Accepted Manuscript

A scalable block-preconditioning strategy for divergence-conforming B-spline discretizations of the Stokes problem

A.M.A. Côrtes, L. Dalcin, A. Sarmiento, N. Collier, V.M. Calo

PII: S0045-7825(16)31309-3

DOI: <http://dx.doi.org/10.1016/j.cma.2016.10.014>

Reference: CMA 11175

To appear in: *Comput. Methods Appl. Mech. Engrg.*

Please cite this article as: A.M.A. Côrtes, L. Dalcin, A. Sarmiento, N. Collier, V.M. Calo, A scalable block-preconditioning strategy for divergence-conforming B-spline discretizations of the Stokes problem, *Comput. Methods Appl. Mech. Engrg.* (2016), <http://dx.doi.org/10.1016/j.cma.2016.10.014>

This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

A scalable block-preconditioning strategy for divergence-conforming B-spline discretizations of the Stokes problem

A.M.A. Côrtes^{a,d,*}, L. Dalcin^{a,e}, A. Sarmiento^{a,b}, N. Collier^f, V.M. Calo^{a,b,c}

^aCenter for Numerical Porous Media (NumPor)
King Abdullah University of Science and Technology (KAUST)
Thuwal, Saudi Arabia

^bApplied Mathematics and Computational Science (AMCS)
King Abdullah University of Science and Technology (KAUST)
Thuwal, Saudi Arabia

^cEarth Science and Engineering (ErSE)
King Abdullah University of Science and Technology (KAUST)
Thuwal, Saudi Arabia

^dDepartment of Mathematics and Statistics
Federal State University of Rio de Janeiro
Rio de Janeiro, Brazil

^eCentro de Investigación de Métodos Computacionales (CIMEC)
Consejo Nacional de Investigaciones Científicas y Técnicas (CONICET)
Santa Fe, Argentina

^fComputer Science and Mathematics Division
Oak Ridge National Laboratory
Oak Ridge, TN, USA

Abstract

The recently introduced divergence-conforming B-spline discretizations allow the construction of smooth discrete velocity-pressure pairs for viscous incompressible flows that are at the same time inf – sup stable and pointwise divergence-free. When applied to the discretized Stokes problem, these spaces generate a symmetric and indefinite saddle-point linear system. The iterative method of choice to solve such system is the Generalized Minimum Residual Method. This method lacks robustness, and one remedy is to use preconditioners. For linear systems of saddle-point type, a large family of preconditioners can be obtained by using a block factorization of the system. In this paper, we show how the nesting of “black-box” solvers and preconditioners can be put together in a block triangular strategy to build a scalable block preconditioner for the Stokes system discretized by divergence-conforming B-splines. Besides the known cavity flow problem, we used for benchmark flows defined on complex geometries: an eccentric annulus and hollow torus of an eccentric annular cross-section.

Keywords: isogeometric analysis, B-spline compatible vector field discretization, Krylov subspace methods, block preconditioners, Stokes flow

*Corresponding author

Email addresses: adrimacortes@gmail.com (A.M.A. Côrtes), dalcinl@gmail.com (L. Dalcin), adel.sarmientorodriguez@kaust.edu.sa (A. Sarmiento), nathaniel.collier@gmail.com (N. Collier), vmcalo@gmail.com (V.M. Calo)

1. Introduction

In [1], Buffa et al. introduced a discretization based on the isogeometric analysis (IGA) concept that satisfies a de Rham diagram and delivers good results for electromagnetic simulations. At the heart of their construction are the simple control of the polynomial order and smoothness that B-splines functions allow. The generalization for three-dimensional vector fields and the mathematical theory of such discretization appeared in [2]. Their approach, called isogeometric discrete differential forms, was inspired by the theory of finite element exterior calculus of Arnold et al. [3].

Following the developments of Buffa et al. [2], Evans [4] further explored these discrete spaces in the context of Hilbert complexes. Indeed, by using the stable projectors of [2], a divergence-preserving transformation (Piola transformation) of the velocity vector field and an integral-preserving transformation of the pressure field, Evans devised a Stokes complex with a compatible sub-complex that furnishes a discretization scheme, which is both $\inf - \sup$ stable and divergence-free. Evans [4], and Evans and Hughes [5–7] applied this discretization scheme to several viscous incompressible flows, and also started its mathematical analysis. This discretization scheme is overviewed in section 3.

The discretization of the Stokes equations by $\inf - \sup$ stable mixed elements requires the solution of a symmetric indefinite linear system, called the (discrete) Stokes system, with a block coefficient matrix of saddle-point type. Several strategies for solving the Stokes linear system have appeared in the literature, the most popular being variants of Uzawa method [8, 9], like the *inexact* Uzawa method, and Krylov subspace methods such as the Minimum Residual Method (MINRES) [10] and the Generalized Minimum Residual Method (GMRES) [11–14]. The Krylov subspace methods have better convergence properties than the Uzawa variants, and both lack robustness with respect to discretization parameters and equation coefficients. Thus, the use of a preconditioner is indispensable. For an excellent survey on the topic see [15]. One efficient preconditioning strategy for such problems, comes from a block factorization of the saddle-point system, that when used with other preconditioners as their building blocks renders powerful and scalable solution methods. For example, in [16], an Earth’s mantle convection flow solver based on high-order continuous/discontinuous finite elements with octree-based adaptive mesh refinement and coarsening (AMR/C) is reported to scale up to 1.5 million processor cores such strategy. The referred paper won the prestigious Gordon Bell prize in 2015.

To the best of our knowledge, the first study of block preconditioning strategies for B-splines divergence-conforming spaces was performed by the authors in [17]. In that study, the spectral equivalence of the block factorization strategy (thus, the invariance in the number of iterations of the iterative solver with mesh refinement) was shown possible, even though we used non-scalable direct methods to solve for the blocks. The question of a scalable choice for the blocks was only possible to address with our parallel implementation of divergence-conforming spaces [18] based on the high-performance isogeometric framework,

named PetIGA [19], built on top of PETSc [20], that releases a whole set of solvers and preconditioners. Here we investigate an approach built from *state-of-the-art* and *black-box* preconditioning strategies that furnish an h -independent preconditioner for the Stokes problem discretized with divergence-conforming B-spline spaces, with good performance for both sequential and parallel setups.

The rest of the paper is organized as follows. In section 2, we review some isogeometric analysis definitions to introduce the nomenclature related to the discretization. Section 3 is dedicated to overviewing the divergence-conforming B-spline spaces of [4, 5]. First, we present the discrete velocity-pressure pair defined on the parametric domain, and how it is pushed-forward to general domains described by a geometric mapping, and using a proper transformation. We also emphasize the inf – sup stability and the divergence-free property of the discrete scheme. We formulate the discrete variational problem and how Nitsche’s method is used to impose Dirichlet boundary conditions weakly.

In section 4, we review solution methods and block preconditioning strategies for the Stokes saddle-point systems. In section 5 we present a practical strategy based on composing black-box solver kernels to build a preconditioner for the discrete Stokes linear system. In section 6, we present our numerical results in two and three dimensions showing the good performance of the strategy. Finally, section 7 summarizes our conclusions and sketches possible extensions of this work.

2. Isogeometric analysis concepts

With the purpose of defining the divergence-conforming discretization of [2, 5] in the next section, we recall notation and definitions and that are standard in the IGA literature. We closely follow the expositions in [2, 4, 5, 21].

2.1. Univariate B-Splines

A set of univariate B-spline basis functions is defined by specifying the number n of basis functions, the polynomial degree p , and a finite nondecreasing sequence of real numbers $\Xi = \{\xi_1, \dots, \xi_{n+p+1}\}$ called *knot vector*. Since the knots can be repeated, we introduce the sequences $\zeta = \{\zeta_1, \dots, \zeta_m\}$ of breakpoints (knots without repetitions) and $\mathbf{r} = \{r_1, \dots, r_m\}$ of corresponding multiplicities.

The B-spline basis functions are p -degree piecewise polynomials on ζ , and the subintervals $[\zeta_j, \zeta_{j+1}]$ define a set of non-empty knot spans, akin to one-dimensional elements in the finite element parlance. A practical and efficient algorithm for evaluating B-spline basis functions on computers uses the Cox–de Boor recurrence formula, see [22].

Knot multiplicity is an important ingredient in spline theory, since it allows the control of smoothness. Indeed, if a breakpoint ζ_j has multiplicity r_j , then the basis functions have at least $\alpha_j := p - r_j$ continuous derivatives at ζ_j . Hence, the maximum multiplicity allowed for ζ_j is $r_j = p + 1$, in this case $\alpha_j = -1$ and

the basis is discontinuous at ζ_j . We restrict ourselves to open knot vectors, in this case $r_1 = r_m = p + 1$ and $1 \leq r_k \leq p$ for $k = 2, \dots, m - 1$, which implies $n \geq p + 1$ and $\alpha_1 = \alpha_m = -1$. The sequence collecting regularities is denoted by $\alpha := \{\alpha_1, \dots, \alpha_m\}$, and we define $\alpha - 1 = \{-1, \alpha_2 - 1, \dots, \alpha_{m-1} - 1, -1\}$, where $\alpha_k \geq 0$ for $k = 2, \dots, m - 1$, and $|\alpha| = \min\{\alpha_2, \dots, \alpha_{m-1}\}$.

The set $\{B_i^p\}_{i=1}^n$ defines a linear independent set of basis functions with all the properties required for analysis purposes [23]. The B-spline space spanned by them is denoted by

$$\mathcal{S}_\alpha^p := \text{span} \{B_i^p\}_{i=1}^n. \quad (1)$$

For univariate spaces, when $p \geq 1$ and $\alpha_k \geq 0$ for $k = 2, \dots, m - 1$, the derivative of a spline is also a spline. Moreover, the derivative is a surjective operator, that is,

$$\left\{ \frac{d}{dx} u : u \in \mathcal{S}_\alpha^p \right\} \equiv \mathcal{S}_{\alpha-1}^{p-1}. \quad (2)$$

Besides these properties, the refinement of B-spline spaces by knot insertion and degree elevation are also relevant features of B-spline technology resembling h -refinement and p -refinement in classical finite element methods. These procedures do not commute, and the application of degree elevation followed by knot insertion yields a new refinement procedure called k -refinement. For more details and examples, see [22–24].

