

An overlapping Schwarz method for spectral element solution of the incompressible Navier-Stokes equations

P. F. Fischer *

Abstract

Efficient solution of the Navier-Stokes equations in complex domains is dependent upon the availability of fast solvers for sparse linear systems. For unsteady incompressible flows, the pressure operator is the leading contributor to stiffness, as the characteristic propagation speed is infinite. In the context of operator splitting formulations, it is the pressure solve which is the most computationally challenging, despite its elliptic origins. We examine several preconditioners for the consistent \mathcal{L}^2 Poisson operator arising in the $\mathbb{P}_N - \mathbb{P}_{N-2}$ spectral element formulation of the incompressible Navier-Stokes equations. We develop a finite element based additive Schwarz preconditioner using overlapping subdomains plus a coarse grid projection operator which is applied directly to the pressure on the interior Gauss points. For large two-dimensional problems this approach can yield as much as a five-fold reduction in simulation time over previously employed methods based upon deflation.

To appear in *J. of Comp. Phys.*, 1997.

*Present address: Division of Applied Mathematics, Brown University, Providence, RI 02912, USA.
E-mail: pff@cfm.brown.edu

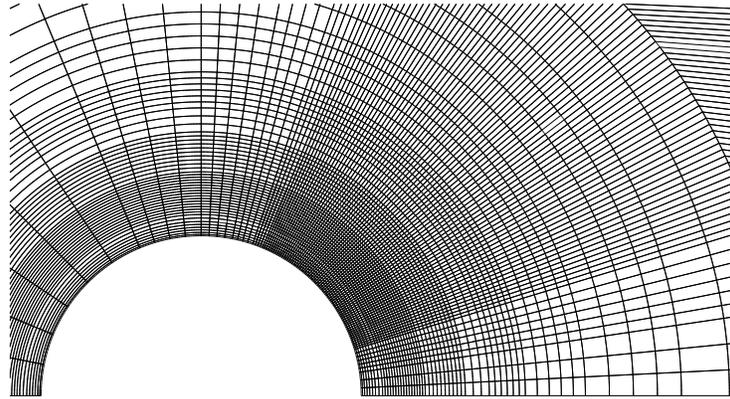
1 Introduction

We consider the problems encountered in large-scale spectral element simulations of unsteady incompressible flows. Accurate simulation of even two-dimensional flows can require hundreds of thousands of grid points when the Reynolds number is on the order of 10^4 . In the spectral element method, this elevated resolution can be attained by either increasing K , the number of elements, or increasing N , the order of approximation within each element. In practice, it is common to keep the order at a moderate level, i.e., in the range $N = 4$ to 15, and increase the number of elements to capture increasing physical and geometrical complexity.

We have followed this approach in a number of recent high-Reynolds number simulations of start-up flow past a cylinder using a time-splitting procedure which decouples the CFL limited convection steps, the linear viscous step, and the divergence-free projection into independent subproblems to be solved at each time step. Fig. 1a shows an example of a mesh used to compute the early evolution of wake vortices at $Re_D = U_\infty D/\nu = 9500$. Fig. 1b shows the vortex structure at a non-dimensional time of $\tau = tU_\infty/D = 3.1$. The drag history, shown in (c), agrees well with the results of [24] which were based upon an adaptive vortex method using up to 10^6 elements. The present calculation used a total of $K = 6112$ spectral elements, with the order varying from $N = 4$ at early times to $N = 9$ at later times.

At elevated resolutions, the linear system which imposes the pressure/divergence-free constraint at each time step can become very ill-conditioned and consequently tends to be the computational bottleneck when iterative solvers are employed. This problem can be exacerbated by the presence of high-aspect ratio elements and/or widely varying scales of resolution which are frequently encountered in practice, but often not present in model problems. Consequently, all of our recent iterative development work has focused upon a suite of cylinder problems of the type shown in Fig. 1.

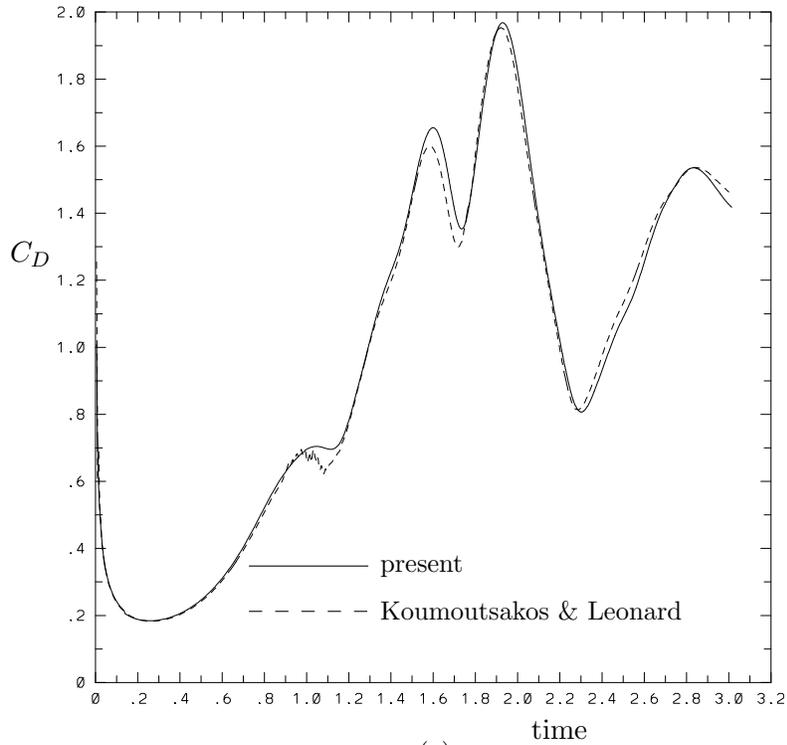
In this paper, we present a preconditioner for the pressure problem which is derived from a low-order finite element Laplacian with appropriate boundary conditions. The low-order operator defines a system to which additive overlapping Schwarz methods can be readily applied, as proposed by Dryja and Widlund, e.g., [12]. The combination of spectral methods and finite element preconditioning was first proposed by Orszag [33] and has been studied extensively by Deville, Mund, and coworkers, e.g., [9, 10]. The combination of spectral methods, finite element preconditioning, and additive Schwarz methods has been investigated by Pahl [34], Pavarino and Widlund [37], and Casarin [5] for the case of the discrete Laplacian. Rønquist [40] and Casarin [6] have studied iterative substructuring methods for spectral element solution of the fully-coupled steady Navier-Stokes equations. Rønquist also proposed a block-Jacobi/deflation based scheme applied to the consistent



(a)



(b)



(c)

Figure 1: (a) close-up of $K = 6112$ spectral element mesh for computation of start-up flow past a cylinder at $Re = 9500$ (b) vorticity contours at a convective time of $t = 3.10$ (c) drag coefficient C_D vs. non-dimensional time (Computation by G.W. Kruse).

Poisson operator governing the pressure for the unsteady case [17, 39]. To our knowledge, this is the first application of additive Schwarz methods directly to the consistent Poisson operator which governs the pressure in the $\mathbb{P}_N - \mathbb{P}_{N-2}$ spectral element method. In several large production runs, we have found this new method to yield a five-fold reduction in Navier-Stokes solution times over the deflation based scheme used previously [17, 18, 39]. Moreover, in some of these runs, the conditioning is improved to the point where 32-bit arithmetic suffices in cases previously requiring 64-bit precision, resulting in a two-fold savings in storage.

The outline of the remainder of the paper is as follows. In Section 2, we review the spectral element formulation for the unsteady Navier-Stokes equations, derive the system governing the pressure and demonstrate numerical convergence of the method. In Section 3, we examine the potential of several finite element based Laplacian operators as a basis for the pressure preconditioner. In Section 4, we couple an additive Schwarz method with the low-order Laplacian to develop a new preconditioner and demonstrate a technique to overcome some of the difficulties presented by high-aspect ratio subdomains. We draw some conclusions about this approach in Section 5, and present arguments favoring the choice of linear finite element preconditioning in the Appendix.

2 Navier-Stokes Discretization

As the nature of the pressure operator is quite different from discrete Laplacians based upon standard weighted residual techniques, we carefully review the temporal and spatial discretization for the spectral element method. We consider solution of the incompressible Navier-Stokes equations in \mathbb{R}^d , $d = 2$ or 3 :

$$\begin{aligned} \frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} &= -\nabla p + \frac{1}{Re} \nabla^2 \mathbf{u} & \text{in } \Omega, \\ \nabla \cdot \mathbf{u} &= 0 & \text{in } \Omega, \end{aligned} \tag{1}$$

where $\mathbf{u} = (u_1, \dots, u_d)$ is the velocity vector, p the pressure, and $Re = \frac{UL}{\nu}$ the Reynolds number based on a characteristic velocity, length scale, and kinematic viscosity. We have associated initial and boundary conditions

$$\mathbf{u}(\mathbf{x}, 0) = \mathbf{u}^0(\mathbf{x}), \quad \mathbf{u} = \mathbf{u}_v \text{ on } \partial\Omega_v, \quad \nabla u_i \cdot \hat{\mathbf{n}} = 0 \text{ on } \partial\Omega_o, \tag{2}$$

where $\hat{\mathbf{n}}$ is the outward pointing normal on the boundary, and subscripts v and o refer to boundary regions where either “velocity” or “outflow” boundary conditions are specified.

2.1 Temporal discretization

Temporal discretization is based upon an operator splitting approach in which the nonlinear convective terms are decoupled from the viscous and divergence operators via an operation-

integration-factor (OIF) technique developed by Maday, Patera, and Rønquist [26] and studied in detail in the thesis of Couzy [8]. The OIF scheme leads to a Stokes problem of the form:

$$\begin{aligned} \mathcal{H} \mathbf{u}^n + \nabla p^n &= \beta_1 \tilde{\mathbf{u}}_1^n + \beta_2 \tilde{\mathbf{u}}_2^n & \text{in } \Omega, \\ \nabla \cdot \mathbf{u}^n &= 0 & \text{in } \Omega. \end{aligned} \quad (3)$$

Here \mathcal{H} is the Helmholtz operator $\mathcal{H} = \left(-\frac{1}{Re} \nabla^2 + \frac{\beta_0}{\Delta t}\right)$, and $\beta_0 = \frac{3}{2}$, $\beta_1 = 2$, and $\beta_2 = -\frac{1}{2}$ are coefficients associated with second-order backwards differentiation (BDF2). The inhomogeneous terms, $\tilde{\mathbf{u}}_l^n \equiv \tilde{\mathbf{u}}_l(\mathbf{x}, t^n)$, are computed as solutions to the pure convection problem:

$$\begin{aligned} \frac{\partial \tilde{\mathbf{u}}_l}{\partial t} + \mathbf{u} \cdot \nabla \tilde{\mathbf{u}} &= 0, \\ \tilde{\mathbf{u}}_l(\mathbf{x}, t^{n-l}) &= \mathbf{u}(\mathbf{x}, t^{n-l}). \end{aligned} \quad (4)$$

The initial value problem (4) is solved using an explicit fourth-order Runge-Kutta scheme with step size $\Delta s \leq \Delta t$ which satisfies appropriate CFL criteria. The values of \mathbf{u} in (4) are interpolated/extrapolated from the previous velocity fields, $(\mathbf{u}^{n-1}, \dots, \mathbf{u}^{n-l})$. Note that the motivation for the OIF/BDF2 formulation is to decouple the CFL-limited convection steps from the implicit Stokes problem in order to permit a much larger step size, Δt , between (expensive) Stokes solves. Typical values of the CFL number, $\max_{\Omega} \frac{U \Delta t}{\Delta x}$, for our spectral element calculations are on the order of two to five. Numerous numerical tests have verified that the accuracy of this approach is indeed $O(\Delta t^2)$. In the sequel, we drop the superscript n in (3) and define the forcing function $\mathbf{f} = \beta_1 \tilde{\mathbf{u}}_1 + \beta_2 \tilde{\mathbf{u}}_2$.

