-6} if $\mu = 0$. The outer FGMRES solver's absolute tolerance is set to 10^{-6} if $\mu > 0$ and 10^{-7} if $\mu = 0$, and its relative tolerance is set to 10^{-5} . We apply an augmented Lagrangian two-grid multigrid preconditioner (aL2) to the linear systems with five FGMRES iterations for the relaxation of the fine level and $\gamma_d = 10^5$. The augmented momentum block is solved to an absolute or relative tolerance of 10^{-4} or 10^{-7} , respectively. Of the original 14 solutions, the nonlinear solver successfully converges to seven fine-grid solutions.

The resulting iteration counts are given in Table 5, the resulting fine-mesh solutions are shown in Figure 9, and their cross sections are shown in Figure 10. A further mesh refinement of branch 0 is shown in Figure 11 as computed using the augmented Lagrangian (aL2) with a three-grid multigrid cycle. As expected, the five




FIG. 9. The material distribution of seven solutions to the 3D five-holes quadruple-pipe optimization problem discretized with 2,014,113 degrees of freedom as discovered by the deflated barrier method and grid-sequenced once. The resulting power dissipation values for branches 0, 1, 2, 3, 7, 11, and 13 are $J_h = 55.02, 54.73, 62.27, 62.22, 59.56, 55.28$, and 62.78 , respectively.

holes obstruct the channels and prevent a large channel from passing through the center. The best solutions found are branches 0, 1, and 11 where the channels form one large channel and either move to the left, upwards, or downwards to avoid the middle internal hole. As in the five-holes double-pipe example, there are remaining solutions that we have not yet computed by symmetry.

In Table 6 we give timings for the solvers as run on a machine with 512 GB of RAM and 32 core Intel Xeon CPU E5-4627 v2 @ 3.30 GHz. We note that a direct LU factorization of the full BM system (4.18) fails to converge due to the ill-conditioning of the system (even with $\gamma_d = 0$). Thus the preconditioning strategies are necessary to find the BM updates. On the coarse mesh, the (aL1) strategy is faster than the (aL2)




FIG. 10. The crinkled cross sections at $x=3/4$ for the discovered solutions of the 3D five-holes quadruple-pipe. The grey regions are part of the five cuboid holes in the box domain. The material distribution has a value of one in the red regions and zero in the blue regions, with intermediate values for the intermediate colored regions.




FIG. 11. The material distribution of branch 0 after grid-sequencing to a second uniform mesh refinement resulting in 15,953,537 degrees of freedom ($J_h = 39.11$). The augmented Lagrangian three-grid multigrid preconditioner (aL2) was used.

strategy. Moreover, we show that the average time taken per BM iteration drops from 277 seconds with 1 CPU to 35 seconds when using 32 CPUs. On the fine mesh, the assembly of the full fine-grid augmented block requires more memory than what is available on the workstation (512 GB). Hence only the (aL2) strategy succeeded in grid-sequencing the solutions on this machine.

TABLE 6

Time measurements of the linear solves of the BM system. On the coarse mesh we measure the number of BM iterations and overall time taken to find the first branch at $\mu = 200$. On the fine mesh we measure the number of BM iterations and overall time taken to grid-sequence the first solution starting at $\mu = 10^{-5}$.

Strategy	BM	Time taken (s)	Avg. time taken (s)
Coarse mesh			
(aL1) 1 CPU	6	1664.02	277.33
(aL1) 4 CPUs	44	4425.71	100.58
(aL1) 8 CPUs	37	2310.22	62.44
(aL1) 16 CPUs	42	1844.64	43.92
(aL1) 32 CPUs	6	208.64	34.77
(aL2) 32 CPUs	7	2150.24	307.18
Fine mesh			
(aL2) 32 CPUs	41	53,356.56	1301.38

Code availability: For reproducibility, the implementation of the deflated barrier method used in this work, as well as scripts to generate the solutions, can be found at <https://github.com/ioannisPApapadopoulos/fir3dab/>. The version of the software used in this paper is archived on Zenodo [50].