2.2. Trivariate B-splines

Given $p_1, p_2, p_3, n_1, n_2, n_3 \in \mathbb{N}$, and the knot vectors Ξ_1, Ξ_2, Ξ_3 , we construct a set of univariate B-spline basis functions in each direction, that is, $\{B_{i_d,d}^{p_d}\}_{i_d=1}^{n_d}$ for $d = 1, 2, 3$. Trivariate B-spline basis functions are defined by tensor products of the univariate ones as

$$B_{i_1,i_2,i_3}^{p_1,p_2,p_3} := B_{i_1,1}^{p_1} \otimes B_{i_2,2}^{p_2} \otimes B_{i_3,3}^{p_3}, \quad (3)$$

where $i_1 = 1, \dots, n_1$; $i_2 = 1, \dots, n_2$; $i_3 = 1, \dots, n_3$. The bivariate case is defined similarly.

As in the univariate case, the sets of breakpoints $\zeta_d = \{\zeta_{1,d}, \dots, \zeta_{m_d,d}\}$ in each direction $d = 1, 2, 3$ define a structured mesh

$$\mathcal{M}_h = \{Q = (\zeta_{i_1,1}, \zeta_{i_1+1,1}) \times (\zeta_{i_2,2}, \zeta_{i_2+1,2}) \times (\zeta_{i_3,3}, \zeta_{i_3+1,3}) : 1 \leq i_d \leq m_d - 1, d = 1, 2, 3\}, \quad (4)$$

on the parametric domain $\hat{\Omega}$, also called *parametric mesh*, where $h := \max_{Q \in \mathcal{M}_h} h_Q$ is the global mesh size and h_Q is a characteristic size of element Q . The trivariate B-spline space is defined as

$$\mathcal{S}_{\alpha_1, \alpha_2, \alpha_3}^{p_1, p_2, p_3} \equiv \mathcal{S}_{\alpha_1, \alpha_2, \alpha_3}^{p_1, p_2, p_3}(\mathcal{M}_h) := \text{span} \{B_{i_1, i_2, i_3}^{p_1, p_2, p_3}\}_{i_1, i_2, i_3=1}^{n_1, n_2, n_3}, \quad (5)$$

and its global regularity is $\alpha := \min\{|\alpha_1|, |\alpha_2|, |\alpha_3|\}$.

2.3. Geometric mapping

We assume the physical domain Ω is defined by a geometric mapping $\mathbf{F} : \widehat{\Omega} \rightarrow \Omega$ that is piecewise smooth over $\mathcal{M}_h$ and with piecewise smooth inverse. Within the isogeometric framework, $\mathbf{F}$ is given as a B-splines or NURBS parametrization defined on a coarse mesh $\mathcal{M}_{h_0}$; however, other choices are possible, see section 6.2. We call the *physical mesh* the image of the parametric mesh $\mathcal{M}_h$ by the geometric mapping $\mathbf{F}$.

3. Divergence-conforming B-spline discretization for the Stokes problem

The Stokes equations in strong form are

$$-\nabla \cdot (2\nu \nabla^s \mathbf{u}) + \nabla p = \mathbf{f} \text{ in } \Omega, \quad (6a)$$

$$\nabla \cdot \mathbf{u} = 0 \text{ in } \Omega, \quad (6b)$$

with $\Omega \subset \mathbb{R}^{2,3}$ a bounded simply connected Lipschitz open set, $\mathbf{u}$ is the flow velocity, p is the pressure, $\nu > 0$ the kinematic viscosity and $\mathbf{f}$ denotes a body force acting on the fluid. These equations must be augmented with appropriate boundary conditions to ensure well-posedness. To simplify the presentation we consider here the no-slip case, that is, homogeneous Dirichlet boundary conditions $\mathbf{u} = 0$ in $\partial\Omega$. As usual in the finite element framework, the strong form is recast as a weak formulation stated as: find $\mathbf{u} \in \mathbf{H}_0^1(\Omega)$ and $p \in L_0^2(\Omega)$ such that

$$a(\mathbf{u}, \mathbf{v}) + b(\mathbf{v}, p) + b(\mathbf{u}, q) = (\mathbf{f}, \mathbf{v})_{\mathbf{L}^2(\Omega)}, \quad (7)$$

for all $\mathbf{v} \in \mathbf{H}_0^1(\Omega)$ and $q \in L_0^2(\Omega)$ where

$$a(\mathbf{w}, \mathbf{v}) = (2\nu \nabla^s \mathbf{w}, \nabla^s \mathbf{v})_{\mathbf{L}^2(\Omega)}, \quad (8)$$

$$b(\mathbf{v}, q) = -(\operatorname{div} \mathbf{v}, q)_{L^2(\Omega)}. \quad (9)$$

From Brezzi [25], it is known that a solution $(\mathbf{u}, p) \in \mathbf{H}_0^1(\Omega) \times L_0^2(\Omega)$ of (7) exists given that the following conditions hold: the continuity of the bilinear forms $a(\cdot, \cdot)$ and $b(\cdot, \cdot)$, the coercivity of $a(\cdot, \cdot)$, and the inf – sup condition

$$\inf_{\substack{q \in L_0^2(\Omega) \\ q \neq 0}} \sup_{\substack{\mathbf{v} \in \mathbf{H}_0^1(\Omega) \\ \mathbf{v} \neq 0}} \frac{b(\mathbf{v}, q)}{\|\mathbf{v}\|_{\mathbf{H}^1(\Omega)} \|q\|_{L^2(\Omega)}} \geq \beta > 0, \quad (10)$$

where β is known as the inf – sup constant.

In the following, we review definitions and results for the divergence-conforming B-spline discretization for incompressible flows from [4], which first appeared in [2]. Assuming a global regularity $\alpha \geq 0$, the discrete velocity space on the parametric domain $\widehat{\Omega}$ is defined as

$$\widehat{\mathcal{V}}_h := \mathcal{S}_{\alpha_1+1, \alpha_2, \alpha_3}^{p_1+1, p_2, p_3} \times \mathcal{S}_{\alpha_1, \alpha_2+1, \alpha_3}^{p_1, p_2+1, p_3} \times \mathcal{S}_{\alpha_1, \alpha_2, \alpha_3+1}^{p_1, p_2, p_3+1} \quad (11)$$

and the discrete pressure space on the parametric domain $\widehat{\Omega}$ as

$$\widehat{\mathcal{Q}}_h := \mathcal{S}_{\alpha_1, \alpha_2, \alpha_3}^{p_1, p_2, p_3}. \quad (12)$$

The discrete velocity space with no-penetration boundary conditions is defined as

$$\widehat{\mathcal{V}}_{0,h} := \left\{ \widehat{\mathbf{v}}_h \in \widehat{\mathbf{V}}_h : \widehat{\mathbf{v}}_h \cdot \widehat{\mathbf{n}} = 0 \text{ on } \partial\widehat{\Omega} \right\} \subset \mathbf{H}_0(\widehat{\text{div}}; \widehat{\Omega}), \quad (13)$$

where $\widehat{\mathbf{n}}$ denotes the unit outward normal to $\partial\widehat{\Omega}$ and $\widehat{\text{div}}$ is the divergence operator in parametric coordinates.

With this choice for the velocity space, a constrained discrete pressure space is defined as

$$\widehat{\mathcal{Q}}_{0,h} := \left\{ \widehat{q}_h \in \widehat{\mathcal{Q}}_h : \int_{\widehat{\Omega}} \widehat{q}_h = 0 \right\} \subset L_0^2(\widehat{\Omega}). \quad (14)$$

Such choices of constrained spaces guarantee the surjectivity of the divergence operator at the discrete level, thus ensuring divergence-free velocity fields as well as an inf – sup-stable velocity-pressure pair. Indeed, using the surjectivity of the derivative between univariate B-spline spaces (2), it can be proved that

$$\widehat{\text{div}}(\widehat{\mathcal{V}}_{0,h}) = \widehat{\mathcal{Q}}_{0,h}. \quad (15)$$

As the incompressibility condition is weakly satisfied, that is,

$$(\widehat{\text{div}} \widehat{\mathbf{v}}_h, \widehat{q}_h)_{L^2(\Omega)} = 0 \text{ for all } \widehat{q}_h \in \widehat{\mathcal{Q}}_{0,h}, \quad (16)$$

taking $\widehat{q}_h = \widehat{\text{div}} \widehat{\mathbf{v}}_h$ implies $\|\widehat{\text{div}} \widehat{\mathbf{v}}_h\|_{L^2(\Omega)} = 0$. Since $\widehat{\mathbf{v}}_h$ is continuous, then $\widehat{\text{div}} \widehat{\mathbf{v}}_h = 0$ pointwise. With the growing popularity and successful application of Nitsche's method to impose boundary conditions weakly [26–29], the above velocity-pressure pair choice is not a limitation, as the imposition of no-slip conditions is made weakly by augmenting the variational formulation with additional terms, as we discuss in the next section.

The discrete velocity and pressure spaces on the physical domain Ω are defined by appropriate transformations induced by the geometric mapping $\mathbf{F}$ through the pullback operators

$$\iota_{\mathbf{u}}(\mathbf{v}) = \det(D\mathbf{F})(D\mathbf{F})^{-1}(\mathbf{v} \circ \mathbf{F}), \quad \mathbf{v} \in \mathbf{H}_0(\text{div}; \Omega) \quad (17)$$

$$\iota_p(q) = \det(D\mathbf{F})(q \circ \mathbf{F}), \quad q \in L_0^2(\Omega) \quad (18)$$

where $D\mathbf{F}$ is the Jacobian of the geometrical mapping $\mathbf{F}$. Equation (17) is known as the Piola transform, which is a standard choice to build approximation spaces in $\mathbf{H}(\text{div}; \Omega)$ within the context of mixed finite elements [30]. The Piola transform preserves divergence and the normal component of the transformed vector field. Equation (18), being integral-preserving, guarantees satisfaction of the zero-mean pressure constraint on the physical domain Ω . To preserve the surjectivity of the divergence operator as before, however now in physical coordinates, the discrete velocity and pressure spaces on the physical domain are defined as

$$\mathcal{V}_{0,h} := \left\{ \mathbf{v} \in \mathbf{H}_0(\text{div}; \Omega) : \widehat{\mathbf{v}} = \iota_{\mathbf{u}}(\mathbf{v}) \in \widehat{\mathcal{V}}_{0,h} \right\}, \quad (19)$$

$$\mathcal{Q}_{0,h} := \left\{ q \in L_0^2(\Omega) : \widehat{q} = \iota_p(q) \in \widehat{\mathcal{Q}}_{0,h} \right\}. \quad (20)$$

The last required ingredient within the framework of isogeometric differential forms is the existence of suitable projectors commuting with the divergence operator, see [2]. This commutativity property is key to prove the velocity-pressure pair is inf – sup stable in the discrete sense, see [5].

