

Preconditioned Mixed Spectral Element Methods for Elasticity and Stokes Problems

Luca F. Pavarino*
Department of Mathematics
Università di Pavia, Via Abbiategrasso 209
27100 Pavia, ITALY
pavarino@dragon.ian.pv.cnr.it

Abstract

Preconditioned iterative methods for the indefinite systems obtained by discretizing the linear elasticity and Stokes problems with mixed spectral elements in three dimensions are introduced and analyzed. The resulting stiffness matrices have the structure of saddle point problems with a penalty term, which is associated with the Poisson ratio for elasticity problems or with stabilization techniques for Stokes problems. The main results of this paper show that the convergence rate of the resulting algorithms is independent of the penalty parameter, the number of spectral elements N and mildly dependent on the spectral degree n via the inf-sup constant. The preconditioners proposed for the whole indefinite system are block-diagonal and block-triangular. Numerical experiments presented in the final section show that these algorithms are a practical and efficient strategy for the iterative solution of the indefinite problems arising from mixed spectral element discretizations of elliptic systems.

*This work was supported by the National Science Foundation under Grant NSF-CCR-9503408 and by the National Aeronautics and Space Administration under NASA Contract No. NAS1-19480 while the author was in residence at the Institute for Computer Applications in Science and Engineering (ICASE), NASA Langley Research Center, Hampton, VA 23681-0001.

1 Introduction

The goal of this paper is to introduce and analyze some preconditioned iterative methods for the large indefinite linear systems arising from the mixed spectral discretization of the linear elasticity system and the limiting Stokes problem in three dimensions. Standard finite element discretizations of elasticity problems can suffer from the phenomenon of locking when the Poisson ratio tends to $1/2$ (almost incompressible case); see Babuška and Suri [3]. This means that the convergence rate of the finite element method deteriorates when ν approaches $1/2$. Moreover, the resulting linear system, even though symmetric and positive definite, has a condition number that goes to infinity when the Poisson ratio tends to $1/2$. Both problems can be overcome by using a mixed finite element formulation and rewriting the problem as a saddle point problem with a penalty term; see Brezzi and Fortin [12]. The penalty term depends on the Poisson ratio for elasticity problems or on stabilization parameters for Stokes problems. By carefully choosing the finite element spaces in order to satisfy the inf-sup condition, we obtain a convergent method. The stiffness matrix is symmetric and indefinite.

In recent years, several iterative methods have been proposed and studied in the case of low-order h -version finite elements, such as Uzawa's algorithm (see Elman and Golub [17], Bramble, Pasciak and Vassilev [10]), multigrid (see Verfürth [37], Wittum [39], Braess and Blömer [8], Brenner [11]), preconditioned conjugate gradient (see Bramble and Pasciak [9]), and preconditioned conjugate residuals (PCR) (see Rusten and Winther [34], Silvester and Wathen [35], [38], Klawonn [22], [24]). Elman [16] has carried out a careful comparison of the performance of four of these methods applied to Stokes problems in two dimensions.

Here we consider instead spectral element discretizations. For a general introduction to spectral methods, we refer to Canuto, Hussaini, Quarteroni, and Zang [13], Bernardi and Maday [7], and Funaro [20]. See also Babuška and Suri [4] for the related p -version of the finite element method. Already for scalar problems, the stiffness matrices obtained by spectral and p -version finite elements are less sparse and more ill-conditioned than those obtained with h -version finite elements. The construction and analysis of efficient preconditioned iterative methods is therefore more challenging. We refer to Pavarino and Widlund [31], Casarin [14] and to the references therein for an overview of recent results based on domain decomposition techniques for elliptic scalar problems. In the context of spectral elements for Stokes and Navier-Stokes problems, iterative methods have been studied in Maday, Meiron, Patera and Rønquist [25], Maday, Patera and Rønquist [26], Fischer and Rønquist [18], and Rønquist [33]. The methods proposed by these authors are based on conjugate gradient iterations on the reduced Schur complement of the discrete Stokes matrix involving only the pressure unknowns. In the context of linear elasticity and p -version finite elements, iterative methods have been studied by Mandel; see [28], [29] and the references therein. These works are based on the pure displacement formulation and are concerned mainly with compressible materials.

In this paper, we propose solving the whole indefinite system arising from the mixed spectral element discretization using the results in Klawonn [24], [23] and extending his h -version study to spectral elements. We will consider both block-diagonal and block-

triangular preconditioners. In the first case, the preconditioned operator is symmetric indefinite and we can use the PCR method. In the second case, the preconditioned operator is no longer symmetric and we will consider three iterative methods: GMRES without restart, Bi-CGSTAB and QMR; see Barret et al. [5] and Freund, Golub, and Nachtigal [19] for an introduction to these methods.

The main result of this paper is that the convergence rate of the proposed algorithms is independent of the penalty parameter ν , the number of spectral elements N and mildly dependent on the spectral degree n via the inf-sup constant. This is due to the dependence on n of the inf-sup constant for our choices of spectral element spaces in the discretization. We will consider two choices of mixed spectral spaces, known as the $Q_n - Q_{n-2}$ and $Q_n - P_{n-1}$ methods; see Maday, Patera and Rønquist [26] and Stenberg and Suri [36]. Several numerical experiments reported in the final section confirm this result and show that the number of iterations required by the triangular preconditioner is much smaller than the number of iterations required by the block-diagonal preconditioner, while its cost is only marginally higher.

The organization of the paper is as follows. Section 2 introduces the elasticity and Stokes systems in both the pure displacement and mixed formulation. In Section 3, two mixed spectral element discretizations are introduced. These are based on the Gauss-Legendre-Lobatto (GLL) quadrature, briefly reviewed in Section 3.1. The known results for the associated inf-sup constants are reported in Section 3.2. The preconditioned iterative methods and the main convergence results are introduced in Section 4, with the block-diagonal preconditioner in 4.1 and the triangular preconditioner in 4.2. In Section 5, we report the results of several numerical experiments in three dimensions, both with block-diagonal and triangular preconditioners, with one and many spectral elements.

2 The linear elasticity and Stokes systems

We consider a polyhedral domain $\Omega \subset R^3$, fixed along a subset of its boundary Γ_0 , subject to a surface force of density \mathbf{g} along $\Gamma_1 = \partial\Omega - \Gamma_0$ and subject to an external force \mathbf{f} . Let \mathbf{V} be the Sobolev space $\mathbf{V} = \{\mathbf{v} \in H^1(\Omega)^3 : \mathbf{v}|_{\Gamma_0} = 0\}$. The linear elasticity problem (pure displacement model) consists in finding the displacement $\mathbf{u} \in \mathbf{V}$ of the domain Ω such that:

$$2\mu \int_{\Omega} \epsilon(\mathbf{u}) : \epsilon(\mathbf{v}) \, dx + \lambda \int_{\Omega} \operatorname{div} \mathbf{u} \operatorname{div} \mathbf{v} \, dx = \langle \mathbf{F}, \mathbf{v} \rangle \quad \forall \mathbf{v} \in \mathbf{V}, \quad (1)$$

where λ and μ are the Lamè constants, $\epsilon_{ij}(\mathbf{u}) = \frac{1}{2}(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i})$ is the linearized stress tensor, and the inner products are defined as

$$\epsilon(\mathbf{u}) : \epsilon(\mathbf{v}) = \sum_{i=1}^3 \sum_{j=1}^3 \epsilon_{ij}(\mathbf{u}) \epsilon_{ij}(\mathbf{v}), \quad \langle \mathbf{F}, \mathbf{v} \rangle = \int_{\Omega} \sum_{i=1}^3 f_i v_i \, dx + \int_{\Gamma_1} \sum_{i=1}^3 g_i v_i \, ds.$$

Almost incompressible materials are characterized by very large values of λ , or, in terms of the Poisson ratio $\nu = \frac{\lambda}{2(\lambda + \mu)}$, by ν close to 1/2. When low order h -version finite elements are used in the discretization of (1), the locking phenomenon causes a

deterioration of the convergence rate as $h \rightarrow 0$; see Babuška and Suri [3]. If the p -version is used instead, locking in \mathbf{u} is eliminated, but it could still be present in quantities of interest such as $\lambda \operatorname{div} \mathbf{u}$. Moreover, the stiffness matrix obtained by discretizing the pure displacement model (1) has a condition number that goes to infinity when $\nu \rightarrow 1/2$. Therefore, the convergence rate of iterative methods deteriorates rapidly as the material becomes almost incompressible.

