Accepted Manuscript

Title: Performance evaluation of block-diagonal preconditioners for the divergence-conforming B-Spline discretization of the Stokes system

Author: A.M.A. Côrtes A.L.G.A. Coutinho L. Dalcin V.M. Calo

PII: S1877-7503(15)00009-5
DOI: <http://dx.doi.org/doi:10.1016/j.jocs.2015.01.005>
Reference: JOCS 322

To appear in:

Received date: 2-11-2014
Revised date: 23-12-2014
Accepted date: 28-1-2015

Please cite this article as: A.M.A. Côrtes, 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), <http://dx.doi.org/10.1016/j.jocs.2015.01.005>

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.

Performance evaluation of block-diagonal preconditioners for the divergence-conforming B-Spline discretization of the Stokes system

A.M.A. Côrtes^{a,b,*}, A.L.G.A. Coutinho^b, L. Dalcin^{a,c}, V.M. Calo^{a,d}

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

^bDepartment of Civil Engineering, NACAD, COPPE, Federal University of Rio de Janeiro, Rio de Janeiro, Brazil

^cConsejo Nacional de Investigaciones Científicas y Técnicas (CONICET)
Santa Fe, Argentina

^dPhysical Sciences and Engineering Division (PSE)
Earth Science and Engineering (ErSE)
King Abdullah University of Science and Technology (KAUST)
Thuwal, Saudi Arabia

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 discretized Stokes equations, these spaces generate a symmetric and indefinite saddle-point linear system. Krylov subspace methods are usually the most efficient procedures to solve such systems. One of such methods, for symmetric systems, is the Minimum Residual Method (MINRES). However, the efficiency and robustness of Krylov subspace methods is closely tied to appropriate preconditioning strategies. For the discrete Stokes system, in particular, block-diagonal strategies provide efficient preconditioners. In this paper, we compare the performance of block-diagonal preconditioners for several block choices. We verify how the eigenvalue clustering promoted by the preconditioning strategies affects MINRES convergence. We also compare the number of iterations and wall-clock timings. We conclude that among the building blocks we tested, the strategy with relaxed inner conjugate gradients preconditioned with incomplete Cholesky provided the best results.

Keywords: Isogeometric analysis, Divergence-conforming B-Spline spaces, Krylov subspace method, Block-diagonal preconditioner, Stokes problem

1. Introduction

The concept of isogeometric analysis (IGA) first appeared in [1], and since then several papers followed, either exploring their mathematical theory, for example [2, 3], or showing their potential

*Correspondence to: A.M.A. Côrtes, 4700 King Abdullah University of Science and Technology, al-Khawarizmi Bldg (Bldg 1), Office #4319-CU11, Thuwal 23955-6900, Kingdom of Saudi Arabia
Email address: adriano.cortes@kaust.edu.sa (A.M.A. Côrtes)

in engineering applications, to mention some [4–12]. In [13], the IGA concept is used to discretize vector fields of electromagnetic problems. For such problems, it is known that the function spaces satisfy a de Rham diagram at the continuous level, and for a discretization to be successfully applied to them, the finite dimensional spaces should also satisfy the de Rham diagram at the discrete level. Exploring one of the main features of spline basis functions, that is the easy control of the basis polynomial degree and regularity, and by a suitable choice of B-spline spaces of each component of the two-dimensional vector field, Buffa et al. [13] introduced an IGA discretization satisfying a de Rham diagram. They have shown that the technique can be viewed as a smooth generalization of Nédélec elements, and thus good results were reported.

The generalization for three-dimensional vector fields and the mathematical theory of such discretization appeared in [14]. Their approach, called Isogeometric Discrete Differential Forms, was inspired by the theory of finite element exterior calculus of Arnold et al. [15].

In [16], Buffa et al. introduced three similar vector field discretizations for the Stokes problem. By a proper choice of the polynomial degrees and the regularity of the components of the discrete velocity field and the discrete pressure field, these discretizations can be interpreted as smooth generalizations of Nédélec, Taylor-Hood and Raviart-Thomas elements. Because of the smoothness of the basis functions used, the discrete velocity spaces of these elements are $\mathbf{H}^1$ -conforming, which make them suitable to discretize the Stokes system. Furthermore, in the case of the Raviart-Thomas element type, Buffa et al. [16] characterize the image of the divergence operator from the discrete velocity space (with and without boundary conditions) onto the discrete pressure space, guaranteeing this way a point-wise divergence-free discrete vector field, a condition that is generally only satisfied weakly by classical mixed finite elements.

Following the developments of Buffa et al. [14], Evans and Hughes [17] further developed the Raviart-Thomas element type in the context of Hilbert complexes. Indeed, by using the stable projectors of [14], a divergence-preserving transformation (Piola transformation) of the velocity field and an integral-preserving transformation of the pressure field, Evans and Hughes devised a Stokes complex with a compatible sub-complex that furnishes a discretization scheme, that is at the same time *inf-sup* stable and divergence free. In [17–19], Evans and Hughes applied this discretization scheme to several viscous incompressible flows, and also started its mathematical theory as well.

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 ([20–23]), the most popular being variants of Uzawa’s method, such as the *inexact* Uzawa method, and the Minimum Residual Method (MINRES) [24]. The latter is a member from the family of Krylov subspace methods, and as such, its robustness and performance is very

dependent on the preconditioning strategy.

For example, MINRES is being used to solve large-scale problems in science, such as Earth's mantle convection flows in parallel by finite elements with octree-based adaptive mesh refinement and coarsening (AMR/C), demonstrating scalability up to 122880 cores [25].

The rest of the paper is organized as follows. In section 2, we review some isogeometric analysis definitions, in order to setup the nomenclature for the divergence-conforming discretization. Section 3 reviews the results of [17] with respect to Stokes flow. First, we present the discrete velocity-pressure pair on the parametric domain, and how it is mapped to general geometries by means of proper transformations. Also, the inf-sup stability and the divergence-free property of the divergence-conforming discrete velocity-pressure pair is presented. The next section deals with the discrete variational problem, and how Nitsche's method is used to impose Dirichlet boundary conditions weakly. In section 5, we review the Minimum Residual Method. We discuss its convergence properties and how to precondition it. We also present the block-diagonal preconditioning strategy introduced by Wathen and Silvester in [26, 27], and the choices we made for the preconditioners blocks. Section 6 describes our numerical results. We present the results for three examples: two manufactured analytical solutions for different geometries and the lid-driven cavity flow benchmark. For the lid-driven problem we analyze the preconditioners performance.

2. Isogeometric notation: Spline spaces and the geometrical mapping

We recall some spline spaces definition and related notation to describe the divergence-conforming spaces introduced in [17]. Here, we follow closely [14, 16, 17].

2.1. Univariate B-Splines

To define a univariate B-spline basis we specify the number n of basis functions wanted, the polynomial degree p of the basis and a knot vector Ξ . A knot vector Ξ is a finite nondecreasing sequence $\Xi = \{0 = \xi_1, \dots, \xi_{n+p+1} = 1\}$. The sequence may have repeated knots, in this case one says that the knot has multiplicity greater than one. Introducing the vector $\zeta = \{\zeta_1, \dots, \zeta_m\}$ of knots without repetitions, also called breakpoints, and the vector $\{r_1, \dots, r_m\}$ of their corresponding multiplicities, one has that,

$$\Xi = \underbrace{\{\zeta_1, \dots, \zeta_1\}}_{r_1 \text{ times}}, \underbrace{\{\zeta_2, \dots, \zeta_2\}}_{r_2 \text{ times}}, \dots, \underbrace{\{\zeta_m, \dots, \zeta_m\}}_{r_m \text{ times}}, \quad (1)$$

with $\sum_{i=1}^m r_i = n + p + 1$.

The B-spline basis functions are p -degree piecewise polynomials on the subdivision $\{\zeta_1, \dots, \zeta_m\}$. A stable way of generating them is by using the Cox-de Boor recursion algorithm [28], which receives as inputs p and Ξ . Knot multiplicity is an essential ingredient in spline theory, which

controls the basis 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 $r_1 = r_m = p + 1$, which implies $n \geq p + 1$ and $\alpha_1 = \alpha_m = -1$. The vector $\boldsymbol{\alpha} := \{\alpha_1, \dots, \alpha_m\}$ collects the basis regularity. We define $\boldsymbol{\alpha} - 1 = \{-1, \alpha_2 - 1, \dots, \alpha_{m-1} - 1, -1\}$, when $\alpha_j \geq 0$ for $2 \leq j \leq m - 1$, and $|\boldsymbol{\alpha}| = \min\{\alpha_2, \dots, \alpha_{m-1}\}$.

The set $\{B_i^p\}_{i=1}^n$ defines a linearly independent set of functions with all the good properties wanted for analysis purposes [29]. The space of B-splines spanned by them is denoted by,

$$\mathcal{S}_{\boldsymbol{\alpha}}^p := \text{span} \{B_i^p\}_{i=1}^n. \quad (2)$$

For univariate spline spaces, when $p \geq 1$ and $\alpha_j \geq 0$ for $2 \leq j \leq m - 1$, the derivative of a spline is a spline too, indeed the derivative is a surjective operator, that is,

$$\left\{ \frac{d}{dx} u : u \in \mathcal{S}_{\boldsymbol{\alpha}}^p \right\} \equiv \mathcal{S}_{\boldsymbol{\alpha}-1}^{p-1}. \quad (3)$$

2.2. Bivariate B-splines

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

$$B_{i_1,i_2}^{p_1,p_2} := B_{i_1,1}^{p_1} \otimes B_{i_2,2}^{p_2}, \quad i_1 = 1, \dots, n_1; \quad i_2 = 1, \dots, n_2. \quad (4)$$

The breakpoints $\zeta_d = \{\zeta_{1,d}, \dots, \zeta_{m_d,d}\}$ in each direction $d = 1, 2$ define a mesh

$$\mathcal{M}_h = \{Q = (\zeta_{i_1,1}, \zeta_{i_1+1,1}) \times (\zeta_{i_2,2}, \zeta_{i_2+1,2}) : 1 \leq i_1 \leq m_1 - 1, 1 \leq i_2 \leq m_2 - 1\}, \quad (5)$$

called the parametric mesh, on the parametric domain $\hat{\Omega} = (0, 1)^2$. The subscript h stands for the global mesh size, and is defined as $h := \max_{Q \in \mathcal{M}_h} h_Q$, where $h_Q := \text{diam}(Q)$. To guarantee theoretical convergence estimates, the mesh $\mathcal{M}_h$ should satisfy a shape-regularity condition [2],

$$\lambda^{-1} \leq \frac{h_{Q,\min}}{h_Q} \leq \lambda, \quad \forall Q \in \mathcal{M}_h, \quad (6)$$

for constant $\lambda > 0$, where $h_{Q,\min}$ is the length of the smallest edge of Q . If the same λ holds for a sequence of nested refined meshes $\{\mathcal{M}_h\}_{h \leq h_0}$, this sequence is said to be locally quasi-uniform, which we assume hereafter.

Using the notation $\boldsymbol{\alpha}_1 = \{\alpha_{1,1}, \dots, \alpha_{m_1,1}\}$ and $\boldsymbol{\alpha}_2 = \{\alpha_{1,2}, \dots, \alpha_{m_2,2}\}$ for the regularity vectors in each direction, the bivariate B-spline space is defined as

$$\mathcal{S}_{\boldsymbol{\alpha}_1, \boldsymbol{\alpha}_2}^{p_1, p_2} \equiv \mathcal{S}_{\boldsymbol{\alpha}_1, \boldsymbol{\alpha}_2}^{p_1, p_2}(\mathcal{M}_h) := \text{span} \{B_{i_1, i_2}^{p_1, p_2}\}_{i_1, i_2=1}^{n_1, n_2}. \quad (7)$$

The global regularity of the space is defined as $\alpha := \min\{|\boldsymbol{\alpha}_1|, |\boldsymbol{\alpha}_2|\}$.