7. Conclusions. In this work we extended the application of the deflated barrier method to discovering multiple three-dimensional solutions of the fluid topology optimization model of Borrvall and Petersson [14]. This was achieved by developing preconditioners for the linear systems that arise in the deflated barrier method. The preconditioning strategy reduces the discretized 4×4 -block matrix solve to an outer FGMRES method and the solve of one diagonal matrix, one block-diagonal matrix (that can be factorized once or quickly solved with multigrid), and an augmented momentum block. Moreover, we develop a geometric multigrid cycle for the augmented momentum block that consists of a specialized relaxation method and a characterization of the active set on coarser levels. We found that the preconditioner is robust to the mesh size and successfully computed three solutions of a 3D cross-channel problem and seven solutions of a 3D five-holes quadruple-pipe problem.

REFERENCES

- [1] N. AAGE AND B. S. LAZAROV, *Parallel framework for topology optimization using the method of moving asymptotes*, Struct. Multidiscip. Optim., 47 (2013), pp. 493–505, <https://doi.org/10.1007/s00158-012-0869-2>.
- [2] N. AAGE, T. H. POULSEN, A. GERSBORG-HANSEN, AND O. SIGMUND, *Topology optimization of large scale Stokes flow problems*, Struct. Multidiscip. Optim., 35 (2007), pp. 175–180, <https://doi.org/10.1007/s00158-007-0128-0>.
- [3] J. ALEXANDERSEN AND C. S. ANDREASEN, *A review of topology optimisation for fluid-based problems*, Fluids, 5 (2020), 29, <https://doi.org/10.3390/fluids5010029>.
- [4] F. P. ALI BEIK AND M. BENZI, *Iterative methods for double saddle point systems*, SIAM J. Matrix Anal. Appl., 39 (2018), pp. 902–921, <https://doi.org/10.1137/17M1121226>.
- [5] P. R. AMESTOY, I. S. DUFF, J.-Y. L'EXCELLENT, AND J. KOSTER, *A fully asynchronous multifrontal solver using distributed dynamic scheduling*, SIAM J. Matrix Anal. Appl., 23 (2001), pp. 15–41, <https://doi.org/10.1137/S0895479899358194>.
- [6] 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, <https://doi.org/10.1007/PL00005386>.
- [7] F. R. A. AZNARAN, P. E. FARRELL, AND R. C. KIRBY, *Transformations for Piola-mapped elements*, SMAI J. Comput. Math., 8 (2022), pp. 399–437, <https://doi.org/10.5802/smai-jcm.91>.