2.2 Spatial discretization

The Stokes problem (3) can be recast in an equivalent variational form:

Find $\mathbf{u} \in X, p \in Y$ such that:

$$\begin{aligned} \frac{1}{Re} (\nabla \mathbf{u}, \nabla \mathbf{v}) + \frac{\beta_0}{\Delta t} (\mathbf{u}, \mathbf{v}) - (p, \nabla \cdot \mathbf{v}) &= (\mathbf{f}, \mathbf{v}) \quad \forall \mathbf{v} \in X_0 \\ - (q, \nabla \cdot \mathbf{u}) &= 0 \quad \forall q \in Y \end{aligned} \quad (5)$$

where

$$\forall \phi, \psi \in \mathcal{L}^2(\Omega), \quad (\phi, \psi) \equiv \int_{\Omega} \phi(\mathbf{x}) \psi(\mathbf{x}) d\mathbf{x}. \quad (6)$$

The proper subspaces for \mathbf{u}, \mathbf{v} and p, q are [20]:

$$\begin{aligned} X &= \{\mathbf{v} : v_i \in \mathcal{H}^1(\Omega), i = 1, \dots, d, \mathbf{v} = \mathbf{u}_v \text{ on } \partial\Omega_v\} \\ X_0 &= \{\mathbf{v} : v_i \in \mathcal{H}^1(\Omega), i = 1, \dots, d, \mathbf{v} = 0 \text{ on } \partial\Omega_v\} \\ Y &= \mathcal{L}^2(\Omega). \end{aligned}$$

Here $\mathcal{L}^2(\Omega)$ is the space of all functions which are square integrable over Ω and $\mathcal{H}^1(\Omega)$ is the space of all functions belonging to $\mathcal{L}^2(\Omega)$ whose first derivatives are also in $\mathcal{L}^2(\Omega)$.

Spatial discretization proceeds by restricting \mathbf{u} , \mathbf{v} , p and q to compatible finite dimensional velocity and pressure subspaces, $X^N \subset X$ and $Y^N \subset Y$, respectively, and using appropriate quadrature to approximate the inner products in (5):

Find $\mathbf{u} \in X^N$, $p \in Y^N$ such that:

$$\begin{aligned} \frac{1}{Re} (\nabla \mathbf{u}, \nabla \mathbf{v})_{GL} + \frac{\beta_0}{\Delta t} (\mathbf{u}, \mathbf{v})_{GL} - (p, \nabla \cdot \mathbf{v})_G &= (\mathbf{f}, \mathbf{v})_{GL} \quad \forall \mathbf{v} \in X_0^N \\ - (q, \nabla \cdot \mathbf{u})_G &= 0 \quad \forall q \in Y^N \end{aligned} \quad (7)$$

where the quadrature rules $(\cdot, \cdot)_{GL}$ and $(\cdot, \cdot)_G$ will be related to the spaces X^N and Y^N . To simplify subsequent operator definition, we symmetrize the problem (7) by introducing the splitting $\mathbf{u} = \mathbf{u}_0 + \mathbf{u}_b$ where $\mathbf{u}_0 = 0$ on Ω_v and \mathbf{u}_b is any known function in X^N . This yields the following restatement of the Stokes problem:

Find $\mathbf{u}_0 \in X_0^N$, $p \in Y^N$ such that:

$$\begin{aligned} \frac{1}{Re} (\nabla \mathbf{u}_0, \nabla \mathbf{v})_{GL} + \frac{\beta_0}{\Delta t} (\mathbf{u}_0, \mathbf{v})_{GL} - (p, \nabla \cdot \mathbf{v})_G &= (\mathbf{f}_v, \mathbf{v})_{GL} \quad \forall \mathbf{v} \in X_0^N \\ - (q, \nabla \cdot \mathbf{u}_0)_G &= (q, \nabla \cdot \mathbf{u}_b)_G \quad \forall q \in Y^N, \end{aligned} \quad (8)$$

where \mathbf{f}_v is the augmented inhomogeneity which incorporates the boundary terms.

In the spectral element method [36, 25] the bases for X^N and Y^N are defined by tessellating the domain into K non-overlapping subdomains, $\Omega = \cup_{k=1}^K \Omega^k$, and representing functions within each subdomain in terms of tensor-product polynomials on a reference subdomain $\hat{\Omega} = [-1, +1]^d$. (We will refer to the Ω^k 's as *subdomains* to distinguish them from *elements* which will be defined in the context of finite element preconditioners in the next section.) Each Ω^k is the image of the reference subdomain under a mapping: $\mathbf{x}^k(\mathbf{r}) \in \Omega^k \implies \mathbf{r} \in \hat{\Omega}$, with well defined inverse: $\mathbf{r}^k(\mathbf{x}) \in \hat{\Omega} \implies \mathbf{x} \in \Omega^k$. Thus, each subdomain is a deformed quadrilateral in \mathbb{R}^2 or deformed parallelepiped in \mathbb{R}^3 . The intersection of the closure of any two subdomains is either void, a vertex, an entire edge, or an entire face.

To avoid spurious pressure modes, Maday, Patera and Rønquist [27] and Bernardi and Maday [2] suggest the following approximation spaces for the velocity and pressure:

$$\begin{aligned} X^N &= X \cap \mathbb{P}_{N,K}^d(\Omega) \\ Y^N &= Y \cap \mathbb{P}_{N-2,K}(\Omega) \end{aligned} \quad (9)$$

where

$$\mathbb{P}_{N,K}(\Omega) = \left\{ v(\mathbf{x}^k(\mathbf{r})) \Big|_{\Omega^k} \in \mathbb{P}_{\mathcal{N}}(r_1) \otimes \dots \otimes \mathbb{P}_{\mathcal{N}}(r_d), k = 1, \dots, K \right\} \quad (10)$$

and $\mathbb{P}_{\mathcal{N}}(r)$ is the space of all polynomials of degree less than or equal to \mathcal{N} in the argument. Note that the dimension of Y^N is $K(N-1)^d$ since function continuity is not enforced for

functions in Y^N . However, the dimension of X^N is less than $dK(N+1)^d$ due to the restriction that functions in X^N must be continuous across subdomain interfaces and must also satisfy Dirichlet boundary conditions on $\partial\Omega_v$.

For the velocity space, we choose as a basis for $\mathbb{P}_N(r)$ the set of Lagrangian interpolants on the Gauss-Lobatto Legendre (GL) quadrature points in the reference domain: $\xi_i \in [-1, +1]$, $i = 0, \dots, N$, whereas for the pressure space, the basis for $\mathbb{P}_{N-2}(r)$ is the set of Lagrangian interpolants on the Gauss Legendre (G) quadrature points $\eta_i \in]-1, +1[$, $i = 1, \dots, N-1$. Fig 2 shows the nodal points for both the velocity (GL) and pressure (G) meshes for a regular subdomain configuration. Note that the basis for velocity is continuous across subdomain interfaces, while the basis for the pressure is not.

The Lagrangian bases permit convenient implementation of the quadrature rules which we now define. Let $f^k(\mathbf{r}) \equiv f(\mathbf{x}^k(\mathbf{r}))$, $\mathbf{r} \in \hat{\Omega}$. In \mathbb{R}^2 we have:

$$(f, g)_{GL} \equiv \sum_k \sum_{i=0}^N \sum_{j=0}^N f^k(\xi_i, \xi_j) \cdot g^k(\xi_i, \xi_j) \cdot |J^k(\xi_i, \xi_j)| \cdot \rho_i \rho_j \quad (11)$$

$$(f, g)_G \equiv \sum_k \sum_{i=1}^{N-1} \sum_{j=1}^{N-1} f^k(\eta_i, \eta_j) \cdot g^k(\eta_i, \eta_j) \cdot |J^k(\eta_i, \eta_j)| \cdot \sigma_i \sigma_j \quad (12)$$

where $J^k(\mathbf{r})$ is the Jacobian arising from the transformation $\mathbf{x}^k(\mathbf{r})$, ρ_i is the GL quadrature weight associated with ξ_i , and σ_i is the G quadrature weight associated with η_i . The extension to \mathbb{R}^3 follows readily from the tensor product forms. It has been shown both

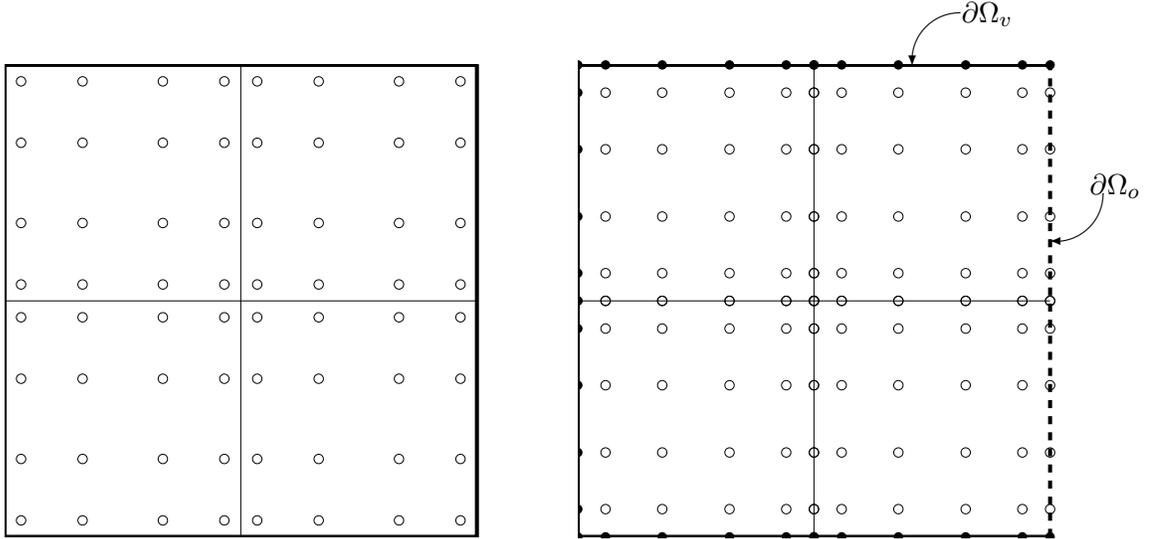


Figure 2: Spectral element configuration ($K = 4$, $N = 5$) showing Lagrange interpolation points for the pressure (Gauss) mesh on the left, and for the velocity (Gauss-Lobatto) mesh on the right. Open circles denote true degrees-of-freedom. Solid circles denote Dirichlet boundary nodes for velocity.

numerically and theoretically that the spectral element solution to the Stokes problem converges exponentially as $N \rightarrow \infty$ for problems having smooth solutions [25].

2.3 Spectral element operators

The locally-structured/globally-unstructured bases of the spectral element method naturally define a two-level operator and data hierarchy which we now describe. Our notation will be two-dimensional, restricted to the case of affine mappings: $\mathbf{x}^k(r_1, r_2) = (x_{0,1}^k + \frac{L_1^k}{2}r_1, x_{0,2}^k + \frac{L_2^k}{2}r_2)$, where $x_{0,i}^k$ and L_i^k represent local translation and dilation constants.

We first define the bases and operators associated with the velocity space. Every scalar field in $\mathbb{P}_{N,K}(\Omega)$ is represented in the form:

$$f(\mathbf{x})|_{\Omega^k} = \sum_{i=0}^N \sum_{j=0}^N f_{ij}^k h_i(r_1) h_j(r_2) \quad r_1, r_2 \in [-1, 1]^2, \quad (13)$$

where $h_i(r) \in \mathbb{P}_N(r)$ is the Lagrange polynomial satisfying $h_i(\xi_j) = \delta_{ij}$, and δ_{ij} is the Kronecker delta function. For each subdomain, we associate a natural ordering of the nodal values f_{ij}^k , $i, j \in \{0, \dots, N\}^2$ with the vector \underline{f}^k and, in turn, associate a natural ordering of the vectors \underline{f}^k , $k \in \{1, \dots, K\}$ with the $K(N+1)^2 \times 1$ vector \underline{f}_L .