2.3. Geometrical Mapping and the physical mesh

The great potential of IGA concept stems from the possibility of working on geometries of varied complexities. This is achieved by the introduction of a geometrical mapping $\mathbf{F} : \widehat{\Omega} \rightarrow \Omega$, from the parametric domain $\widehat{\Omega} = (0,1)^2$ to the general physical domain Ω . We assume that $\mathbf{F}$ is a piecewise smooth mapping over $\mathcal{M}_h$, with piecewise smooth inverse. Moreover, $\mathbf{F}$ is generally given by B-splines or NURBS basis defined on the coarsest mesh $\mathcal{M}_{h_0}$. The advent of the IGA concept started with the observation that many CAD systems provide $\mathbf{F}$.

Implicitly, we have a notion of physical mesh. Indeed, the image of a parametric mesh $\mathcal{M}_h$ induces a mesh on the physical domain Ω , denoted by $\mathcal{K}_h$. Also, the images of the elements boundaries by $\mathbf{F}$ are denoted by $\mathcal{F}_h$, and the boundaries that are contained in $\partial\Omega$ define the boundary mesh, denoted by Γ_h .

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

3.1. The Stokes problem

The Stokes system in its strong form is

$$-\operatorname{div}(2\nu\nabla^s \mathbf{u}) + \nabla p = \mathbf{f} \text{ in } \Omega, \quad (8a)$$

$$\operatorname{div} \mathbf{u} = 0 \text{ in } \Omega, \quad (8b)$$

with $\Omega \subset \mathbb{R}^2$ a bounded simply connected Lipschitz open set, ∇^s the symmetric gradient operator, $\mathbf{u}$ is the flow velocity, p is the pressure, $\nu > 0$ is the kinematic viscosity and $\mathbf{f}$ denotes a body force acting on the fluid. For the system to be well posed, it must be augmented with appropriate boundary conditions. To simplify the presentation, we consider here the case of homogeneous Dirichlet boundary conditions that is, the no-slip case $\mathbf{u} = \mathbf{0}$. Then, as usual in the finite element framework, the strong form is recast in a weak formulation given by:

Find $\mathbf{u} \in \mathbf{H}_0^1(\Omega)$ and $p \in L_0^2(\Omega)$ such that

$$a(\mathbf{u}, \mathbf{v}) + b(\mathbf{v}, p) + b(\mathbf{u}, q) = (\mathbf{f}, \mathbf{v})_{\mathbf{L}^2(\Omega)}, \quad (9)$$

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

$$a(\mathbf{w}, \mathbf{v}) = (2\nu\nabla^s \mathbf{w}, \nabla^s \mathbf{v})_{\mathbf{L}^2(\Omega)}, \quad (10)$$

$$b(\mathbf{v}, q) = -(\operatorname{div} \mathbf{v}, q)_{L^2(\Omega)}. \quad (11)$$

From Brezzi's theory [30], it is known that (9) is an optimality condition for a saddle-point $(\mathbf{u}, p)$ of a Lagrangian functional, and that a solution $(\mathbf{u}, p) \in \mathbf{H}_0^1(\Omega) \times L_0^2(\Omega)$ 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

$$\inf_{q \in L_0^2(\Omega), q \neq 0} \sup_{\mathbf{v} \in \mathbf{H}_0^1(\Omega)} \frac{b(\mathbf{v}, q)}{\|\mathbf{v}\|_{\mathbf{H}^1(\Omega)} \|q\|_{L^2(\Omega)}} \geq \beta, \quad (12)$$

with the constant $\beta > 0$.

3.2. Divergence-conforming B-spline discretization

In this section, we review the definitions and results of the divergence-conforming B-spline discretization of the Stokes problem, as elaborated for general incompressible flows in [17], which first appeared in [14].

Assuming the global regularity $\alpha \geq 0$, the discrete velocity space on the parametric domain $\widehat{\Omega}$ is defined as

$$\widehat{\mathcal{V}}_h := \mathcal{S}_{\alpha_1+1, \alpha_2}^{p_1+1, p_2} \times \mathcal{S}_{\alpha_1, \alpha_2+1}^{p_1, p_2+1} \quad (13)$$

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

$$\widehat{\mathcal{Q}}_h := \mathcal{S}_{\alpha_1, \alpha_2}^{p_1, p_2}. \quad (14)$$

The space $\widehat{\mathcal{V}}_h$ is as smooth generalization of the Raviart-Thomas element, indeed when $\alpha = -1$ it coincides with the classical Raviart-Thomas elements on quadrilaterals. However, for $\alpha \geq 0$, (13) is $\mathbf{H}^1$ -conforming (patch-wise), which makes it appropriate as discrete velocity approximations for the Stokes and Navier-Stokes equations. We adopt the convention that everything referring to the parametric space is denoted with a hat superscript.

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

$$\widehat{\mathcal{V}}_{0,h} := \left\{ \widehat{\mathbf{v}}_h \in \widehat{\mathcal{V}}_h : \widehat{\mathbf{v}}_h \cdot \widehat{\mathbf{n}} = 0 \text{ on } \partial\widehat{\Omega} \right\} \subset \mathbf{H}_0(\widehat{\text{div}}; \widehat{\Omega}), \quad (15)$$

where $\widehat{\mathbf{n}}$ denotes the outward normal to $\partial\widehat{\Omega}$ and $\widehat{\text{div}}$ the divergence operator in parametric coordinates. With this choice for the velocity space, a constrained discrete pressure space

$$\widehat{\mathcal{Q}}_{0,h} := \left\{ \widehat{q}_h \in \widehat{\mathcal{Q}}_h : \int_{\widehat{\Omega}} \widehat{q}_h = 0 \right\} \subset L_0^2(\widehat{\Omega}). \quad (16)$$

must be defined.

The rationale for choosing these constrained spaces is that, in order to guarantee a divergence-free velocity field that does not conflict with the *inf-sup* stability of the velocity-pressure pair, we must guarantee the surjectivity of the divergence operator at the discrete level. Indeed, together with the surjectivity of the derivative between B-spline spaces (3), one can prove that

$$\widehat{\mathcal{V}}_{0,h} \xrightarrow{\widehat{\text{div}}} \widehat{\mathcal{Q}}_{0,h} \quad (17)$$

forms a cochain complex. Then, if we have the incompressibility condition weakly satisfied, that is,

$$(\widehat{\operatorname{div}} \widehat{\mathbf{v}}_h, \widehat{q}_h)_{L^2(\Omega)} = 0 \text{ for all } \widehat{q}_h \in \widehat{\mathcal{Q}}_{0,h}, \quad (18)$$

we can take $\widehat{q}_h = \widehat{\operatorname{div}} \widehat{\mathbf{v}}_h$ above, which implies $\|\widehat{\operatorname{div}} \widehat{\mathbf{v}}_h\|_{L^2(\Omega)} = 0$, and since $\widehat{\operatorname{div}} \widehat{\mathbf{v}}_h$ is at least continuous, we have that $\operatorname{div} \widehat{\mathbf{v}}_h = 0$ pointwise.

However, if the discrete velocity space is constrained by no-slip conditions on $\partial\widehat{\Omega}$, the discrete pressure space would also be constrained on the corners of $\widehat{\Omega}$, which renders the pressure approximation less accurate. For a complete discussion, see [16].

With the growing popularity and successful application of Nitsche's method to impose boundary conditions weakly [11, 31–33], the above velocity-pressure pair choice are not a limitation, as no-slip conditions are imposed weakly by augmenting the variational formulation with additional terms, as we describe in the next section.

Up to now we worked on the parametric domain $\widehat{\Omega}$. The definition of the discrete velocity and pressure spaces on the physical domain Ω are established by considering appropriate transformations induced by $\mathbf{F}$. These transformations are due to the pullback operators:

$$\iota_{\mathbf{u}}(\mathbf{v}) = \det(D\mathbf{F})(D\mathbf{F})^{-1}(\mathbf{v} \circ \mathbf{F}), \quad \mathbf{v} \in \mathbf{H}_0(\operatorname{div}; \Omega) \quad (19)$$

$$\iota_p(q) = \det(D\mathbf{F})(q \circ \mathbf{F}), \quad q \in L_0^2(\Omega) \quad (20)$$

where $D\mathbf{F}$ is the Jacobian matrix of the geometrical mapping $\mathbf{F}$. The first one, the Piola transform, is a standard choice to build approximation spaces in $\mathbf{H}(\operatorname{div}; \Omega)$, mainly in the context of mixed finite elements, since it is divergence-preserving and also preserves the normal component of the transformed vector field. The second one is necessary to preserve the zero mean pressure constraint on the physical domain Ω .

With the goal of preserving the surjectivity of the divergence operator, but now on physical coordinates, the discrete velocity and pressure spaces on the physical domain are defined by

$$\mathcal{V}_{0,h} := \left\{ \mathbf{v} \in \mathbf{H}_0(\operatorname{div}; \Omega) : \widehat{\mathbf{v}} = \iota_{\mathbf{u}}(\mathbf{v}) \in \widehat{\mathcal{V}}_{0,h} \right\}, \quad (21)$$

$$\mathcal{Q}_{0,h} := \left\{ q \in L_0^2(\Omega) : \widehat{q} = \iota_p(q) \in \widehat{\mathcal{Q}}_{0,h} \right\} \quad (22)$$

The last ingredient necessary by the framework of isogeometric differential forms is the existence of suitable projectors. In [14], Buffa et al. introduced L^2 -stable projection operators $\Pi_{\mathcal{V}_h}^0 : \mathbf{H}_0(\operatorname{div}; \Omega) \rightarrow \mathcal{V}_{0,h}$ and $\Pi_{\mathcal{Q}_h}^0 : L_0^2(\Omega) \rightarrow \mathcal{Q}_{0,h}$ that commute with the divergence operator. This commutativity property is the key ingredient to prove that the velocity-pressure pair is *inf-sup* stable in the discrete sense (for the proof see [17]).

In the sections that follow, we denote the polynomial degree of the velocity-pressure pair by $k' = \min\{p_1, p_2\}$. In our numerical examples, we used $p_1 = p_2 = p$, so in this case $k' = p$ denotes the polynomial degree of the pressure space $\mathcal{Q}_{0,h}$ of the pair.

4. Discrete variational formulation

With the discrete divergence-conforming velocity-pressure spaces pair properly defined, we consider the discrete formulation of (9). Since the discrete velocity space $\mathcal{V}_{0,h}$ only satisfies the no-penetration ($\mathbf{u} \cdot \mathbf{n} = 0$) constraint, the no-slip condition has to be imposed weakly, that is, by modifying the variational formulation properly. Following [17], Nitsche's method is applied. It works as a penalty method by adding variationally consistent terms to the bilinear form $a(\cdot, \cdot)$. Indeed, defining the new bilinear form

$$a_h(\mathbf{u}_h, \mathbf{v}_h) = a(\mathbf{u}_h, \mathbf{v}_h) - \eta_h(\mathbf{u}_h, \mathbf{v}_h), \quad (23)$$

where

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

with $C_{pen} > 0$ the Nitsche's penalization parameter and h_F the characteristic length of the face F , the discrete formulation for the no-slip boundary condition is written as:

Find $\mathbf{u}_h \in \mathcal{V}_{0,h}$ and $p_h \in \mathcal{Q}_{0,h}$ such that

$$a_h(\mathbf{u}_h, \mathbf{v}_h) + b(p_h, \mathbf{v}_h) + b(q_h, \mathbf{u}_h) = (\mathbf{f}, \mathbf{v}_h)_{\mathbf{L}^2(\Omega)}, \quad (25)$$

for all $\mathbf{v}_h \in \mathcal{V}_{0,h}$ and $q_h \in \mathcal{Q}_{0,h}$.