These locking problems can be overcome by introducing the new variable $p = -\lambda \operatorname{div} \mathbf{u} \in L^2(\Omega) = W$ and by rewriting the pure displacement problem in the following mixed formulation (see Brezzi and Fortin [12]). Find $(\mathbf{u}, p) \in \mathbf{V} \times W$ such that

$$\begin{cases} 2\mu \int_{\Omega} \epsilon(\mathbf{u}) : \epsilon(\mathbf{v}) \, dx - \int_{\Omega} \operatorname{div} \mathbf{v} \, p \, dx = \langle \mathbf{F}, \mathbf{v} \rangle & \forall \mathbf{v} \in \mathbf{V} \\ - \int_{\Omega} \operatorname{div} \mathbf{u} \, q \, dx - \frac{1}{\lambda} \int_{\Omega} p q \, dx = 0 & \forall q \in W. \end{cases} \quad (2)$$

With the standard definitions $a^\epsilon(\mathbf{u}, \mathbf{v}) = 2\mu \int_{\Omega} \epsilon(\mathbf{u}) : \epsilon(\mathbf{v}) \, dx$, $b(\mathbf{v}, q) = - \int_{\Omega} \operatorname{div} \mathbf{v} q \, dx$, and $c(p, q) = \int_{\Omega} p q \, dx$, this problem takes the form: find $(\mathbf{u}, p) \in \mathbf{V} \times W$ such that

$$\begin{cases} a^\epsilon(\mathbf{u}, \mathbf{v}) + b(\mathbf{v}, p) = \langle \mathbf{F}, \mathbf{v} \rangle & \forall \mathbf{v} \in \mathbf{V} \\ b(\mathbf{u}, q) - \frac{1}{\lambda} c(p, q) = 0 & \forall q \in W. \end{cases} \quad (3)$$

When $\lambda \rightarrow \infty$ (or, equivalently, $\nu \rightarrow 1/2$), we obtain from (2) the limiting problem for incompressible elasticity:

$$\begin{cases} a^\epsilon(\mathbf{u}, \mathbf{v}) + b(\mathbf{v}, p) = \langle \mathbf{F}, \mathbf{v} \rangle & \forall \mathbf{v} \in \mathbf{V} \\ b(\mathbf{u}, q) = 0 & \forall q \in W. \end{cases} \quad (4)$$

In case of homogeneous Dirichlet boundary conditions on the whole boundary $\partial\Omega$, the pressure will have zero mean value, so we define $W = L_0^2(\Omega)$. In this case, problem (2) can equivalently be written in the following way (see Brezzi and Fortin [12]):

$$\begin{cases} a(\mathbf{u}, \mathbf{v}) + b(\mathbf{v}, p) = \langle \mathbf{F}, \mathbf{v} \rangle & \forall \mathbf{v} \in \mathbf{V} \\ b(\mathbf{u}, q) - \frac{1}{\lambda + \mu} c(p, q) = 0 & \forall q \in W, \end{cases} \quad (5)$$

where here $a(\mathbf{u}, \mathbf{v}) = \mu \int_{\Omega} \nabla \mathbf{u} : \nabla \mathbf{v} \, dx$. The limiting problem when $\lambda \rightarrow \infty$ is the Stokes system describing the velocity \mathbf{u} and pressure p of a fluid of viscosity μ :

$$\begin{cases} a(\mathbf{u}, \mathbf{v}) + b(\mathbf{v}, p) = \langle \mathbf{F}, \mathbf{v} \rangle & \forall \mathbf{v} \in \mathbf{V} \\ b(\mathbf{u}, q) = 0 & \forall q \in W. \end{cases} \quad (6)$$

The penalty term in (5) might be present due to stabilization techniques.

3 Mixed spectral element methods

For an introduction to spectral elements see Patera [30], Maday and Patera [27], Maday, Patera and Rønquist [26] and the references therein.

Let Ω_{ref} be the reference cube $[-1, 1]^3$ and let $Q_n(\Omega_{ref})$ be the set of polynomials on Ω_{ref} of degree n in each variable and $P_n(\Omega_{ref})$ be the set of polynomials on Ω_{ref} of total degree n . Let the domain Ω be decomposed into a finite element triangulation $\bigcup_{i=1}^N \Omega_i$

of nonoverlapping elements. Each Ω_i is the affine image of the reference cube $\Omega_i = F_i(\Omega_{ref})$, where F_i is an affine mapping. We discretize each displacement component by conforming spectral elements, i.e. by continuous, piecewise polynomials of degree n :

$$\mathbf{V}^n = \{\mathbf{v} \in \mathbf{V} : v_k|_{\Omega_i} \circ F_i \in Q_n(\Omega_{ref}), i = 1, \dots, N, k = 1, 2, 3\}.$$

We consider two choices for the discrete pressure space W^n :

$$W_1^n = \{q \in W : q_i \circ F_i \in Q_{n-2}(\Omega_{ref}), i = 1, \dots, N\},$$

$$W_2^n = \{q \in W : q_i \circ F_i \in P_{n-1}(\Omega_{ref}), i = 1, \dots, N\}.$$

The choice W_1^n has been proposed for the Stokes system by Maday, Patera and Rønquist [26] and it is known as the $Q_n - Q_{n-2}$ method. A basis for W_1^n can be constructed by using the tensor-product Lagrangian interpolants associated with the internal GLL nodes, described in the next section in more details.

The second choice corresponds to Method 2 analyzed in Stenberg and Suri [36]. We will call this method $Q_n - P_{n-1}$. For P_{n-1} it is not possible to have a tensorial basis, but other standard bases, common in the p -version finite element literature, can be used.

3.1 Gauss-Lobatto-Legendre (GLL) quadrature and the discrete problem

The efficient evaluation of the multiple integrals of polynomials, involved in our model problem, is based on numerical quadrature at the GLL points. Let $\{\xi_i, \xi_j, \xi_k\}_{i,j,k=0}^n$ be the set of GLL points on the reference cube $[-1, 1]^3$, and let σ_i be the weight associated with ξ_i . Let $l_i(x)$ be the Lagrange interpolating polynomial vanishing at all the GLL nodes except at ξ_i , where it equals one. By tensor product, the basis functions on the reference cube are then defined by

$$l_i(x)l_j(y)l_k(z), \quad 0 \leq i, j, k \leq n.$$

Since every polynomial in $Q_n(\Omega_{ref})$ can be written as

$$u(x, y, z) = \sum_{i=0}^n \sum_{j=0}^n \sum_{k=0}^n u(\xi_i, \xi_j, \xi_k) l_i(x) l_j(y) l_k(z),$$

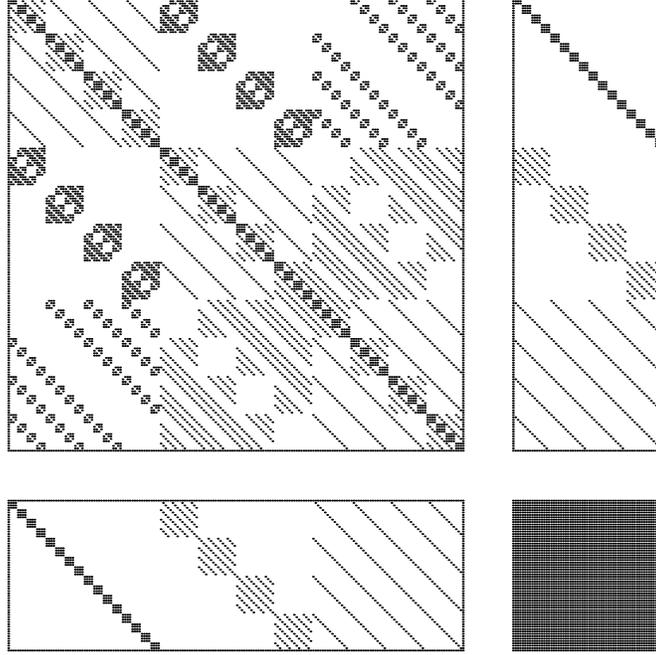
these basis functions form a nodal basis. Each integral of the continuous model (2) is then replaced by GLL quadrature sums. On Ω_{ref}

$$(u, v)_{Q, \Omega_{ref}} = \sum_{i=0}^n \sum_{j=0}^n \sum_{k=0}^n u(\xi_i, \xi_j, \xi_k) v(\xi_i, \xi_j, \xi_k) \sigma_i \sigma_j \sigma_k,$$

and in general on Ω

$$(u, v)_{Q, \Omega} = \sum_{s=1}^N \sum_{i,j,k=0}^n (u \circ F_s)(v \circ F_s) |J_s|(\xi_i, \xi_j, \xi_k) \sigma_i \sigma_j \sigma_k,$$

Figure 1: Sparsity pattern of the stiffness matrix K for model problem (7) discretized with method $Q_n - Q_{n-2}$ on one element, $n = 5$, $\nu = 0.3$



where $|J_s|$ is the determinant of the Jacobian of F_s . The analysis of this discretization technique can be found in Bernardi and Maday [7] and Maday, Patera and Rønquist [26].