With the discrete divergence-conforming velocity-pressure spaces pair properly defined, we consider the discrete formulation of (7). Recall that the discrete velocity space $\mathcal{V}_{0,h}$ only satisfies the no-penetration constraint $\mathbf{u} \cdot \mathbf{n} = 0$. To impose no-slip boundary conditions or, more generally, non-homogeneous tangential Dirichlet boundary conditions, we resort to Nitsche's method and add the boundary terms

$$\sum_{F \in \Gamma_h} \int_F -((2\nu \nabla^s \mathbf{u}) \cdot \mathbf{n}) \cdot \mathbf{v} + 2\nu \frac{C_{pen}}{h_F} (\mathbf{u} - \mathbf{g}) \cdot \mathbf{v} - ((2\nu \nabla^s \mathbf{v}) \cdot \mathbf{n}) \cdot (\mathbf{u} - \mathbf{g}) d\Gamma, \quad (21)$$

to (7), where $\mathbf{g}$ is a function defined on $\partial\Omega$ that corresponds to the prescribed tangential component of $\mathbf{u}$ on $\partial\Omega$, h_F is the mesh size in the direction normal to the face F , and $C_{pen} > 0$ is a penalization parameter. The last term in (21) guarantees adjoint consistency, which implies a symmetric linear system. Defining the bilinear form

$$a_h(\mathbf{u}_h, \mathbf{v}_h) = a(\mathbf{u}_h, \mathbf{v}_h) - \eta_h(\mathbf{u}_h, \mathbf{v}_h), \quad (22)$$

where

$$\eta_h(\mathbf{u}_h, \mathbf{v}_h) = \sum_{F \in \Gamma_h} \int_F 2\nu \left(((\nabla^s \mathbf{v}_h) \cdot \mathbf{n}) \cdot \mathbf{u}_h + ((\nabla^s \mathbf{u}_h) \cdot \mathbf{n}) \cdot \mathbf{v}_h - \frac{C_{pen}}{h_F} \mathbf{u}_h \cdot \mathbf{v}_h \right) dS, \quad (23)$$

and the linear form

$$l_h(\mathbf{v}_h) = (\mathbf{f}, \mathbf{v}_h)_{\mathbf{L}^2(\Omega)} + \sum_{F \in \Gamma_h} \int_F 2\nu \left(-((\nabla^s \mathbf{v}_h) \cdot \mathbf{n}) \cdot \mathbf{g} + \frac{C_{pen}}{h_F} \mathbf{g} \cdot \mathbf{v}_h \right) dS \quad (24)$$

the discrete formulation for the Stokes problem is written as: find $\mathbf{u}_h \in \mathcal{V}_{0,h}$ and $p_h \in \mathcal{Q}_{0,h}$ such that

$$a_h(\mathbf{u}_h, \mathbf{v}_h) + b(p_h, \mathbf{v}_h) + b(q_h, \mathbf{u}_h) = l_h(\mathbf{v}_h), \quad (25)$$

for all $\mathbf{v}_h \in \mathcal{V}_{0,h}$ and $q \in \mathcal{Q}_{0,h}$. Results about existence and uniqueness of the discrete solution, as well as the mathematical theory of *a priori* error estimates for the generalized Stokes problem can be found in [5]. Also, all numerical results obtained in this paper used the Nitsche's penalization constant $C_{pen} = 5(\alpha + 1)$ as proposed in [5].

We assume hereafter that discrete velocity and pressure spaces are both defined on the same parametric mesh, with N_i elements in i -direction ($i = 1, 2, 3$), and that $n_{\mathbf{u}} = \dim \widehat{\mathcal{V}}_{h,0} = n_{u_1} + n_{u_2} + n_{u_3}$, where $n_{u_i} = \dim(\widehat{\mathcal{V}}_{h,0})_i$ ($i = 1, 2, 3$) and $n_p = \dim \widehat{\mathcal{Q}}_{h,0}$. We also assume that $p_1 = p_2 = p_3 = p$ and that $\alpha_1 = \alpha_2 = \alpha_3$ are constant sequences with constant component α . We denote $\{\mathbf{N}_1, \dots, \mathbf{N}_{n_{\mathbf{u}}}\} = \{\mathbf{N}_1^1, \dots, \mathbf{N}_{n_{u_1}}^1, \mathbf{N}_1^2, \dots, \mathbf{N}_{n_{u_2}}^2, \mathbf{N}_1^3, \dots, \mathbf{N}_{n_{u_3}}^3\}$ the basis of $\mathcal{V}_{0,h}$, where $\mathbf{N}_{A_i}^i = \iota_{\mathbf{u}}^{-1}(\widehat{N}_{A_i}^i \mathbf{e}_i)$, $i = 1, 2, 3$, $A_i = 1, \dots, n_i$; and $\{N_1, \dots, N_{n_p}\}$ the basis of $\mathcal{Q}_{0,h}$, where $N_A = \iota_p^{-1}(\widehat{N}_A)$, $A = 1, \dots, n_p$. Then (25) reduces to

the solution of the algebraic linear system

$$\begin{bmatrix} \mathbf{A} & \mathbf{B}^T \\ \mathbf{B} & \mathbf{0} \end{bmatrix} \begin{pmatrix} \mathbf{U} \\ \mathbf{P} \end{pmatrix} = \begin{pmatrix} \mathbf{F} \\ \mathbf{0} \end{pmatrix}, \quad (26)$$

where $\mathbf{A} \in \mathbb{R}^{n_u \times n_u}$, $\mathbf{B} \in \mathbb{R}^{n_p \times n_u}$, and $\mathbf{F} \in \mathbb{R}^{n_u}$ are defined as

$$[\mathbf{A}]_{i,j} := a_h(\mathbf{N}_j, \mathbf{N}_i) = a(\mathbf{N}_j, \mathbf{N}_i) - \eta_h(\mathbf{N}_j, \mathbf{N}_i), \quad (27)$$

$$[\mathbf{B}]_{k,j} := b(\mathbf{N}_j, N_k), \quad (28)$$

$$[\mathbf{F}]_i := l_h(\mathbf{N}_i), \quad (29)$$

where $i, j = 1, \dots, n_u$, $k = 1, \dots, n_p$, $\mathbf{U} \in \mathbb{R}^{n_u}$ is the coefficient vector of the discrete velocity $\mathbf{u}_h \in \mathcal{V}_{0,h}$, and $\mathbf{P} \in \mathbb{R}^{n_p}$ is the coefficient vector of the discrete pressure $p_h \in \mathcal{Q}_{0,h}$. Note that $\mathbf{A}$ is symmetric and positive definite. For future reference, we detail the block structure of the coefficient matrix in (26) and write it as

$$\mathcal{A} := \begin{bmatrix} \mathbf{A} & \mathbf{B}^T \\ \mathbf{B} & \mathbf{0} \end{bmatrix} = \begin{bmatrix} A_{u_1 u_1} & A_{u_1 u_2} & A_{u_1 u_3} & B_{u_1 p}^T \\ A_{u_2 u_1} & A_{u_2 u_2} & A_{u_2 u_3} & B_{u_2 p}^T \\ A_{u_3 u_1} & A_{u_3 u_2} & A_{u_3 u_3} & B_{u_3 p}^T \\ B_{p u_1} & B_{p u_2} & B_{p u_3} & 0 \end{bmatrix}. \quad (30)$$

4. Block preconditioners for the Stokes linear system

In this section we briefly recall the most commonly used iterative solution methods for the Stokes linear system obtained from mixed, inf – sup stable finite element discretizations, and the block preconditioning strategies used to increase the robustness of such methods. These block strategies have their rationale motivated by the block $\mathcal{LDU}$ factorization of the coefficient matrix $\mathcal{A}$ (eq. 31).

4.1. Iterative linear solvers

The linear system (26) is symmetric and indefinite. Its indefiniteness follows from Sylvester's law of inertia applied to the block factorization

$$\mathcal{A} = \begin{bmatrix} \mathbf{A} & \mathbf{B}^T \\ \mathbf{B} & \mathbf{0} \end{bmatrix} = \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{B}\mathbf{A}^{-1} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{A} & \mathbf{0} \\ \mathbf{0} & -\mathbf{B}\mathbf{A}^{-1}\mathbf{B}^T \end{bmatrix} \begin{bmatrix} \mathbf{I} & \mathbf{A}^{-1}\mathbf{B}^T \\ \mathbf{0} & \mathbf{I} \end{bmatrix} = \mathcal{LDU}. \quad (31)$$

Indeed, since the block $\mathcal{LDU}$ factorization above defines a congruence we conclude that $\mathcal{A}$ has n_u positive eigenvalues (as $\mathbf{A}$ is positive definite), and has at least $n_p - 1$ negative eigenvalues, since for enclosed flows we have $\text{Ker}(\mathbf{B}^T) \neq \mathbf{0}$ and then the matrix $\mathbf{S} = \mathbf{B}\mathbf{A}^{-1}\mathbf{B}^T$, called the *pressure Schur complement*, is positive semidefinite.