We define the unassembled mass matrix to be the block-diagonal matrix $B_L \equiv \text{diag}(B^k)$, where each local mass matrix is expressed as a tensor product of one-dimensional operators:

$$B^k = \left(\frac{L_1^k L_2^k}{4} \right) \hat{B} \otimes \hat{B}. \quad (14)$$

Here $\hat{B} = \text{diag}(\rho_i)$, $i = 0, \dots, N$. The inner-product (11) is then expressed in terms of the mass matrices as:

$$\forall f, g \in \mathbb{P}_{N,K}(\Omega), \quad (f, g)_{GL} = \sum_k (\underline{f}^k)^T B^k \underline{g}^k = \underline{f}_L^T B_L \underline{g}_L. \quad (15)$$

In a similar fashion, we develop an expression for the bilinear form $(\nabla f, \nabla g)$:

$$\forall f, g \in \mathbb{P}_{N,K}(\Omega), \quad (\nabla f, \nabla g)_{GL} = \sum_k (\underline{f}^k)^T A^k \underline{g}^k = \underline{f}_L^T A_L \underline{g}_L, \quad (16)$$

Here, $A_L = \text{diag}(A^k)$ is the unassembled stiffness matrix and A^k is the local stiffness matrix:

$$A^k = \left(\frac{L_2^k}{L_1^k} \right) \hat{B} \otimes \hat{A} + \left(\frac{L_1^k}{L_2^k} \right) \hat{A} \otimes \hat{B}. \quad (17)$$

The one-dimensional stiffness matrix, \hat{A} , is defined in terms of the spectral differentiation matrix, \hat{D} :

$$\hat{A}_{ij} = \sum_{l=0}^N \hat{D}_{li} \rho_l \hat{D}_{lj} \quad i, j \in \{0, \dots, N\}^2 \quad (18)$$

with

$$\hat{D}_{ij} \equiv \left. \frac{dh_j}{dr} \right|_{r=\xi_i} \quad . \quad i, j \in \{0, \dots, N\}^2 \quad (19)$$

Deville and Mund [10] noted that, whereas \hat{A} is full, the two-dimensional operator, A^k , is sparse due to the use of the diagonal mass matrix, \hat{B} . In fact, the computational “stencil” of A^k is a cross, much like a high-order finite difference stencil. For deformed subdomains A^k is generally full with $(N+1)^d$ nonzero entries. Nonetheless, it was pointed out by Orszag [33] the action of A^k upon a vector can be efficiently computed in $O(N^{d+1})$ operations if one retains its tensor product form in favor of its explicit formation. Implementation details for fully deformed three-dimensional geometries can be found in [14].

The local subdomain operators are incorporated into global $n_v \times n_v$ system matrices by defining an index set $q_{ijk} \in \{1, \dots, n_v\}$ which maps vectors from their local representation, \underline{f}_L , to their global form, \underline{f} . The index set has repeated entries for any node (i, j, k) which is physically coincident with another node (i', j', k') , i.e.,

$$q_{ijk} = q_{i'j'k'} \quad \text{iff} \quad \mathbf{x}^k(r_i, r_j) = \mathbf{x}^{k'}(r_{i'}, r_{j'}) . \quad (20)$$

The index map can be represented in matrix form as a prolongation operator Q which maps from the set of global indices to the local index set. Q is a $K(N+1)^d \times n_v$ Boolean matrix with a single “1” in each row and zeros elsewhere. If $m = (k-1) \cdot (N+1)^2 + j \cdot (N+1) + i + 1$ is the position of f_{ij}^k in \underline{f}_L , and $q = q_{ijk}$ is the corresponding global index, then the m th column of Q^T is the unit vector \hat{e}_q , i.e., the q th column of the identity matrix. Application of Q to a vector implies the distribution of information, whereas application of Q^T to a vector implies summation, or gathering of information. Q^T is sometimes referred to as the “direct-stiffness-summation” operator.

A direct consequence of the unique mapping property (20) and the use of Lagrangian basis is that:

$$\forall f, g \in \mathbb{P}_{N,K}(\Omega) \cap \mathcal{H}^1, \quad (\nabla f, \nabla g)_{GL} = \underline{f}^T Q^T A_L Q \underline{g} . \quad (21)$$

We define $Q^T A_L Q$ as the Neumann Laplacian operator – it has a null-space of dimension unity corresponding to the constant mode. We define the associated Dirichlet operator as $\mathcal{M} Q^T A_L Q \mathcal{M}$, where \mathcal{M} is the diagonal mask matrix having ones on the diagonal at points $q_{ijk} : \mathbf{x}_{ij}^k \in \Omega \cup \partial\Omega_o$, and zeros elsewhere. With the operators Q and \mathcal{M} , the following problems are equivalent: For $f \in \mathbb{P}_{N,K}(\Omega)$,

Find $u \in X_0^N$, such that:

$$(\nabla v, \nabla u)_{GL} = (v, f)_{GL} \quad \forall v \in X_0^N \quad (22)$$

Find $\underline{u} \in \mathcal{R}(\mathcal{M})$, such that:

$$\underline{v}^T \mathcal{M} Q^T A_L Q \mathcal{M} \underline{u} = \underline{v}^T \mathcal{M} Q^T B_L \underline{f}_L \quad \forall \underline{v} \in \mathcal{R}(\mathcal{M}). \quad (23)$$

Here, $\mathcal{R}()$ is the range of the argument, and \underline{f}_L is the vector of nodal values of $f(\mathbf{x})$. The direct-stiffness-summation operator Q ensures that the solution will lie in \mathcal{H}^1 , while the mask \mathcal{M} enforces the homogeneous Dirichlet boundary condition: $u = 0$ on $\partial\Omega_v$. We define the discrete Laplacian and mass matrices as:

$$A \equiv \mathcal{M}Q^T A_L Q \mathcal{M}, \quad (24)$$

$$B \equiv \mathcal{M}Q^T B_L Q \mathcal{M}, \quad (25)$$

respectively. We will treat them as being both invertible and symmetric positive definite (SPD), although this is not strictly true due to the null space associated with the boundaries.

2.4 Stokes operators

To complete the description of the Stokes operator we need to consider the bilinear form:

$$(q, \nabla \cdot \mathbf{u})_G = \sum_{l=1}^d \left(q, \frac{\partial u_l}{\partial x_l} \right)_G. \quad (26)$$

Using the definition (12), the contribution to (26) from a single element in \mathbb{R}^2 is given by:

$$\sum_{l=1}^d \sum_{i=1}^{N-1} \sum_{j=1}^{N-1} q^k(\eta_i, \eta_j) \cdot \frac{\partial u_l^k}{\partial x_l}(\eta_i, \eta_j) \cdot |J^k(\eta_i, \eta_j)| \cdot \sigma_i \sigma_j. \quad (27)$$

The contribution from q presents no difficulty, as it is represented by Lagrangian interpolants on the Gauss points, i.e., $q^k(\eta_i, \eta_j) = q_{ij}^k$. The derivative of the velocity must be interpolated, giving rise to the following matrix form:

$$(q, \nabla \cdot \mathbf{u})_G = \sum_{k=1}^K (\underline{q}^k)^T \left(D_1^k \underline{u}_1^k + D_2^k \underline{u}_2^k \right). \quad (28)$$

For the case of the affine mappings defined above, the local derivative matrices are defined as

$$D_1^k = \left(\frac{L_2^k}{2} \right) \tilde{I} \otimes \tilde{D}, \quad D_2^k = \left(\frac{L_1^k}{2} \right) \tilde{D} \otimes \tilde{I}, \quad (29)$$

where

$$\tilde{I}_{ij} = \sigma_i h_j(\eta_i) \quad (30)$$

is the weighted one-dimensional interpolation matrix mapping from the Gauss-Lobatto points to the Gauss points, and

$$\tilde{D}_{ij} = \sigma_i \frac{dh_j}{dr} \Big|_{r=\eta_i} \quad (31)$$

is the weighted one-dimensional differentiation matrix, interpolated onto the Gauss points.

The extension from the local operator to the global operator proceeds exactly as in the previous section. The space of admissible functions in X_0^N is limited by the constraints that the velocity must be continuous at the subdomain interfaces and must satisfy the homogeneous boundary conditions, enforced by the action of the operators Q and \mathcal{M} , respectively.

Let $D_i \equiv D_{L,i}Q\mathcal{M}$, $i = 1, \dots, d$, with $D_{L,i} \equiv \text{diag}(D_i^k)$. In \mathbb{R}^2 , the matrix form of the Stokes problem (8) is then:

$$\begin{bmatrix} H & & -D_1^T \\ & H & -D_2^T \\ -D_1 & -D_2 & 0 \end{bmatrix} \begin{pmatrix} \underline{u}_1 \\ \underline{u}_2 \\ \underline{p} \end{pmatrix} = \begin{pmatrix} \underline{f}_1 \\ \underline{f}_2 \\ \underline{f}_p \end{pmatrix}, \quad (32)$$

where $H = \frac{1}{Re}A + \frac{\beta_0}{\Delta t}B$ is the discrete Helmholtz operator.

2.5 Stokes solvers

We now consider solution of the unsteady Stokes problem represented by the linear system (32). As noted in the previous section, explicit formation of the spectral operators will lead to $O(KN^{2d})$ nonzeros in the operators A and D_i , whereas their action upon a vector can be computed in $O(KN^{d+1})$ operations. This, coupled with the large system bandwidth which is inevitable in any three-dimensional calculation mandates the use of iterative solvers. The choice of Legendre based quadrature (vs., say, Chebyshev) results in symmetric positive definite operators for which preconditioned conjugate gradient iteration is a natural choice.

A common approach to solving the Stokes problem is to decouple the velocity and pressure by formally carrying out block LU factorization (Uzawa decoupling) on the system (32). This yields the Schur complement system for the pressure:

$$S\underline{p} = \underline{g}, \quad (33)$$

which is solved iteratively. Here, $S = \sum_i D_i H^{-1} D_i^T$ and \underline{g} is the corresponding inhomogeneity. Once the pressure is known, d Helmholtz solves serve to compute the velocity and complete the solution at time level n .

As it stands, the Uzawa approach requires a set of d Helmholtz solves for *each* iteration, since H^{-1} is embedded in S . An effective means to circumvent this difficulty is to decouple the viscous and pressure terms via an additional time splitting. Such an approach was suggested by Maday, Patera and Rønquist [26] and clearly analyzed in the articles of Blair Perot [3] and Couzy [8]. The approach has a common foundation with classical splitting approaches (e.g., [19, 32]) which lead to a Poisson equation for the pressure save that, in the present case, the splitting is effected in the discrete form of the equations. The correct boundary conditions are preserved and no steady-state temporal errors are introduced.

Following [8], we rewrite the Stokes system (32) by introducing an auxiliary system \mathbf{Q} :

$$\begin{bmatrix} \mathbf{H} & -\mathbf{H}\mathbf{Q}\mathbf{D}^T \\ -\mathbf{D} & 0 \end{bmatrix} \begin{pmatrix} \underline{\mathbf{u}}^n \\ \underline{p}^n - \underline{p}^{n-1} \end{pmatrix} = \begin{pmatrix} \mathbf{B}\underline{\mathbf{f}} + \mathbf{D}^T \underline{p}^{n-1} \\ \underline{f}_p \end{pmatrix} + \begin{pmatrix} \underline{\mathbf{r}} \\ \underline{0} \end{pmatrix}. \quad (34)$$

Here, boldface indicates the d -dimensional vector form of the previously defined operators, and the residual term is:

$$\underline{\mathbf{r}} = (\mathbf{H}\mathbf{Q} - \mathbf{I})\mathbf{D}^T(\underline{p}^{n-1} - \underline{p}^n). \quad (35)$$

Couzy [8] examines several choices for \mathbf{Q} . The choice $\mathbf{Q} = \mathbf{H}^{-1}$ yields no splitting error ($\underline{\mathbf{r}} = 0$), while the more computationally convenient choice $\mathbf{Q} = \frac{\Delta t}{\beta_0}\mathbf{B}^{-1}$ leads to

$$\underline{\mathbf{r}} = \frac{\Delta t}{\beta_0 Re} \mathbf{A}\mathbf{B}^{-1}\mathbf{D}^T(\underline{p}^n - \underline{p}^{n-1}) = O(\Delta t^2). \quad (36)$$

Since there is a factor of Δt^{-1} in front of the velocity in (34) the local truncation error incurred by neglecting $\underline{\mathbf{r}}$ is $O(\Delta t^3)$, and the method is formally second-order accurate in time. This is borne out by numerical results at the end of this section.

Dropping the residual in (34) and formally carrying out a single round of block Gaussian elimination leads to the reformulated Stokes problem to be solved at each time step:

$$\begin{bmatrix} \mathbf{H} & -\frac{\Delta t}{\beta_0}\mathbf{H}\mathbf{B}^{-1}\mathbf{D}^T \\ \mathbf{0} & E \end{bmatrix} \begin{pmatrix} \underline{\mathbf{u}}^n \\ \underline{p}^n - \underline{p}^{n-1} \end{pmatrix} = \begin{pmatrix} \mathbf{B}\underline{\mathbf{f}} + \mathbf{D}^T \underline{p}^{n-1} \\ \underline{g} \end{pmatrix}, \quad (37)$$

where

$$E \equiv \frac{\Delta t}{\beta_0} \sum_{i=1}^d D_i B^{-1} D_i^T, \quad (38)$$

and \underline{g} is the modified inhomogeneity arising from Gaussian elimination. The advantage of the splitting procedure is that matrix vector products involving E can be computed *without* system solves, since B is diagonal. It can be shown that E is SPD, save a possible one-dimensional null space associated with the hydrostatic pressure mode in cases where $\partial\Omega_o = \emptyset$.