The discrete problem obtained from pb. (3) is:

$$\begin{cases} a_Q^\epsilon(\mathbf{u}, \mathbf{v}) + b_Q(\mathbf{v}, p) = \langle \mathbf{F}, \mathbf{v} \rangle_{Q, \Omega} & \forall \mathbf{v} \in \mathbf{V}^n \\ b_Q(\mathbf{u}, q) - \frac{1}{\lambda} c_Q(p, q) = 0 & \forall q \in W^n, \end{cases} \quad (7)$$

where $a_Q^\epsilon(\mathbf{u}, \mathbf{v}) = 2\mu(\epsilon(\mathbf{u}) : \epsilon(\mathbf{v}))_{Q, \Omega}$, $b_Q(\mathbf{v}, q) = -(\text{div} \mathbf{v}, q)_{Q, \Omega}$, $c(p, q) = (p, q)_{Q, \Omega}$. The bilinear forms $b(\cdot, \cdot)$ and $c(\cdot, \cdot)$ are computed exactly by GLL quadrature since the Ω_i are affine images of the reference cube. This system is a saddle point problem with a penalty term and has the following matrix form:

$$Kx = \begin{bmatrix} A & B^T \\ B & -\frac{1}{\lambda}C \end{bmatrix} x = b. \quad (8)$$

The stiffness matrix K is symmetric and indefinite. It is less sparse than the one obtained by low-order finite elements, but is still well-structured. See Figure 1 for the sparsity structure of K . An analogous discrete problem with $C = 0$ is obtained in the incompressible case. For the Stokes problem, the discretization of the equivalent formulations (5) and (6) lead to an analogous block structure, with A consisting of three uncoupled discrete laplacians and with the penalty term in (5) scaled by $1/(\lambda + \mu)$.

3.2 Estimates of the inf-sup constant for spectral elements

The convergence of mixed methods depends not only on the approximation properties of the discrete spaces \mathbf{V}^n and W^n , but also on a stability condition known as the inf-sup (or LBB) condition; see Brezzi and Fortin [12]. For numerical studies of the inf-sup constant of various h -version finite elements, see Bathe and Chapelle [6] and Aristov and Chizhonkov [1]. While many important h -version finite elements for Stokes problems satisfy the inf-sup condition with a constant independent of h , the important spectral elements proposed for Stokes problems, such as the $Q_n - Q_{n-2}$ and $Q_n - P_{n-1}$ methods, have an inf-sup constant that approaches zero as $n^{-(d-1)/2}$ ($d = 2, 3$). This result has been proven for the $Q_n - Q_{n-2}$ method by Maday, Patera and Rønquist [26], where an example is constructed showing that this estimate is sharp. Stenberg and Suri [36] proved the following, more general, result covering both methods.

Theorem 1 (Stenberg and Suri [36]) *Let the spaces \mathbf{V}^n and W^n satisfy assumptions (A1)-(A4) of [36] (satisfied by both our methods). Then for $d = 2, 3$*

$$\sup_{\mathbf{v} \in \mathbf{V}^n \setminus \{0\}} \frac{(\operatorname{div} \mathbf{v}, q)}{\|\mathbf{v}\|_{H^1}} \geq C n^{-(\frac{d-1}{2})} \|q\|_{L^2} \quad \forall q \in W^n,$$

where the constant C is independent of n and q .

For the $Q_n - P_{n-1}$ method, no example is known regarding the sharpness of this estimate. We can rewrite the inf-sup condition in matrix form as

$$q^t B A^{-1} B^t q \geq \beta_0^2 q^t C q \quad \forall q \in W^n, \quad (9)$$

where $\beta_0 = C n^{-(\frac{d-1}{2})}$ is the inf-sup constant of the method; see Brezzi and Fortin [12]. Therefore β_0^2 scales as $\lambda_{\min}(C^{-1} B A^{-1} B^t)$. If β_1 is the continuity constant of B , we have

$$\mathbf{v}^t B^t q \leq \beta_1 (q^t C q)^{1/2} (\mathbf{v}^t A \mathbf{v})^{1/2} \quad \forall \mathbf{v} \in \mathbf{V}^n, \forall q \in W^n. \quad (10)$$

From (9) and (10) it follows that the

$$\beta_0^2 \leq \frac{q^t B A^{-1} B^t q}{q^t C q} \leq \beta_1^2 \quad \forall q \in W^n,$$

for positive constant β_0 and β_1 . For stable h -version finite elements, both β_0 and β_1 are independent of h . Theorem 1 shows that this is no longer the case when spectral elements are used. However, numerical experiments by Maday, Meiron, Patera and Rønquist [25] and [26], have shown that for the $Q_n - Q_{n-2}$ method, for practical values of n (e.g. $n \leq 16$), the dependence of β_0 on n is much weaker. In our numerical experiments in Section 5, we show that the situation is even better for the $Q_n - P_{n-1}$ method. Of course, the trade-off in this case is the loss of a tensorial basis.

4 Preconditioned iterative methods

The indefinite system $Kx = b$ obtained from our spectral element discretization (8), will be solved iteratively by preconditioned Krylov methods for indefinite systems. Two classes of preconditioners will be considered: block-diagonal and triangular.

4.1 Block-diagonal preconditioners

We first consider a block-diagonal preconditioner for K with positive definite blocks \hat{A} and \hat{C} :

$$\hat{D} = \begin{bmatrix} \hat{A} & 0 \\ 0 & \hat{C} \end{bmatrix}. \quad (11)$$

We will denote by D the case with exact blocks $\hat{A} = A$ and $\hat{C} = C$. Interesting choices for \hat{A} are given by h -version finite element discretizations on the GLL mesh or by substructuring domain decomposition methods, where a_0 and a_1 have a polylogarithmic dependence on the spectral degree n (for the scalar case, see Pavarino and Widlund [32] and Casarin [14]). Since the resulting preconditioned system is symmetric, we can use the Preconditioned Conjugate Residual Method (PCR); see Ashby, Manteuffel and Saylor [2] and Hackbusch [21]. See Elman [16] for a short description of the ORTHMIN version of PCR for symmetric indefinite systems.

In his thesis [24], Klawonn considered low-order finite elements and proved an estimate for the condition number of $\hat{D}^{-1}K$, under the following assumptions that \hat{A} and \hat{C} are good preconditioners for A and C respectively:

i) $\exists a_0, a_1 > 0$ such that

$$a_0^2 \mathbf{v}^t \hat{A} \mathbf{v} \leq \mathbf{v}^t A \mathbf{v} \leq a_1^2 \mathbf{v}^t \hat{A} \mathbf{v} \quad \forall \mathbf{v} \in \mathbf{V}^n;$$

ii) \hat{C} is spectrally equivalent to the pressure mass matrix C : $\exists c_0, c_1 > 0$ such that

$$c_0^2 q^t \hat{C} q \leq q^t C q \leq c_1^2 q^t \hat{C} q \quad \forall q \in W^n.$$

Theorem 2 (Klawonn [24], pp. 46-47)

$$\text{cond}(\hat{D}^{-1}K) \leq \frac{\max\{a_1^2, c_1^2\}}{\min\{a_0^2, c_0^2\}} \text{cond}(D^{-1}K)$$

and

$$\text{cond}(D^{-1}K) \leq \frac{1/2 + \sqrt{\beta_1^2 + 1/4}}{-1/2 + \sqrt{\beta_0^2 + 1/4}},$$

where β_0 is the inf-sup constant of the method and β_1 is the continuity constant of B .

Clearly, this abstract result can also be applied to high-order elements. Combining Theorems 1 and 2, we obtain convergence estimates for both methods we have proposed.

Theorem 3 *If K is the stiffness matrix of the discrete system (7) obtained with either the $Q_n - Q_{n-2}$ or the $Q_n - P_{n-1}$ method and \hat{D} is the block-diagonal preconditioner (11), then*

$$\text{cond}(\hat{D}^{-1}K) \leq C\beta_0^{-1} = Cn^{(\frac{d-1}{2})}, \quad d = 2, 3.$$

We remark that the number of iterations of the PCR algorithm applied to an indefinite system is bounded by the condition number of the system (see Hackbusch [21]). This is different from the bounds for conjugate gradient algorithms, where the number of iterations is bounded by the square root of the condition number of the system. Therefore, the number of iterations of our preconditioned algorithm is bounded by $Cn^{(\frac{d-1}{2})}$.

4.2 Triangular preconditioners

An alternative way to precondition the saddle-point problem (8) is provided by the lower and upper triangular preconditioners

$$\hat{T}_L = \begin{bmatrix} \hat{A} & 0 \\ B & \hat{C} \end{bmatrix}, \quad \hat{T}_U = \begin{bmatrix} \hat{A} & B^T \\ 0 & \hat{C} \end{bmatrix}, \quad (12)$$

where \hat{A} and \hat{C} are positive definite matrices. Again, we will denote by T_L and T_U the case with exact blocks $\hat{A} = A$ and $\hat{C} = C$. Since the resulting preconditioned system is no longer symmetric or positive definite, we need to use Krylov methods for general nonsymmetric systems. In particular, we will consider three relatively recent methods: GMRES, Bi-CGSTAB and QMR; see Barret et al. [5] and Freund, Golub and Nachtigal [19]. We remark that each application of the inverse of the triangular preconditioners \hat{T}_L or \hat{T}_U is only marginally more expensive than the block-diagonal preconditioner. In fact, both preconditioners require the solution of a system for \hat{A} and one for \hat{C} . In addition, the triangular preconditioner requires only one application of B (or B^t):

$$\begin{bmatrix} \hat{A} & 0 \\ B & \hat{C} \end{bmatrix}^{-1} \begin{bmatrix} u \\ p \end{bmatrix} = \begin{bmatrix} \hat{A}^{-1} & 0 \\ -\hat{C}^{-1}B\hat{A}^{-1} & \hat{C}^{-1} \end{bmatrix} \begin{bmatrix} u \\ p \end{bmatrix} = \begin{bmatrix} \hat{A}^{-1}u \\ \hat{C}^{-1}(-B\hat{A}^{-1}u + p) \end{bmatrix}. \quad (13)$$

In Klawonn [24], it is first proved a bound for the spectrum σ of the preconditioned operator with exact blocks. The surprising result is that such spectrum is a subset of the positive real axis.