- [8] S. BALAY, S. ABHYANKAR, M. ADAMS, J. BROWN, P. BRUNE, K. BUSCHELMAN, L. DALCIN, A. DENER, V. ELJKOUT, W. GROPP, R. TRAN MILLS, T. MUNSON, K. RUPP, P. SANA, B. SMITH, S. ZAMPINI, H. ZHANG, AND H. ZHANG, *PETSc Users Manual*, Technical Report ANL-95/11 - Revision 3.11, Argonne National Laboratory, 2019, <http://www.mcs.anl.gov/petsc>.
- [9] R. BEHROU, R. RANJAN, AND J. K. GUEST, *Adaptive topology optimization for incompressible laminar flow problems with mass flow constraints*, *Comput. Methods Appl. Mech. Engrg.*, 346 (2019), pp. 612–641, <https://doi.org/10.1016/j.cma.2018.11.037>.
- [10] M. P. BENDSØE AND O. SIGMUND, *Topology Optimization: Theory, Methods, and Applications*, Springer, 2004, <https://doi.org/10.1007/978-3-662-05086-6>.
- [11] S. J. BENSON AND T. S. MUNSON, *Flexible complementarity solvers for large-scale applications*, *Optim. Methods Softw.*, 21 (2006), pp. 155–168, <https://doi.org/10.1080/10556780500065382>.
- [12] M. BENZI, G. H. GOLUB, AND J. LIESEN, *Numerical solution of saddle point problems*, *Acta Numer.*, 14 (2005), pp. 1–137, <https://doi.org/10.1017/S0962492904000212>.
- [13] M. BENZI AND M. A. OLSHANSKII, *An augmented Lagrangian-based approach to the Oseen problem*, *SIAM J. Sci. Comput.*, 28 (2006), pp. 2095–2113, <https://doi.org/10.1137/050646421>.
- [14] T. BORRVALL AND J. PETERSSON, *Topology optimization of fluids in Stokes flow*, *Internat. J. Numer. Methods Fluids*, 41 (2003), pp. 77–107, <https://doi.org/10.1002/fld.426>.
- [15] F. BREZZI, J. DOUGLAS, JR., R. DURÁN, AND M. FORTIN, *Mixed finite elements for second order elliptic problems in three variables*, *Numer. Math.*, 51 (1987), pp. 237–250, <https://doi.org/10.1007/BF01396752>.
- [16] F. BREZZI, J. DOUGLAS, JR., AND L. D. MARINI, *Two families of mixed finite elements for second order elliptic problems*, *Numer. Math.*, 47 (1985), pp. 217–235, <https://doi.org/10.1007/BF01389710>.
- [17] P. R. BRUNE, M. G. KNEPLEY, B. F. SMITH, AND X. TU, *Composing scalable nonlinear algebraic solvers*, *SIAM Rev.*, 57 (2015), pp. 535–565, <https://doi.org/10.1137/130936725>.
- [18] J. CAHOUEU AND J.-P. CHABARD, *Some fast 3D finite element solvers for the generalized Stokes problem*, *Internat. J. Numer. Methods Fluids*, 8 (1988), pp. 869–895, <https://doi.org/10.1002/fld.1650080802>.
- [19] V. J. CHALLIS AND J. K. GUEST, *Level set topology optimization of fluids in Stokes flow*, *Internat. J. Numer. Methods Engrg.*, 79 (2009), pp. 1284–1308, <https://doi.org/10.1002/nme.2616>.
- [20] Y. DENG, Z. LIU, J. WU, AND Y. WU, *Topology optimization of steady Navier-Stokes flow with body force*, *Comput. Methods Appl. Mech. Engrg.*, 255 (2013), pp. 306–321, <https://doi.org/10.1016/j.cma.2012.11.015>.
- [21] H. C. ELMAN, D. J. SILVESTER, AND A. J. WATHEN, *Finite Elements and Fast Iterative Solvers: With Applications in Incompressible Fluid Dynamics*, 2nd ed., Oxford University Press, 2014, <https://doi.org/10.1093/acprof:oso/9780199678792.001.0001>.
- [22] M. ENGEL AND M. GRIEBEL, *A multigrid method for constrained optimal control problems*, *J. Comput. Appl. Math.*, 235 (2011), pp. 4368–4388, <https://doi.org/10.1016/j.cam.2011.04.002>.
- [23] L. C. EVANS, *Partial Differential Equations*, 2nd ed., American Mathematical Society, 2010.
- [24] A. EVGRAFOV, *State space Newton’s method for topology optimization*, *Comput. Methods Appl. Mech. Engrg.*, 278 (2014), pp. 272–290, <https://doi.org/10.1016/j.cma.2014.06.005>.
- [25] P. E. FARRELL, Å. BIRKISSON, AND S. W. FUNKE, *Deflation techniques for finding distinct solutions of nonlinear partial differential equations*, *SIAM J. Sci. Comput.*, 37 (2015), pp. A2026–A2045, <https://doi.org/10.1137/140984798>.
- [26] P. E. FARRELL, M. CROCI, AND T. M. SUROWIEC, *Deflation for semismooth equations*, *Optim. Methods Softw.*, 35 (2020), pp. 1248–1271, <https://doi.org/10.1080/10556788.2019.1613655>.
- [27] P. E. FARRELL AND P. A. GAZCA-OROZCO, *An augmented Lagrangian preconditioner for implicitly constituted non-Newtonian incompressible flow*, *SIAM J. Sci. Comput.*, 42 (2020), pp. B1329–B1349, <https://doi.org/10.1137/20M1336618>.
- [28] P. E. FARRELL, M. G. KNEPLEY, L. MITCHELL, AND F. WECHSUNG, *PCPATCH: Software for the topological construction of multigrid relaxation methods*, *ACM Trans. Math. Softw. (TOMS)*, 47 (2021), 25, <https://doi.org/10.1145/3445791>.
- [29] P. E. FARRELL, L. MITCHELL, L. R. SCOTT, AND F. WECHSUNG, *A Reynolds-robust preconditioner for the Scott–Vogelius discretization of the stationary incompressible Navier–Stokes equations*, *SMAI J. Comput. Math.*, 7 (2021), pp. 75–96, <https://doi.org/10.5802/smai-jcm.72>.