To summarize, time advancement of the Navier-Stokes equations involves: advancing the convective terms through the solution of (4), solving for the viscous contribution in the construction of g (37), solving for the pressure (37-38), and finally computing the divergence-free solution, $\underline{\mathbf{u}}^n$ (37). The systems involving H and E are solved iteratively. Typically, $\Delta t/Re$ is small, implying that $H = \frac{1}{Re}A + \frac{\beta_0}{\Delta t}B$ is strongly diagonally dominant and readily amenable to solution via Jacobi preconditioned conjugate gradients. E is the more challenging operator. As shown in [3, 39] it has properties similar to a Poisson operator, and is often referred to as a ‘‘consistent Poisson operator.’’ Its condition number is independent of Δt and Re . In the remaining sections of this paper, we address preconditioning strategies for the E operator.

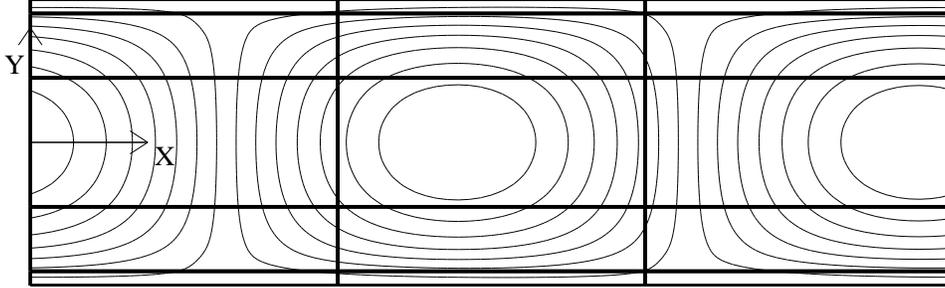


Figure 3: Spectral element mesh ($K = 15$) and perturbation streamfunction for Orr-Sommerfeld problem.

2.6 Numerical study of the Orr-Sommerfeld problem

To close this section, we demonstrate the accuracy of the numerical method outlined above by computing the growth rates of small amplitude two-dimensional Tollmien-Schlichting (TS) waves in plane Poiseuille flow, following [38, 28]. This benchmark is useful for verifying both temporal and spatial convergence in that it is an unforced time-dependent Navier-Stokes problem for which accurate solutions are available from linear stability theory.

The geometry consists of two walls separated by a distance $2h$ with periodic boundary conditions in the streamwise direction at $x = 0$ and $x = 2\pi h$. The initial condition is

$$\begin{aligned} u^0(x, y) &= 1 - \left(\frac{y}{h}\right)^2 + \epsilon \hat{u} \\ v^0(x, y) &= \epsilon \hat{v}, \end{aligned} \quad (39)$$

where (\hat{u}, \hat{v}) corresponds to the only unstable eigensolution (wave number unity) of the Orr-Sommerfeld equation. The Reynolds number is $Re = U_c h / \nu = 7500$, based upon the centerline velocity, $U_c = 1$. A constant body force is applied to sustain the mean flow. The perturbation velocity is normalized to $|\hat{\mathbf{u}}| = 1$ and ϵ is set to .00001. The spectral element discretization consists of $K = 15$ subdomains in the configuration shown in Fig. 3. The calculations are performed in 64-bit arithmetic with iteration tolerances set to 10^{-13} in order to observe high-order spatial and second-order temporal convergence rates.

According to linear theory, the energy of the perturbation:

$$E(t) = \int_0^{2\pi} \int_{-1}^1 \left\{ (1 - y^2 - u)^2 + v^2 \right\} dy dx \quad (40)$$

should grow as $e^{2\omega_i t}$, where $\omega_i = .002234976$. Following [28], we take as a measure of error the quantity: $error(t) = e^{2\omega_i t} - E(t_i)/E(0)$, where $E(t_i)$ is derived from our computed Navier-Stokes solution at times $t_1 = 25.1437$ and $t_2 = 50.2873$, corresponding to one and

Table 1: Temporal Convergence, O-S problem: $K = 15$, $N = 17$

Δt	$E(t_1)$	$error_1$	$E(t_2)$	$error_2$	$error_g$
0.20000	1.13215830	0.013208	1.28480818	0.032758	0.11213
0.10000	1.12278058	0.003830	1.26120270	0.009153	0.03347
0.05000	1.11997704	0.001026	1.25446507	0.002415	0.00896
0.02500	1.11921579	0.000265	1.25266990	0.000620	0.00230
0.01250	1.11901773	0.000067	1.25220678	0.000157	0.00058
0.00625	1.11896723	0.000017	1.25208917	0.000040	0.00015

Table 2: Spatial Convergence, O-S problem: $K = 15$, $\Delta t = .003125$

N	$E(t_1)$	$error_1$	$E(t_2)$	$error_2$	$error_g$
7	1.11498657	0.003963	1.21465285	0.037396	0.313602
9	1.11519192	0.003758	1.24838788	0.003661	0.001820
11	1.11910382	0.000153	1.25303597	0.000986	0.004407
13	1.11896714	0.000016	1.25205855	0.000009	0.000097
15	1.11895646	0.000006	1.25206398	0.000014	0.000041

two periods of oscillation for the TS waves, respectively. In addition, we compute the error in the growth rate at time $t = 50$ according to:

$$error_g = \frac{1}{\omega_i} \left| \omega_i - \frac{1}{2\Delta t} \ln \left(\frac{E(50)}{E(50 - \Delta t)} \right) \right| \quad (41)$$

Table 1 shows the computed energies and the error for varying Δt with $N = 17$. The error is clearly $O(\Delta t^2)$. Table 2 shows the results for a spatial convergence study carried out using a third-order accurate (though less stable) implementation of the above time stepping algorithms, with $\Delta t = .003125$. The convergence is non-monotonic due to the fact that the growth rates are oscillating about the analytical value. However, spectral convergence is clearly attained.

3 Laplacian based preconditioning for the pressure

Since E (38) is symmetric positive definite, it is natural to consider conjugate gradient iteration as an iterative solver [22]. The development of a fast solver requires finding an SPD preconditioning matrix, M^{-1} , which can be inexpensively applied and which renders the condition number of $M^{-1}E$ as close to unity as possible. In this section, we examine the potential of three discrete Laplace operators as alternatives to working with the more cumbersome matrix E in the ultimate development of our additive Schwarz preconditioner. The first two are based upon the Gauss-Lobatto (\mathcal{H}^1) mesh, for which a natural triangulation

of the domain exists. The third is based upon the Gauss points and is more difficult to develop. However, it is surprisingly effective, and has proven to be the key to developing a fast solver for this problem.

3.1 Preconditioners

One of the fundamental difficulties in developing a preconditioner for E is that the inter-domain coupling leads to a fully connected graph between adjacent subdomains, i.e., every degree-of-freedom in a given subdomain is coupled to every degree-of-freedom in its adjacent subdomains. If formed explicitly, E would correspondingly have approximately $3^d(N-1)^d$ nonzeros per row for a geometry consisting of a regularly structured array of subdomains. Thus, E has no natural interface structure, or separator, such as commonly found in operators derived from finite element bases of compact support. Consequently, E cannot be readily treated by substructuring, or Schur complement, approaches as in [6, 37]. In addition, there are no boundary conditions associated directly with the pressure, as boundary conditions for the Stokes problem are applied in the velocity space.

Despite the \mathcal{L}^2 approximation space used for the pressure it is nonetheless clear that E is related to a Laplacian on Ω . This can be seen most readily from the continuous analog by taking the divergence of the momentum equation in (1) to derive a Poisson equation for the pressure. The article of Blair Perot [3] explores this relationship in detail.

As the starting point, we consider the variational formulation of the Poisson problem:
Find $\tilde{p} \in Z_p^h(\Omega)$, such that:

$$(\nabla v, \nabla \tilde{p})_* = (v, f)_* \quad \forall v \in Z_p^h(\Omega) \quad (42)$$

where the space Z_p^h and quadrature rules $(\cdot, \cdot)_*$ will be defined implicitly in the subsequent choice of basis for the discrete problem. The boundary conditions for \tilde{p} are:

$$\nabla \tilde{p} \cdot \hat{\mathbf{n}} = 0 \text{ on } \partial\Omega_v, \quad \tilde{p} = 0 \text{ on } \partial\Omega_o, \quad (43)$$

which are derived from standard pressure-splitting formulations of the Navier-Stokes equations, e.g. [3, 19, 32]. Once the basis and quadrature formulae are chosen, the preconditioner is constructed by combining the resultant stiffness matrix, A_* , with a suitable restriction operator, R , to yield:

$$M_*^{-1} = R^T A_*^{-1} R. \quad (44)$$

In all the cases we consider, R^T is simply an interpolant from the nodal basis of Z_p^h onto the nodal points for the Lagrangian basis of Y^N . In general, exact solution of the system A_* will be too expensive for problems of practical interest and an appropriate surrogate must be developed. However, to illustrate the potential of the Laplacian based preconditioning

strategy, we make a brief preliminary study of three choices for (A_*, R^T) when exact solvers are used for A_*^{-1} .

The first choice is:

$$M_s^{-1} = R_{vp}^T A_s^{-1} R_{vp} \quad (45)$$

where R_{vp}^T is the spectral interpolation operator from the velocity mesh to the pressure mesh, and A_s is the spectral element Laplacian on the velocity mesh, with associated boundary conditions of homogeneous Dirichlet at outflow and homogeneous Neumann elsewhere. Other than a change in the mask, $A_s = \mathcal{M}_p Q^T A_L Q \mathcal{M}_p$ is identical to A . In particular, the inter-domain connectivity effected by Q is preserved and software development costs are consequently minimal. We note that A_s is a viable preconditioner as it is roughly a factor of d less costly to apply than E since E first maps a scalar field to a vector field, then maps back to a scalar (38), whereas A_s is simply a scalar-scalar mapping. Moreover, the Laplacian is readily available in the context of any Navier-Stokes solver as it is an essential building block of any such code.

As a second choice, we consider

$$M_t^{-1} = R_{vp}^T A_t^{-1} R_{vp} \quad (46)$$

where A_t is the discrete Laplacian derived from linear finite elements based upon a triangulation of the *Gauss-Lobatto* points of the originating spectral element discretization, and R_{vp}^T is the spectral interpolant described above. The technique of using low-order operators as preconditioners for spectral methods was first proposed by Orszag [33], and has been studied by many others, including numerous articles by Deville, Mund and coworkers, e.g., [9, 10, 11]. The principal advantage of this approach is that the degradation in convergence rate is offset by the increased sparsity of A_t over A_s . The connectivity of A_t can be effected in the same manner as A_s and one therefore only needs to develop a triangulation of the Gauss-Lobatto points on the reference subdomain.

Finally, we consider a third choice:

$$M_g^{-1} = R_{\text{aug}}^T A_g^{-1} R_{\text{aug}} \quad (47)$$

where A_g is the discrete Laplacian derived from linear finite elements based upon a triangulation of the *Gauss* (pressure) points, augmented with additional points to be described further in the next section. Since the pressure is already represented by nodal values on the Gauss points, R_{aug}^T is a Boolean operator which is essentially the identity matrix, save for the addition of some zero rows to account for the additional vertices. Consequently, the cost of applying R_{aug}^T is nil, in contrast to R_{vp}^T for which the cost scales as $O(KN^{d+1})$. Moreover, the dimension of A_g is approximately $K(N-1)^d$, whereas the dimension of A_s and A_t is approximately KN^d (not $K(N+1)^d$, due to the shared interface variables). Note

that it is not possible to interpolate onto a lower-order grid (e.g., a polynomial of degree $N - 1$) in developing M_s as the boundary conditions would reduce the dimension of A_s to less than $K(N - 1)^d$, thereby making M_s^{-1} singular.

3.2 Tiling the Gauss points

One of the attractive features of triangulating the velocity (Gauss-Lobatto) mesh is that the points which reside on $\partial\Omega_k$ are shared by adjacent subdomains and thus serve to define the connectivity between those domains. In the case of the Gauss points, there are no shared points between domains and the triangulation used to interconnect subdomains can be rather arbitrary, particularly at the corners where an unspecified number of subdomains must be joined together.

In order to provide a mechanism for imposing Dirichlet boundary conditions and avoid development of a general purpose triangulation routine, we propose the following tiling scheme for connecting adjacent sets of Gauss points together. The “tile” is defined in a self-contained manner that reduces the global triangulation problem to a local one which is then solved for each subdomain. Within each subdomain, the triangulation proceeds in three steps. First the $(N - 2)^2$ rectangles associated with the nodes (η_i, η_j) are divided into two triangles each. Then, the edges and boundary regions are triangulated. Finally, the regions in the vicinity of the vertices are triangulated.