Non-homogeneous tangential Dirichlet boundary conditions are also treated by Nitsche's method. In this case, we add the linear form

$$l_h(\mathbf{v}_h) = \sum_{F \in \Gamma_h} \int_F 2\nu \left(-((\nabla^s \mathbf{v}_h) \mathbf{n}) \cdot \mathbf{g} + \frac{C_{pen}}{h_F} \mathbf{g} \cdot \mathbf{v}_h \right) dS \quad (26)$$

to the right hand side of (25), where $\mathbf{g}$ is a function defined on $\partial\Omega$ that corresponds to the prescribed tangential component of $\mathbf{u}$ on $\partial\Omega$.

Denoting by $\{\Phi_1, \dots, \Phi_{n_u}\}$ the basis of $\mathcal{V}_{0,h}$ and by $\{\phi_1, \dots, \phi_{n_p}\}$ the basis of $\mathcal{Q}_{0,h}$, then the solution of (25)-(26) resumes to the solution of the discrete Stokes system

$$\begin{bmatrix} \mathbf{A}_h & B^T \\ B & 0 \end{bmatrix} \begin{bmatrix} \mathbf{U} \\ \mathbf{P} \end{bmatrix} = \begin{bmatrix} \mathbf{F} \\ \mathbf{0} \end{bmatrix} \quad (27)$$

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

$$[\mathbf{A}_h]_{i,j} = a_h(\Phi_j, \Phi_i) = a(\Phi_j, \Phi_i) - \eta_h(\Phi_j, \Phi_i), \quad (28)$$

$$[B]_{k,j} = b(\phi_k, \Phi_j), \quad (29)$$

$$[\mathbf{F}]_i = (\mathbf{f}, \Phi_i)_{\mathbf{L}^2(\Omega)} + l_h(\Phi_i), \quad (30)$$

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 $p_h \in \mathcal{Q}_{0,h}$.

By Lemma 6.3.2 of [34] we have that, if $C_{pen} > 0$ is sufficiently large (see Chapter 6, [34] for details), the bilinear form $a_h(\cdot, \cdot)$ satisfies

$$\langle \mathbf{A}_h \mathbf{U}, \mathbf{U} \rangle = a_h(\mathbf{u}_h, \mathbf{u}_h) \geq \frac{\nu}{C_{Korn}} |\mathbf{u}_h|_{\mathbf{H}^1(\Omega)}^2 + \sum_{F \in \Gamma_h} \frac{\nu C_{pen}}{2h_F} \|\mathbf{u}_h\|_{\mathbf{L}^2(F)}^2, \quad (31)$$

for all $\mathbf{u}_h \in \mathcal{V}_{0,h}$, where C_{Korn} is the constant of Korn's inequality and $|\cdot|_{\mathbf{H}^1}$ is the $\mathbf{H}^1$ -seminorm. In particular, this estimate implies that the symmetric matrix $\mathbf{A}_h$ is positive definite, and that $\|\mathbf{u}_h\|_{\mathcal{V}} = a_h(\mathbf{u}_h, \mathbf{u}_h)^{1/2} = \langle \mathbf{A}_h \mathbf{U}, \mathbf{U} \rangle^{1/2}$ defines a norm.

Combining Lemmas 6.3.1 and 6.3.3 of [34], we have the stronger discrete *inf-sup* stability condition

$$\inf_{p_h \in \mathcal{Q}_{0,h}, p_h \neq 0} \sup_{\mathbf{u}_h \in \mathcal{V}_{0,h}} \frac{b(p_h, \mathbf{u}_h)}{\|\mathbf{u}_h\|_{\mathcal{V}} \|p_h\|_{\mathcal{Q}}} \geq \beta_0, \quad (32)$$

where $\|p_h\|_{\mathcal{Q}} := \frac{1}{2\nu} \|p_h\|_{L^2(\Omega)}$ and $\beta_0 > 0$ is the stability constant that is independent of h , ν , and scales as $O(C_{pen}^{-1/2})$. Tables 1 and 2 show the dependency of the stability constant β_0 on the penalization parameter C_{pen} for $h = 1/16$, $\Omega = (0, 1)^2$, $k' = 2$, and $h = 1/16$, $\Omega = (0, 1)^2$, $k' = 3$, respectively, confirming the dependency numerically.

C_{pen}	4	8	16	32	64	128	256	1024
β_0	0.6015	0.5543	0.4266	0.3106	0.2227	0.1584	0.1122	0.0566
order	—	-0.1178	-0.3779	-0.4577	-0.4801	-0.4913	-0.4971	—

Table 1: Inf-sup stability constant β_0 for $h = 1/16$, $\Omega = (0, 1)^2$, $k' = 2$ and $\nu = 1$.

C_{pen}	4	8	16	32	64	128	256	1024
β_0	0.6001	0.5853	0.4759	0.3445	0.2456	0.1744	0.1233	0.0616
order	—	-0.0359	-0.2985	-0.4661	-0.4885	-0.4940	-0.5000	—

Table 2: Inf-sup stability constant β_0 for $h = 1/16$, $\Omega = (0, 1)^2$, $k' = 3$ and $\nu = 1$.

Thus, as remarked by Evans [34], Nitsche's penalization parameter C_{pen} should be taken as small as possible to guarantee the coercivity of a_h , and as we discuss in the next section, the inf-sup stability constant β_0 also plays a role in the convergence analysis of the preconditioned MINRES, and it is also desirable to keep it as large as possible (keeping C_{pen} small) for numerical reasons. Finally, in [17, 34] one can find the results for stability, existence and uniqueness of the discrete solution, as well as the mathematical theory of a priori error estimates for the generalized Stokes problem.

5. Linear Solver

Here, we discuss the solution of the resulting Stokes system discretized by inf-sup stable mixed elements. That is, the solution of the symmetric indefinite linear system with a block coefficient matrix as in equation (27).

The symmetry of $\mathbf{A}$ follows from equation (28). The indefinite property follows from the block LDU factorization,

$$\mathcal{A} = \begin{bmatrix} \mathbf{A} & B^T \\ B & 0 \end{bmatrix} = \begin{bmatrix} I & 0 \\ B\mathbf{A}^{-1} & I \end{bmatrix} \begin{bmatrix} \mathbf{A} & 0 \\ 0 & -B\mathbf{A}^{-1}B^T \end{bmatrix} \begin{bmatrix} I & \mathbf{A}^{-1}B^T \\ 0 & I \end{bmatrix}. \quad (33)$$

Indeed, from (33) we conclude that $\mathcal{A}$ has n_u positive eigenvalues ($\mathbf{A}$ is positive definite), and it has at least $n_p - 1$ negative eigenvalues, since for enclosed flows we have $\mathbf{1} \in \text{Ker}(B^T)$, where $\mathbf{1} \in \mathbb{R}^{n_p}$ is the vector with all the components equal to one. The positive semidefinite matrix $S = B\mathbf{A}^{-1}B^T$ is called the *pressure Schur complement* and it is fundamental in devising good preconditioners.

Although being symmetric, the indefiniteness of $\mathcal{A}$ precludes the application of the Conjugate Gradient method to solve (27). Another Krylov subspace method is better suited for this task: the Minimum Residual Method (MINRES) [22, 24], which requires only the symmetry of $\mathcal{A}$. For completeness, we review the main characteristics and the convergence results of MINRES and its preconditioned version.

Consider the linear system $\mathcal{A}\mathbf{x} = \mathbf{b}$, with $\mathcal{A}$ symmetric and possibly indefinite. Given an initial guess $\mathbf{x}_0$, the MINRES method generates a sequence of approximate solutions $\mathbf{x}_k$, $k = 1, 2, \dots$ with the property

$$\mathbf{x}_k \in \mathbf{x}_0 + \mathcal{K}_k(\mathcal{A}, \mathbf{r}_0), \quad (34)$$

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

$$\begin{aligned} \|\mathbf{r}_k\|_2 &= \min_{\mathbf{x} \in \mathbf{x}_0 + \mathcal{K}_k(\mathcal{A}, \mathbf{r}_0)} \|\mathbf{b} - \mathcal{A}\mathbf{x}\|_2 = \min_{p \in \Pi_k} \|p(\mathcal{A})\mathbf{r}_0\|_2 \\ &\leq \min_{p \in \Pi_k} \max_{\lambda \in \sigma(\mathcal{A})} |p(\lambda)| \|\mathbf{r}_0\|_2. \end{aligned} \quad (35)$$

where Π_k is the set of polynomials of degree at most k with $p(0) = 1$, and $\sigma(\mathcal{A})$ is the spectrum of $\mathcal{A}$ (see [22, 23]). The inequality (35) follows from the symmetry of $\mathcal{A}$. Additionally, in terms of implementation, the symmetry ensures the ultimate memory efficiency goal of a *short-term recurrence* to generate an orthogonal basis for the Krylov subspace $\mathcal{K}_k(\mathcal{A}, \mathbf{r}_0)$, which is achieved by the Lanczos method [35].

It is clear from (35) the dependence of MINRES convergence on the spectrum $\sigma(\mathcal{A})$. Since, in this case, $\mathcal{A}$ has positive and negative eigenvalues, clusterizing such eigenvalues is the goal of a good preconditioner for (27).

In order to preserve the symmetry of the preconditioned system, we assume a symmetric and positive definite preconditioner $\mathcal{M}$. In this case, by using the Cholesky decomposition $\mathcal{M} = \mathcal{L}\mathcal{L}^T$ one can write the symmetric preconditioned version of the linear system, that is,

$$(\mathcal{L}^{-1}\mathcal{A}\mathcal{L}^{-T})\mathbf{y} = \mathcal{L}^{-1}\mathbf{b}, \quad \mathbf{y} = \mathcal{L}^T\mathbf{x}. \quad (36)$$

Then, MINRES is applied to the preconditioned system. The Euclidian norm of the preconditioned residual $\tilde{\mathbf{r}}$ is related to the residual of the original system by,

$$\|\tilde{\mathbf{r}}_k\|_2 = \langle \mathcal{L}^{-1}\mathbf{r}_k, \mathcal{L}^{-1}\mathbf{r}_k \rangle^{1/2} = \langle \mathcal{L}^{-T}\mathcal{L}^{-1}\mathbf{r}_k, \mathbf{r}_k \rangle^{1/2} = \langle \mathcal{M}^{-1}\mathbf{r}_k, \mathbf{r}_k \rangle^{1/2} = \|\mathbf{r}_k\|_{\mathcal{M}^{-1}}, \quad (37)$$

and the convergence bound (35),

$$\|\tilde{\mathbf{r}}_k\|_2 \leq \min_{p \in \Pi_k} \max_{\lambda \in \sigma(\mathcal{L}^{-1}\mathcal{A}\mathcal{L}^{-T})} |p(\lambda)| \|\tilde{\mathbf{r}}_0\|_2, \quad (38)$$

for the preconditioned system by

$$\|\mathbf{r}_k\|_{\mathcal{M}^{-1}} \leq \min_{p \in \Pi_k} \max_{\lambda \in \sigma(\mathcal{M}^{-1}\mathcal{A})} |p(\lambda)| \|\mathbf{r}_0\|_{\mathcal{M}^{-1}}, \quad (39)$$

observing that $\sigma(\mathcal{L}^{-1}\mathcal{A}\mathcal{L}^{-T}) = \sigma(\mathcal{M}^{-1}\mathcal{A})$ since $\mathcal{L}^{-1}\mathcal{A}\mathcal{L}^{-T}$ and $\mathcal{M}^{-1}\mathcal{A}$ are similar matrices.