- [30] P. E. FARRELL, L. MITCHELL, AND F. WECHSUNG, *An augmented Lagrangian preconditioner for the 3D stationary incompressible Navier–Stokes equations at high Reynolds number*, SIAM J. Sci. Comput., 41 (2019), pp. A3073–A3096, <https://doi.org/10.1137/18M1219370>.
- [31] P. E. FARRELL, P. A. G. OROZCO, AND E. SÜLI, *Finite element approximation and preconditioning for anisothermal implicitly-constituted non-Newtonian flow*, Math. Comp., 91 (2022), pp. 659–697, <https://doi.org/10.1090/mcom/3703>.
- [32] A. V. FIACCO AND G. P. MCCORMICK, *Nonlinear Programming*, Classics Appl. Math. 4, SIAM, 1990, <https://doi.org/10.1137/1.9781611971316>.
- [33] A. FORSGREN, P. E. GILL, AND M. H. WRIGHT, *Interior methods for nonlinear optimization*, SIAM Rev., 44 (2002), pp. 525–597, <https://doi.org/10.1137/S0036144502414942>.
- [34] M. FORTIN AND R. GLOWINSKI, *Augmented Lagrangian Methods: Applications to the Numerical Solution of Boundary-Value Problems*, Stud. Math. Its Appl. 15, Elsevier, 1983.
- [35] K. R. FRISCH, *The Logarithmic Potential Method of Convex Programming*, Technical report, University Institute of Economics, Oslo, Norway, 1955.
- [36] N. R. GAUGER, A. LINKE, AND P. W. SCHROEDER, *On high-order pressure-robust space discretisations, their advantages for incompressible high Reynolds number generalised Beltrami flows and beyond*, SMAI J. Comput. Math., 5 (2019), pp. 89–129, <https://doi.org/10.5802/smai-jcm.44>.
- [37] C. GEUZAIN AND J.-F. REMACLE, *Gmsh: A 3-D finite element mesh generator with built-in pre- and post-processing facilities*, Internat. J. Numer. Methods Engrg., 79 (2009), pp. 1309–1331, <https://doi.org/10.1002/nme.2579>.
- [38] M. HINTERMÜLLER, K. ITO, AND K. KUNISCH, *The primal-dual active set strategy as a semismooth Newton method*, SIAM J. Optim., 13 (2002), pp. 865–888, <https://doi.org/10.1137/S1052623401383558>.
- [39] Q. HONG, J. KRAUS, J. XU, AND L. ZIKATANOV, *A robust multigrid method for discontinuous Galerkin discretizations of Stokes and linear elasticity equations*, Numer. Math., 132 (2016), pp. 23–49, <https://doi.org/10.1007/s00211-015-0712-y>.
- [40] R. H. HOPPE, *Multigrid algorithms for variational inequalities*, SIAM J. Numer. Anal., 24 (1987), pp. 1046–1065, <https://doi.org/10.1137/0724069>.
- [41] K. E. JENSEN, *Topology optimization of Stokes flow on dynamic meshes using simple optimizers*, Comput. & Fluids, 174 (2018), pp. 66–77, <https://doi.org/10.1016/j.compfluid.2018.07.011>.
- [42] V. JOHN, A. LINKE, C. MERDON, M. NEILAN, AND L. G. REBHOLZ, *On the divergence constraint in mixed finite element methods for incompressible flows*, SIAM Rev., 59 (2017), pp. 492–544, <https://doi.org/10.1137/15M1047696>.
- [43] J. KÖNNÖ AND R. STENBERG, *Numerical computations with $H(\text{div})$ -finite elements for the Brinkman problem*, Comput. Geosci., 16 (2012), pp. 139–158, <https://doi.org/10.1007/s10596-011-9259-x>.
- [44] F. LAAKMANN, P. E. FARRELL, AND L. MITCHELL, *An augmented Lagrangian preconditioner for the magnetohydrodynamics equations at high Reynolds and coupling numbers*, SIAM J. Sci. Comput., 44 (2022), pp. B1018–B1044, <https://doi.org/10.1137/21M1416539>.
- [45] A. LINKE AND L. G. REBHOLZ, *Pressure-induced locking in mixed methods for time-dependent (Navier–) Stokes equations*, J. Comput. Phys., 388 (2019), pp. 350–356, <https://doi.org/10.1016/j.jcp.2019.03.010>.
- [46] K. A. MARDAL, X.-C. TAI, AND R. WINTHER, *A robust finite element method for Darcy–Stokes flow*, SIAM J. Numer. Anal., 40 (2002), pp. 1605–1631, <https://doi.org/10.1137/S0036142901383910>.
- [47] M. F. MURPHY, G. H. GOLUB, AND A. J. WATHEN, *A note on preconditioning for indefinite linear systems*, SIAM J. Sci. Comput., 21 (2000), pp. 1969–1972, <https://doi.org/10.1137/S1064827599355153>.
- [48] I. P. PAPADOPOULOS, *Numerical analysis of a discontinuous Galerkin method for the Borrvall–Pettersson topology optimization problem*, SIAM J. Numer. Anal., 60 (2022), pp. 2538–2564, <https://doi.org/10.1137/21M1438943>.
- [49] I. P. A. PAPADOPOULOS, *Computing Multiple Solutions of Topology Optimization Problems*, Ph.D. thesis, University of Oxford, 2021, <https://ora.ox.ac.uk/objects/uuid:455f5465-af68-49c7-8e11-96909cdb82a1>.
- [50] I. P. A. PAPADOPOULOS AND P. E. FARRELL, *– fir3 dab Code Used in ‘Preconditioners for Computing Multiple Solutions in Three-Dimensional Fluid Topology Optimization,’ 2022*, <https://doi.org/10.5281/zenodo.5899634>.
- [51] I. P. A. PAPADOPOULOS, P. E. FARRELL, AND T. M. SUROWIEC, *Computing multiple solutions of topology optimization problems*, SIAM J. Sci. Comput., 43 (2021), pp. A1555–A1582, <https://doi.org/10.1137/20M1326209>.