Theorem 4 (Klawonn [24], p. 56)

$$\sigma(T_L^{-1}K) = \sigma(KT_U^{-1}) \subset [\beta_0^2, \beta_1^2 + 1] \cup \{1\}.$$

Combining this result and the estimate of β_0 given in Theorem 1 for our spectral element spaces, we obtain the following result.

Theorem 5 *If K is the stiffness matrix of the discrete system (7) obtained with either $Q_n - Q_{n-2}$ or $Q_n - P_{n-1}$ spectral elements and T is the lower or upper triangular preconditioner (12) with exact blocks, then*

$$\text{cond}(T^{-1}K) \leq C\beta_0^{-2} = Cn^{(d-1)}, \quad d = 2, 3.$$

If GMRES is used to solve our problem, it is possible to prove that the number of iterations required is bounded by the square root of the condition number of $T^{-1}K$; see Klawonn [24], Theorems 5.3 and 5.4. Therefore, we have the same bound $Cn^{(\frac{d-1}{2})}$ as for PCR with block-diagonal preconditioner.

The case of a triangular preconditioner with inexact blocks is studied in Theorem 5.2 in Klawonn [24], pg. 59, under the standard assumptions i) and ii) of the previous section. The estimate provided is analog to the case with exact blocks, but it is more complicated and we refer to [24] for the details.

5 Numerical results

The numerical results are divided into a preliminary section regarding the inf-sup constant and into four sections corresponding to block-diagonal and triangular preconditioners, each divided into single-element and multi-element case. The iterative methods considered are PCR for the block-diagonal preconditioner and GMRES (without restart), Bi-CGSTAB and QMR for the triangular preconditioner. All the computations were performed in MATLAB 4.2 on a SPARCcenter 2000. The model problem considered is (2), discretized with the $Q_n - Q_{n-2}$ or $Q_n - P_{n-1}$ spectral element methods. The resulting discrete systems have a structure as in (8). The implementations of GMRES, Bi-CGSTAB and QMR are the Matlab templates from [5], while the implementation of PCR is the same as in [16]. Except Table 11 showing the discretization errors in the L^2 -norm, all the results reported are iteration counts for the iterative methods considered. The initial guess is always zero and the right-hand side f consists of uniformly distributed random numbers in $[-1,1]$. The stopping criterion is $\|r_i\|_2/\|r_0\|_2 \leq 10^{-6}$, where r_i is the i -th residual. We did not try to optimize any of these routines and in each table, the size of the largest problem we were able to run was determined by the size of the available memory. This was particularly limiting in the multi-element case, where already with four elements, we could run only cases up to $n = 6$. We considered only preconditioners with exact blocks, in order to study the algorithms under the best of circumstances. For the single-element block-diagonal case, we considered also preconditioners with inexact blocks based on piecewise linear finite elements on the GLL mesh.

5.1 The inf-sup constant

We first report in Table 1 a comparison of the spectrum of the matrices $C^{-1}S = C^{-1}BA^{-1}B^t$ associated with the two methods $Q_n - Q_{n-2}$ and $Q_n - P_{n-1}$. Since the inf-sup constant β_0 scales like $\sqrt{\lambda_{min}}$, these results give an indication on the performance of the PCR method reported in the following tables. The first set of results for the $Q_n - Q_{n-2}$ method agree with the 3-d results of Maday, Patera and Rønquist [26]. For these relatively low values of n , λ_{min} scales like n^{-1} and therefore β_0 scales like $n^{-0.5}$, which is better than the value predicted by the theory (n^{-1}). 2-d numerical results in Maday, Meiron, Patera and Rønquist [25] for higher values of n ($16 \leq n \leq 36$) show that the decay of β_0 approaches the theoretical bound, but is still better than the value predicted by the theory ($n^{-0.5}$). The case $n = 10$ could not be run due to memory limitations. The second set of results in Table 1 show that the $Q_n - P_{n-1}$ method has a much better inf-sup constant. From so few values of n , it might look like λ_{min} is bounded away from zero. However, a closer look shows that λ_{min} has now a zig-zag behavior. By separating odd and even values of n , we found that λ_{min} scales approximately like $n^{-0.1}$. Higher values of n are needed in order to understand the asymptotic behavior of λ_{min} for this method. The maximum eigenvalue quickly approaches the same value $\lambda_{max} = 0.65$ for both methods.

We remark that for h -version finite elements, a numerical study of the inf-sup constant seems simpler. In Bathe and Chapelle [6], only three or four values of h are needed to predict the asymptotic behavior of λ_{min} .

Table 1: Inf-sup constant $\beta_0 = \lambda_{min}^{1/2}$

n	$Q_n - Q_{n-2}$			$Q_n - P_{n-1}$		
	$cond(C^{-1}S)$	λ_{max}	λ_{min}	$cond(C^{-1}S)$	λ_{max}	λ_{min}
3	2.1084	0.2284	0.1083	2.1527	0.3611	0.1677
4	7.3040	0.6334	0.0867	2.6771	0.4570	0.1707
5	9.2670	0.6447	0.0695	3.4258	0.5973	0.1743
6	11.2829	0.6500	0.0576	3.7291	0.6097	0.1635
7	13.4537	0.6500	0.0483	3.9161	0.6499	0.1659
8	15.8016	0.6500	0.0411	4.0713	0.6499	0.1596
9	18.3445	0.6500	0.0354	4.0446	0.6500	0.1607
10	-	-	-	4.1653	0.6500	0.1560

Table 2: Condition numbers for $\nu = 0.5$; exact block-diagonal preconditioner

n	$Q_n - Q_{n-2}$		$Q_n - P_{n-1}$	
	$cond(D^{-1}K)$	$cond(K)$	$cond(D^{-1}K)$	$cond(K)$
3	10.827	29.338	7.931	69.613
4	16.259	130.52	8.196	214.63
5	20.040	404.29	8.556	869.92
6	23.977	1,042.1	9.096	3,656.6
7	28.332	2,371.1	9.121	18,427
8	33.038	4,609.2	9.437	86,673
9	38.132	8,845.7	9.382	452,767

Table 2 reports the condition numbers of the preconditioned system $D^{-1}K$: here K is the Stokes matrix obtained for $\nu = 1/2$ and D is the preconditioner with exact blocks. By Theorem 3, these condition numbers scale like the inverse of the respective inf-sup constant. In fact, the results for $Q_n - Q_{n-2}$ clearly show a linear growth with n . The results for $Q_n - P_{n-1}$ are much better and, in comparison, they almost look bounded by a constant. However, by again separating odd and even values of n , the growth still appears linear.

5.2 Block-diagonal preconditioner: single-element case

In Table 3, we report the PCR iteration counts for both methods with exact preconditioner. We followed the PCR implementation of Elman [16], which switches from the ORTHOMIN to the ORTHODIR version to avoid breakdown. In our experiments, this switch often took place for ν near and equal to $1/2$. As in Klawonn [22], the results are uniform in the Poisson ratio ν : for each fixed degree n , the number of PCR iterations is bounded by a constant independent of ν . As the material becomes almost incompressible, the number of iterations tends to a constant which is the number of iterations required by the limiting Stokes problem. As the spectral degree n increases, the number of iterations increases, in agreement with Theorem 3. This effect is less pronounced for compressible materials (for $\nu = 0.3$ and 0.4 the number of iterations stays practically constant), but becomes more important near or at the incompressible limit. This is particularly true for the $Q_n - Q_{n-2}$ method, where the growth of the number of iterations for $\nu = 1/2$ is clearly linear. The results for $Q_n - P_{n-1}$ are better, as expected from the better inf-sup constant of this method. In this case, it is even hard to read a linear growth from the table, which has large constant blocks. Graphs showing the convergence history of both methods for $n = 8$ and $\nu = 1/2$ can be found in Figure 2. In Table 4, the same results are reported for the equivalent formulation (5) instead of (2). This implies that block A in K now consists of three uncoupled discrete laplacians, one for each component of \mathbf{u} . The problem is somewhat harder to solve and PCR takes more iterations than in each corresponding case of the previous table (except $n = 3$ for $Q_n - P_{n-1}$). Again, the results for $Q_n - P_{n-1}$ are consistently better than those for $Q_n - Q_{n-2}$. Now a linear growth with n is clear for both methods (for $Q_n - P_{n-1}$ the odd and even values of n have to be separated).