We see that the convergence property of the preconditioned MINRES not only depends on $\sigma(\mathcal{M}^{-1}\mathcal{A})$, but is measured in a norm induced by the preconditioner. The factorization $\mathcal{M} = \mathcal{L}\mathcal{L}^T$ is only needed for theoretical purposes and is never used in practice. Practical implementations only need the action of $\mathcal{M}^{-1}$, or equivalently, the solution of a linear system with $\mathcal{M}$ as the coefficient matrix. Consequently, besides clustering the spectrum of $\mathcal{A}$, a good preconditioner $\mathcal{M}$ should result in a linear system that is easy to solve.

5.1. Block-diagonal Preconditioning Strategy

We now discuss a preconditioning strategy for the discrete Stokes system (27) well established in the literature [22, 23, 26, 27]. First, consider the block factorization of (33). A possible preconditioner in this case is the positive definite block diagonal matrix

$$\mathcal{M} = \begin{bmatrix} \mathbf{A} & 0 \\ 0 & S \end{bmatrix}, \quad (40)$$

with $S = B\mathbf{A}^{-1}B^T$ (pressure Schur complement). Indeed, one can prove that $\sigma(\mathcal{M}^{-1}\mathcal{A}) = \{(1 - \sqrt{5})/2, 1, (1 + \sqrt{5})/2\}$ (Chapter 6, [22]), and the minimax polynomial convergence estimate (39) guarantees the convergence of the preconditioned MINRES to the exact solution after at

most three iterations. Clearly, this preconditioner does not fulfill the requirement of being easily solvable, because solving a linear system with the pressure Schur complement is not an easy task since it is generally a dense matrix and we do not have $\mathbf{A}^{-1}$ at hand.

A fundamental concept for deriving scalable preconditioning strategies with respect to the system size or equivalently reducing mesh size h is that of *spectral equivalence*. Two matrices K_1 and K_2 are said to be spectrally equivalent, if there are constants $c, C > 0$, both independent of h , such that,

$$c \leq \frac{\langle K_1 \mathbf{x}, \mathbf{x} \rangle}{\langle K_2 \mathbf{x}, \mathbf{x} \rangle} \leq C, \quad \forall \mathbf{x} \neq \mathbf{0}. \quad (41)$$

For general inf-sup stable and conforming mixed discretizations, the discrete inf-sup stability condition and the boundedness of the bilinear form $b(\mathbf{u}, p) = -(\operatorname{div} \mathbf{u}, p)_{L^2(\Omega)}$ (Chapter 5, [22]) imply that

$$\beta_0^2 \leq \frac{\langle B\mathbf{A}^{-1}B^T \mathbf{P}, \mathbf{P} \rangle}{\langle Q\mathbf{P}, \mathbf{P} \rangle} \leq C_b^2, \quad \forall \mathbf{P} \in \mathbb{R}^{n_p} \setminus \operatorname{Ker}(B^T), \quad (42)$$

where $\beta_0 > 0$ is the inf-sup constant and $C_b > 0$ is the continuity constant, both independent of h . The above inequality implies the spectral equivalence of the pressure Schur complement $S = B\mathbf{A}^{-1}B^T$ and the pressure mass matrix Q , that is, the matrix whose coefficients are $[Q]_{i,j} = (\phi_j, \phi_i)_{L^2(\Omega)}$, for $i, j = 1, \dots, n_p$.

Then, a more viable preconditioner, in the case of inf-sup stable and conforming discretization, is

$$\mathcal{M} = \begin{bmatrix} \mathbf{A} & 0 \\ 0 & Q \end{bmatrix}. \quad (43)$$

One can prove in this case [22, 27] that the spectrum $\sigma(\mathcal{M}^{-1}\mathcal{A})$ is included in the union

$$\left[\frac{1 - \sqrt{1 + 4C_b^2}}{2}, \frac{1 - \sqrt{1 + 4\beta_0^2}}{2} \right] \cup \{1\} \cup \left[\frac{1 + \sqrt{1 + 4\beta_0^2}}{2}, \frac{1 + \sqrt{1 + 4C_b^2}}{2} \right]. \quad (44)$$

In our case, the divergence-conforming discretization, the stronger inf-sup stability condition (32) and the boundedness of the bilinear form b imply that

$$\beta_0^2 \leq \frac{\langle B\mathbf{A}_h^{-1}B^T \mathbf{P}, \mathbf{P} \rangle}{\langle Q_\nu \mathbf{P}, \mathbf{P} \rangle} \leq C_b^2, \quad \forall \mathbf{P} \in \mathbb{R}^{n_p} \setminus \operatorname{Ker}(B^T), \quad (45)$$

where $Q_\nu := \frac{1}{2\nu}Q$ is the properly scaled pressure mass matrix, since it takes into account the viscosity parameter $\nu > 0$. Then using the preconditioner (43) with $\mathbf{A} = \mathbf{A}_h$ and $Q = Q_\nu$ we have the inclusion (44) for the spectrum $\sigma(\mathcal{M}^{-1}\mathcal{A})$. Also remember that, in our case, $\beta_0^2 = O(C_{pen}^{-1})$.

When the preconditioner (43) is applied using a direct solver for each block, it is denoted as an *ideal preconditioner*. In this case, by the eigenvalues bounds (44), the inclusion intervals are independent of the mesh-size parameter h , and invariance on the number of iterations for the preconditioned MINRES to converge to a prescribed tolerance, for fixed ν and C_{pen} , is expected.

In the next section, we verify numerically that these bounds are indeed sharp, thus, achieving an optimal clustering of the eigenvalues of the preconditioned matrix.

Solving a linear system with $\mathbf{A}$ and Q may not be an easy task, either. A more general and practical preconditioning strategy is to consider approximations $\mathbf{A} \approx M_{\mathbf{A}}$ and $Q \approx M_Q$ and the preconditioner

$$\mathcal{M} = \begin{bmatrix} M_{\mathbf{A}} & 0 \\ 0 & M_Q \end{bmatrix}, \quad (46)$$

where $M_{\mathbf{A}} \in \mathbb{R}^{n_u \times n_u}$ and $M_Q \in \mathbb{R}^{n_p \times n_p}$ are symmetric and positive-definite. The effectiveness of such strategy is given by the following spectral bounds: let $\gamma_{\mathbf{A}}, \Gamma_{\mathbf{A}} > 0$ and $M_{\mathbf{A}}$ be such that

$$\gamma_{\mathbf{A}} \leq \frac{\langle \mathbf{A}\mathbf{U}, \mathbf{U} \rangle}{\langle M_{\mathbf{A}}\mathbf{U}, \mathbf{U} \rangle} \leq \Gamma_{\mathbf{A}}, \quad \forall \mathbf{U} \in \mathbb{R}^{n_u} \setminus \{0\}, \quad (47)$$

and $\gamma_Q, \Gamma_Q > 0$ and M_Q be such that

$$\gamma_Q \leq \frac{\langle Q\mathbf{P}, \mathbf{P} \rangle}{\langle M_Q\mathbf{P}, \mathbf{P} \rangle} \leq \Gamma_Q, \quad \forall \mathbf{P} \in \mathbb{R}^{n_p} \setminus \{0\}. \quad (48)$$

Moreover, the eigenvalues bound for the spectrum $\sigma(\mathcal{M}^{-1}\mathcal{A})$ are given by the theorem,

Theorem 1 (Wathen and Silvester [22, 27]). *For an inf-sup stable and conforming mixed discretization and a block diagonal preconditioner $\mathcal{M}$ of the form (46) satisfying (47) and (48), the eigenvalues of the preconditioned Stokes matrix is contained in the union*

$$\left[\frac{\gamma_{\mathbf{A}} - \sqrt{\gamma_{\mathbf{A}}^2 + 4C_b^2\Gamma_{\mathbf{A}}\Gamma_Q}}{2}, \frac{\gamma_{\mathbf{A}} + \sqrt{\gamma_{\mathbf{A}}^2 + 4\beta_0^2\gamma_{\mathbf{A}}\gamma_Q}}{2} \right] \cup \left[\gamma_{\mathbf{A}}, \frac{\Gamma_{\mathbf{A}} + \sqrt{\Gamma_{\mathbf{A}}^2 + 4C_b^2\Gamma_{\mathbf{A}}\Gamma_Q}}{2} \right]. \quad (49)$$

Obviously, several choices are possible for $M_{\mathbf{A}}$ and M_Q , but clearly the most robust ones are those where the spectral bounds (47) and (48) are indeed spectral equivalences, that is, do not depend on the mesh-size parameter h . In our numerical tests, we consider several possible combinations for $M_{\mathbf{A}}$ and M_Q that we call:

- *Ideal($\mathbf{A}, Q$) preconditioning*, where $M_{\mathbf{A}} = \mathbf{A}$ and $M_Q = Q$ and the preconditioner is solved by a direct solver. In this case, the stronger eigenvalues inclusion estimate (44) holds;
- *PCG($\mathbf{A}, Q$) preconditioning*, where $M_{\mathbf{A}}$ approximates $\mathbf{A}$ by solving a system with the coefficient matrix $\mathbf{A}$ by a diagonally preconditioned conjugate gradient method. The same is used for the preconditioner block M_Q ;
- *Ideal($\mathbf{A}$), Diag(Q) preconditioning*, where $M_{\mathbf{A}} = \mathbf{A}$ is solved by a direct solver and $M_Q = \text{Diag}(Q)$. For classical Lagrangian finite element bases functions, it is known [26] that $\text{Diag}(Q)$ is spectrally equivalent to Q ;
- *PCG($\mathbf{A}$), Diag(Q) preconditioning*. Is a combination of the two choices above;

- $\text{Diag}(\mathbf{A}, Q)$ preconditioning, where $M_{\mathbf{A}} = \text{Diag}(\mathbf{A})$ and $M_Q = \text{Diag}(Q)$. For classical Lagrangian finite element bases functions, it is known [26] that when $\mathbf{A}$ is a discrete vectorial Laplacian, the spectral bound (47) holds with $\gamma_{\mathbf{A}} = O(h^2)$ and $\Gamma_{\mathbf{A}} = O(1)$.

6. Numerical Results

Our numerical results were performed on the MATLAB®\Octave toolbox GeoPDEs [36]. We had to make some minor modifications on the divergence-conforming B-spline discretization already implemented in GeoPDEs. Also, we added the Nitsche's method to it. Our development was incorporated to the official GeoPDEs release, and is now available in the package distribution, which makes our results reproducible by anyone.

We used three different test cases: two manufactured solutions in two different geometries, a square and a 1/8th of the annulus, and the lid-driven cavity flow benchmark.

6.1. Square domain

This example was used for validation and verification in [16, 17]. Here we will use it with the same goal, namely, to validate our implementation on GeoPDEs and to understand the preconditioners performance.