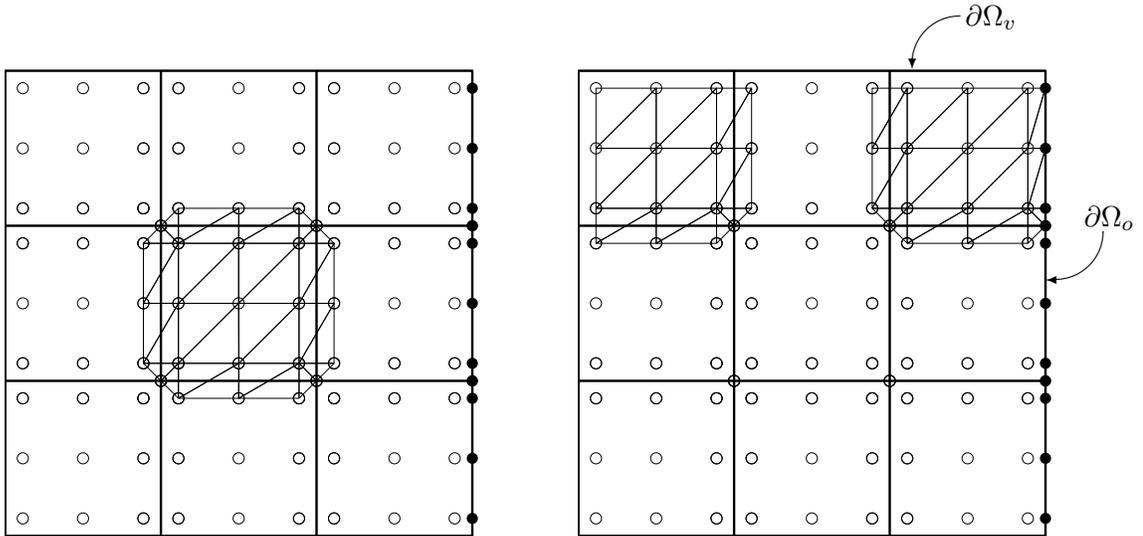


Figure 4: Spectral element configuration ($K = 9$, $N = 4$) showing standard “tiling” of the Gauss points for a subdomain with four neighboring subdomains (left); a subdomain with two Neumann boundaries (right, upper left); and a subdomain with one Neumann and one Dirichlet boundary (right, upper right). Solid circles denote Dirichlet boundary nodes.

To triangulate the edge regions, we begin by noting that, for the conforming spectral element discretizations considered here, each domain has either a neighbor or a boundary condition associated with each of $2d$ edges (faces). In the case where the edge separates two subdomains, we connect corresponding nodes on either side of the separating edge to form $N - 2$ quadrilaterals. In the case where the edge corresponds to Dirichlet boundary conditions (for \tilde{p}), we add “ghost” points along the edge at locations $(-1, \eta_j)$ for a left boundary, $(+1, \eta_j)$ for a right boundary, $(\eta_i, -1)$ for a lower boundary, and $(\eta_i, +1)$ for an upper boundary, all defined with respect to the reference subdomain. Because these points coincide with the Dirichlet boundary, no new degrees-of-freedom are introduced by this procedure. If the boundary is a Neumann boundary, the tile is not extended beyond the Gauss points; the boundary condition is applied directly to these points even though they are actually separated from the boundary by a distance of $O(N^{-2})$. Once the edges are extended, the resultant quadrilaterals are subdivided into triangles. The dividing line extends from the highest globally numbered vertex in each quadrilateral to ensure that the triangulation is consistent from one tile to the next.

The vertex regions are a potential source of difficulty in developing a standardized subdomain interconnection scheme because an unspecified number of subdomains may be joined at each vertex. While triangulation of this region presents no particular difficulty in \mathbb{R}^2 , finding a tetrahedral decomposition subject to conformity with the surrounding region in \mathbb{R}^3 is not an easy task. This difficulty can be eliminated in either \mathbb{R}^2 or \mathbb{R}^3 by augmenting the Gauss points with the set of points comprising the subdomain vertices. In \mathbb{R}^2 , a line is then extended from each of the four vertices to the corresponding (closest) Gauss point interior to Ω^k . At each vertex, additional lines are extended to the nearest point on the extended edges defined in the preceding paragraph. If either of the subdomain edges corresponds to a Neumann boundary (no extension) then the augmented vertex is eliminated and the corresponding region is not triangulated.

Examples of the basic tiles are shown in Fig. 4 for a $K = 9$ subdomain configuration. The union of all such tilings defines the triangulation for the finite element operator A_g . Note the addition of four “ghost” vertices coincident with the four vertices of the central subdomain. These are degrees-of-freedom in the A_g system, but are not enumerated in the range of the prolongation operator R_{aug}^T , i.e., the values of the residual are allowed to float at these points, but they make no direct contribution to the output of the preconditioner. A more expensive alternative would be to have the restriction operator, R_{aug} , take the average of the surrounding values in computing the right hand side at the ghost points.

After the tile has been defined for each domain, a clean-up routine removes all multiply defined triangles. The entire procedure can be defined in the context of the (extended) reference element without having to consider numerous special cases, and is readily extended to \mathbb{R}^3 .

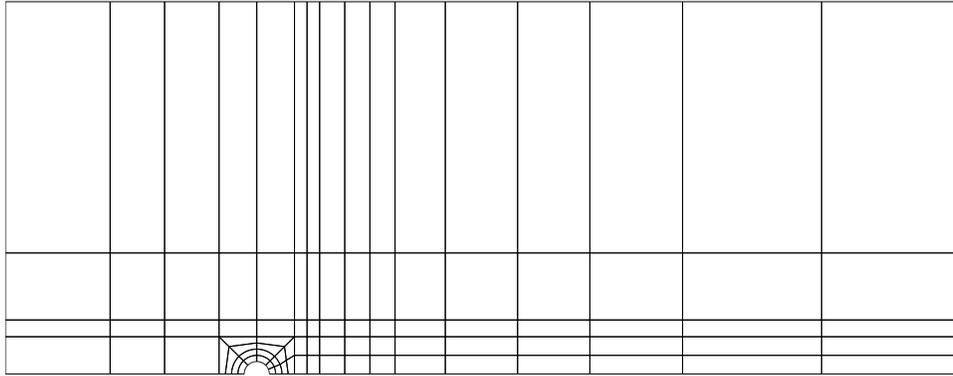


Figure 5: Spectral element mesh ($K = 93$) for iterative convergence study.

3.3 Numerical results for Laplacian preconditioning

We now test the preconditioners M_s , M_t , and M_g for the problem of start-up flow past a cylinder in the half-domain $\Omega = [-10, 28] \times [0, 15]$. Fig. 5 shows the subdomain configuration for the baseline case of $K = 93$. A cylinder of diameter $D = 1$ is centered at the origin. The Reynolds number is $Re = DU/\nu = 5000$, where $(U, 0)$ is the free-stream velocity taken as the initial condition and inflow boundary condition at $x = -10$. Symmetry boundary conditions are imposed at $y = 0$ and $y = 15$ with Neumann-velocity (outflow) boundary conditions at $x = 28$. The free stream velocity is $U = 1$ and the time step is $\Delta t = .025$.

While this problem is hardly canonical, it is typical of the types of domains encountered in practice. It contains a spectrum of subdomain shapes and sizes which, combined, have a pronounced effect on the convergence behavior of many iterative solvers. A comparison of this particular problem in contrast to a regular tessellation of a square has been carried out in [18]. It was noted there, and in the thesis of Couzy [8] that irregular domains can lead to significant degradation of the convergence rate of the deflation based solver we have been using up to the present [17, 39]. We remark that the first time step is the most difficult as the full spectrum of the pressure must be computed, rather than just the perturbation from the previous step. Moreover, the initial condition, \mathbf{u}^0 , is not even divergence free which makes the projection onto a divergence free space all the more difficult. In practice, in subsequent time steps, the right hand side will be largely devoid of low-wave number components as they can be eliminated by projecting onto the space of previous solutions [15].

Table 3 shows the number of iterations required to reduce the residual of the E system on the first step by 10^{-5} for the three Laplacian based preconditioners just described. The $K = 372$ and $K = 1488$ meshes are obtained through successive quarterings of the elements in the base configuration. The dimension of the E system is $K(N - 1)^2$. Note that the

Table 3: Laplacian preconditioners for cylinder flow.

	$N = 7$			$N = 9$		
K	M_s	M_t	M_g	M_s	M_t	M_g
93	15	24	16	12	27	16
372	12	26	18	12	28	17
1488	12	26	18	12	29	18

spectral preconditioner exhibits no K dependence in the iteration count for this problem, while the finite element preconditioners exhibit mild degradation in convergence rate.

Of the two finite element preconditioners, it is clear that M_g is to be preferred over M_t . Not only does it yield a lower iteration count, it has a much less expensive restriction operator, and a lower dimensional system to be solved with each iteration. The only drawback of M_g is that it is more complicated to triangulate the Gauss points than the Gauss-Lobatto points, something which is of particular consequence in \mathbb{R}^3 .

In the preceding examples, the choice of finite element spaces was not arbitrary. Given the locally tensor-product spaces of the spectral element method, the natural choice for a low-order finite element preconditioner would be bilinear (trilinear in \mathbb{R}^3) finite elements. However, linear elements based upon triangles are not only easier to implement in complex geometries, but also have favorable conditioning properties over their bilinear counterparts. This point is addressed in the Appendix.

4 Overlapping Schwarz Method

One of the principal aims of employing the Laplacian based preconditioner is to derive a preconditioner for the E system which leads to a bounded iteration count. As developed in the previous section, the Laplacian based preconditioner requires the solution of a system in A_g for *each* iteration. In practice, this is too expensive from the standpoint of both storage and operation count, particularly in the three-dimensional case. However, the Laplacian preconditioning strategy effectively reduces the problem of solving the E system to one of solving a low-order \mathcal{H}^1 operator, a problem which is well understood. In this section, we develop an overlapping additive Schwarz preconditioner, M_o^{-1} , based upon linear finite element discretizations of Poisson's equation. M_o^{-1} can either be used as a preconditioner for solving the system A_g in (47), resulting in a nested iteration, or directly as a preconditioner for E . We have found the latter approach to significantly outperform the former and therefore consider it as the only option.

Though a fairly recent development [12, 13, 30], the literature on additive Schwarz preconditioners is quite vast and the subject is covered in depth in at least four recent books

[1, 23, 41, 42]. The overlapping methods are based upon a decomposition of the domain into overlapping subdomains and solving local Poisson problems within each subdomain. In addition, a coarse grid problem is solved which involves few degrees-of-freedom, but covers the entire domain. The solutions from the subproblems are then added together to form the output of the preconditioner.

For finite element solution of Poisson's equation, Dryja and Widlund [12] have shown that overlapping Schwarz procedures with a coarse grid space lead to bounded iteration counts in the case of quasi-uniform subdomains. Cai [4] has derived a uniform bound in which this restriction is lifted, provided that the coarse grid triangulation is quasi-uniform. Presently, we are insisting that our coarse grid space coincide with the subdomains which are predetermined by the original spectral element tessellation of Ω , and hence do not have control over the uniformity of either the subdomains or the coarse grid space. We investigate the consequences of this in the numerical results at the end of this section and suggest some remedies.

The additive Schwarz preconditioner is expressed as (e.g. [1, 42]):

$$M_o^{-1} = R_0^T A_0^{-1} R_0 + \sum_{k=1}^K R_k^T A_k^{-1} R_k \quad , \quad (48)$$

which is the sum of the global coarse grid operator (subscript 0) and local subdomain operators (subscript k). The method has a natural parallel aspect in that the subdomain problems can be solved independently. While the coarse grid problem is not trivially parallelized it is usually of sufficiently modest size to admit fast direct solution methods, even in parallel (e.g. [16, 18, 21]).

The coarse grid operator, A_0 , is constructed as the finite element Laplacian derived from linear elements whose vertices are coincident with the subdomain vertices, i.e., in \mathbb{R}^2 , each subdomain is partitioned into two triangles. The boundary conditions are inherited from the variational problem (42). The prolongation operator, R_0^T , is defined somewhat differently from standard (purely finite element based) domain decomposition methods. R_0^T is the operator which interpolates the coarse grid solution onto the tensor product array of Gauss points *in the reference element* $\hat{\Omega}$. Thus, even in situations where some points lie outside of the coarse grid triangulation in physical space, there is a natural interpolation operator to transfer the coarse grid solution to the fine grid. This approach, coupled with our treatment of Neumann boundary conditions for the fine grid, ensures that the coarse grid covers the fine grid, which is required for stability according to Chan, Smith, and Zou [7]

To define the subdomain operators, let V_k , $k = 1, \dots, K$ be subsets of the augmented vertex set defined by the tiling procedure of the previous section, and let \hat{R}_k be the Boolean restriction matrix defined such that $\hat{R}_k \underline{u}$ returns the nodal values of \underline{u} corresponding to the

vertex set V_k . The subdomain operators are defined as

$$A_k \equiv \hat{R}_k A_g \hat{R}_k^T . \quad (49)$$

A_k is consequently a principal submatrix of A_g , the linear finite element Laplacian on the augmented Gauss points. To transfer the solution to the Gauss points, we define the operator $R_k^T \equiv R_{\text{aug}}^T \hat{R}_k^T$, which is essentially a map from V_k to the Gauss points with extension by zero to points not in V_k .