Next, we consider an inexact preconditioner by choosing as u-block \hat{A} the Q_1 finite element stiffness matrix obtained by discretizing the term $\int_{\Omega} \nabla \mathbf{u} : \nabla \mathbf{v} \, dx$ on the GLL grid. In the scalar case, it is well-known that such matrix is spectrally equivalent to the stiffness matrix obtained by spectral discretization; see Deville and Mund [15]. In Table 5, we study numerically the quality of such preconditioner in three dimensions for a Poisson problem on the reference cube with homogeneous Dirichlet boundary conditions. In the first column, we report the condition number of $F_{Q_1}^{-1}K$. It is not obvious that the values are bounded by a constant, but they appear to grow slower than any power of n or $\log(n)$, as results from log-log plots. In any case, these values are larger than the corresponding ones reported by Rønquist [33] for F_{P_1} (the P_1 finite element stiffness matrix obtained by dividing each element of the GLL mesh into tetrahedra). The values reported by Rønquist are all from 2 to 2.65, for values of n between 4 and

Table 3: Iteration counts with exact block-diagonal preconditioner

$Q_n - Q_{n-2}$								
n	ν							
	0.3	0.4	0.49	0.499	0.4999	0.49999	0.499999	0.5
2	1	1	1	1	1	1	1	1
3	7	7	7	7	7	7	7	7
4	11	13	21	21	21	21	21	21
5	11	15	27	31	31	31	31	31
6	11	15	29	33	35	35	35	35
7	11	15	29	35	35	35	37	37
8	11	15	31	37	39	39	39	39
9	11	15	31	39	41	41	41	41

$Q_n - P_{n-1}$								
n	ν							
	0.3	0.4	0.49	0.499	0.4999	0.49999	0.499999	0.5
2	2	2	2	2	2	2	2	2
3	9	9	9	9	9	9	9	9
4	9	11	15	15	15	15	15	15
5	9	13	19	21	21	21	21	21
6	9	13	21	21	21	21	21	21
7	11	13	21	21	21	21	21	21
8	11	13	23	25	25	25	25	25
9	11	13	23	25	25	25	25	25
10	11	13	23	25	25	25	25	25

Table 4: Equivalent formulation: iteration counts with exact block-diagonal preconditioner

$Q_n - Q_{n-2}$								
n	ν							
	0.3	0.4	0.49	0.499	0.4999	0.49999	0.499999	0.5
2	1	1	1	1	1	1	1	1
3	7	7	7	7	7	7	7	7
4	15	17	21	23	23	23	23	23
5	19	21	33	37	37	37	37	37
6	17	21	37	41	43	43	43	43
7	19	21	39	45	45	45	45	45
8	19	23	41	47	47	47	47	47
9	19	21	41	47	49	49	49	49

$Q_n - P_{n-1}$								
n	ν							
	0.3	0.4	0.49	0.499	0.4999	0.49999	0.499999	0.5
2	2	2	2	2	2	2	2	2
3	7	7	7	7	7	7	7	7
4	10	13	15	15	15	15	15	15
5	11	15	23	25	25	25	25	25
6	13	17	25	27	27	27	27	27
7	13	17	27	29	29	29	29	29
8	13	17	31	33	33	33	33	33
9	13	17	29	31	31	31	31	31
10	13	18	31	33	35	35	35	35

Table 5: Q_1 finite element preconditioner on the GLL mesh for Poisson equation: condition numbers and relative errors with known exact solution

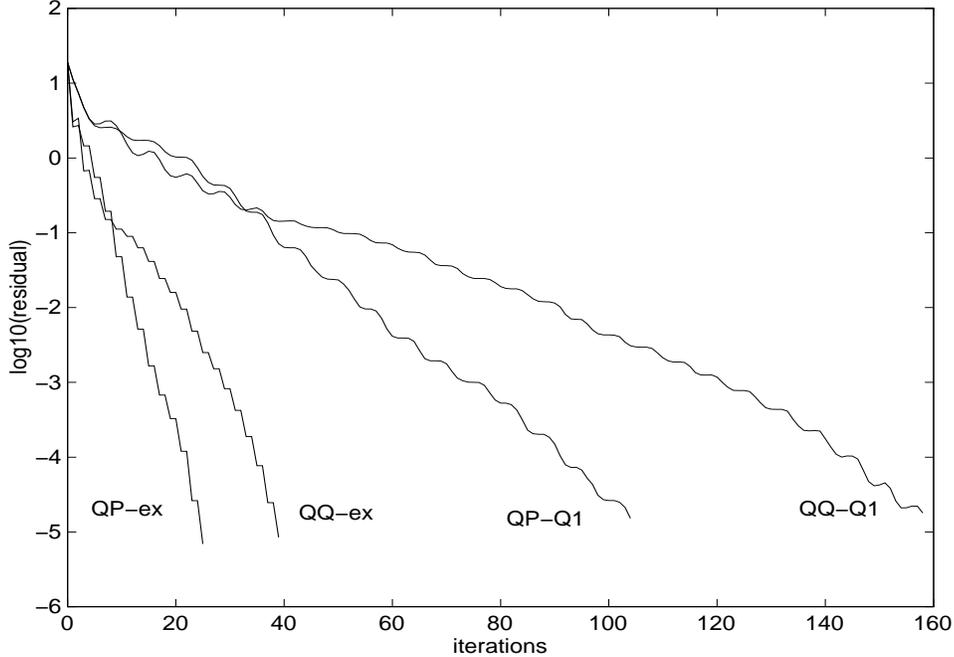
n	$cond(F_{Q_1}^{-1}K)$	$\frac{\ u_{ex} - u_{Q_1}\ _2}{\ u_{ex}\ _2}$	$\frac{\ u_{ex} - u_n\ _2}{\ u_{ex}\ _2}$
3	4.8150	0.9794	0.0130
4	8.4566	0.5892	0.0020
5	11.1569	0.3481	5.1163e-5
6	13.0747	0.2352	1.4581e-5
7	14.4623	0.1704	2.1302e-7
8	15.4977	0.1296	7.4257e-8
9	16.2954	0.1021	7.0685e-10
10	16.9275	0.0826	2.7414e-10
11	17.4406	0.0682	1.8431e-12
12	17.8653	-	-

Table 6: Iteration counts with Q_1 u-block preconditioner and exact p-block

$Q_n - Q_{n-2}$								
n	ν							
	0.3	0.4	0.49	0.499	0.4999	0.49999	0.499999	0.5
2	1	1	1	1	1	1	1	1
3	13	13	13	13	13	13	13	13
4	28	32	45	49	48	48	48	48
5	36	46	75	84	84	84	84	84
6	44	52	94	110	111	111	111	111
7	49	60	110	133	134	133	134	134
8	55	68	128	154	158	159	158	158
9	57	72	139	173	179	179	179	179

$Q_n - P_{n-1}$								
n	ν							
	0.3	0.4	0.49	0.499	0.4999	0.49999	0.499999	0.5
2	2	2	2	2	2	2	2	2
3	14	14	14	14	14	14	14	14
4	27	30	36	37	37	37	37	37
5	34	40	56	61	61	61	61	61
6	42	49	68	75	75	75	75	75
7	46	54	80	87	87	87	87	87
8	52	61	92	102	103	103	103	104
9	55	65	97	109	109	109	110	109
10	57	69	107	121	121	121	122	122

Figure 2: Convergence history for $n = 8$ and $\nu = 1/2$: $QP - ex = Q_n - P_{n-1}$ method with exact preconditioner, $QQ - ex = Q_n - Q_{n-2}$ method with exact preconditioner, $QP - Q1 = Q_n - P_{n-1}$ method with inexact Q_1 preconditioner, $QQ - Q1 = Q_n - Q_{n-2}$ method with inexact Q_1 preconditioner



12. In the second and third columns of Table 5, we report the relative errors in the discrete l_2 -norm between the exact solution $u_{ex} = \sin \frac{\pi}{2}(x+1) \sin \frac{\pi}{2}(y+1) \sin \frac{\pi}{2}(z+1)$, obtained by computing the appropriate right-hand side $f = -\Delta u$, and the discrete solution u_n (spectral) or u_{Q_1} (Q_1 fem on the GLL mesh). The difference between spectral and h -version finite element accuracy is very clear.

In Table 6, we report the iteration counts for the model problem (2) when the inexact block $\hat{A} = \text{diag}(F_{Q_1}, F_{Q_1}, F_{Q_1})$ is used in the preconditioner. Even if the uniformity in ν is preserved, the number of iterations grows considerably, especially for higher values of n . Therefore, it does not appear that this inexact preconditioner is effective for PCR methods applied to mixed spectral systems. Results for the equivalent model problem (5) were similar and are not reported.

Figure 2 shows the convergence history of the $Q_n - Q_{n-2}$ and $Q_n - P_{n-1}$ methods with $n = 8$ and with exact and inexact Q_1 preconditioners for the Stokes problem. The resulting graphs are similar to the ones reported in Elman [16].