Let $\hat{\Omega} = \Omega = (0, 1)^2$. The analytical solution is given by

$$\mathbf{u} = \begin{pmatrix} 2e^x(-1+x)^2x^2(y^2-y)(-1+2y) \\ e^x(-1+x)x(-2+x(3+x))(-1+y)^2y^2 \end{pmatrix} \quad (50)$$

and

$$p = -424 + 156e + (y^2 - y)(-456 + e^x(456 + x^2(228 - 5(y^2 - y))) \quad (51)$$

$$+ 2x(-228 + (y^2 - y)) + 2x^3(-36 + (y^2 - y)) + x^4(12 + (y^2 - y))). \quad (52)$$

The boundary condition is $\mathbf{u} = 0$ on $\partial\Omega$, the geometric mapping $\mathbf{F}$ is the identity mapping and the body force is

$$\mathbf{f} = -\nabla \cdot (2\nu\nabla^s \mathbf{u}) + \nabla p. \quad (53)$$

Following [17], we select in all computations the Nitsche's penalization constant as $C_{pen} = 5(k' + 1)$. The approximation errors for the velocity and the pressure, and the convergence orders for the polynomials degrees $k' = 2$ and $k' = 3$ are shown at Tables 3 and 4, respectively. As predicted by the a priori convergence estimates of [17], the order of convergence of the velocity in $\mathbf{H}^1$ -seminorm is k' , and in the $\mathbf{L}^2$ -norm is $k' + 1$. Although the computed order of convergence for the pressure in the L^2 -norm is $k' + 1$ in this case, we observed from tests where the meshes are distorted, that the a priori convergence order estimate holds, that is, k' .

Figure 1: Velocity magnitude and some streamlines of the manufactured solution for the square domain.

h	1/8	1/16	1/32	1/64
$\ \mathbf{u} - \mathbf{u}_h\ _{\mathbf{H}^1}$	2.26e-03	5.58e-04	1.39e-04	3.46e-05
order	—	2.02	2.01	2.00
$\ \mathbf{u} - \mathbf{u}_h\ _{\mathbf{L}^2}$	4.10e-05	5.18e-06	6.55e-07	8.26e-08
order	—	2.98	2.98	2.99
$\ p - p_h\ _{\mathbf{L}^2}$	8.96e-05	8.58e-06	9.24e-07	1.07e-07
order	—	3.38	3.21	3.11

Table 3: Errors and convergence orders for $k' = 2$ (Square domain).

h	1/8	1/16	1/32	1/64
$\ \mathbf{u} - \mathbf{u}_h\ _{\mathbf{H}^1}$	1.24e-04	1.62e-05	2.09e-06	2.65e-07
order	—	2.93	2.96	2.98
$\ \mathbf{u} - \mathbf{u}_h\ _{\mathbf{L}^2}$	2.33e-06	1.58e-07	1.03e-08	6.53e-10
order	—	3.88	3.95	3.97
$\ p - p_h\ _{\mathbf{L}^2}$	4.84e-06	3.30e-07	2.19e-08	1.41e-09
order	—	3.87	3.92	3.96

Table 4: Errors and convergence orders for $k' = 3$ (Square domain).

	Ideal($\mathbf{A}$, Q)	PCG($\mathbf{A}$, Q)	Ideal($\mathbf{A}$),Diag(Q)	PCG($\mathbf{A}$),Diag(Q)	Diag($\mathbf{A}$, Q)
h	<i>iter(s)</i>	<i>iter(s)</i>	<i>iter(s)</i>	<i>iter(s)</i>	<i>iter(s)</i>
1/8	33 (0.79)	50 (0.28)	105 (0.33)	138 (0.40)	280 (0.07)
1/16	35 (0.55)	45 (0.43)	125 (1.46)	152 (0.94)	526 (0.16)
1/32	35 (2.72)	41 (1.50)	135 (8.53)	164 (4.51)	681 (0.76)
1/64	37 (14.17)	43 (8.67)	135 (42.60)	172 (28.43)	1248 (3.89)
1/128	35 (79.74)	45 (53.05)	131 (252.68)	172 (179.45)	2753 (32.79)

Table 5: Number of iterations and times for $k' = 2$ (Square domain).

	Ideal($\mathbf{A}$, Q)	PCG($\mathbf{A}$, Q)	Ideal($\mathbf{A}$),Diag(Q)	PCG($\mathbf{A}$),Diag(Q)	Diag($\mathbf{A}$, Q)
h	<i>iter(s)</i>	<i>iter(s)</i>	<i>iter(s)</i>	<i>iter(s)</i>	<i>iter(s)</i>
1/8	33 (0.44)	269 (2.08)	195 (0.94)	341 (1.20)	336 (0.08)
1/16	35 (0.90)	60 (0.71)	243 (5.14)	315 (2.50)	1045 (0.48)
1/32	35 (4.49)	44 (2.46)	291 (30.92)	336 (14.53)	2493 (3.76)
1/64	37 (34.27)	44 (13.02)	273 (387.19)	346 (77.33)	2819 (14.48)
1/128	35 (190.75)	46 (82.38)	263 (1058.93)	348 (519.76)	4431 (83.16)

Table 6: Number of iterations and times for $k' = 3$ (Square domain).

Tables 5 and 6 show the number of iterations and the time in seconds of P-MINRES for all preconditioners discussed in section 5.1, with a tolerance $tol_1 = 10^{-12}$ for the relative residual. In the cases that PCG was used, its relative residual tolerance was $tol_2 = 10^{-6}$, and the preconditioner is simply the diagonal of either $\mathbf{A}$ and Q . Observe that in all cases, the number of iterations is almost constant, except for Diag($\mathbf{A}$, Q). Since the mesh is uniformly refined, these results indicate that the preconditioning strategies are spectrally equivalent.

6.2. 1/8th annulus domain

This example, also available in the GeoPDEs package [36], is used for validation and verification for the Stokes equations discretizations as in [16].

For this example $\Omega \neq \hat{\Omega}$ is one eighth of an annulus and it is parameterized by two types of geometric mappings $\mathbf{F}$, the first one is a NURBS parameterization and the second one is a polar parameterization. We used both parameterizations to test if there is any influence of the geometric mapping on the preconditioner and the solver behavior.

The boundary conditions for this problem is no-slip over $\partial\Omega$, and $\mathbf{f}$ is given (53) for an analytical solution $(\mathbf{u}, p)$ given a priori, that is, a manufactured solution. Figure 2 shows the domain, the velocity magnitude and some streamlines of the analytical velocity field for this example.

Figure 2: Velocity magnitude and some streamlines of the manufactured solution for the 1/8th annulus domain.

h	1/8	1/16	1/32	1/64
$\ \mathbf{u} - \mathbf{u}_h\ _{\mathbf{H}^1}$	3.41e+00	8.38e-01	2.07e-01	5.15e-02
order	—	2.02	2.02	2.01
$\ \mathbf{u} - \mathbf{u}_h\ _{\mathbf{L}^2}$	6.19e-02	7.52e-03	9.52e-04	1.21e-04
order	—	3.04	2.98	2.98
$\ p - p_h\ _{\mathbf{L}^2}$	9.00e-02	7.71e-03	6.62e-04	5.75e-05
order	—	3.55	3.54	3.53

Table 7: Errors and convergence orders for $k' = 2$ (1/8th annulus domain).

Tables 7 and 8 show the approximation errors for the velocity and the pressure of the manufactured solution, and the convergence orders for the polynomial degrees $k' = 2$ and $k' = 3$, respectively, where the geometric mapping $\mathbf{F}$ is a NURBS parametrization and for Nitsche's penalization constant $C_{pen} = 5(k' + 1)$. Again, as predicted by the a priori convergence estimates of [17], the order of convergence of the velocity in $\mathbf{H}^1$ -seminorm is k' , and in the $\mathbf{L}^2$ -norm is $k' + 1$, since the meshes used are not distorted as we already observed in the square case. The approximation errors for the case of the polar parameterization are not shown since up to the second decimal digit all the errors and rates are identical.

Tables 9 and 10 show the number of iterations and the time in seconds of the P-MINRES for all preconditioners discussed in section 5.1, with a tolerance $tol_1 = 10^{-12}$ for the relative residual. As in the last example, the relative residual tolerance for PCG was $tol_2 = 10^{-6}$. Similar to the square domain example, the number of iterations of all strategies, except $\text{Diag}(\mathbf{A}, \mathbf{Q})$, are

h	1/8	1/16	1/32	1/64
$\ \mathbf{u} - \mathbf{u}_h\ _{\mathbf{H}^1}$	3.84e-01	5.20e-02	6.88e-03	8.87e-04
order	—	2.89	2.92	2.95
$\ \mathbf{u} - \mathbf{u}_h\ _{\mathbf{L}^2}$	6.42e-03	4.87e-04	3.34e-05	2.18e-06
order	—	3.72	3.86	3.94
$\ p - p_h\ _{\mathbf{L}^2}$	3.61e-03	1.02e-04	4.28e-06	2.06e-07
order	—	5.15	4.57	4.38

Table 8: Errors and convergence orders for $k' = 3$ (1/8th annulus domain).

	Ideal($\mathbf{A}, Q$)	PCG($\mathbf{A}, Q$)	Ideal($\mathbf{A}$),Diag(Q)	PCG($\mathbf{A}$),Diag(Q)	Diag($\mathbf{A}, Q$)
h	<i>iter(s)</i>	<i>iter(s)</i>	<i>iter(s)</i>	<i>iter(s)</i>	<i>iter(s)</i>
1/8	29 (0.31)	44 (0.27)	89 (0.29)	113 (0.36)	251 (0.05)
1/16	31 (0.47)	38 (0.38)	101 (1.22)	130 (0.82)	525 (0.17)
1/32	31 (2.62)	38 (1.43)	105 (6.80)	140 (3.89)	728 (0.81)
1/64	33 (13.12)	39 (8.02)	103 (32.80)	142 (23.79)	1216 (4.17)
1/128	33 (77.17)	41 (49.83)	97 (190.12)	144 (153.60)	2814 (33.70)

Table 9: Number of iterations and times for $k' = 2$ (1/8th annulus domain).

almost constant, indicating the spectral equivalence of these preconditioners. Besides its increasing number of iterations, as the mesh is refined, for $k' = 2$, Diag($\mathbf{A}, Q$) gave the best time results. Clearly, we do not expect this behavior with more refined meshes. Indeed, for $k' = 3$ and for $h \leq 1/32$, Diag($\mathbf{A}, Q$) is already slower than PCG($\mathbf{A}, Q$).

6.3. Lid-driven cavity flow

For the lid-driven cavity, the computational domain is the unit square, thus, $\widehat{\Omega} = \Omega = (0, 1)^2$. The boundary conditions are no-slip at the bottom, and at the sides. At the top, the tangential

	Ideal($\mathbf{A}, Q$)	PCG($\mathbf{A}, Q$)	Ideal($\mathbf{A}$),Diag(Q)	PCG($\mathbf{A}$),Diag(Q)	Diag($\mathbf{A}, Q$)
h	<i>iter(s)</i>	<i>iter(s)</i>	<i>iter(s)</i>	<i>iter(s)</i>	<i>iter(s)</i>
1/8	31 (0.33)	69 (0.75)	113 (0.60)	203 (0.71)	341 (0.11)
1/16	31 (0.81)	48 (0.60)	175 (3.82)	227 (1.82)	1044 (0.49)
1/32	31 (4.09)	40 (2.20)	191 (20.52)	246 (10.93)	2328 (3.60)
1/64	33 (41.53)	40 (11.94)	175 (231.46)	244 (57.92)	2882 (14.73)
1/128	33 (175.33)	41 (76.24)	163 (668.10)	248 (389.04)	4371 (82.27)

Table 10: Number of iterations and times for $k' = 3$ (1/8th annulus domain).

h	$size(\mathcal{A})$	$nnz(\mathcal{A})$	$size(\mathbf{A}) = n_{\mathbf{u}}$	$nnz(\mathbf{A})$	$size(Q) = n_p$	$nnz(Q)$
1/8	282	17654	182	9170	100	1936
1/16	938	68774	614	36482	324	7056
1/32	3402	271622	2246	145634	1156	26896
1/64	12938	1079750	8582	582050	4356	104976
1/128	50442	4305734	33542	2327330	16900	414736

Table 11: Sizes and numbers of non-zeros for the Stokes systems $\mathcal{A}$, $\mathbf{A}$ and Q for $k' = 2$.

h	$size(\mathcal{A})$	$nnz(\mathcal{A})$	$size(\mathbf{A}) = n_{\mathbf{u}}$	$nnz(\mathbf{A})$	$size(Q) = n_p$	$nnz(Q)$
1/8	343	36222	222	18478	121	4225
1/16	1047	133294	686	69342	361	14641
1/32	3607	510990	2382	268606	1225	54289
1/64	13335	2000590	8846	1057278	4489	208849
1/128	51223	7916622	34062	4195198	17161	819025

Table 12: Sizes and numbers of non-zeros for the Stokes systems $\mathcal{A}$, $\mathbf{A}$ and Q for $k' = 3$.

velocity component is constant, and in our test it is equal to 1. Finally, the kinematic viscosity is $\nu = 1$.