For each subdomain, overlap is generated by starting with the vertex set V_k^0 comprising the Gauss points interior to Ω^k . Let N_o be the desired amount of overlap. Then, for $i = 1, \dots, N_o$, set:

$$V_k^i = V_k^{i-1} \cup \mathcal{N}_k^{i-1} , \quad (50)$$

and finally, set $V_k = V_k^{N_o}$. Here \mathcal{N}_k^i is the set of vertices neighboring V_k^i . A vertex is defined to be neighboring V_k^i if it is connected by at least one edge in the triangulation to any vertex in V_k^i . Note that $N_o = 0$ is admissible and corresponds to block-Jacobi preconditioning.

4.1 Numerical results for the additive Schwarz preconditioner

We reconsider the cylinder problem ($N = 7$, $K = 93$, 372 , and 1488) of the previous section to illustrate several aspects of the additive Schwarz scheme. As a baseline, we compare the additive Schwarz procedure with the deflation based solver we have been using to the present.

Briefly, for the baseline results, the consistent Poisson problem (37) is solved via the two-level iteration scheme developed by Rønquist [39] in which a coarse-grid operator is folded into a global conjugate-gradient iteration through deflation [29, 31]. A coarse (subscript c) and fine (subscript f) decomposition is effected through a subdomain-motivated prolongation operator $J \in \mathbb{R}^{n_p \times m}$, where $n_p = K(N-1)^d$ is the number of pressure degrees-of-freedom, and m is the dimension of the coarse-grid approximation space. The column space of the prolongation operator J is intended to approximate the span of the low eigenmodes of the E system. The pressure is then expressed as $\underline{p} = J\underline{p}_c + \underline{p}_f$, leading to an algebraic reformulation of the original problem as solvable fine and coarse subproblems,

$$E_f \underline{p}_f = \underline{g} - J E_c^{-1} J^T \underline{g} , \quad (51)$$

$$E_c \underline{p}_c = J^T \underline{g} - J^T E \underline{p}_f , \quad (52)$$

respectively. Here $E_f = E - E J E_c^{-1} J^T E$, and $E_c = J^T E J$. Each application of the fine grid operator requires two multiplications by E , plus the solution of the relatively small ($m \times m$) system, E_c . The fine system (51) is solved by conjugate-gradient iteration restricted to the complement of $\mathcal{R}(J)$, where $\mathcal{R}(\cdot)$ denotes the column space of the argument.

Table 4: Iteration count and CPU time for the deflation scheme

K	constant		linear		quadratic	
	iter	CPU	iter	CPU	iter	CPU
93	126	17.	80	12.	60	10.
372	216	125.	120	75.	84	67.
1488	327	845.	159	467.	108	437.

Once \underline{p}_f is established, the coarse-grid problem is solved (directly) for \underline{p}_c , and the procedure is complete. With appropriate application of a local, element-based preconditioner to E_f , the condition number of the fine system is significantly reduced relative to the originating E matrix.

Since we are comparing CPU times as well as iteration counts, a few remarks regarding implementation are in order. In spectral element methods, seventy to ninety percent of the operation count is devoted to matrix-matrix products ($m \times m$'s) invoked in the application of one-dimensional spectral operators (e.g., (18-19)). We have found that an $m \times m$ implemented with a completely unrolled inner-product loop is roughly three times faster than most implementations *dgemm* for the values of N considered. The subdomain and coarse grid solves are computed using a version of the banded SPD linpack solver, *dgesl*, modified to multiply by the inverse of the diagonal elements, rather than divide, for each solve. Each system is reordered using reverse Cuthill-McKee to minimize bandwidth. All CPU times (in seconds) are for an SGI Onyx with sufficient RAM to ensure that disk swapping is not required.

Table 4 shows the number of iterations and CPU time required to reduce the pressure residual for the first time step by five orders of magnitude for the deflation scheme using piecewise (discontinuous) constant, bilinear, and biquadratic coarse grid spaces. Our parallel production code [16] has been based upon the piecewise constant prolongation operator. The higher-order coarse grid spaces were studied in [18] in an attempt to improve this scheme. Although the results of Table 4 show a two-fold reduction in CPU time for the first time step, these extensions yielded only a thirty percent reduction in subsequent steps, and would be even less effective in \mathbb{R}^3 due to the rapid increase in the dimension of the coarse grid problem. These two considerations motivated the present study.

Table 5 shows the iteration count and CPU times for the additive Schwarz procedure with varying amounts of overlap. The $N_o = 0$ column corresponds to block-Jacobi preconditioning (no overlap). Introducing a minimal amount of overlap ($N_o = 1$) reduces the iteration count almost two-fold and the CPU time slightly less than two-fold. Increasing the overlap to $N_o = 3$ does not yield much further improvement, and memory considerations consequently favor the choice $N_o = 1$. The importance of the coarse grid is demonstrated

Table 5: Performance of the additive Schwarz algorithm

K	$N_o = 0$		$N_o = 1$		$N_o = 3$		$A_0 = 0$		$N_o - var.$	
	iter	CPU	iter	CPU	iter	CPU	iter	CPU	iter	CPU
93	121	10.	64	5.9	49	5.6	169	19.	45	5.4
372	203	74.	106	43.	73	39.	364	193.	75	32.
1488	303	470.	158	274.	107	242.	802	1798.	102	183.
1744	183	329	97	199.	68	180.	801	2089.	78	163.

in column 4 for the case $N_o = 3$; in the absence of a coarse grid the iteration count roughly doubles for each successive quad-level refinement from $K = 93$ to 1488.

Although the results of the overlapping Schwarz scheme are impressive, the fact that the iteration counts in the first three columns are not bounded with K is somewhat disappointing. Fig. 6 shows the residual history for $M_o^{-1}E$ with $N_o = 1$ for $K = 93, 372,$ and 1488. All three curves show the same rapid initial convergence. However, after the 25th iteration, there appears to be a persistent mode which prevents the convergence from otherwise being order-independent. This mode can be viewed by taking the difference between the 25th and final iterates, as shown in Fig. 7 for $K = 1488$. The mode is clearly centered about the high-aspect ratio (HAR) subdomains and suggests that elimination of such domains might improve convergence.

An easy way to alleviate the aspect ratio problem is to simply subdivide the HAR subdomains along the appropriate axis. Since we are presently considering only conforming discretizations, this implies splitting subdomains along the entire length of the domain. The

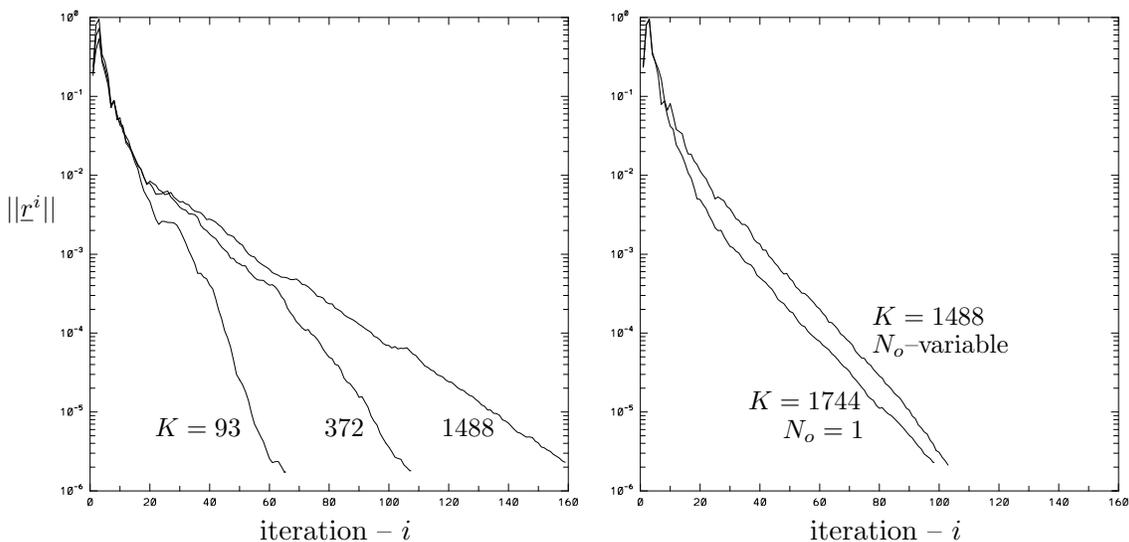


Figure 6: Residual history showing deterioration of convergence rate for $K = 93 \rightarrow 1488$ (left). Convergence is improved for regularized geometry and/or variable overlap (right).

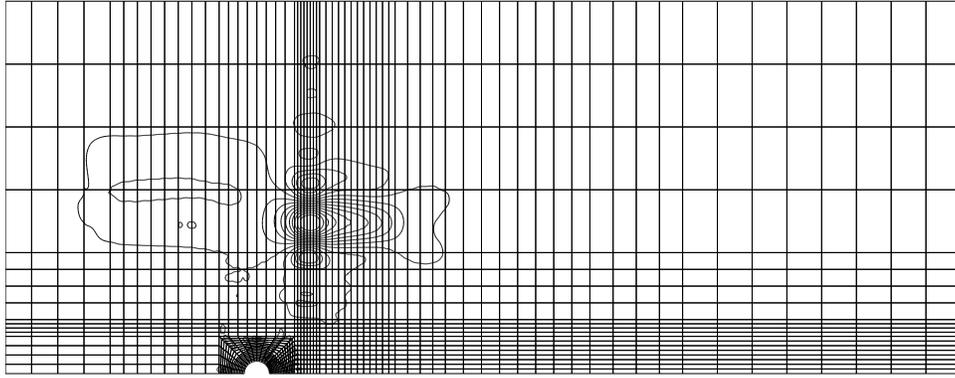


Figure 7: Spectral element mesh ($K = 1488$) showing resistant pressure mode centered about high-aspect ratio subdomains.

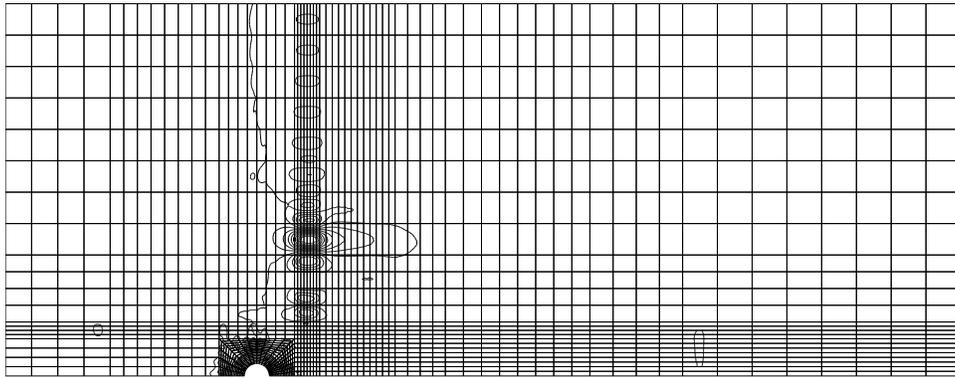


Figure 8: Regularized spectral element mesh ($K = 1744$) showing resistant pressure mode for flow past a cylinder.

resultant configuration is shown in Fig. 8, along with contours of the mode $\underline{p}^{final} - \underline{p}^{25}$. Despite the fact that the $K = 1744$ problem has more degrees-of-freedom, the convergence rate is superior to the $K = 1488$ case, as seen in Fig. 6 and in the last row of Table 5.