5.3 Block-diagonal preconditioner: multi-element case

Tables 7 and 8 report the iteration counts for $Q_n - Q_{n-2}$ and $Q_n - P_{n-1}$ respectively. Here we study the dependence of the number of iterations on the number of elements N for a fixed spectral degree n . This is analog to studying the dependence on h for a hp -finite element method. We divide the domain Ω into $N = N_x \times N_y \times N_z$ subcubes and

Table 7: Iteration count for $Q_n - Q_{n-2}$ with exact block-diagonal preconditioner: changing N for fixed n

n	$N = N_x \times N_y \times N_z$	ν							
		0.3	0.4	0.49	0.499	0.4999	0.49999	0.499999	0.5
2	$1 = 1 \times 1 \times 1$	1	1	1	1	1	1	1	1
2	$8 = 2 \times 2 \times 2$	6	7	7	7	7	7	7	7
2	$27 = 3 \times 3 \times 3$	9	11	17	17	17	17	17	17
2	$64 = 4 \times 4 \times 4$	10	13	19	21	21	21	21	21
2	$125 = 5 \times 5 \times 5$	10	13	21	23	23	23	23	23
3	$1 = 1 \times 1 \times 1$	7	7	7	7	7	7	7	7
3	$8 = 2 \times 2 \times 2$	11	13	21	23	23	23	23	23
3	$27 = 3 \times 3 \times 3$	11	13	21	23	23	23	23	23
3	$64 = 4 \times 4 \times 4$	11	13	21	23	25	25	25	25
4	$1 = 1 \times 1 \times 1$	11	13	21	21	21	21	21	21
4	$8 = 2 \times 2 \times 2$	11	15	23	27	27	27	27	27

Table 8: Iteration count for $Q_n - P_{n-1}$ with exact block-diagonal preconditioner: changing N for fixed n

n	$N = N_x \times N_y \times N_z$	ν							
		0.3	0.4	0.49	0.499	0.4999	0.49999	0.499999	0.5
2	$1 = 1 \times 1 \times 1$	2	2	2	2	2	2	2	2
2	$8 = 2 \times 2 \times 2$	9	13	19	21	21	21	21	21
2	$27 = 3 \times 3 \times 3$	11	14	23	27	27	27	27	27
2	$64 = 4 \times 4 \times 4$	11	15	23	27	27	27	27	27
2	$125 = 5 \times 5 \times 5$	11	15	23	27	27	27	27	27
3	$1 = 1 \times 1 \times 1$	9	9	9	9	9	9	9	9
3	$8 = 2 \times 2 \times 2$	11	13	21	23	23	23	23	23
3	$27 = 3 \times 3 \times 3$	11	13	23	25	25	25	25	25
4	$1 = 1 \times 1 \times 1$	9	11	15	15	15	15	15	15
4	$8 = 2 \times 2 \times 2$	11	13	23	25	25	25	25	25

Table 9: Iteration count for $Q_n - Q_{n-2}$ with exact block-diagonal preconditioner: changing n for fixed $N = 4 = 2 \times 2 \times 1$

n	ν							
	0.3	0.4	0.49	0.499	0.4999	0.49999	0.499999	0.5
2	5	5	5	5	5	5	5	5
3	10	13	21	23	23	23	23	23
4	11	15	25	29	29	29	29	29
5	11	15	27	31	33	33	33	33
6	11	15	27	35	35	35	35	35

Table 10: Iteration count for $Q_n - P_{n-1}$ with exact block-diagonal preconditioner: changing n for fixed $N = 4 = 2 \times 2 \times 1$

n	ν							
	0.3	0.4	0.49	0.499	0.4999	0.49999	0.499999	0.5
2	9	11	17	19	19	19	19	19
3	10	13	19	23	23	23	23	23
4	11	14	23	25	27	27	27	27
5	11	14	23	27	27	27	27	27
6	11	14	23	27	27	27	27	27

we take $N_x = N_y = N_z$ in order to always have a cubic domain (and avoid comparing problems with different aspect ratios). Due to the cubic growth of N , we could only run cases with a low value of $n = 2, 3, 4$. The results of the tables seem to indicate a bound on the number of iterations independent of N , in agreement with the theory. This is particularly evident for the $Q_n - P_{n-1}$ method and for $n = 2$, which allows us to run with sufficiently many elements. As before, the results are uniform in ν and the incompressible limit is the hardest case, for each n and N fixed.

In tables 9 and 10, we study the dependence of the number of iterations on n , for a small fixed number of elements $N = 4$, with $N_x = N_y = 2$ and $N_z = 1$. The linear growth of the number of iterations with n is clearly visible in the incompressible limit for $Q_n - Q_{n-2}$ (Table 9), while for compressible materials ($\nu \approx 0.3 - 0.4$) the number of iterations seems insensitive to n and bounded by a constant. For the $Q_n - P_{n-1}$ method, the results of Table 10 are better and the number of iterations seems bounded by a constant also in the incompressible limit. However, as was shown in the single-element case, higher values of n might reveal a growth which is still linear, just with a better constant in front of the linear term.

Table 11: L^2 -errors with known exact solution for $\nu = 0.3$; $Q_n - Q_{n-2}$ with exact lower-triangular preconditioner; G=GMRES, B=Bi-CGSTAB, Q=QMR

$N = N_x \times N_y \times N_z$	n	iter.			$\frac{\ \mathbf{u}_{gmres} - \mathbf{u}_{ex}\ _2}{\ \mathbf{u}_{ex}\ _2}$	$\frac{\ p_{gmres} - p_{ex}\ _2}{\ p_{ex}\ _2}$
		G	B	Q		
$1 = 1 \times 1 \times 1$	2	1	1	1	4.6502e-1	1.0000e+1
	3	2	1	2	1.0027e-1	4.0346e-1
	4	4	3	4	1.8246e-2	7.6835e-2
	5	6	4	6	3.7748e-3	2.5745e-2
	6	7	4	7	3.6803e-4	3.0630e-3
	7	7	4	7	5.9270e-5	6.8210e-4
	8	7	4	7	4.0743e-6	5.0991e-5
	9	7	4	7	5.3355e-7	8.7179e-6
$8 = 2 \times 2 \times 2$	2	2	1	2	9.4308e-2	2.5707e-1
	3	6	4	6	1.2310e-2	1.1750e-1
	4	7	4	7	1.1785e-3	1.3814e-2
	5	7	4	7	9.1187e-5	1.6598e-3

5.4 Triangular preconditioner: single-element case

In the following tables, we have used the convention G = GMRES, B = Bi-CGSTAB, Q = QMR. In all cases, we have used the (left) lower-triangular preconditioner T_L with exact blocks.

In the first part of Table 11, we report the errors in the L^2 -norm between the $Q_n - Q_{n-2}$ spectral element solution and the known exact solution $u_1 = u_2 = u_3 = \sin(\frac{\pi}{2N_x}x) \sin(\frac{\pi}{2N_y}y) \sin(\frac{\pi}{2N_z}z)$, $p = \lambda \operatorname{div}(\mathbf{u})$ for the elasticity problem with $\nu = 0.3$ on the reference cube (i.e. $N_x = N_y = N_z = 1$). On each row, corresponding to each value of n , we report the iteration counts for the three methods and the errors for the displacement \mathbf{u} and for p . The results clearly show the spectral convergence of the discrete solution to the exact solution. The second part of the table shows the same results for a multi-element case with 8 elements. Tables 12 and 13 report the iteration counts for $Q_n - Q_{n-2}$ and $Q_n - P_{n-1}$ respectively, on the reference element. For each value of n and ν , the results for GMRES, Bi-CGSTAB and QMR are reported. As in the block-diagonal case, the results are uniform in ν , i.e. for each fixed n , the number of iterations tends to the number of iterations of the limiting incompressible case. Moreover, for each fixed ν , the number of iterations grows at worst linearly with n , in agreement with the theory. This is clear at the incompressible limit, while away from it the results are much better and in practice bounded independently of n . Among the three iterative methods, Bi-CGSTAB requires the least number of iterations, in some cases half of those required by GMRES, but it requires twice as many applications of the matrix and the preconditioner. Moreover, Bi-CGSTAB shows a more irregular convergence behavior than the other two methods. QMR has iterations counts in between Bi-CGSTAB and GMRES, often closer to the last one. QMR also

Table 12: Iteration count for $Q_n - Q_{n-2}$ on one element, with exact lower-triangular preconditioner; G=GMRES, B=Bi-CGSTAB, Q=QMR

n	ν							
	0.3 G B Q	0.4 G B Q	0.49 G B Q	0.499 G B Q	0.4999 G B Q	0.49999 G B Q	0.499999 G B Q	0.5 G B Q
2	1 1 1	1 1 1	1 1 1	1 1 1	1 1 1	1 1 1	1 1 1	1 1 1
3	4 2 4	4 3 4	4 3 4	4 3 4	4 3 4	4 3 4	4 3 4	4 3 4
4	6 3 6	7 4 7	11 7 11	11 8 11	11 10 12	11 10 12	11 10 12	11 10 12
5	7 4 6	8 5 8	15 9 14	18 17 18	19 17 18	19 14 18	19 18 18	19 14 18
6	7 4 6	9 5 8	16 12 15	20 17 19	21 15 20	21 15 20	21 15 20	21 15 20
7	7 4 6	9 5 8	17 14 16	22 17 20	22 17 21	22 16 21	22 16 21	22 20 21
8	7 4 7	9 5 8	20 14 18	24 20 21	24 24 22	24 20 22	24 20 22	24 20 22
9	7 4 7	9 5 8	20 13 18	25 19 22	26 20 23	26 21 23	26 19 23	26 17 23