One of the first methods to solve (26) is the (exact) Uzawa method [8, 31], which is a Richardson iteration in the pressure variables $\mathbf{P}$ with the pressure Schur complement $\mathbf{S}$ as the coefficient matrix. Its optimal convergence rate is $\rho := (\kappa - 1)/(\kappa + 1)$, where $\kappa := \kappa(\mathbf{S})$ is the spectral condition number of $\mathbf{S}$. The application of $\mathbf{A}^{-1}$ at each iteration, as required by the exact version of Uzawa method, can be relaxed by the introduction of an inexact variant which employs an approximation to $\mathbf{A}^{-1}$. In the quest of better convergence rates, preconditioned versions of the inexact Uzawa method are also considered in the literature [9]. Ultimately, all such variants can be recast as preconditioned Richardson iterations

$$\begin{bmatrix} \mathbf{M}_A & \mathbf{0} \\ \mathbf{B} & \mathbf{M}_S \end{bmatrix} \begin{pmatrix} \mathbf{U}_{k+1} - \mathbf{U}_k \\ \mathbf{P}_{k+1} - \mathbf{P}_k \end{pmatrix} = \begin{bmatrix} \mathbf{F} \\ \mathbf{0} \end{bmatrix} - \begin{bmatrix} \mathbf{A} & \mathbf{B}^T \\ \mathbf{B} & \mathbf{0} \end{bmatrix} \begin{pmatrix} \mathbf{U}_k \\ \mathbf{P}_k \end{pmatrix}, \quad (32)$$

where $\mathbf{M}_A$ and $\mathbf{M}_S$ are preconditioners for $\mathbf{A}$ and $\mathbf{S}$, respectively.

Richardson iterations are typically not optimal as they do not guarantee the minimization of the linear residual at each iteration. A more effective alternative to stationary iterative methods like Richardson is to use Krylov Subspace Methods, also called Krylov accelerators. For the Stokes linear system, two options are prevalent: the Minimum Residual Method (MINRES) [10, 14], which requires $\mathcal{A}$ to be symmetric, and the Generalized Minimum Residual Method (GMRES) [11, 14]. Given an initial guess $\mathbf{x}_0$, both methods generate a sequence of approximate solutions $\mathbf{x}_k$, $k = 1, 2, \dots$ with

$$\mathbf{x}_k \in \mathbf{x}_0 + \mathcal{K}_k(\mathcal{A}, \mathbf{r}_0), \quad (33)$$

where $\mathbf{r}_0 = \mathbf{b} - \mathcal{A}\mathbf{x}_0$ is the initial residual and $\mathcal{K}_k(\mathcal{A}, \mathbf{r}_0) := \text{span}\{\mathbf{r}_0, \mathcal{A}\mathbf{r}_0, \dots, \mathcal{A}^{k-1}\mathbf{r}_0\}$ is the k th dimensional Krylov subspace generated by $\mathcal{A}$ and $\mathbf{r}_0$. The iterate $\mathbf{x}_k$ in MINRES and GMRES is sought to satisfy the optimality condition

$$\|\mathbf{r}_k\|_2 = \|\mathbf{b} - \mathcal{A}\mathbf{x}_k\|_2 = \min_{\mathbf{x} \in \mathbf{x}_0 + \mathcal{K}_k(\mathcal{A}, \mathbf{r}_0)} \|\mathbf{b} - \mathcal{A}\mathbf{x}\|_2. \quad (34)$$

The main difference between them is the procedure used to construct an orthogonal basis for the Krylov subspace $\mathcal{K}_k(\mathcal{A}, \mathbf{r}_0)$. For MINRES, the symmetric case, Lanczos iterations are sufficient. For GMRES, the general case, Arnoldi iterations are needed. In practice, this implies different memory requirements. MINRES can afford optimality and memory efficiency with a three-term recurrence. However, optimality in GMRES comes at the expense of larger memory requirements since a basis for the Krylov subspace generated during the iterative process needs to be stored.

As known from the literature [15], the convergence rate attained by (34) is influenced by the spectrum of $\mathcal{A}$, denoted by $\sigma(\mathcal{A})$. For matrices coming from discretizations of partial differential equations, $\sigma(\mathcal{A})$ depends on both equation coefficients and discretization parameters in such a way that methods can fail to converge. To increase the robustness of those methods preconditioning strategies must be devised with the goal of clustering the eigenvalues of $\mathcal{A}$.

4.2. Preconditioning strategies

The block $\mathcal{LDU}$ factorization (31) provides important insights on how to build efficient preconditioners for $\mathcal{A}$. In the following, we present optimal choices and the rationale to build practical versions of these. A thorough review on preconditioners for saddle-point linear systems can be found in [15].

4.2.1. Block diagonal variant

A simple choice is to use the diagonal factor $\mathcal{D}$ as preconditioner, that is,

$$\mathcal{M} = \begin{bmatrix} \mathbf{A} & \mathbf{0} \\ \mathbf{0} & \mathbf{S} \end{bmatrix}, \quad (35)$$

It can be proven [14, Chapter 6] that $\sigma(\mathcal{M}^{-1}\mathcal{A}) = \{(1 - \sqrt{5})/2, 1, (1 + \sqrt{5})/2\}$. In this case (34) guarantees convergence to the exact solution in at most three iterations of either MINRES or GMRES. The application of $\mathcal{M}^{-1}$ requires the exact solution of two linear systems, one with $\mathbf{A}$ and another with $\mathbf{S}$. As the pressure Schur complement $\mathbf{S}$ usually is a dense matrix that is expensive to compute, store, and further solve for the linear system, a more practical choice is

$$\mathcal{M}_D = \begin{bmatrix} \mathbf{M}_A & \mathbf{0} \\ \mathbf{0} & \mathbf{M}_S \end{bmatrix}, \quad (36)$$

where $\mathbf{M}_A$ and $\mathbf{M}_S$ are approximations for $\mathbf{A}$ and $\mathbf{S}$, respectively.

For a general inf – sup stable and conforming mixed discretization of the Stokes problem, one has [14, Chapter 5]

$$\beta_0^2 \leq \frac{\langle \mathbf{B}\mathbf{A}^{-1}\mathbf{B}^T\mathbf{p}, \mathbf{p} \rangle}{\langle \mathbf{Q}\mathbf{p}, \mathbf{p} \rangle} \leq C_b^2, \quad \forall \mathbf{p} \in \mathbb{R}^{n_p} \setminus \text{Ker}(\mathbf{B}^T), \quad (37)$$

where $\beta_0 > 0$ is the discrete counterpart of the continuous inf – sup constant in (10) and $C_b > 0$ is the continuity constant of the bilinear form $b(\cdot, \cdot)$, both independent of the mesh size h . The above inequality implies the spectral equivalence of the pressure Schur complement $\mathbf{S}$ and the scaled pressure matrix $\mathbf{Q} := (2\nu)^{-1}\mathbf{M}_p$, where $\mathbf{M}_p$ is the matrix corresponding to the pressure space. Therefore a reasonable choice is $\mathbf{M}_S$ is $\mathbf{Q}$ or, alternatively, any spectrally equivalent approximation to $\mathbf{Q}$. For example, in traditional finite elements it is known that $\mathbf{diag}(\mathbf{Q})$, the main diagonal of $\mathbf{Q}$, and $\mathbf{lump}(\mathbf{Q})$ the row sum lumping of $\mathbf{Q}$, are both spectrally equivalent to $\mathbf{Q}$, see [32].

Several choices are possible for $\mathbf{M}_A$, the most robust being the spectrally equivalent ones, that is, when the bound constants do not depend on the mesh-size parameter h . For the case of the Stokes problem written in vector-Laplacian form with $\mathbf{H}^1$ -conforming basis, two robust options for $\mathbf{M}_A$ are multigrid and domain decomposition methods on the sub-blocks $A_{u_i u_i}$, $i = 1, 2, 3$, since $A_{u_i u_j} = 0$ for $i \neq j$ (eq. (30)), see [14]. However, such an approach does not straightforwardly fit in our setting. We have a symmetric

formulation of the viscous term, and the push-forward of the discrete velocity space to the physical space is done through the Piola transform. As a result, the momentum block $\mathbf{A}$ is not block-diagonal but fully coupled instead. We detail on our choice for $\mathbf{M}_\mathbf{A}$ in section 5.

Finally, the block diagonal preconditioner $\mathcal{M}_D$ is the only sensible choice for MINRES, as it requires a symmetric positive-definite preconditioner. This is no longer the case for the block triangular variant discussed in the following section.

4.2.2. Block triangular variants

The $\mathcal{LDU}$ factorization (31) induces two block triangular preconditioner variants approximating either the lower triangular matrix $\mathcal{LD}$ or the upper triangular matrix $\mathcal{DU}$. These two variants are the lower block triangular preconditioner

$$\mathcal{M}_L = \begin{bmatrix} \mathbf{M}_\mathbf{A} & \mathbf{0} \\ \mathbf{B} & -\mathbf{M}_\mathbf{S} \end{bmatrix}, \quad (38)$$

and the upper block triangular preconditioner

$$\mathcal{M}_U = \begin{bmatrix} \mathbf{M}_\mathbf{A} & \mathbf{B}^T \\ \mathbf{0} & -\mathbf{M}_\mathbf{S} \end{bmatrix}. \quad (39)$$

In these variants, the optimal choice corresponds to $\mathbf{M}_\mathbf{A} = \mathbf{A}$ and $\mathbf{M}_\mathbf{S} = \mathbf{S}$. With such choices, $\sigma(\mathcal{A}\mathcal{M}_U^{-1}) = \{1\}$ and $\sigma(\mathcal{M}_L^{-1}\mathcal{A}) = \{1\}$, and right-preconditioned GMRES with $\mathcal{M}_U$ or left-preconditioned GMRES with $\mathcal{M}_L$ are optimal and converge in at most two iterations. Comparing with the diagonal variant $\mathcal{M}_D$, the additional cost incurred in applying any of the block triangular variants $\mathcal{M}_L$ and $\mathcal{M}_U$ is just an additional matrix-vector product with either $\mathbf{B}$ or $\mathbf{B}^T$. Estimates for the clustering of the spectrum on the complex plane for the linear system (26) preconditioned with the block triangular variants (38) and (39) under several choices for $\mathbf{M}_\mathbf{A}$ and $\mathbf{M}_\mathbf{S}$ can be found in [33, 34].