A more viable approach to curing the aspect ratio problem is to change the preconditioner. One possibility is to regularize the coarse grid space by adding additional nodes to the coarse grid triangulation, as suggested by Cai [4]. However, this requires retriangulating the coarse grid and is inherently nonlocal. A more readily implemented scheme, proposed by Widlund [43], is to increase the overlap only for subdomains having high aspect ratio. After a brief inspection of the distribution of aspect ratios for several of our meshes, we developed the following, quite arbitrary, heuristic for determining the amount of overlap for

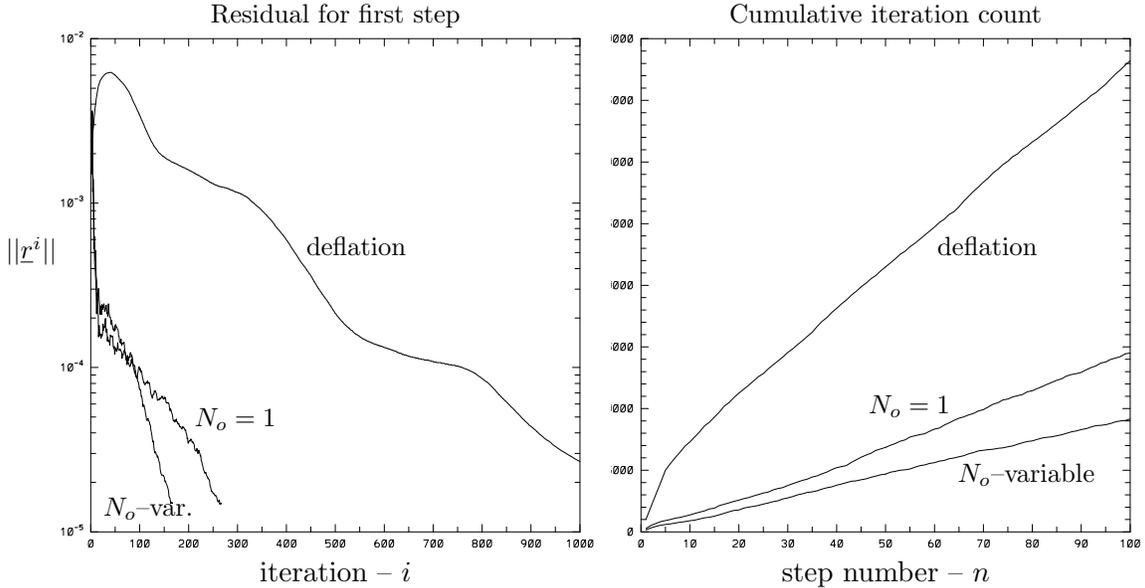


Figure 9: Residual history for the first step (left) and cumulative pressure iteration count for 100 steps (right) for $K = 6112$ cylinder calculation of Fig. 1, restarted from $t = 3.0$.

each domain, k :

$$N_o^k = \begin{cases} 3 & \text{if } AR \geq 10 \\ 2 & \text{if } 10 > AR \geq 5 \\ 1 & \text{otherwise} \end{cases} \quad (53)$$

The convergence behavior for $K = 1488$ using variable overlap is shown on the right in Fig. 6. It can be seen that convergence rates for the variable overlap case and the regularized geometry ($K = 1744$) do not exhibit the degradation associated with HAR subdomains. The last column of Table 5 shows the marked improvement of variable overlap over minimal overlap ($N_o = 1$). The iteration counts are roughly equal to those for $N_o = 3$, while the CPU time is lower due to the reduced work in the subdomain solves. Note that variable overlap does not provide significant gain (163s vs. 199s) for the regularized mesh.

Finally we remark that, in some cases, the overlapping Schwarz method improves the conditioning to the point where 32-bit arithmetic suffices where 64-bit arithmetic had been previously required for convergence of the pressure solve. One such example is the $Re = 9500$ cylinder calculation of Fig. 1. The number of pressure degrees-of-freedom for this ($K = 6112$, $N = 9$) problem is 390,000. Fig. 9 (left) shows the first time step convergence behavior for the overlapping Schwarz preconditioner with $N_o = 1$ and N_o varying according to (53). Also shown is the convergence behavior for the deflation scheme with piecewise constant prolongation. On the right is the cumulative number of pressure iterations for the first one hundred time steps. Note that the initial time steps are the most costly, as

subsequent solutions are computed from initial guesses based upon previous solutions [15]. Also recall that the deflation based scheme requires *two* matrix vector products with E for each iteration.

5 Conclusions

We have developed an additive overlapping Schwarz preconditioner for the computationally challenging pressure operator which arises when an Uzawa decoupling procedure is applied to the $\mathbb{P}_N - \mathbb{P}_{N-2}$ spectral element formulation of the incompressible Navier-Stokes equations. The pressure preconditioner is derived from local finite element Laplacians based upon a triangulation of the Gauss (pressure) points, coupled with a global coarse grid operator based upon a triangulation of the spectral element vertices. The Schwarz procedure yielded significantly improved convergence rates over previously employed deflation/block-Jacobi based schemes. It was also found that the convergence of the Schwarz procedure deteriorated significantly in the presence of high-aspect ratio subdomains. For the case where the number of such subdomains is small, it was found that the problem could be easily remedied by introducing more overlap for high-aspect ratio domains. The overall Navier-Stokes solution times for several production runs has been reduced by a factor of five with the development of this preconditioner.

Future research efforts include extending the preconditioner to three dimensions and implementing it on distributed memory architectures. In addition, the need for more sophisticated coarse grid spaces which address the aspect ratio problem should be assessed.

Acknowledgements

The author would like to thank Dr. Einar Rønquist and Prof. Olof Widlund for valuable discussions in the course of this work, and Gerald W. Kruse for providing the results for Fig. 1. This work was supported by the NSF under Grant ASC-9405403 and by the AFOSR under Grant F49620-95-1-0074. Computational support was provided on the Intel Delta by Caltech under NSF Cooperative agreement CCR-8809615 and by AFOSR Grant F49620-95-1-0100.

Appendix: Choice of finite element spaces

In the development of the low-order Laplacian in Section 3, the choice of finite element spaces was not arbitrary. Given the locally tensor-product spaces of the spectral element method, the natural choice for a low-order finite element preconditioner would be bilinear (trilinear in \mathbb{R}^3) finite elements. However, *linear* elements based upon a triangulation of

the Gauss or Gauss-Lobatto points yield a lower condition number in the mono-domain case. This is seemingly at odds with the results of Deville and Mund, who in [10] found the contrary, so a careful study of this issue is pertinent.

As a point of reference, we first consider finite element preconditioning of the spectral element Poisson operator, A , derived for the mono-domain case on the reference subdomain, $\Omega = \hat{\Omega}$, with homogeneous Dirichlet boundary conditions on $\partial\Omega$. This case was studied in detail by Deville and Mund in [10] for the spectral collocation method.

In addition to the linear spaces based on triangles, we will consider as preconditioners the Laplacian based upon bilinear finite elements, A_f , and its lumped-mass counterpart, A_l . The easiest way to define these operators is as a tensor product of one-dimensional finite element operators. Consider the space of piecewise linear functions, $\phi_i(\xi)$, $\xi \in [-1, +1]$:

$$\phi_i(\xi) = \begin{cases} \frac{\xi - \xi_{i-1}}{\xi_i - \xi_{i-1}} & \xi_{i-1} \leq \xi < \xi_i \\ \frac{\xi - \xi_{i+1}}{\xi_i - \xi_{i+1}} & \xi_i \leq \xi < \xi_{i+1} \\ 0 & \text{otherwise} \end{cases}, \quad i \in \{1, \dots, N-1\}, \quad (54)$$

where the ξ_i 's are the Gauss-Lobatto points. The bilinear forms for the homogeneous Dirichlet problem give rise to the tridiagonal stiffness and mass matrices:

$$\tilde{A}_{ij} = \left(\frac{d\phi_i}{d\xi}, \frac{d\phi_j}{d\xi} \right), \quad \tilde{B}_{ij} = (\phi_i, \phi_j), \quad i, j \in \{1, \dots, N-1\}^2, \quad (55)$$

where $(., .)$ is the one-dimensional counterpart of (6). In addition, we define the lumped mass matrix for the one dimensional case:

$$\bar{B}_{ij} = \delta_{ij} \int_{-1}^{+1} \phi_j(\xi) d\xi \quad i, j \in \{1, \dots, N-1\}^2, \quad (56)$$

which corresponds to a diagonal matrix with nonzeros equal to the row sums of \tilde{B} .

If the interior grid points on the reference domain are numbered lexicographically, then the two-dimensional stiffness and mass matrices for the bilinear finite elements are:

$$A_f = \tilde{B} \otimes \tilde{A} + \tilde{A} \otimes \tilde{B}, \quad B_f = \tilde{B} \otimes \tilde{B}, \quad (57)$$

respectively. The lumped-mass stiffness and mass matrices in two dimensions are:

$$A_l = \bar{B} \otimes \tilde{A} + \tilde{A} \otimes \bar{B}, \quad B_l = \bar{B} \otimes \bar{B}. \quad (58)$$

The stiffness matrix A_l is identical to the Laplacian derived from linear triangles, A_t , or the classic five-point finite difference stencil. However, whereas the lumped-mass matrix, B_l , is diagonal, the linear finite element mass matrix, B_t , has six nonzeros per row for a regular array of triangles. Since both A_f and B_f result from Kronecker products of tridiagonal matrices, they yield a full nine-point stencil.

Table 6 shows the condition number and extreme eigenvalues for the preconditioned system $A_*^{-1}A$. The values of $\kappa(A_t^{-1}A)$ and $\kappa(A_l^{-1}A)$ appear to be approaching the bound

Table 6: Symmetric preconditioner for A , mono-domain case.

N	$A_f^{-1}A$			$A_l^{-1}A$			$A_t^{-1}A$		
	κ	λ_{\max}	λ_{\min}	κ	λ_{\max}	λ_{\min}	κ	λ_{\max}	λ_{\min}
4	3.62	4.63	1.28	1.55	1.84	1.18	1.55	1.84	1.18
6	4.84	5.44	1.12	1.80	1.95	1.08	1.80	1.95	1.08
8	5.47	5.86	1.07	1.95	2.04	1.05	1.95	2.04	1.05
10	5.86	6.12	1.05	2.04	2.10	1.03	2.04	2.10	1.03
20	6.64	6.71	1.01	2.24	2.26	1.01	2.24	2.26	1.01
40	7.02	7.04	1.00	2.35	2.35	1.00	2.35	2.35	1.00

Table 7: Nonsymmetric preconditioner for A , mono-domain case.

N	$A_f^{-1}B_fB^{-1}A$			$A_l^{-1}B_lB^{-1}A$			$A_t^{-1}B_tB^{-1}A$		
	κ	λ_{\max}	λ_{\min}	κ	λ_{\max}	λ_{\min}	κ	λ_{\max}	λ_{\min}
4	1.46	0.92	0.63	1.55	1.55	1.00	2.18	0.90	0.41
6	1.41	0.96	0.68	1.80	1.80	1.00	2.49	0.96	0.39
8	1.35	0.98	0.72	1.95	1.95	1.00	2.63	0.98	0.37
10	1.36	0.99	0.72	2.04	2.04	1.00	2.69	0.99	0.37
20	1.37	1.00	0.73	2.24	2.24	1.00	2.73	1.00	0.37
40	1.40	1.00	0.71	2.35	2.34	1.00	2.79	1.00	0.36

$\pi^2/4$ which has been found by several authors [10, 35]. Note that preconditioners having a “star” type of stencil yield a lower condition number than the seemingly natural choice of bilinear finite elements. It was noted earlier that, for undeformed geometries, the spectral Laplacian also has a star-shaped stencil due to the diagonal mass matrix, and this appears to be the reason for the superior conditioning by the linear and lumped mass operators in this symmetric case.

By contrast, if one preconditions the original *differential operator* (42) by taking into account the mass matrices as done by Deville and Mund [10], then bilinear finite elements are superior to linear finite elements. Table 7 shows the extreme *real* part of the eigenvalues for $A_*^{-1}B_*B^{-1}A$ and associated “condition number,” $\kappa = \max \operatorname{Re}(\lambda)/\min \operatorname{Re}(\lambda)$. (We note that the spectrum of several of the operators was complex. However, the maximum imaginary magnitude was less than ten percent of the minimum real part the spectrum.) In this case, the asymptotic value of $\kappa = 1.4$ for the bilinear element is significantly lower than found for the symmetric preconditioners, as noted in [10], and suggests that non-symmetric iteration schemes such as GMRES might be of interest in conjunction with this preconditioner.