In order to check the correctness of the implementation in Figure 3 we show the streamlines for the lid-driven cavity flow problem. Also, we check an assertion of Evans and Hughes [17], that with a uniform refined mesh with mesh-size $h \leq 1/256$, the discretization can capture the second Moffatt eddy on the lower corners [37].

Since the cavity flow is an established benchmark problem, we decided to take a closer look on the analysis of the block-diagonal preconditioning strategy. First, we considered the sizes and numbers of non-zeros of each matrix involved on the linear solver: the coefficient matrix $\mathcal{A}$ of the Stokes system, the viscosity matrix $\mathbf{A}$ and the pressure mass matrix Q , and for several discretizations, namely, for degrees $k' = 2$ and $k' = 3$, and for five uniform mesh refinements levels. Table 11 shows the values for the case $k' = 2$, and Table 12 for the case $k' = 3$.

The sizes of the matrices for both cases $k' = 2$ and $k' = 3$ are almost the same, but the numbers of non-zeros components almost doubles. Obviously this has an impact on the matrix-vector products, but not on vector-vector operations like a dot product or a vector update. A comparison of the cost of the matrix-vector operation of a scalar Laplacian for continuous and more regular B-spline bases can be found on Collier et al. [38].

Tables 13 and 14 show the numbers of iterations and the times (in seconds) for P-MINRES for the five preconditioning strategies discussed in section 5.1. The notation of the column PCG($\mathbf{A}, Q$)

(a) Velocity magnitude and some streamlines for lid-driven cavity flow.

(b) Zoom in at the left corner showing the primary and the secondary Moffatt eddy.

Figure 3: Lid-driven cavity problem with an identity geometrical map.

	Ideal($\mathbf{A}, Q$)	PCG($\mathbf{A}, Q$)	Ideal($\mathbf{A}$),Diag(Q)	PCG($\mathbf{A}$),Diag(Q)	Diag($\mathbf{A}, Q$)
h	$iter(s)$	$iter(s)$	$iter(s)$	$iter(s)$	$iter(s)$
1/8	29 (0.27)	48/14.29/10 (0.31)	101 (0.32)	136 (0.74)	243 (0.05)
1/16	29 (0.46)	37/25.68/18 (0.36)	135 (1.59)	160 (0.99)	518 (0.16)
1/32	29 (2.34)	37/48.81/34 (1.36)	135 (9.08)	166 (4.54)	696 (0.77)
1/64	29 (11.09)	37/93.65/47.92 (7.51)	131 (40.97)	166 (27.16)	1111 (3.45)
1/128	27 (62.15)	37/183.65/47.19 (43.92)	123 (237.23)	164 (171.19)	2357 (27.55)

Table 13: Number of iterations and times for $k' = 2$ (Lid-driven cavity).

	Ideal($\mathbf{A}, Q$)	PCG($\mathbf{A}, Q$)	Ideal($\mathbf{A}$),Diag(Q)	PCG($\mathbf{A}$),Diag(Q)	Diag($\mathbf{A}, Q$)
h	$iter(s)$	$iter(s)$	$iter(s)$	$iter(s)$	$iter(s)$
1/8	31 (0.31)	92/15.11/9.80 (0.58)	151 (0.75)	303 (1.14)	343 (0.09)
1/16	31 (0.80)	50/27.54/18.98 (0.63)	277 (5.96)	325 (2.59)	1046 (0.49)
1/32	31 (4.05)	38/50.13/35 (2.16)	315 (33.85)	364 (15.75)	2671 (4.04)
1/64	29 (57.06)	38/85.36/67 (10.98)	297 (451.46)	366 (82.07)	2986 (15.37)
1/128	29 (167.07)	38/161.42/116.53 (66.23)	279 (1121.41)	362 (521.79)	4172 (73.64)

Table 14: Number of iterations and times for $k' = 3$ (Lid-driven cavity).

is as follows: the first number refers to the number of iterations of P-MINRES, while the last two are the mean number of iterations for the diagonally preconditioned PCG applied to $\mathbf{A}$ and Q , respectively, where the mean is taken with respect to the number of P-MINRES iterations. In all cases, we used a tolerance of $tol_1 = 10^{-12}$ for the relative residual of P-MINRES, and on the cases that we used the diagonally preconditioned PCG the tolerance for the relative residual of PCG was $tol_2 = 10^{-6}$.

We discuss the $k' = 2$ case first. Except for Diag($\mathbf{A}, Q$), all others yield an almost constant number of iterations with respect to mesh refinement. This is an indication that the preconditioners are spectrally equivalent. Clearly, the Ideal($\mathbf{A}, Q$) strategy gave the best results in terms of number of iterations for the P-MINRES, since we solved the preconditioner systems up to machine precision by the backslash command of MATLAB®, but this is a costly strategy [38]. Despite the significant increase in the number of iterations, as h is decreased, the Diag($\mathbf{A}, Q$) strategy has the best times for all mesh refinement levels. Obviously, we attributed it to the small cost of solving the preconditioner systems. The second best time is that corresponding to PCG($\mathbf{A}, Q$).

It is also interesting to notice that the worst time corresponds to Ideal($\mathbf{A}$),Diag(Q) followed by PCG($\mathbf{A}$),Diag(Q). Also note the increase in the number of iterations for these cases as compared to Ideal($\mathbf{A}, Q$) and PCG($\mathbf{A}, Q$), respectively is significant. We conclude that, while Q and Diag(Q)

Figure 4: Spectrum of the discrete Stokes system and the preconditioned systems for $k' = 2$ and $h = 1/32$. (Lid-driven cavity flow)

are spectrally equivalent, they do not perform efficiently in P-MINRES. Indeed, we computed the spectral bounds $\gamma_Q \approx 0.058$, and $\Gamma_Q \approx 3.33$ numerically, where it is clear that the lower bound is too small, signaling the deficiency of the diagonal approach. Below, we see this is even worse for the case $k' = 3$.

For the Nitsche's penalization constant $C_{pen} = 5(k' + 1) = 15$ in this case, the squared inf-sup constant is approximately $\beta_0^2 \approx 0.1924$ and the continuity constant is approximately $C_b^2 \approx 1.0134$. Computing the upper bound for the negative eigenvalues given by the inclusion estimate (49) of Theorem 1 for $\text{Ideal}(\mathbf{A}), \text{Diag}(Q)$ we obtain $\frac{1 - \sqrt{1 + 4(0.1924)(0.058)}}{2} \approx -0.011$. This indicates that the negative part of the spectrum is tending to zero, which is undesirable for the minmax convergence estimate of MINRES, and indeed this happens (see Table 15).

To corroborate the results of Table 13, we computed numerically $\sigma(\mathcal{A})$ and $\sigma(\mathcal{M}^{-1}\mathcal{A})$ for the preconditioned cases: $\text{Ideal}(\mathbf{A}, Q)$, $\text{Ideal}(\mathbf{A}), \text{Diag}(Q)$ and $\text{Diag}(\mathbf{A}, Q)$. A picture of the spectrum of all cases is shown in Figure 4, and some limiting eigenvalues on Table 15. We disregarded the eigenvalues 0 and 1 because we imposed the zero mean pressure constraint after the solver step by filtering the solution. In our case the Stokes system matrix is singular and has 0 as an eigenvalue of multiplicity one, which is the dimension of its kernel. On the other hand, 1 is always an eigenvalue of multiplicity at least $n_u - n_p$ of $\mathcal{M}^{-1}\mathcal{A}$ when $M_{\mathbf{A}} = \mathbf{A}$.

For the cases $\text{Ideal}(\mathbf{A}, Q)$ and $\text{Ideal}(\mathbf{A}), \text{Diag}(Q)$, the clustered spectra is symmetric around the value $1/2$, as predicted by the inclusion sets estimates (44) and (49). Also, the bounds for (44) are sharp as can be seen by computing its values with $\beta_0^2 \approx 0.1924$ and $C_b^2 \approx 1.0134$, and comparing with the eigenvalues of Table 15.

Preconditioner	$\lambda_{\min}^-$	$\lambda_{\max}^-$	$\lambda_{\min}^+$	$\lambda_{\max}^+$
No precondition.	-4.7522e-4	-1.3939e-6	0.0284	27.2898
Ideal($\mathbf{A}$, Q)	-0.6240	-0.1651	1.1651	1.6240
Ideal($\mathbf{A}$),Diag(Q)	-1.3773	-0.0417	1.0417	2.3773
Diag($\mathbf{A}$, Q)	-1.2698	-0.0375	0.0416	3.0817

Table 15: Limiting eigenvalues of $\sigma(\mathcal{M}^{-1}\mathcal{A})$ ($k' = 2$).

Now, we analyze the case $k' = 3$. Our numerical evidence indicates the spectral equivalence of the preconditioners, except for the case Diag($\mathbf{A}$, Q). Comparing the columns Ideal($\mathbf{A}$, Q) and PCG($\mathbf{A}$, Q) of Table 14 with the same columns of Table 13, we observe that the number of iterations for both discretizations degrees are almost the same, but clearly the case $k' = 3$ takes more time. This is expected since, for this case, the matrix-vector operations take twice the time $k' = 2$ case takes, because of the increase on the number of non-zeros of the systems matrices for $k' = 3$.

In terms of time comparison, we have another picture in this case. Here the strategy PCG($\mathbf{A}$, Q) has the best times for mesh-sizes $h \leq 1/32$, followed by Diag($\mathbf{A}$, Q). Comparing the number of iterations for Diag($\mathbf{A}$, Q) for both degrees, one can find an increase (for some meshes iteration more than double) of the numbers of iterations, that with the additional cost of the matrix-vector operations, led to this result.

The worst time continues to be Ideal($\mathbf{A}$),Diag(Q), followed by PCG($\mathbf{A}$),Diag(Q), that for most mesh-sizes took half the time of the former, indicating again that besides being spectrally equivalent, Diag(Q) misses a lot of information about Q . Indeed, the numerically computed spectral bounds are $\gamma_Q \approx 0.012$ and $\Gamma_Q \approx 4.46$. It is clear that the lower bound is too small, and even smaller than that for $k' = 2$ (as we already anticipated), causing the negative part of the spectrum to become closer to zero, as it is indicated by computing the upper bound $\frac{1 - \sqrt{1 + 4(0.1852)(0.012)}}{2} \approx -0.0022$ for the negative eigenvalues given by the inclusion estimate (49) of Theorem 1 for Ideal($\mathbf{A}$),Diag(Q), where the squared inf-sup constant is approximately $\beta_0^2 \approx 0.1852$, for the Nitsche's penalization constant $C_{pen} = 5(k' + 1) = 20$ in this case (see Figure 5 and Table 16).