- [52] I. P. A. PAPADOPOULOS AND E. SÜLI, *Numerical analysis of a topology optimization problem for Stokes flow*, J. Comput. Appl. Math., 412 (2022), 114295, <https://doi.org/10.1016/j.cam.2022.114295>.
- [53] G. PINGEN, A. EVGRAFOV, AND K. MAUTE, *Topology optimization of flow domains using the lattice Boltzmann method*, Struct. Multidiscip. Optim., 34 (2007), pp. 507–524, <https://doi.org/10.1007/s00158-007-0105-7>.
- [54] P. POPOV, *Preconditioning of linear systems arising in finite element discretizations of the Brinkman equation*, in Large-Scale Scientific Computing (LSSC 2011), Lecture Notes in Comput. Sci. 7116, Springer, 2012, pp. 381–389, https://doi.org/10.1007/978-3-642-29843-1_43.
- [55] L. QI, *Convergence analysis of some algorithms for solving nonsmooth equations*, Math. Oper. Res., 18 (1993), pp. 227–244, <https://doi.org/10.1287/moor.18.1.227>.
- [56] L. QI AND J. SUN, *A nonsmooth version of Newton's method*, Math. Program., 58 (1993), pp. 353–367, <https://doi.org/10.1007/BF01581275>.
- [57] F. RATHGEBER, D. A. HAM, L. MITCHELL, M. LANGE, F. LUPORINI, A. T. MCRAE, G.-T. BERCEA, G. R. MARKALL, AND P. H. KELLY, *Firedrake: Automating the finite element method by composing abstractions*, ACM Trans. Math. Softw., 43 (2016), pp. 1–27, <https://doi.org/10.1145/2998441>.
- [58] L. F. N. SÁ, R. C. R. AMIGO, A. A. NOVOTNY, AND E. C. N. SILVA, *Topological derivatives applied to fluid flow channel design optimization problems*, Struct. Multidiscip. Optim., 54 (2016), pp. 249–264, <https://doi.org/10.1007/s00158-016-1399-0>.
- [59] Y. SAAD, *A flexible inner-outer preconditioned GMRES algorithm*, SIAM J. Sci. Comput., 14 (1993), pp. 461–469, <https://doi.org/10.1137/0914028>.
- [60] A. SCHIELA, *Barrier methods for optimal control problems with state constraints*, SIAM J. Optim., 20 (2009), pp. 1002–1031, <https://doi.org/10.1137/070692789>.
- [61] A. SCHIELA AND M. WEISER, *Superlinear convergence of the control reduced interior point method for PDE constrained optimization*, Comput. Optim. Appl., 39 (2008), pp. 369–393, <https://doi.org/10.1007/s10589-007-9057-5>.
- [62] J. SCHÖBERL, *Multigrid methods for a parameter dependent problem in primal variables*, Numer. Math., 84 (1999), pp. 97–119, <https://doi.org/10.1007/s002110050465>.
- [63] D. SILVESTER AND A. WATHEN, *Fast iterative solution of stabilised Stokes systems part II: Using general block preconditioners*, SIAM J. Numer. Anal., 31 (1994), pp. 1352–1367, <https://doi.org/10.1137/0731070>.