5. Algorithmic building blocks

In this section, we discuss how a small set of “black-box” algorithms, which are available in a variety of software frameworks, can be combined to yield a robust and scalable preconditioner. First, we show results for the optimal preconditioners $\mathcal{M}_D$ and $\mathcal{M}_U$, when applied to the Krylov subspace methods MINRES and GMRES. Despite the fact that both block preconditioning strategies deliver an almost constant number of iterations with respect to h - and k -refinement in the optimal case, they differ substantially in the total solve time. In other words, these methods are optimal but not practical. As discussed in section 4.2, the definition of a practical preconditioner $\mathcal{M}_U$ involves constructing spectrally equivalent approximations to $\mathbf{A}$ and $\mathbf{S}$. Next, we present a practical preconditioning setup by combining classical solvers and preconditioners, and

method	preconditioner	number of elements				
		64^2	128^2	256^2	512^2	1024^2
MINRES	Optimal $\mathcal{M}_D$	23	23	21	21	21
GMRES	Optimal $\mathcal{M}_D$	23	23	21	21	21
	Optimal $\mathcal{M}_U$	12	12	11	11	11

Table 1: Number of iterations for the optimal preconditioners. As benchmark we used the bi-dimensional lid-driven cavity flow problem with uniform mesh refinement. GMRES is right-preconditioned.

then we discuss in detail how to build $\mathbf{M}_A$ and $\mathbf{M}_S$ in the approximate block factorization. Finally, we report a set of cost measurements (timings) for the two most time-consuming kernels of our strategy, namely, the matrix-vector product of the outer iteration and the multigrid sweep.

5.1. Optimal preconditioner

To indirectly measure the performance of our block choices in terms of spectrum clustering, we computed the number of iterations for the Krylov subspace methods MINRES and GMRES preconditioned by optimal preconditioners, here denoted by Optimal $\mathcal{M}_D$ and Optimal $\mathcal{M}_U$. That is, taking $\mathbf{M}_A = \mathbf{A}$ and $\mathbf{M}_S = \mathbf{Q}$ in both cases and applying them by using a direct solver (LU factorization), in our case MUMPS [35, 36]. We used an absolute residual tolerance of 10^{-8} as stopping criteria. The results are shown in Table 1. For each method we only show one row, since the number of iterations are almost the same (differing by no more than two iterations) for the different discretization regularities $p = 1, \alpha = 0$; $p = 2, \alpha = 1$ and $p = 3, \alpha = 2$. Hereafter, discretization regularity refers to the regularity of the discrete pressure space on the parametric domain $\hat{\mathcal{Q}}_h = \mathcal{S}_{\alpha, \alpha}^{p, p}$ for space dimension two and $\hat{\mathcal{Q}}_h = \mathcal{S}_{\alpha, \alpha, \alpha}^{p, p, p}$ for space dimension three. We limited our numerical investigations to those regularities.

From Table 1 we can verify that both the upper block triangular variant $\mathcal{M}_U$ and the block diagonal variant $\mathcal{M}_D$ are h -scalable in terms of the number of iterations. Also, this table indicates the known trend that the number of iterations using a block triangular preconditioner is approximately half of that using the block diagonal one. Despite the almost constant number of iterations with mesh refinement (h -refinement) and k -refinement, the same is not true for the total solve time, which increases with a rate dependent on p . The reason for such an increase is directly related to the use of a multi-frontal solver to perform the LU factorizations of $\mathbf{A}$ and $\mathbf{Q}$. The p -dependence on the cost of a multi-frontal algorithm when applied to the solution of a linear system with coefficient matrix $\mathbf{Q}$ was investigated by Collier et al. [37]. We verified that the same p -dependent time complexity of $O(p^3 N^{3/2})$ [37] holds for the LU factorization of matrix $\mathbf{A}$,

obtained using spaces with mixed regularities. The authors conducted an investigation closely related to this one [17], using specifically MINRES and applying different approximations for $\mathbf{A}$ and $\mathbf{Q}$, to compose the block diagonal preconditioner variant (36).

5.2. Practical preconditioner

Instead of using direct solvers, a natural choice for a practical preconditioner is to use iterative methods to approximate the application of $\mathbf{A}^{-1}$ and $\mathbf{Q}^{-1}$. The problem of such methods is the lack of robustness; however, that can be somehow remedied using a preconditioned version of the iterative method of choice, which depends on the matrix at hand. In particular, for the momentum block $\mathbf{A}$, it is known that its spectrum depends on the mesh size being used, such that the first “black-box” choice for $\mathbf{M}_{\mathbf{A}}^{-1}$ is to apply a few cycles of an algebraic multigrid method (AMG) to solve for $\mathbf{A}$. Initially developed for scalar elliptic problems, AMG methods try to achieve the optimality of a geometric multigrid method using as unique information the coefficient matrix given as input. In this case, a coarsening algorithm works directly with this matrix, building a family of coarser *grids* and a set of transfer operators (restriction and interpolation) between each consecutive pair of grids [38].

“Black-box” implementations of AMG for systems (e.g. elasticity) assume standard equal-order discretizations with interpolatory basis functions and degrees of freedoms numbered consecutively within each grid node. This way, in a pre-process step, AMG algorithms for systems are able to account for the coupling between fields by blocking the degrees of freedom, building a connectivity graph and promoting the coarsening on it [39]. However, the sparsity pattern of B-spline divergence-conforming isogeometric discrete operators is akin to the ones obtained with staggered finite differences using different stencil widths in each parametric direction. The AMG implementations we explored cannot account for this very special structure of the coupled system, thus we had to resort to a simpler approach exploiting the block structure of $\mathbf{A}$ and using “black-box” AMG methods for scalar problems.

We build $\mathbf{M}_{\mathbf{A}}$, more precisely the application of $\mathbf{M}_{\mathbf{A}}^{-1}$, using a preconditioned Krylov subspace method applied to $\mathbf{A}$ (see eq. (30)). For the Krylov subspace method, we choose the Flexible Generalized Minimum Residual Method (FGMRES,[40]): a variant of the Generalized Minimum Residual Method (GMRES,[11]) allowing a variable preconditioner at each iteration. We fixed the number of iterations of this inner FGMRES to five. As a preconditioner for this inner FGMRES we used an Inexact Block Gauss–Seidel (BGS) approach applied to $\mathbf{A}$ combined with AMG sweeps. That is, at each BGS sweep the inverses of the diagonal blocks are approximated by an AMG method, that is, $A_{u_i u_i}^{-1} \sim \text{AMG}(A_{u_i u_i})$, for $i = 1, 2, 3$ (see eq. (30)). We use such definition of $\mathbf{M}_{\mathbf{A}}^{-1}$ for both sequential and parallel runs. Additionally, since we build on PETSc (version 3.6), we were able to try different algebraic multigrid implementations, namely, the ML library (version 6.2) [41] and the BoomerAMG library (*hypra*, version 2.10.0b) [42]. We experimented with tuning of AMG parameters. Even though we do not achieve perfect weak scaling, the AMG parameters we chose

(and kept fixed for all runs) are the ones that gave the smallest time to solution for the range of problem sizes and number of processes that we used.

Our experience [17, 43] suggests building the approximation $\mathbf{M}_S$ using an incomplete LU factorization of $\mathbf{Q}$ with no fill-in, hereafter denoted $\text{ILU}(\mathbf{Q})$. For example, in [17] the block diagonal preconditioning variant of (36) delivered h - and k -independence in the number of iterations. In that work, we defined the action of $\mathbf{M}_S^{-1}$ using a preconditioned conjugate gradient method (PCG) applied to $\mathbf{Q}$ using an incomplete Cholesky factorization with no fill-in as a preconditioner. In what follows, we use $\mathbf{M}_S^{-1} \sim \text{ILU}(\mathbf{Q})$ for the serial runs. For the parallel runs, we used five preconditioned conjugate gradients iterations applied to $\mathbf{Q}$, preconditioned by an inexact Block Jacobi strategy, where submatrices of $\mathbf{Q}$ corresponding to the sets of unknowns local to each process are approximately inverted using an incomplete LU factorization with no fill-in, i.e., $\mathbf{M}_S^{-1} \sim (\text{PCG}(5) + \text{Block-Jacobi} + \text{ILU})(\mathbf{Q})$.

To clarify the nesting of iterative solvers and preconditioners that our strategy promotes we summarize it as follows: an outer flexible GMRES loop right preconditioned with $\mathcal{M}_U$, where the application of $\mathbf{M}_A^{-1}$ is realized through five iterations of flexible GMRES right preconditioned by an Inexact Block Gauss-Seidel sweep with the inverse of the diagonal blocks approximated by $\text{AMG}(A_{u_i u_i})$ for $i = 1, 2, 3$, denoted $\mathbf{M}_A^{-1} \sim (\text{FGMRES}(5) + \text{BGS} + \text{AMG})(\mathbf{A})$. Finally, for $\mathbf{M}_S$ we have two possibilities: $\mathbf{M}_S^{-1} \sim \text{ILU}(\mathbf{Q})$ in serial, or $\mathbf{M}_S^{-1} \sim (\text{PCG}(5) + \text{Block-Jacobi} + \text{ILU})(\mathbf{Q})$ in parallel. Inspecting the methods used in our strategy the most intensive kernels used are the matrix-vector product, mainly with the matrix $\mathcal{A}$ (see eq. (30)) and the multigrid sweeps. Before we discuss the performance of our strategy applied to the benchmarks problems, in the next two subsections we report the behavior of both kernels independently.

5.2.1. Matrix-vector product

To investigate the cost of the matrix-vector product with the matrix $\mathcal{A}$ from (30), resulting from the discretization of a two-dimensional problem with increasing regularity in the discrete spaces (k -refinement), we fix the number of elements to 512^2 . Table 2 and Figure 1 summarize the results. Despite the mild increase in the number of rows of $\mathcal{A}$ with k -refinement, the number of non-zero entries in the matrix grows at a faster pace due to the increasing support of basis functions. Figure 1a shows the increase in the mean number of non-zero entries per row. The wall-clock time required to perform a matrix-vector product matches closely the increase of non-zero entries in the matrix, as shown in Figure 1b.