For the cases Ideal($\mathbf{A}$, Q) and Ideal($\mathbf{A}$),Diag(Q), the clustered spectra is symmetric around the value $1/2$, as predicted by the inclusion sets estimates (44) and (49). Also, the bounds for (44) are sharp, as can be seen by computing its values with $\beta_0^2 \approx 0.1852$ and $C_b^2 \approx 1.0367$, and comparing with the eigenvalues of Table 16.

For a more comprehensive comparison, we also tested two global strategies to solve the Stokes system. The first one using also an iterative solver, in this case, the Generalized Minimum Residual

Figure 5: Spectrum of the discrete Stokes system and the preconditioned systems for $k' = 3$ and $h = 1/32$. (Lid-driven cavity flow)

Preconditioner	$\lambda_{\min}^-$	$\lambda_{\max}^-$	$\lambda_{\min}^+$	$\lambda_{\max}^+$
No precondition.	-4.7124e-4	-3.3468e-7	0.0221	35.7423
Ideal($\mathbf{A}$, Q)	-0.6343	-0.1597	1.1597	1.6343
Ideal($\mathbf{A}$),Diag(Q)	-1.6357	-0.0107	1.0107	2.6357
Diag($\mathbf{A}$, Q)	-1.6320	-0.0093	0.0299	3.8923

Table 16: Limiting eigenvalues of $\sigma(\mathcal{M}^{-1}\mathcal{A})$ ($k' = 3$).

k'	h	$nnz(L)$	$nnz(U)$	$iter(outer)$	$iter(inner)$	$time(s)$
2	1/64	4617086	8593117	1	7	21.26
	1/128	30400992	51347840	1	9	230.02
3	1/64	8380970	13280191	1	9	64.84
	1/128	47298593	77150221	1	11	501.33

Table 17: Results for P-GMRES(50) (Lid-driven cavity flow).

k'	h	$time(s)$
2	1/64	1.70
	1/128	12.94
	1/256	435.00
3	1/64	5.77
	1/128	36.36
	1/256	2289.70

Table 18: Results for UMFPACK (Lid-driven cavity flow).

Method (GMRES) and the second one using a sparse direct solver, namely the Unsymmetric Multi-frontal Sparse LU Factorization Package (UMFPACK), called in MATLAB® by the backslash command.

We first discuss the iterative solver. Primarily, a reordering of the unknowns was performed using the column approximate minimum degree permutation (COLAMD) algorithm, followed by an ILUT(τ) factorization with $\tau = 10^{-6}$ because, in this case, we could not use an ILU(0) factorization since the algorithm breaks down with a zero diagonal element. Then, the factors L and U were used with a left preconditioned GMRES(50), since, in this case, the preconditioner is not symmetric and positive definite. The tolerance used was 10^{-12} . The results for $k' = 2$ and $k' = 3$ are shown at Table 17. The total time presented on the last column incorporates the time to set up the preconditioner that is, the factorization time, that corresponds approximately to 90% of the total time.

The second global strategy, using UMFPACK, gave the time results shown in Table 18. For the case $k' = 3$, mesh-size $h = 1/256$, UMFPACK used a standard partial pivoting factorization because the problem is ill-conditioned.

For real problems, some computational techniques may be used in advance. One such technique is matrix reordering that in addition to possibly improving data locality, also helps to improve the quality of the preconditioner for purely algebraic strategies, such as incomplete factorizations. For a brief review see [39], Section 2. Also, Collier et al. [38] showed that incomplete factorization

IC(0)-PCG($\mathbf{A}$, Q)							
h	$R + F(s)$	$nnz(IC(\mathbf{A}))$	$nnz(IC(Q))$	$tol_2 = 10^{-3}$		$tol_2 = 10^{-6}$	
				$iter$	$time(s)$	$iter$	$time(s)$
1/8	9.64e-4	4676	1018	39/2.10/1	0.10	34/4.18/1	0.17
1/16	0.0033	18548	3690	40/3.35/1	0.16	36/6.89/1	0.18
1/32	0.0125	73940	14026	42/6.14/1	0.51	37/12.65/1	0.71
1/64	0.0471	295316	54666	43/11.05/1	2.60	36/24.06/1	4.16
1/128	0.1824	1180436	215818	44/19.48/1	18.11	37/46.70/1	33.46

Table 19: Number of iterations of P-MINRES, mean number of iterations for the PCG($\mathbf{A}$) and PCG(Q), and total time for $k' = 2$ (Lid-driven cavity).

IC(0)-PCG($\mathbf{A}$, Q)							
h	$R + F(s)$	$nnz(IC(\mathbf{A}))$	$nnz(IC(Q))$	$tol_2 = 10^{-3}$		$tol_2 = 10^{-6}$	
				$iter$	$time(s)$	$iter$	$time(s)$
1/8	0.0024	9350	2173	42/1.95/1	0.18	34/3.56/1	0.18
1/16	0.0074	35014	7501	40/2.55/1	0.18	37/5.11/1	0.21
1/32	0.0283	135494	27757	40/4.48/1	0.61	37/9.14/1	0.85
1/64	0.1053	533062	106669	43/7.63/1	2.96	37/17.41/1	4.68
1/128	0.4038	2114630	418093	44/13.80/1	21.06	38/33.66/1	37.81

Table 20: Number of iterations of P-MINRES, mean number of iterations for the PCG($\mathbf{A}$) and PCG(Q), and total time for $k' = 3$ (Lid-driven cavity).

with zero fill-in performs well as a preconditioner for the conjugate gradient method for isogeometric discretizations of a Laplace problem. Moreover, the incomplete factorization preconditioning presented p -scalability, that is, scalability under polynomial refinement for C^{p-1} bases, but no spectral equivalence with mesh refinement.

In this context, our last numerical experiment is as follows: first we performed a separate reordering of the matrices $\mathbf{A}$ and Q using the reverse Cuthill-McKee (RCM) algorithm, then we perform an incomplete Cholesky factorization with zero fill-in, IC(0), of both $\mathbf{A}$ and Q . Finally, the factors were used as preconditioners for the PCG method to solve the preconditioner systems with $\mathbf{A}$ and Q for P-MINRES. We also relaxed the relative residual tolerance of PCG, hence we tested with $tol_2 = 10^{-6}$ as above, and with $tol_2 = 10^{-3}$. The results are shown on Tables 19 and 20.

As can be seen by the second column, $R + F(s)$, of Tables 19 and 20, that measures the time spent on the reordering and factorization steps, these procedures worked extremely fast in both

cases, when compared to the overall solver time. With respect to the number of non-zeros of the factors for $\mathbf{A}$ and Q , one can see that, for both cases and all mesh-sizes, they have almost half the number of non-zeros of $\mathbf{A}$ and Q themselves. Again we see an almost constant number of iteration for P-MINRES for both $k' = 2$ and $k' = 3$, and here also for both $tol_2 = 10^{-3}$ and $tol_2 = 10^{-6}$, indicating a spectral equivalence with mesh refinement.

Moreover, comparing the number of P-MINRES iterations of Table 13 and Table 16, column $PCG(\mathbf{A}, Q)$, with those for $IC(0)-PCG(\mathbf{A}, Q)$ in Tables 19 and 20, they are almost the same, but obviously for the incomplete Cholesky case, that offers a better preconditioning, the number of internal iterations of $PCG(\mathbf{A})$ and $PCG(Q)$ were considerably reduced. Surprisingly, $PCG(Q)$ converged in one iteration for all mesh-sizes and tolerances, showing that reordering by RCM and incomplete Cholesky offers a very good preconditioning for the pressure mass matrix Q . Comparing the total time for $PCG(\mathbf{A}, Q)$ and $IC(0)-PCG(\mathbf{A}, Q)$ ($tol_2 = 10^{-6}$) we observe a reduction of more than 10 seconds for $k' = 2, h = 1/128$ and almost 20 seconds for $k' = 3, h = 1/128$.

The relaxation of the relative residual tolerance for PCG also improved the total time for both $k' = 2$ to $k' = 3$. We can observe that the number of iterations of P-MINRES increased a little, on the other hand, the mean number of inner iterations of $PCG(\mathbf{A})$ decreased, causing an overall decrease in time as compared to the case where $tol_2 = 10^{-6}$.

Finally, we observe that the $IC(0)-PCG(\mathbf{A}, Q)$ preconditioning strategy gave the best time results of all strategies, losing only to the sparse direct solver when applied to $k' = 2$, and the mesh-sizes $h = 1/64$ and $h = 1/128$. That is why we tested additionally the mesh-size $h = 1/256$, for both $IC(0)-PCG(\mathbf{A}, Q)$ and the direct solver, where we see that $IC(0)-PCG(\mathbf{A}, Q)$ with PCG tolerance of 10^{-3} performed almost 4 times faster. Also, motivated by the excellent results of [40], where relaxed PCG tolerances were used, we ran the case $k' = 3, h = 1/256$ with a PCG tolerance of 10^{-2} and the total time was 153.97 seconds, which is bigger than our best time, 133.48 seconds with PCG tolerance of 10^{-3} .

7. Conclusion and future work

Divergence-conforming B-spline discretizations are based on the Isogeometric discrete differential forms concept. In addition to being inf-sup stable, they are also pointwise divergence-free, a feature that is not easily achieved by mixed inf-sup stable elements, nor for stabilized ones. Their mathematical properties, presented by Evans and Hughes in [17–19], highlight their usefulness for viscous incompressible flows analyses. As usual, divergence-conforming discretizations require the solution of a linear system, and the efficient solution of such systems is of fundamental importance. In this paper, we analyze the performance of block-diagonal preconditioners, as introduced by Wathen and Silvester in [22, 26, 27], for divergence-conforming discretizations, applied to the Stokes problem.

We have shown our block selection that the theoretical bounds for the spectra, derived by Wathen and Silvester [26, 27] in the context of classical elements, also hold for divergence-conforming discretizations. One of the ingredients, in the block-diagonal preconditioning strategy, is the proper approximation of the pressure mass matrix $\mathbf{Q}$. We have shown that, for divergence-conforming discretizations, the approximation of $\mathbf{Q}$ by taking its diagonal entries generates a poor preconditioner which degrades as the polynomial order of approximation grows.

Another ingredient is the approximation of the viscosity matrix $\mathbf{A}$. As the lid-cavity flow results have shown, the use of iterative solvers to approximate $\mathbf{A}$, such as preconditioned conjugate gradients, performed well for both polynomial degrees $k' = 2$ and $k' = 3$, yielding the smallest time results for fine meshes. Also, we have shown that reordering the unknowns with zero fill-in Incomplete Cholesky factorization as a preconditioner, for both $\mathbf{A}$ and $\mathbf{Q}$, with relaxed inner relative residual tolerances, yields very good preconditioners. Nevertheless, there is still room for improvements, since the mean number of iterations of the inner preconditioned conjugate gradients applied to $\mathbf{A}$ is not spectrally equivalent with mesh refinement.

Several investigations may unfold. Since we only tested two-dimensional problems, we feel that the performance evaluation of Krylov solvers and block preconditioning strategies for divergence-conforming discretizations applied to larger problems, followed by a scalability analysis, must be performed. Another aspect that we will pursue is the coupling of incompressible flows, discretized by divergence-conforming spaces, with transport problems, to evaluate the importance of the pointwise divergence-free property.