Table 13: Iteration count for $Q_n - P_{n-1}$ on one element, with exact lower-triangular preconditioner; G=GMRES, B=Bi-CGSTAB, Q=QMR

n	ν							
	0.3 G B Q	0.4 G B Q	0.49 G B Q	0.499 G B Q	0.4999 G B Q	0.49999 G B Q	0.499999 G B Q	0.5 G B Q
2	2 1 2	2 1 2	2 1 2	2 1 2	2 1 2	2 1 2	2 1 2	2* 2
3	5 3 5	5 3 5	5 4 5	5 4 5	5 4 5	5 4 5	5 4 5	5 4 5
4	5 3 5	6 4 6	8 6 8	8 6 8	8 6 8	8 6 8	8 6 8	8 6 8
5	6 3 6	7 4 7	11 6 10	11 7 11	12 7 11	12 7 11	12 7 11	12 7 11
6	7 3 6	8 4 7	12 7 11	13 8 11	13 8 11	13 8 11	13 8 11	13 8 11
7	7 4 6	8 4 7	13 6 11	14 7 12	14 7 12	14 7 12	14 7 12	14 7 12
8	7 3 6	9 4 8	14 7 12	15 8 13	15 8 13	15 8 13	15 8 13	15 8 13
9	8 4 6	9 4 8	14 7 12	15 8 13	15 8 13	15 8 13	15 8 13	15 8 13
10	8 4 6	10 5 8	15 7 12	18 8 13	18 8 13	18 8 13	18 8 13	18 8 13

requires twice as many matrix and preconditioner applications compared with GMRES. Of course, GMRES without restart requires much more memory than the other two methods. In comparison with the block-diagonal results, the triangular preconditioner requires many less iterations, sometimes half of those required by PCR.

5.5 Triangular preconditioner: multi-element case

Tables 14 and 15 are the analog for triangular preconditioners of Tables 7 and 8 for block-diagonal preconditioners. Here, we verify numerically the independence of the iteration counts on N while keeping n fixed. Again, we could run with many elements only for small values of $n = 2, 3, 4$, and we used decompositions of the cubic domain into cubic powers of subdomains. The results indicate in all cases an upper bound independent of N . Regarding the different convergence performance of the three methods, the same considerations as for the single-element case apply.

Table 14: Iteration count for $Q_n - Q_{n-2}$ with exact lower-triangular preconditioner: changing N for fixed n

n	N	ν							
		0.3 G B Q	0.4 G B Q	0.49 G B Q	0.499 G B Q	0.4999 G B Q	0.49999 G B Q	0.499999 G B Q	0.5 G B Q
2	1^3	1 1 1	1 1 1	1 1 1	1 1 1	1 1 1	1 1 1	1 1 1	1 1 1
2	2^3	3 2 4	4 2 4	4 3 4	4 3 4	4 3 4	4 3 4	4 3 4	4 3 4
2	3^3	5 3 6	6 4 7	8 5 9	8 6 9	8 6 9	8 6 9	8 6 9	8 6 9
2	4^3	5 4 6	6 4 7	9 6 10	10 7 11	10 7 11	10 7 11	10 7 11	10 7 11
2	5^3	5 4 6	7 5 7	10 7 11	10 7 12	11 7 12	11 7 12	11 7 12	11 7 12
3	1^3	4 2 4	4 3 4	4 3 4	4 3 4	4 3 4	4 3 4	4 3 4	4 3 4
3	2^3	6 3 6	7 4 8	11 7 11	12 8 12	12 8 12	12 8 12	12 8 12	12 8 12
3	3^3	6 4 6	7 5 8	11 7 12	12 8 12	12 8 13	12 9 13	12 9 13	12 9 13
4	1^3	6 3 6	7 4 7	11 7 11	11 8 11	11 10 12	11 10 12	11 10 12	11 10 12
4	2^3	6 4 6	8 5 8	12 8 13	14 9 14	14 10 15	14 10 15	14 10 15	14 10 15

Tables 16 and 17 are the analog for triangular preconditioners of Tables 9 and 10. Here, we fix a small number of elements $N = 4$ and we study the iteration counts by increasing n and ν . Again, the results are uniform in ν and linear (at worst) in n , with the incompressible case being the hardest one. In comparison with the block-diagonal preconditioner, the triangular preconditioner considerably decreases the number of iterations, sometimes by as much as one half.

6 Conclusions

We have proposed and analyzed iterative methods for the sparse indefinite systems arising from the mixed spectral element discretization of elasticity and Stokes problems. These systems are solved with a preconditioned conjugate residual method when a block-diagonal preconditioner is used or with Krylov methods for nonsymmetric systems such as GMRES, Bi-CGSTAB and QMR when a triangular preconditioner is used. We have proven and have numerically shown that such algorithms have convergence rates bounded by the inverse of the inf-sup constant and independent of the penalty parameter in the saddle point formulation (the Poisson ratio for elasticity or a stabilization parameter for Stokes). The two mixed spectral methods considered, $Q_n - Q_{n-2}$ and $Q_n - P_{n-1}$, have equivalent theoretical convergence bounds, but we have numerically shown that the latter one has a better inf-sup constant and gives better iteration counts. On the other hand, P_{n-1} does not have a tensorial basis. The exact blocks in the preconditioners could be replaced by appropriate preconditioners based on low-order discretizations on the GLL mesh and/or domain decomposition techniques. The inexact preconditioner based on Q_1 finite elements on the GLL mesh largely increases the iteration counts. Future work should address other inexact preconditioners, such as P_1 finite elements on the GLL mesh, multigrid or domain decomposition methods and preconditioners for the mass matrix C .

Table 15: Iteration count for $Q_n - P_{n-1}$ with exact lower-triangular preconditioner: changing N for fixed n

n	N	ν							
		0.3 G B Q	0.4 G B Q	0.49 G B Q	0.499 G B Q	0.4999 G B Q	0.49999 G B Q	0.499999 G B Q	0.5 G B Q
2	1^3	2 1 2	2 1 2	2 1 2	2 1 2	2 1 2	2 1 2	2 1 2	2 * 2
2	2^3	5 3 6	6 4 7	9 6 10	10 7 11	10 7 11	10 7 11	10 7 11	10 7 11
2	3^3	6 4 6	7 5 8	11 8 12	13 9 14	13 9 14	13 9 14	13 9 14	13 9 14
2	4^3	6 4 6	7 5 8	11 8 12	13 9 14	13 9 14	13 9 14	13 9 14	13 9 14
2	5^3	6 4 6	7 5 8	12 8 13	13 9 14	13 10 14	13 10 14	13 10 14	13 10 14
3	1^3	5 3 5	5 3 5	5 4 5	5 4 5	5 4 5	5 4 5	5 4 5	5 4 5
3	2^3	6 4 6	7 5 7	10 7 11	11 7 12	11 7 12	11 7 12	11 7 12	11 7 12
3	3^3	6 4 6	7 5 8	11 7 12	12 8 13	12 8 13	12 8 13	12 8 13	12 8 13
4	1^3	5 3 5	6 4 6	8 6 8	8 6 8	8 6 8	8 6 8	8 6 8	8 6 8
4	2^3	6 4 6	8 5 8	12 7 12	12 8 13	12 8 13	12 8 13	12 8 13	12 8 13

Table 16: Iteration count for $Q_n - Q_{n-2}$ with exact lower-triangular preconditioner: changing n for fixed $N = 4 = 2 \times 2 \times 1$

n	ν							
	0.3 G B Q	0.4 G B Q	0.49 G B Q	0.499 G B Q	0.4999 G B Q	0.49999 G B Q	0.499999 G B Q	0.5 G B Q
2	3 2 3	3 2 3	3 2 3	3 2 3	3 2 3	3 2 3	3 2 3	3 2 3
3	6 3 6	6 4 7	10 7 11	11 8 12	11 8 12	12 8 12	12 8 12	12 8 12
4	6 4 6	8 5 8	13 9 13	15 11 15	15 11 16	15 13 16	15 13 16	15 13 16
5	7 4 6	8 5 8	15 9 14	18 14 18	19 14 19	19 14 19	19 14 19	19 14 19
6	7 4 6	8 5 8	16 10 15	20 14 19	21 16 20	21 15 20	21 15 20	21 15 20

Table 17: Iteration count for $Q_n - P_{n-1}$ with exact lower-triangular preconditioner: changing n for fixed $N = 4 = 2 \times 2 \times 1$

n	ν							
	0.3 G B Q	0.4 G B Q	0.49 G B Q	0.499 G B Q	0.4999 G B Q	0.49999 G B Q	0.499999 G B Q	0.5 G B Q
2	5 3 5	6 4 6	8 6 9	9 7 10	9 7 10	9 7 10	9 7 10	9 7 10
3	6 4 6	7 4 7	10 7 11	11 8 12	11 8 12	11 8 12	11 8 12	11 8 12
4	6 4 6	8 5 8	12 8 12	14 8 14	14 8 14	14 9 14	14 8 14	14 9 14
5	6 4 6	8 5 8	13 9 13	15 12 14	15 14 14	15 12 14	15 13 14	15 14 14
6	7 4 6	8 5 8	14 9 13	15 11 14	15 11 14	15 11 14	15 11 14	15 11 14

Acknowledgments. The author wishes to thank Olof Widlund, Franco Brezzi, Howard Elman and Axel Klawonn for many helpful discussions and suggestions.