For a comprehensive analysis of the cost of matrix-vector products and other computational kernels common in iterative methods see [43]. Our results are in close agreement with the theoretical estimates and numerical evidence presented in that work.

regularity	rows	nnz	time
$p = 1, \alpha = 0$	790,533	28,890,136	0.037
$p = 2, \alpha = 1$	793,616	68,941,040	0.087
$p = 3, \alpha = 2$	796,705	126,019,568	0.158
$p = 4, \alpha = 3$	799,800	200,224,576	0.251

Table 2: Effect of k -refinement on the cost of matrix-vector products for the matrix $\mathcal{A}$ in (30). The number of matrix rows increases mildly with k -refinement. The total number of non-zero entries (nnz) grows at a faster pace, thus increasing the time cost of matrix-vector products. Wall-clock time correspond to the mean of 50 runs and is measured in seconds.

(a) Mean number of non-zero entries per row.

(b) Time cost of matrix-vector products.

Figure 1: Effect of k -refinement on matrix sparsity. The mean number of non-zero entries per row increases with k -refinement. The time cost of matrix-vector products increases proportionally to the mean number of non-zero entries per row.

Figure 2: Effect of k -refinement on the cost of multigrid sweeps for the matrix block $A_{u_1 u_1}$ in (30).

5.2.2. Multigrid sweeps

For each outer FGMRES iteration, we apply at least $5d$ ($d = 2, 3$ the space dimension) multigrid sweeps, corresponding to the approximate solution of $A_{u_i u_i}$, that is, $A_{u_i u_i}^{-1} \sim \text{AMG}(A_{u_i u_i})$. Thus, we seek to quantify appropriately the cost of this important building block of our strategy. Here we report the time to apply one multigrid cycle for the first matrix block $A_{u_1 u_1}$. The time reported does not consider the setup time, that is, the time required to build all the algebraic multigrid infra-structure that includes: selection of a hierarchic of coarse grids (coarsening algorithm), building the transfer operators (restriction and interpolation) between them, and the construction of coarse-level matrices. A multigrid cycle involves sparse matrix-vector products which accounts for the transfer operators, and a pre- and a post-smoother at each level of the hierarchy. The actual cost to perform a multigrid cycle is dominated in many cases by the cost of the pre- and post-smoother on the finest grid.

As before we investigate how the cost increases with k -refinement, to this end we fixed the number of elements to 512^2 . Figure 2 shows the results of applying a Full-Multigrid cycle (FMG) to the matrix $A_{u_1 u_1}$, where the maximum number of grid levels is fixed at five, the Symmetric Successive Overrelaxation method (SSOR) method is used as a pre- and post-smoother in every level and the coarse-grid solver is the LU factorization method. As expected the time is proportional to the mean number of non-zero entries per row, as one can verify from Figures 1 and 2. Also, Figures 1 and 2 follow a similar qualitative trend.

6. Performance results

In this section, we report the performance of our upper block triangular preconditioning strategy applied to the flexible variant of the GMRES algorithm without restart to solve the Stokes linear system (26). The strategy was implemented in PetIGA-MF [18], a toolbox built on top of PetIGA [19] extending it to support flexible and parallel multi-field discretizations based on structure-preserving B-spline spaces, in the sense of [2]. For examples of the Navier-Stokes-Cahn-Hilliard model using PetIGA-MF see [44, 45].

As benchmark problems, we use the flows on a square and a cubic cavity, the flow in an eccentric annulus, and the flow in a hollow torus of eccentric annular cross-section. The last two geometries entail the use of nonlinear geometric mappings and considerably graded physical meshes. After defining each problem, we discuss the specific set of results for it. The main performance results we report are wall-clock time and number of iterations for different levels of mesh refinements and space regularities. All performance results from the two-dimensional benchmarks were obtained with sequential runs, whereas the three-dimensional ones were obtained with parallel runs. All benchmarks ran on Stampede [46], a distributed memory supercomputing system featuring compute nodes with 2×8 -core Intel Xeon E5-2680 (Sandy Bridge) processors, 32 GB of memory per node, and an InfiniBand Mellanox interconnect.

6.1. Lid-driven cavity problem

The lid-driven cavity flow is a standard benchmark problem to test new discretization schemes and preconditioning strategies for saddle-point problem defined by the Stokes system. This benchmark consists of solving a flow problem on a square or cubic domain with homogeneous boundary conditions for the velocity at the bottom and lateral boundaries of the domain, and a constant tangential velocity at the top (lid) boundary. With respect to the geometric mapping $\mathbf{F} : \hat{\Omega} \rightarrow \Omega = [0, 1]^d$ ($d = 2, 3$), we consider here the identity mapping, which is equivalent to working directly on the parametric domain.

We consider uniform h -refinements in the parametric domain, using the same number of elements in each direction. For the square cavity case ($d = 2$), all results were obtained with sequential runs on a single compute node. The number of elements, abbreviated hereafter as nel , are chosen in such a way that the number of degrees of freedom, abbreviated hereafter as $ndof$, approximately double between successive refinements (see Table 3). We run the cubic cavity tests in parallel using multiple compute nodes. To perform a weak scalability analysis, we take as *unit* problem size the largest problem size that fit well on a single node, in our case a problem with 32^3 elements. Then we start our scalability analysis from a problem with 96^3 elements running on 3^3 processes, increasing the number of elements in each direction by 32 and the number of processes accordingly, up to 224^3 elements running on 7^3 processes. In both cases we used the ML library [41] as the algebraic multigrid implementation, with Full-Multigrid cycle (FMG) limiting the number of levels to five, and with two Symmetric Successive Overrelaxation (SSOR) sweeps as pre- and post-smoother and the LU factorization as the coarse grid solver.

Table 3 and Figure 3 present the results for the square cavity. The table shows that the preconditioning strategy provides the spectral equivalence property, reflected by the h -scalability. That is, variations of the mesh sizes maintain the number of iterations almost constant. The table also suggests the k -scalability of the strategy. Finally, from Figure 3, we see that the measured solve time per degree of freedom is approximately constant (asymptotically), which reflects again the scalability property of the strategy. The last mesh size

for the $p = 3, \alpha = 2$ case is off-trend because of memory access effects.¹

¹We ran these tests in a large-memory, two-socket compute node of Stampede [46]. At the largest problem size, the memory requirements were high enough as to force the kernel to allocate physical memory pages out of the remote memory bank (NUMA domain). Because of the architecture of memory subsystem, memory accesses from the remote memory bank have a 20 ~ 25% drop in bandwidth.

regularity	nel	ndof	its	time(s)
$p = 1, \alpha = 0$	64^2	12,805	12	0.27
	90^2	25,025	12	0.37
	128^2	50,181	12	0.68
	180^2	98,645	12	1.28
	256^2	198,661	11	2.40
	360^2	391,685	11	4.94
	512^2	790,533	11	10.74
	720^2	1,560,965	11	20.95
	1024^2	3,153,925	11	43.91
$p = 2, \alpha = 1$	64^2	13,200	12	0.33
	90^2	25,576	12	0.63
	128^2	50,960	12	1.29
	180^2	99,736	12	2.40
	256^2	200,208	11	4.80
	360^2	393,856	11	9.84
	512^2	793,616	11	20.44
	720^2	1,565,296	11	40.92
	1024^2	3,160,080	11	83.31
$p = 3, \alpha = 2$	64^2	13,601	12	0.56
	90^2	26,133	12	1.04
	128^2	51,745	12	2.04
	180^2	100,833	12	4.30
	256^2	201,761	12	9.15
	360^2	396,033	12	18.57
	512^2	796,705	12	38.29
	720^2	1,569,633	12	75.99
	1024^2	3,166,241	12	164.08

Table 3: Number of degrees of freedom ($ndof$), number of iterations (its), and total solve time for the square lid-driven cavity problem using uniform mesh refinements.

Figure 3: Solve time per number of degrees of freedom for the square lid-driven cavity problem using uniform mesh refinement.

Figure 4 shows the results for the cubic cavity benchmark for the spaces with regularities $p = 1, \alpha = 0$ and $p = 2, \alpha = 1$. The wall-clock timings are split in setup, solver and total time (i.e., the sum of the previous two). Solve time is nearly constant. However, setup time is dominated by the time spent on building the algebraic multigrid infrastructure and increases with the number of processes.

6.2. Stokes flow inside an eccentric annulus

Another interesting example we use to test the performance of our preconditioner is the steady Stokes flow in an eccentric annulus. Such domain entails the use of a nonlinear geometric mapping, namely the bipolar parametrization. Although defined on a non-trivial geometry this flow can be solved analytically using bipolar coordinates and the streamline formulation for the Stokes equations. Here we just describe the relevant aspects regarding the example definition; more details about the analytical solution can be found in [47].

The eccentric annulus Ω is described by the radii R_1 and R_2 , assuming $R_1 > R_2$, and the eccentricity ε . In this case, the outer circle is centered in $(l, 0)$ and the inner circle is centered at $(l - \varepsilon, 0)$, where l is computed by $l = \frac{R_1^2 - R_2^2 + \varepsilon^2}{2\varepsilon}$ (see Figure 5a). The bipolar geometric mapping $\mathbf{F} : \hat{\Omega} \rightarrow \Omega$ is defined as

$$(x, y) = \mathbf{F}(\tau, \sigma) = \left(\frac{-a \sinh(\tau)}{\cosh(\tau) - \cos(\sigma)}, \frac{a \sin(\sigma)}{\cosh(\tau) - \cos(\sigma)} \right), \quad (40)$$

where $a = \frac{1}{2\varepsilon}[(R_1^2 + R_2^2 - \varepsilon^2)^2 - 4R_1^2 R_2^2]^{1/2}$, with the parametric space being $\hat{\Omega} = [\tau_2, \tau_1] \times [0, 2\pi)$, with $\tau_2 = -\arcsin(a/R_2)$ and $\tau_1 = -\arcsin(a/R_1)$. We could have also used a NURBS geometric mapping to describe Ω , but since we want to compare against the analytical solution we opted to use the mapping (40).

Figure 4: Total, setup and solve times for the cubic lid-driven cavity problem.

Figure 5: Eccentric annulus and the mesh obtained by uniform refinement of the parametric space $\hat{\Omega}$ mapped by the bipolar mapping.