- [1] T. Hughes, J. Cottrell, Y. Bazilevs, Isogeometric analysis: CAD, finite elements, NURBS, exact geometry, and mesh refinement, *Computer Methods in Applied Mechanics and Engineering* 194 (2005) 4135–4195.
- [2] Y. Bazilevs, L. Beirão Da Veiga, J. Cottrell, T. Hughes, G. Sangalli, Isogeometric analysis: approximation, stability and error estimates for h -refined meshes., *Mathematical Models and Methods in Applied Sciences* 16 (7) (2006) 1031–1090. doi:10.1142/S0218202506001455.
- [3] G. Sangalli, L. Beirão da Veiga, A. Buffa, J. Rivas, Some estimates for h - p - k -refinement in isogeometric analysis, *Numerische Mathematik* 118 (2011) 271–305.
- [4] J. Cottrell, A. Reali, Y. Bazilevs, T. Hughes, Isogeometric analysis of structural vibrations, *Computer Methods in Applied Mechanics and Engineering* 195 (41-43) (2006) 5257–5296. doi:10.1016/j.cma.2005.09.027.
- [5] Y. Bazilevs, V. Calo, Y. Zhang, T. Hughes, Isogeometric fluid–structure interaction analysis with applications to arterial blood flow, *Computational Mechanics* 38 (4-5) (2006) 310–322. doi:10.1007/s00466-006-0084-3.

- [6] Y. Bazilevs, V. Calo, J. Cottrell, T. Hughes, A. Reali, G. Scovazzi, Variational multi-scale residual-based turbulence modeling for large eddy simulation of incompressible flows, *Computer Methods in Applied Mechanics and Engineering* 197 (1-4) (2007) 173–201. doi:10.1016/j.cma.2007.07.016.
- [7] Y. Zhang, Y. Bazilevs, S. Goswami, C. L. Bajaj, T. Hughes, Patient-specific vascular NURBS modeling for isogeometric analysis of blood flow., *Computer Methods in Applied Mechanics and Engineering* 196 (29-30) (2007) 2943–2959. doi:10.1016/j.cma.2007.02.009.
- [8] H. Gómez, V. Calo, Y. Bazilevs, T. Hughes, Isogeometric analysis of the Cahn–Hilliard phase-field model, *Computer Methods in Applied Mechanics and Engineering* 197 (49-50) (2008) 4333–4352. doi:10.1016/j.cma.2008.05.003.
- [9] Y. Bazilevs, V. Calo, T. Hughes, Y. Zhang, Isogeometric fluid-structure interaction: theory, algorithms, and computations, *Computational Mechanics* 43 (1) (2008) 3–37. doi:10.1007/s00466-008-0315-x.
- [10] D. Benson, Y. Bazilevs, M. Hsu, T. Hughes, Isogeometric shell analysis: The Reissner–Mindlin shell, *Computer Methods in Applied Mechanics and Engineering* 199 (5-8) (2010) 276–289. doi:10.1016/j.cma.2009.05.011.
- [11] Y. Bazilevs, C. Michler, V. Calo, T. 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.
- [12] I. Akkerman, Y. Bazilevs, C. Kees, M. Farthing, Isogeometric analysis of free-surface flow, *Journal of Computational Physics* 230 (11) (2011) 4137–4152. doi:10.1016/j.jcp.2010.11.044.
- [13] 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.
- [14] A. Buffa, J. Rivas, G. Sangalli, R. Vázquez, Isogeometric discrete differential forms in three dimensions, *SIAM Journal Numerical Analysis* 49 (2) (2011) 818–844.
- [15] D. Arnold, R. Falk, R. Winther, Finite element exterior calculus: from Hodge theory to numerical stability, *Bulletin American Mathematical Society* 47 (2010) 281–354.
- [16] 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/fld.2337.

- [17] J. Evans, T. 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.
- [18] J. Evans, T. 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.
- [19] J. Evans, T. Hughes, Isogeometric divergence-conforming B-splines for the unsteady Navier-Stokes equations, *Journal of Computational Physics* 241 (0) (2013) 141–167.
- [20] H. Elman, Multigrid and Krylov subspace methods for the discrete Stokes equations, *International Journal for Numerical Methods in Engineering* 22 (1994) 755–770.
- [21] J. Peters, V. Reichelt, A. Reusken, Fast iterative solvers for discrete Stokes equations, *SIAM Journal on Scientific Computing* 27 (2) (2005) 646–666. doi:10.1137/040606028.
- [22] H. Elman, D. Silvester, A. Wathen, *Finite Elements And Fast Iterative Solvers: With Applications in Incompressible Fluid Dynamics*, Numerical Mathematics And Scientific Computation, Oxford University Press, 2005.
- [23] M. Benzi, G. Golub, J. Liesen, Numerical solution of saddle point problems, *Acta Numerica* 14 (2005) 1–137. doi:10.1017/S0962492904000212.
- [24] C. Paige, M. Saunders, Solution of sparse indefinite systems of linear equations, *SIAM Journal on Numerical Analysis* 12 (4) (1975) 617–629. doi:10.1137/0712047.
- [25] C. Burstedde, G. Stadler, L. Alisic, L. C. Wilcox, E. Tan, M. Gurnis, O. Ghattas, Large-scale adaptive mantle convection simulation, *Geophysical Journal International*-doi:10.1093/gji/ggs070.
- [26] 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.
- [27] A. Wathen, D. Silvester, Fast iterative solution of stabilised Stokes systems part II: using general block preconditioners, *SIAM Journal on Numerical Analysis* 31 (5) (1994) 1352–1367. doi:10.1137/0731070.
- [28] L. A. Piegl, W. Tiller, *The NURBS book*, 2nd Edition, Springer, 1996.
- [29] T. Hughes, J. Cottrell, Y. Bazilevs, *Isogeometric Analysis: Toward integration of CAD and FEA*, Wiley, 2009.

- [30] F. Brezzi, On the existence, uniqueness and approximation of saddle-point problems arising from Lagrangian multipliers, *ESAIM: Mathematical Modelling and Numerical Analysis - Modélisation Mathématique et Analyse Numérique* 8 (R2) (1974) 129–151.
- [31] Y. Bazilevs, C. Michler, V. Calo, T. 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.
- [32] Y. Bazilevs, T. 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.
- [33] 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.
- [34] J. Evans, Divergence-free b-spline discretizations for viscous flows, Ph.D. thesis, ICES - UT Texas (2011).
- [35] A. Coutinho, J. Alves, L. Landau, Comparison of Lanczos and conjugate gradients for the element-by-element solution of finite-element equations on the IBM-3090 vector computer, *Computer & Structures* 39 (1-2) (1991) 47–55.
- [36] C. de Falco, A. Reali, R. Vázquez, GeoPDEs: A research tool for Isogeometric Analysis of PDEs, *Advances in Engineering Software* 42 (12) (2011) 1020–1034. doi:10.1016/j.advengsoft.2011.06.010.
- [37] H. Moffatt, Viscous and resistive eddies near a sharp corner, *Journal of Fluid Mechanics* 18 (1964) 1–18.
- [38] N. Collier, D. Pardo, L. Dalcin, M. Paszynski, V. 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.
- [39] J. Camata, A. Rossa, A. Valli, L. Catabriga, G. Carey, A. Coutinho, Reordering and incomplete preconditioning in serial and parallel adaptive mesh refinement and coarsening flow solutions, *International Journal for Numerical Methods in Fluids* 69 (4) (2012) 802–823. doi:10.1002/fld.2614.
- [40] M. Kronbichler, T. Heister, W. Bangerth, High accuracy mantle convection simulation through modern numerical methods, *ACM Trans. Math. Softw.* 38 (2). doi:10.1145/2049673.2049678.

Dr. Adriano Cortes is a post-doctoral fellow at King Abdullah University of Science and Technology (KAUST) since 2013 under the supervision of Prof. Victor M. Calo. He obtained his D.Sc. in computational mechanics from the High Performance Computing Center and the Alberto Luiz Coimbra Institute for Graduate Studies and Research in Engineering (COPPE) - Federal University of Rio de Janeiro (Rio de Janeiro, Brazil) in 2013 under the supervision of Prof. Alvaro L.G.A. Coutinho. His research interests are in computational mechanics and high-performance computing.

Alvaro L.G.A. Coutinho is the Director of the High Performance Computing Center and a Professor at the Department of Civil Engineering in The Alberto Luiz Coimbra Institute for Graduate Studies and Research in Engineering (COPPE), The Federal University of Rio de Janeiro, Brazil; Coordination and participation in over 80 industry projects. Recipient of the IBM Faculty Partnership Award, 2001; Recipient of the Giulio Massarani Academic Award, COPPE, 2007; Organizer of National and International conferences, training workshops and short courses; Recipient of IACM Fellow Award, 2012. Editorial Advisory Board, International Journal for Numerical Methods in Fluids; Associate Editor, Revista Internacional de Métodos Numéricos para Cálculo y Diseño en Ingeniería. Dissertations Directed, 24 Ph.D. and 25 M.Sc., 92 Journal papers, 250 Conference papers. Cited in Web of Science, November 11, 2013: Total citations, 443, h-index: 13. Scopus, November 11, 2013: Total citations, 496, h-index: 11. Google Scholar, November 11, 2103: Total citations, 1145, h-index: 19.

Recent projects: “Research on Simulation of Geological Processes on High Performance Computers”, Network on Basin Modeling, Brazilian Petroleum Agency and PETROBRAS, 2011–2013; “Hoscar-High Performance Computing and Scientific data management driven by highly demanding applications, CNPq, INRIA; “Finite Element Simulator for Complex Free-Surface Problems: Extensions and New Engineering Challenges”, PETROBRAS 2011–2013; “High Performance Computing Infrastructure for the GRADE-BR Node at COPPE/UFRJ”, Thematic Network on Scientific Computing and Visualization, Brazilian Petroleum Agency and PETROBRAS, 2008–2012.

Dr. Lisandro Dalcin is a researcher at King Abdullah University of Science and Technology (KAUST) since 2014 and National Council for Scientific and Technical Research of Argentina (CONICET) since 2010. He obtained his PhD in Engineering from Universidad Nacional del Litoral (Santa Fe, Argentina) in 2008. His research interests are in scientific computing in distributed memory architectures, medium to large-scale numerical simulation, and development of programming tools mixing Python and C/C++/Fortran.

Dr. Victor Manuel Calo is an Associate Professor in Applied Mathematics & Computational Science and Earth Science & Engineering, and is the co-director of the SRI-Center for Numerical Porous Media. Dr. Calo is a highly cited researcher who is actively involved in disseminating knowledge: Dr. Calo has authored over 100 peer-reviewed publications. In addition, in the last two years he has given more than 30 invited presentations and keynotes at conferences and seminars, and organized 12 mini-symposia at international conferences. Dr. Calo holds a professional engineering degree in Civil Engineering from the University of Buenos Aires. He received a master's in Geomechanics and a doctorate in Civil and Environmental Engineering from Stanford University. Dr. Calo's research Interests include modeling and simulation of geomechanics, fluid dynamics, flow in porous media, phase separation, fluid-structure interaction, solid mechanics, and high-performance computing.

Highlights:

1. The paper explores block-diagonal preconditioners for the Stokes (saddle-point) system discretized by B-splines divergence-conforming spaces, evaluating the performance of several options for each block.
2. The paper reviews the inf-sup and the point-wise divergence free properties of the B-splines divergence-conforming velocity/pressure spaces pair.
3. Also, it explains the implementation on GeoPDE (a Matlab/Octave package that implements Isogeometric Analysis) of Nitsche's method is used to impose Dirichlet boundary condition weakly. Our implementation was incorporated by GeoPDE's developers, and is currently part of GeoPDE package.