- [64] M. STOLL AND A. J. WATHEN, *Preconditioning for active set and projected gradient methods as semi-smooth Newton methods for PDE-constrained optimization with control constraints*, 2009.
- [65] M. STOLPE AND K. SVANBERG, *On the trajectories of penalization methods for topology optimization*, Struct. Multidiscip. Optim., 21 (2001), pp. 128–139, <https://doi.org/10.1007/s001580050177>.
- [66] M. ULBRICH, *Semismooth Newton methods for operator equations in function spaces*, SIAM J. Optim., 13 (2002), pp. 805–841, <https://doi.org/10.1137/s1052623400371569>.
- [67] M. ULBRICH AND S. ULBRICH, *Primal-dual interior-point methods for PDE-constrained optimization*, Math. Program., 117 (2009), pp. 435–485, <https://doi.org/10.1007/s10107-007-0168-7>.
- [68] A. WATHEN AND D. SILVESTER, *Fast iterative solution of stabilised Stokes systems. Part I: Using simple diagonal preconditioners*, SIAM J. Numer. Anal., 30 (1993), pp. 630–649, <https://doi.org/10.1137/0730031>.
- [69] A. J. WATHEN, *Preconditioning*, Acta Numer., 24 (2015), pp. 329–376, <https://doi.org/10.1017/S0962492915000021>.
- [70] M. WEISER, T. GÄNZLER, AND A. SCHIELA, *A control reduced primal interior point method for a class of control constrained optimal control problems*, Comput. Optim. Appl., 41 (2008), pp. 127–145, <https://doi.org/10.1007/s10589-007-9088-y>.
- [71] J. XIA, P. E. FARRELL, AND F. WECHSUNG, *Augmented Lagrangian preconditioners for the Oseen–Frank model of nematic and cholesteric liquid crystals*, BIT, 61 (2021), pp. 607–644, <https://doi.org/10.1007/s10543-020-00838-9>.
- [72] J. XU, *Iterative methods by space decomposition and subspace correction*, SIAM Rev., 34 (1992), pp. 581–613, <https://doi.org/10.1137/1034116>.
- [73] J. XU, *The method of subspace corrections*, J. Comput. Appl. Math., 128 (2001), pp. 335–362, [https://doi.org/10.1016/S0377-0427\(00\)00518-5](https://doi.org/10.1016/S0377-0427(00)00518-5).
- [74] K. YAJI, T. YAMADA, M. YOSHINO, T. MATSUMOTO, K. IZUI, AND S. NISHIWAKI, *Topology optimization using the lattice Boltzmann method incorporating level set boundary expressions*, J. Comput. Phys., 274 (2014), pp. 158–181, <https://doi.org/10.1016/j.jcp.2014.06.004>.

- [75] F. ZHANG, *The Schur Complement and Its Applications*, Numer. Methods Algorithms 4, Springer, 2005, <https://doi.org/10.1007/b105056>.
- [76] S. ZHOU AND Q. LI, *A variational level set method for the topology optimization of steady-state Navier-Stokes flow*, J. Comput. Phys., 227 (2008), pp. 10178–10195, <https://doi.org/10.1016/j.jcp.2008.08.022>.