References

- [1] P. ARISTOV AND E. CHIZHONKOV, *On the constant in the LBB condition for rectangular domains*, Tech. Rep. TR 9534, Department of Mathematics, University of Nijmegen, Netherlands, 1995.
- [2] S. ASHBY, T. MANTEUFFEL, AND P. SAYLOR, *A taxonomy for conjugate gradient methods*, SIAM J. Numer. Anal., 27 (1990), pp. 1542–1568.
- [3] I. BABUŠKA AND M. SURI, *Locking effects in the finite element approximation of elasticity problems*, Numer. Math., 62 (1992), pp. 439–463.
- [4] ———, *The p and $h - p$ versions of the finite element method, basic principles and properties*, SIAM Review, 36 (1994), pp. 578–632.
- [5] R. BARRETT, M. BERRY, T. CHAN, J. DEMMEL, J. DONATO, J. DONGARRA, V. EIJKHOUT, R. POZO, C. ROMINE, AND H. VAN DER VORST, *Templates for the Solution of Linear Systems: Building Blocks for Iterative Methods*, SIAM, Philadelphia, 1993.
- [6] K.-J. BATHE AND D. CHAPELLE, *The inf-sup test*, Computers and Structures, 47 (1993), pp. 537–545.
- [7] C. BERNARDI AND Y. MADAY, *Approximations Spectrales de Problèmes aux Limites Elliptiques*, vol. 10 of Mathématiques & Applications, Springer-Verlag France, Paris, 1992.
- [8] D. BRAESS AND C. BLÖMER, *A multigrid method for a parameter dependent problem in solid mechanics*, Numer. Math., 57 (1990), pp. 747–761.
- [9] J. BRAMBLE AND J. PASCIAK, *A preconditioning technique for indefinite systems resulting from mixed approximations of elliptic problems*, Math. Comp., 50 (1988), pp. 1–17.
- [10] J. BRAMBLE, J. PASCIAK, AND A. VASSILEV, *Analysis of the inexact Uzawa algorithm for saddle point problems*, tech. rep., Department of Mathematics, Cornell University, 1994.
- [11] S. BRENNER, *Multigrid methods for parameter dependent problems*, tech. rep., Department of Mathematics, University of South Carolina, 1994. To appear in *M²AN*.
- [12] F. BREZZI AND M. FORTIN, *Mixed and Hybrid Finite Element Methods*, Springer-Verlag, 1991.

- [13] C. CANUTO, M. Y. HUSSAINI, A. QUARTERONI, AND T. A. ZANG, *Spectral Methods in Fluid Dynamics*, Springer-Verlag, 1988.
- [14] M. CASARIN, *Schwarz Preconditioners for Spectral and Mortar Finite Element Methods with Applications to Incompressible Fluids*, PhD thesis, Department of Mathematics, Courant Institute of Mathematical Sciences, New York University, 1996.
- [15] M. DEVILLE AND E. MUND, *Finite element preconditioning for pseudospectral solutions of elliptic problems*, SIAM J. Sci. Stat. Comp., 11 (1990), pp. 311–342.
- [16] H. ELMAN, *Multigrid and Krylov subspace methods for the discrete Stokes equations*, Tech. Rep. UMIA CS-TR-94-76, Department of Computer Science, University of Maryland, 1994. To appear in International Journal for Numerical Methods in Fluids.
- [17] H. ELMAN AND G. GOLUB, *Inexact and preconditioned Uzawa algorithms for saddle point problems*, SIAM J. Numer. Anal., 31 (1994).
- [18] P. FISCHER AND E. RØNQUIST, *Spectral element methods for large scale parallel Navier-Stokes calculations*, Comput. Methods Appl. Mech. Engrg, 116 (1994), pp. 69–76.
- [19] R. FREUND, G. H. GOLUB, AND N. NACHTIGAL, *Iterative Solution of Linear Systems*, Acta Numerica 1992, Cambridge University Press, 1991, pp. 57–100.
- [20] D. FUNARO, *Polynomial Approximation of Differential Equations*, Springer-Verlag, 1992.
- [21] W. HACKBUSCH, *Iterative Solution of Large Sparse Systems of Equations*, Springer-Verlag, 1994.
- [22] A. KLAWONN, *An optimal preconditioner for a class of saddle point problems with a penalty term*, Tech. Rep. TR 19/94, Westfälische Wilhelms-Universität Münster, 1994. To appear in SIAM J. Sci. Comp.
- [23] ———, *Triangular preconditioners for saddle point problems with a penalty term*, tech. rep., Westfälische Wilhelms-Universität Münster, 1995.
- [24] ———, *Preconditioners for Indefinite Problems*, PhD thesis, Westfälische Wilhelms-Universität Münster, Angewandte Mathematik und Informatik, 1996. Tech. Rep. 8/96-N.
- [25] Y. MADAY, D. MEIRON, A. PATERA, AND E. RØNQUIST, *Analysis of iterative methods for the steady and unsteady Stokes problem: application to spectral element discretizations*, SIAM J. Sci. Comp., 14 (1993), pp. 310–337.
- [26] Y. MADAY, A. PATERA, AND E. RØNQUIST, *The $P_N \times P_{N-2}$ method for the approximation of the Stokes problem*, Tech. Rep. 92009, Department of Mechanical Engineering, Massachusetts Institute of Technology, 1992. To appear in Numer. Math.

- [27] Y. MADAY AND A. T. PATERA, *Spectral element methods for the Navier-Stokes equations*, in State of the Art Surveys in Computational Mechanics, A. Noor and J. Oden, eds., New York, 1989, ASME.
- [28] J. MANDEL, *Iterative solvers for p-version finite element method in three dimensions*, Comput. Methods Appl. Mech. Engrg, 116 (1994), pp. 175–183. ICOSAHOM 92, Montpellier, France, June 1992.
- [29] J. MANDEL, *Iterative methods for p-version finite elements: preconditioning thin solids*, Comp. Meths. Appl. Mech. Eng., 133 (1996), pp. 247–257.
- [30] A. T. PATERA, *A spectral element method for fluid dynamics: Laminar flow in a channel expansion*, J. Comp. Phys., 37 (1984), pp. 468–488.
- [31] L. F. PAVARINO AND O. B. WIDLUND, *Preconditioned conjugate gradient solvers for spectral elements in 3D*, in Solution Techniques for Large-Scale CFD Problems, W. G. Habashi, ed., John Wiley & Sons, 1995, pp. 249–270.
- [32] L. F. PAVARINO AND O. B. WIDLUND, *A polylogarithmic bound for an iterative substructuring method for spectral elements in three dimensions*, SIAM J. Numer. Anal., 33 (1996), pp. 1303–1335.
- [33] E. RØNQUIST, *A domain decomposition solver for the steady Navier-Stokes equations*, in Third Int. Conf. on Spectral and High Order Methods for PDE's, 1996. Proceedings of ICOSAHOM 95, a conference held in Houston, TX USA, 1995. To appear in Comput. Meth. Appl. Mech. Eng.
- [34] T. RUSTEN AND R. WINTHER, *A preconditioned iterative method for saddle point problems*, SIAM J. Matr. Anal. Appl., 13 (1992), pp. 887–904.
- [35] D. SILVESTER AND A. WATHEN, *Fast iterative solution of stabilised Stokes systems. part II: Using general block preconditioners*, SIAM J. Numer. Anal., 31 (1994), pp. 1352–1367.
- [36] M. SURI AND R. STENBERG, *Mixed hp finite element methods for problems in elasticity and Stokes flow*, Numer. Math., 72 (1996), pp. 367–390.
- [37] R. VERFÜRTH, *A multilevel algorithm for mixed problems*, SIAM J. Numer. Anal., 21 (1984), pp. 264–271.
- [38] A. WATHEN AND D. SILVESTER, *Fast iterative solution of stabilised Stokes systems. part I: Using simple diagonal preconditioners*, SIAM J. Numer. Anal., 30 (1993), pp. 630–649.
- [39] G. WITTUM, *Multigrid methods for the Stokes and Navier-Stokes equations*, Numer. Math., 54 (1989), pp. 543–564.