The radii R_1 and R_2 , the eccentricity ε , and the angular velocities ω_1 and ω_2 completely describe the qualitative behavior of the flow. In the numerical results that follow, we assume $R_2/R_1 = 0.3$ and $\varepsilon = 0.75(R_1 - R_2)$ to fix the geometric parameters, and consider the ratio $\bar{\omega} = \omega_2/\omega_1$ as a varying parameter. Assuming that $\bar{\omega} > 0$ the analytical investigation in [47] showed that there is a threshold value, $\bar{\omega} = 1.2214$, for the existence of a pair of stagnation points symmetrically located with respect to the x -axis. Figure 6 illustrates the streamlines for three different values of $\bar{\omega}$, namely, for $\bar{\omega} = 1.22$, $\bar{\omega} = 5$ and $\bar{\omega} = 20$.

In the parametric mesh, we used uniform refinements. Due to the use of the bipolar mapping, the resulting physical mesh is curved and graded. An example of such a mesh with 32^2 elements is shown at Figure 5b. A convergence test for the case $\bar{\omega} = 20$ is shown in Figure 7. As expected the convergence order $O(h^{p+1})$ for the L^2 -norm of the velocity is attained.

The performance of our strategy applied to the case with $\bar{\omega} = 20$ is shown in Table 4 and Figure 8. For the AMG implementation we use the BoomerAMG library [42]. BoomerAMG offers to the user several options to customize the algebraic multigrid method. We use most default options, that are: a V-cycle, one sweep of the Symmetric Successive Overrelaxation method (SSOR) method as pre- and post-smoother in every level and LU factorization method as coarse-grid solver. We only customize the parameter known as *strength threshold* related to the coarsening algorithm (see eq. (2) in [42]) and chose a value of 0.75.

6.3. Stokes flow inside an eccentric torus

By rotating the eccentric annulus around a symmetry axis, we obtain the hollow torus of an eccentric annular cross-section. Hereafter, we refer to such domain as eccentric torus. Instead of using an analytical geometric mapping, known as toroidal parametrization, we describe the geometry using a single NURBS patch. The parameters used to define the eccentric torus are the same as for the eccentric annulus, that is, the radii R_1 and R_2 with $R_1 > R_2$, and the eccentricity ε (see Figure 9a).

Boundary conditions are specified through the angular velocities ω_1 and ω_2 of the inner and outer circles of the annular cross-section, respectively; and the angular velocities φ_1 and φ_2 of the inner and outer torus surfaces, respectively (see Figure 9a). We use uniform refinements in the parametric mesh. An example of the resulting physical mesh is shown in Figure 9b. In what follows, we assume $R_2/R_1 = 0.3$ and $\varepsilon = 0.75(R_1 - R_2)$ to fix the geometric parameters, and also fix the ratios $\bar{\omega} = \omega_2/\omega_1 = 6$ and $\bar{\varphi} = \varphi_2/\varphi_1 = 3$. An example of the flow pattern generated using these parameters is shown in Figure 9c.

As our backend algebraic multigrid solver, we use the BoomerAMG [42] implementation with the same settings as for the eccentric annulus: one V-cycle with one SSOR sweep as pre- and post-smoother in every level and LU factorization as the coarse-grid solver. As before, we chose a *strength threshold* coarsening parameter of 0.75. As in the cubic cavity case, the strategy is h -scalable with respect to iteration count and solve time.

Figure 10 shows weak scalability results of our preconditioning strategy for spaces with regularity $p =$

Figure 6: Streamline for different cases varying $\bar{\omega} > 0$.

regularity	nel	ndof	its	time (s)
$p = 1, \alpha = 0$	64^2	12,352	14	0.23
	90^2	24,390	14	0.40
	128^2	49,280	14	0.87
	180^2	97,380	14	1.75
	256^2	196,864	14	3.60
	360^2	389,160	12	5.56
	512^2	786,944	14	16.39
	720^2	1,555,920	15	35.54
$p = 2, \alpha = 1$	64^2	12,544	15	0.39
	90^2	24,660	14	0.80
	128^2	49,664	14	1.64
	180^2	97,920	14	3.35
	256^2	197,632	14	7.35
	360^2	390,240	14	15.79
	512^2	788,480	13	31.86
	720^2	1,558,080	13	65.27
$p = 3, \alpha = 2$	64^2	12,736	17	0.78
	90^2	24,930	16	1.50
	128^2	50,048	16	3.14
	180^2	98,460	16	6.72
	256^2	198,400	16	15.00
	360^2	391,320	16	31.93
	512^2	790,016	16	67.86
	720^2	1,560,240	16	137.82

Table 4: Number of degrees of freedom ($ndof$), number of iterations (its), and total solve time for the eccentric annulus example.

Figure 7: Convergence test for the eccentric annulus.

1, $\alpha = 0$. As in the cubic cavity case, solve time is nearly constant and setup time increases with the number of processes.

Unexpectedly, we were not able to reproduce the previous trends for the $p = 2, \alpha = 1$ regularity case. After investigating the whole stack of nested solvers and how they behave, an eigenvalue analysis of the multigrid-preconditioned operators showed a loss of positive-definiteness. Consequently, a single V-cycle on the diagonal blocks $A_{u_i u_i}$ ($i = 1, \dots, 3$) is not able to sufficiently reduce the residual and degrades considerably the convergence rate of the outer block Gauss-Seidel iteration on the momentum equations. Indeed, by changing the single V-cycle for a preconditioned GMRES method with a prescribed tolerance of 10^{-2} for the reduction of the initial residual norm, we could get a nearly constant number of outer iterations for all levels of refinement, but not h -independence in solve time.

7. Conclusions

Divergence-conforming B-spline discretization is a new compatible scheme developed under the framework of isogeometric analysis. Suitable to approximate viscous flow problems, its inception furnishes a discrete stable velocity-pressure pair that at the same time satisfies the incompressibility condition pointwise. When applied to the Stokes flow problem we are faced with the challenge of solving a symmetric saddle-point linear system. Using a preconditioned Krylov method in conjunction with a block triangular preconditioning strategy built from *black-box* preconditioners, namely algebraic multigrid and incomplete factorization, we built a scalable preconditioner for the Stokes saddle-point linear system. The numeri-

Figure 8: Solve time per number of degrees of freedom for the eccentric annulus flow problem under refinement.

cal results shown the good performance in serial and in parallel of the preconditioner in terms of h - and k -refinement.

For the numerical experiments we used PetIGA-MF, an extension of PetIGA [19], built on top of PETSc [20]. This choice allowed us to experiment with several options of solvers and preconditioners, making our investigation possible, which resulted in an efficient black-box solver that is opening the door for us to investigate several application problems such as Navier-Stokes-Cahn-Hilliard, Magneto-Hydrodynamics, etc.

Figure 9: Geometry, example mesh and Stokes flow pattern for the eccentric torus.

Figure 10: Setup and solve time per number of degrees of freedom for the eccentric torus flow problem under refinement for the regularity $p = 1, \alpha = 0$.

Acknowledgements

This publication was made possible in part by a National Priorities Research Program grant 7-1482-1-278 from the Qatar National Research Fund (a member of The Qatar Foundation), by the European Unions Horizon 2020 Research and Innovation Program under the Marie Skłodowska-Curie grant agreement No 644602 and the Center for Numerical Porous Media at King Abdullah University of Science and Technology (KAUST). The J. Tinsley Oden Faculty Fellowship Research Program at the Institute for Computational Engineering and Sciences (ICES) of the University of Texas at Austin has partially supported the visits of VMC to ICES. The authors acknowledge the Texas Advanced Computing Center (TACC) at The University of Texas at Austin for providing HPC resources that have contributed to the research results reported within this paper.

References

- [1] A. Buffa, G. Sangalli, R. Vázquez, Isogeometric analysis in electromagnetics: B-splines approximation, *Computer Methods in Applied Mechanics and Engineering* 199 (17–20) (2010) 1143–1152. doi:10.1016/j.cma.2009.12.002.
- [2] A. Buffa, J. Rivas, G. Sangalli, R. Vázquez, Isogeometric discrete differential forms in three dimensions, *SIAM Journal Numerical Analysis* 49 (2) (2011) 818–844. doi:10.1137/100786708.
- [3] D.N. Arnold, R.S. Falk, R. Winther, Finite element exterior calculus: from Hodge theory to numerical stability, *Bulletin American Mathematical Society* 47 (2010) 281–354. doi:10.1090/S0273-0979-10-01278-4.
- [4] J.A. Evans, Divergence-free B-spline discretizations for viscous flows, Ph.D. thesis, ICES - UT Texas (2011).
- [5] J.A. Evans, T.J.R. Hughes, Isogeometric divergence-conforming B-splines for the Darcy-Stokes-Brinkman equations, *Mathematical Models and Methods in Applied Sciences* 23 (04) (2013) 671–741. doi:10.1142/S0218202512500583.
- [6] J.A. Evans, T.J.R. Hughes, Isogeometric divergence-conforming B-splines for the steady Navier–Stokes equations, *Mathematical Models and Methods in Applied Sciences* 23 (08) (2013) 1421–1478. doi:10.1142/S0218202513500139.
- [7] J.A. Evans, T.J.R. Hughes, Isogeometric divergence-conforming B-splines for the unsteady Navier–Stokes equations, *Journal of Computational Physics* 241 (0) (2013) 141–167. doi:10.1016/j.jcp.2013.01.006.
- [8] K.J. Arrow, L. Hurwicz, H. Uzawa, *Studies in Linear and Non-linear Programming*, Stanford mathematical studies in the social sciences, Stanford University Press, 1964.
- [9] H.C. Elman, G.H. Golub, Inexact and preconditioned Uzawa algorithms for saddle point problems, *SIAM Journal of Numerical Analysis* 31 (6) (1994) 1645–1661. doi:10.1137/0731085.
- [10] C.C. Paige, M.A. Saunders, Solution of sparse indefinite systems of linear equations, *SIAM Journal on Numerical Analysis* 12 (4) (1975) 617–629. doi:10.1137/0712047.
- [11] Y. Saad, M.H. Schultz, GMRES: a generalized minimal residual algorithm for solving nonsymmetric linear systems, *SIAM Journal Scientific Statistic Computing* 7 (3) (1986) 856–869. doi:10.1137/0907058.
- [12] O. Axelsson, *Iterative Solution Methods*, Cambridge University Press, New York, NY, USA, 1994.
- [13] H.C. Elman, Multigrid and Krylov subspace methods for the discrete Stokes equations, *International Journal for Numerical Methods in Engineering* 22 (8) (1996) 755–770. doi:10.1002/(SICI)1097-0363(19960430)22:8<755::AID-FLD377>3.0.CO;2-1.
- [14] H.C. Elman, D.J. Silvester, A.J. Wathen, *Finite Elements And Fast Iterative Solvers: With Applications in Incompressible Fluid Dynamics*, Numerical Mathematics And Scientific Computation, Oxford University Press, 2005. doi:10.1093/acprof:oso/9780199678792.001.0001.
- [15] M. Benzi, G.H. Golub, J. Liesen, Numerical solution of saddle point problems, *Acta Numerica* 14 (2005) 1–137. doi:10.1017/S0962492904000212.
- [16] J. Rudi, A. C. I. Malossi, T. Isaac, G. Stadler, M. Gurnis, P. W. J. Staar, Y. Ineichen, C. Bekas, A. Curioni, O. Ghattas, An extreme-scale implicit solver for complex pdes: Highly heterogeneous flow in earth’s mantle, in: *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, SC ’15*, ACM, New York, NY, USA, 2015, pp. 5:1–5:12. doi:10.1145/2807591.2807675.
URL <http://doi.acm.org/10.1145/2807591.2807675>
- [17] A.M.A. Cortes, A.L.G.A. Coutinho, L. Dalcin, V.M. Calo, Performance evaluation of block-diagonal preconditioners for the divergence-conforming B-spline discretization of the Stokes system, *Journal of Computational Science* (2015) –doi:10.1016/j.jocs.2015.01.005.
- [18] A. Sarmiento, A.M.A. Cortes, D. Garcia, L. Dalcin, N. Collier, V.M. Calo, PetIGA-MF: a multi-field high-performance toolbox of divergence-conforming B-splines, *Journal of Computational Science* (2016) –doi:10.1016/j.jocs.2016.09.010.
- [19] L. Dalcin, N. Collier, P. Vignal, A.M.A. Cortes, V.M. Calo, PetIGA: A framework for high-performance isogeometric