We now consider the same eigenvalue test for the consistent Poisson operator, E , generated with homogeneous Dirichlet boundary conditions for the velocity. In this case, the finite element preconditioners are Neumann operators (null space of dimension unity) corresponding to a triangulation of the *Gauss* points, (η_i, η_j) . Note that the triangulation does not extend to $\partial\Omega$. The basis for the bilinear finite elements is again a tensor product of one-dimensional functions:

$$\psi_i(\eta) = \begin{cases} \frac{\eta - \eta_{i-1}}{\eta_i - \eta_{i-1}} & \eta_{i-1} \leq \eta < \eta_i \\ \frac{\eta - \eta_{i+1}}{\eta_i - \eta_{i+1}} & \eta_i \leq \eta < \eta_{i+1} \\ 0 & \textit{otherwise} \end{cases}, i \in \{1, \dots, N-1\}, \quad (59)$$

with the restriction that $\eta \in [\eta_1, \eta_{N-1}]$, implying that the basis functions ψ_1 and ψ_{N-1} are non-vanishing at the endpoints. The formation of the two-dimensional operators proceeds exactly as before, yielding bilinear stiffness and mass matrices, E_f and B_f , and lumped-mass counterparts E_l and B_l . The linear triangle based operators are denoted E_t and B_t , respectively. Both E and the preconditioners share a null space containing the constant function, but otherwise have real positive eigenvalues. The condition numbers for the generalized eigenvalue problem were computed by premultiplying E by the pseudoinverse of the preconditioner and then solving the standard eigenvalue problem.

Tables 8 and 9 show the condition numbers for the E operator, analogous to the previous two tables for A . In this case, the condition numbers are not bounded. However, from the results of Table 3 of Section 3, it is clear that finite element preconditioning is effective and apparently only weakly dependent upon N in the multidomain case. We have

Table 8: Symmetric preconditioner for E , mono-domain case.

N	$E_f^{-1}E$			$E_l^{-1}E$			$E_t^{-1}E$		
	κ	λ_{\max}	λ_{\min}	κ	λ_{\max}	λ_{\min}	κ	λ_{\max}	λ_{\min}
4	7.75	10.46	1.35	2.99	3.49	1.17	2.99	3.49	1.17
6	11.23	12.06	1.07	4.08	4.15	1.02	4.08	4.15	1.02
8	13.32	13.76	1.03	5.49	4.74	0.86	5.49	4.74	0.86
10	14.83	15.10	1.02	7.06	5.16	0.73	7.06	5.16	0.73
20	23.98	18.30	0.76	15.94	6.14	0.39	15.94	6.14	0.39
40	50.57	20.14	0.40	35.66	6.72	0.19	35.66	6.72	0.19

Table 9: Nonsymmetric preconditioner for E , mono-domain case.

N	$E_f^{-1}B_fB^{-1}E$			$E_l^{-1}B_lB^{-1}E$			$E_t^{-1}B_tB^{-1}E$		
	κ	λ_{\max}	λ_{\min}	κ	λ_{\max}	λ_{\min}	κ	λ_{\max}	λ_{\min}
4	1.84	0.17	.094	3.93	0.56	.142	3.68	0.81	.219
6	3.81	0.30	.079	5.82	0.91	.156	6.54	1.22	.187
8	5.33	0.36	.068	7.78	1.11	.143	9.21	1.48	.161
10	6.91	0.40	.058	9.81	1.24	.126	11.91	1.65	.139
20	15.78	0.50	.032	21.25	1.52	.072	25.90	2.03	.078
40	35.47	0.55	.016	46.78	1.68	.036	56.05	2.24	.040

also verified that the growth with N is not due to the boundary conditions, by repeating the experiments for $A_*^{-1}A$ using Neumann boundary conditions instead of Dirichlet; the results were largely unchanged from those in Tables 6 and 7. The key information to draw from Tables 8 and 9 is that linear triangles are superior to bilinear elements in the case of the symmetric preconditioner and that, for E , the symmetric linear preconditioner seems to be as good as the nonsymmetric bilinear preconditioner.

We have also investigated the case of deformed elements and found the results for both A and E largely unchanged. As a result of this study we have based our preconditioner upon linear finite elements, which also eases the coding for deformed geometries. Note that E_t is identical to the preconditioner A_g of Sections 3 and 4.

References

- [1] R. Barret, M. Berry, T. Chan, J. Demmel, J. Donato, J. Dongarra, V. Eijkhout, R. Pozo, C. Romine, and H. van der Vorst *Templates*, SIAM, 1994.
- [2] C. Bernardi and Y. Maday, “A collocation method over staggered grids for the Stokes problem,” *Int. J. Numer. Meth. Fluids*, **8** (1988), pp. 537-557.
- [3] J. Blair Perot “An analysis of the fractional step method,” *J. Comput. Phys.*, **108** (1993) pp. 51-58.
- [4] X.-C. Cai, “The use of pointwise interpolation in domain decomposition with non-nested meshes,” *SIAM J. Sci. Comput.*, **14** (1995) pp. 250-256.
- [5] M. Casarin “Quasi-optimal Schwarz methods for the conforming spectral element discretization.” *Tech.Rep. 705*, Dept. Comp. Sci., Courant Inst., NYU (1995).
- [6] M. Casarin “Schwarz preconditioners for spectral and mortar finite element methods with applications to incompressible fluids,” Ph.D. Thesis, Courant Institute of Math. Sci., NYU, 1996.
- [7] T.F. Chan, B.F. Smith and J. Zou, “Overlapping Schwarz methods on unstructured meshes using non-matching coarse grids,” *CAM Report 94-08*, Dept. of Math., UCLA 1994.
- [8] W. Couzy, “Spectral element discretization of the unsteady Navier-Stokes equations and its iterative solution on parallel computers,” Thèse No. 1380, École Polytechnique Fédérale de Lausanne (1995).
- [9] P. Demaret and M.O. Deville, “Chebyshev pseudo-spectral solution of the Stokes equations using finite element preconditioning”, *J. Comput. Phys.*, **83** (1989) pp. 463-484.
- [10] M.O. Deville and E.H. Mund, ”Finite element preconditioning for pseudospectral solutions of elliptic problems”, *SIAM J. Statis. Comput.*, **11**(2), (1990) pp. 311-42.
- [11] M.O. Deville and E.H. Mund, “Fourier analysis of finite element preconditioned collocation schemes”, *SIAM J. Sci. Statist. Comput.*, **13** (1992) pp. 596-610.
- [12] M. Dryja and O.B. Widlund, “An additive variant of the Schwarz alternating method for the case of many subregions,” *Tech. Rep. 339* Dept. Comp. Sci., Courant Inst., NYU(1987).
- [13] M. Dryja, “An additive Schwarz algorithm for two- and three-dimensional finite element elliptic problems,” in *3rd Int. Symp. on Domain Decomposition Methods*, T.F. Chan, R. Glowinski, J. Periaux and O .B. Widlund, eds., SIAM, pp. 168-172 (1989).
- [14] P.F. Fischer, “Spectral Element Solution of the Navier-Stokes Equations on High Performance Distributed-Memory Parallel Processors,” Ph.D. Thesis, Massachusetts Institute of Technology, 1989.
- [15] P.F. Fischer, “Projection techniques for iterative solution of $Ax = \underline{b}$ with successive right-hand sides.” *ICASE Report No. 93-90*, NASA CR-191571.
- [16] P.F. Fischer, “Domain Decomposition Methods for Large Scale Parallel Navier-Stokes Calculations” in *Proceedings of the Sixth International Conference on Domain Decomposition Methods for Partial Differential Equations, Como Italy*, A. Quarteroni, Ed., AMS (1994) p.313-322.

- [17] P.F. Fischer and E.M. Rønquist, “Spectral Element Methods for Large Scale Parallel Navier-Stokes Calculations” *Comp. Meth. Appl. Mech. Engr.*, 1994.
- [18] P.F. Fischer, “Parallel multi-level solvers for spectral element methods” in *Proceedings of Intl. Conf. on Spectral and High-Order Methods '95, Houston, TX*, A.V. Ilin and L.R. Scott, Eds. (1996).
- [19] N.K. Ghaddar, K. Korczak, B.B. Mikic, and A.T. Patera, Numerical investigation of incompressible flow in grooved channels. Part 1: Stability and self-sustained oscillations. *J. Fluid Mech.*, **163**, 1986, 99-127.
- [20] V. Girault and P.A. Raviart, *Finite Element Approximation of the Navier-Stokes Equations*, Springer, 1986.
- [21] W.D. Gropp, *Parallel Computing and Domain Decomposition* in Fifth Conference on Domain Decomposition Methods for Partial Differential Equations T.F. Chan, D.E. Keyes, G.A. Meurant, J.S. Scroggs, and R.G. Voigt, eds., SIAM, Philadelphia, PA, 1992.
- [22] G.H. Golub and C.F. Van Loan, *Matrix Computations*, John Hopkins University Press, Baltimore, Maryland, 1983.
- [23] D.E. Keyes, Y. Saad, and D.G. Truhlar, *Domain-Based Parallelism and Problem Decomposition Methods in Computational Science and Engineering* SIAM, 1995.
- [24] P. Koumoutsakos and A. Leonard “High-resolution simulations of the flow around an impulsively started cylinder using vortex methods” *J. Fluid Mech.* **296** 10 (1995) pp. 1-38.
- [25] Y. Maday and A.T. Patera, Spectral element methods for the Navier-Stokes equations, in *State of the Art Surveys in Computational Mechanics* (Ed. A.K. Noor), ASME, New York, (1989).
- [26] Y. Maday, A.T. Patera, and E.M. Rønquist, *An operator-integration-factor splitting method for time-dependent problems: application to incompressible fluid flow.* *J. Sci. Comput.*, **5**(4), (1990) pp. 310-37.
- [27] Y. Maday, A.T. Patera, and E.M. Rønquist, “The $P_N - P_{N-2}$ method for the approximation of the Stokes problem,” *Numer. Math.*, to appear.
- [28] M.R. Malik, T.A. Zang, and M.Y. Hussaini, “A spectral collocation method for the Navier-Stokes equations,” “Chebyshev pseudo-spectral solution of the Stokes equations using *J. Comput. Phys.*, **61** (1985) pp. 64-88.
- [29] L. Mansfield, *On the use of deflation to improve the convergence of conjugate gradient iteration* *Comm. in Appl. Numer. Meth.*, **4**, (1988) pp. 151-56.
- [30] S.V. Nepomnyaschikh. “Domain decomposition and Schwarz methods in a subspace fo the approximate solution of elliptic boundary value problems,” Ph.D. Thesis, Computing Center of the Siberian Branch of the USSR Acad. of Sci., Novosibirsk, USSR (1986).
- [31] R.A. Nicolaides, *Deflation of conjugate gradients with applications to boundary value problems* *SIAM J. Numer. Anal.*, **24**(2), (1987) pp. 355-65.
- [32] S.A. Orszag and L.C. Kells (1980): Transition to turbulence in plane Poiseuille flow and plane Couette flow. *J. Fluid Mech.* **96**, 159-205

- [33] S.A. Orszag, "Spectral methods for problems in complex geometries" *J. Comput. Phys.* **37**, (1980) pp 70-92.
- [34] S.S. Pahl, "Schwarz type domain decomposition methods for spectral element discretizations." Masters Thesis, Dept. of Comput. and Appl. Mathematics, Univ. of the Witwatersrand, Johannesburg, South Africa, Dec. 1993.
- [35] S.V. Parter and E.E. Rothman, "Preconditioning Legendre spectral collocation approximation to elliptic problems," *SIAM J. Numer. Anal.*, **32**(2), (1995) pp. 333-85.
- [36] A.T. Patera, "A spectral element method for fluid dynamics; Laminar flow in a channel expansion," *J. Comput. Phys.*, **54** (1984), 468-488.
- [37] 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**(4), (1996).
- [38] M.M. Rai and P. Moin, "Direct simulations of turbulent flow using finite difference schemes," presented at the 27th Aerospace Sciences Meeting, Reno NV, AIAA-89-0369, (1989).
- [39] E.M. Rønquist, *A Domain Decomposition Method for Elliptic Boundary Value Problems: Application to Unsteady Incompressible Fluid Flow*, in Fifth Conference on Domain Decomposition Methods for Partial Differential Equations T.F. Chan, D.E. Keyes, G.A. Meurant, J.S. Scroggs, and R.G. Voigt, eds., SIAM, Philadelphia, PA, 1992.
- [40] E.M. Rønquist, "A domain decomposition solver for the steady Navier-Stokes equations" in *Proceedings of Intl. Conf. on Spectral and High-Order Methods '95, Houston, TX*, A.V. Ilin and L.R. Scott, Eds. (1996).
- [41] Y. Saad, *Iterative methods for sparse linear systems*, PWS Publishing, Boston, 1996.
- [42] B. Smith, P. Bjørstad, and W. Gropp, *Domain Decomposition*, Cambridge University Press, 1996.
- [43] O.B. Widlund. Personal communication.