- analysis, *Computer Methods in Applied Mechanics and Engineering* 308 (2016) 151–181. doi:<http://dx.doi.org/10.1016/j.cma.2016.05.011>.
- [20] S. Balay, S. Abhyankar, M.F. Adams, J. Brown, P. Brune, K. Buschelman, V. Eijkhout, W.D. Gropp, D. Kaushik, M.G. Knepley, L.C. McInnes, K. Rupp, B.F. Smith, Hong Zhang, PETSc Web page [online] (2014). URL <http://www.mcs.anl.gov/petsc>
- [21] A. Buffa, C. de Falco, G. Sangalli, Isogeometric analysis: stable elements for the 2D Stokes equation, *International Journal for Numerical Methods in Fluids* 65 (11-12) (2011) 1407–1422. doi:[10.1002/flid.2337](https://doi.org/10.1002/flid.2337).
- [22] Les A. Piegl, W. Tiller, *The NURBS book*, 2nd Edition, Springer, 1996.
- [23] J.A. Cottrell, T.J.R. Hughes, Y. Bazilevs, *Isogeometric Analysis: Toward Integration of CAD and FEA*, Wiley, 2009. doi:[10.1002/9780470749081](https://doi.org/10.1002/9780470749081).
- [24] T.J.R. Hughes, J.A. Cottrell, Y. Bazilevs, Isogeometric analysis: CAD, finite elements, NURBS, exact geometry, and mesh refinement, *cmame* 194 (2005) 4135–4195. doi:[10.1016/j.cma.2004.10.008](https://doi.org/10.1016/j.cma.2004.10.008).
- [25] F. Brezzi, On the existence, uniqueness and approximation of saddle-point problems arising from Lagrangian multipliers, *ESAIM: Mathematical Modelling and Numerical Analysis - Modlisation Mathematique et Analyse Numrique* 8 (R2) (1974) 129–151.
- [26] Y. Bazilevs, C. Michler, V.M. Calo, T.J.R. Hughes, Weak Dirichlet boundary conditions for wall-bounded turbulent flows, *Computer Methods in Applied Mechanics and Engineering* 196 (49-52) (2007) 4853–4862. doi:[10.1016/j.cma.2007.06.026](https://doi.org/10.1016/j.cma.2007.06.026).
- [27] Y. Bazilevs, T.J.R. Hughes, Weak imposition of Dirichlet boundary conditions in fluid mechanics, *Computers & Fluids* 36 (1) (2007) 12–26. doi:[10.1016/j.compfluid.2005.07.012](https://doi.org/10.1016/j.compfluid.2005.07.012).
- [28] Y. Bazilevs, C. Michler, V.M. Calo, T.J.R. Hughes, Isogeometric variational multiscale modeling of wall-bounded turbulent flows with weakly enforced boundary conditions on unstretched meshes, *Computer Methods in Applied Mechanics and Engineering* 199 (13-16) (2010) 780–790. doi:[10.1016/j.cma.2008.11.020](https://doi.org/10.1016/j.cma.2008.11.020).
- [29] A. Embar, J. Dolbow, I. Harari, Imposing Dirichlet boundary conditions with Nitsche’s method and spline-based finite elements, *International Journal for Numerical Methods in Engineering* 83 (7) (2010) 877–898. doi:[10.1002/nme.2863](https://doi.org/10.1002/nme.2863).
- [30] Daniele Boffi, Michel Fortin, Franco Brezzi, *Mixed finite element methods and applications*, Springer series in computational mathematics, Springer, Berlin, Heidelberg, 2013.
- [31] R. Glowinski, M. Fortin, *Augmented Lagrangian Methods: Applications to the Numerical Solution of Boundary-Value Problems*, Studies in Mathematics and its Applications, Elsevier Science, 2000.
- [32] A. Wathen, D. Silvester, Fast iterative solution of stabilised Stokes systems part I: using simple diagonal preconditioners, *SIAM Journal on Numerical Analysis* 30 (3) (1993) 630–649. doi:[10.1137/0730031](https://doi.org/10.1137/0730031).
- [33] V. Simoncini, Block triangular preconditioners for symmetric saddle-point problems, *Applied Numerical Mathematics* 49 (1) (2004) 63–80, numerical Algorithms, Parallelism and Applications. doi:[10.1016/j.apnum.2003.11.012](https://doi.org/10.1016/j.apnum.2003.11.012).
- [34] Y. Notay, A new analysis of block preconditioners for saddle point problems, *SIAM Journal on Matrix Analysis and Applications* 35 (1) (2014) 143–173. doi:[10.1137/130911962](https://doi.org/10.1137/130911962).
- [35] P.R. Amestoy, I.S. Duff, Jean-Yves L’Excellent, J. Koster, A fully asynchronous multifrontal solver using distributed dynamic scheduling, *SIAM J. Matrix Anal. Appl.* 23 (1) (2001) 15–41. doi:[10.1137/S0895479899358194](https://doi.org/10.1137/S0895479899358194).
- [36] P.R. Amestoy, A. Guermouche, Jean-Yves L’Excellent, S. Pralet, Hybrid scheduling for the parallel solution of linear systems, *Parallel Computing* 32 (2) (2006) 136–156. doi:[10.1016/j.parco.2005.07.004](https://doi.org/10.1016/j.parco.2005.07.004).
- [37] N. Collier, D. Pardo, L. Dalcin, M. Paszynski, V.M. Calo, The cost of continuity: A study of the performance of isogeometric finite elements using direct solvers, *Computer Methods in Applied Mechanics and Engineering* 213–216 (0) (2012) 353–361. doi:[10.1016/j.cma.2011.11.002](https://doi.org/10.1016/j.cma.2011.11.002).
- [38] U. Trottenberg, C.W. Oosterlee, A. Schuller, *Multigrid*, Academic Press, Inc., Orlando, FL, USA, 2001.

- [39] M. Adams, Evaluation of three unstructured multigrid methods on 3D finite element problems in solid mechanics, *International Journal for Numerical Methods in Engineering* 55 (5) (2002) 519–534. doi:10.1002/nme.506.
- [40] Y. Saad, A flexible inner-outer preconditioned GMRES algorithm, *SIAM Journal of Scientific Computing* 14 (2) (1993) 461–469. doi:10.1137/0914028.
- [41] M.W. Gee, C.M. Siefert, J.J. Hu, R.S. Tuminaro, M.G. Sala, ML 5.0 smoothed aggregation user’s guide, Tech. Rep. SAND2006-2649, Sandia National Laboratories (2006).
- [42] V.E. Henson, U.M. Yang, BoomerAMG: A parallel algebraic multigrid solver and preconditioner, *Applied Numerical Mathematics* 41 (1) (2002) 155 – 177, developments and Trends in Iterative Methods for Large Systems of Equations - in memorium Rudiger Weiss. doi:10.1016/S0168-9274(01)00115-5.
- [43] N. Collier, L. Dalcin, D. Pardo, V. Calo, The cost of continuity: performance of iterative solvers on isogeometric finite elements, *SIAM Journal on Scientific Computing* 35 (2) (2013) A767–A784. doi:10.1137/120881038.
- [44] P. Vignal, A. Sarmiento, A.M.A. Côrtes, L. Dalcin, V.M. Calo, Coupling Navier–Stokes and Cahn–Hilliard equations in a two-dimensional annular flow configuration, *Procedia Computer Science* 51 (2015) 934–943.
- [45] L. F. R. Espath, A. F. Sarmiento, P. Vignal, B. O. N. Varga, A. M. A. Cortes, L. Dalcin, V. M. Calo, Energy exchange analysis in droplet dynamics via the Navier–Stokes–Cahn–Hilliard model, *Journal of Fluid Mechanics* 797 (2016) 389–430. doi:10.1017/jfm.2016.277.
- [46] Texas Advanced Computing Center [online] (2015).
URL <http://www.tacc.utexas.edu/>
- [47] B.Y. Ballal, R.S. Rivlin, Flow of a Newtonian fluid between eccentric rotating cylinders: Inertial effects, *Archive for Rational Mechanics and Analysis* 62 (3) (1976) 237–294. doi:10.1007/BF00280016.