A post-processing technique for stabilizing the discontinuous pressure projection operator in marginally-resolved incompressible inviscid flow

Sumedh M. Joshi¹

*Center for Applied Mathematics,
657 Rhodes Hall, Cornell University,
Ithaca NY, 14850
smj96@cornell.edu
7135945390*

Peter J. Diamessis

*School of Civil and Environmental Engineering,
105 Hollister Hall, Cornell University,
Ithaca NY, 14850*

Derek T. Steinmoeller

*Department of Applied Mathematics
University of Waterloo
Ontario, CA*

Marek Stastna

*Department of Applied Mathematics
University of Waterloo
Ontario, CA*

Greg N. Thomsen

*Applied Research Laboratories
The University of Texas at Austin
Austin, TX*

Abstract

A method for post-processing the velocity after a pressure projection is developed that helps to maintain stability in an under-resolved, inviscid, discontinuous element-based simulation for use in environmental fluid mechanics process studies. The post-processing method is

¹Corresponding author.

needed because of spurious divergence growth at element interfaces due to the discontinuous nature of the discretization used. This spurious divergence eventually leads to a numerical instability. Previous work has shown that a discontinuous element-local projection onto the space of divergence-free basis functions is capable of stabilizing the projection method, but the discontinuity inherent in this technique may lead to instability in under-resolved simulations. By enforcing inter-element discontinuity and requiring a divergence-free result in the weak sense only, a new post-processing technique is developed that simultaneously improves smoothness and reduces divergence in the pressure-projected velocity field at the same time. When compared against a non-post-processed velocity field, the post-processed velocity field remains stable far longer and exhibits better smoothness and conservation properties.

1. Introduction

Pressure-projection methods are a class of numerical techniques that decouple the solution of pressure from velocity in the numerical simulation of an incompressible flow. These methods, first developed by Chorin [3], overcome the difficult problem of the pressure-velocity coupling through the incompressibility constraint, and are widely used in computational fluid dynamics for simulating time-dependent flow [12, 11]. The fundamental idea in the pressure-projection method is that the momentum equation and the incompressibility constraint are time-integrated separately and in sequence. First the momentum equations are solved to advance the velocity to an interstitial time. Next, a projection operation maps these velocities onto the space of divergence-free functions by way of solving a Poisson equation and advances the velocity to the next time. Thus, projection methods are also sometimes known as fractional-step or time-splitting methods [16, 17, 20] and are widely-used solving viscous, time-dependent, incompressible flows. Because of their ubiquity, much work has been done to construct consistent boundary conditions for the pressure and velocity [17, 16, 24, 2, 19, 26]

and stable spatial discretizations of each constituent operator in the time-splitting [5]. These efforts are largely focused on avoiding the spurious divergence boundary layers that can form when inconsistent boundary conditions are used in the time-splitting within the projection method [16, 26]. A concise review of pressure projection methods can be found in Ref. [11].

While in theory the velocity, once projected, is supposed to be divergence-free, it has recently been observed that in the discontinuous Galerkin (DG) formulation the projection operation may contain non-solenoidal eigenmodes [25]. Thus, the projected velocity fields themselves are not exactly divergence-free, which can lead to numerical instability and inaccuracy. A possible explanation for the cause (which is discussed in section 2.2) of this is that this is due to the discontinuous spatial discretization of the Laplacian operator within the pressure projection method. To remedy this, Steinmoeller et al. [25] construct a post-processing method that explicitly projects the velocity field onto a computed basis set that is exactly the null-space basis of the discrete DG divergence operator, and show that this projection eliminates the spurious divergence due to the non-solenoidal eigenmodes of the pressure projection update operator. While this post-processing technique is effective in eliminating non-solenoidal components of the velocity, by virtue of the discontinuous nature of the spatial discretization, the exact null-space projection (ENP) as in Steinmoeller et al. does not take into account continuity between elements. ENP is a projection that is entirely local to an element.

In this work, we ask whether this locality is ever problematic, and if so what should be done to address it. While this question is broad in scope, it is shown that at least in one instance of a marginally resolved simulation the discontinuity in the ENP can lead to instability. As a remedy, a modified null-space projection technique is used as a post-processing method that explicitly takes into account inter-element continuity in a regularized least-squares sense. It is shown that in this instance the so-called weak null-space projection method appears to yield greater stability as a post-processing technique. In this regard, this line of reasoning is

in parallel to previous efforts which focused on constructing boundary conditions to avoid the spurious divergence boundary layers [16, 26], but focusing instead on the spurious divergence that is observed to form at inter-element boundaries.

To study the relative merits of the exact null-space projection and its modification that captures inter-element continuity, the 2D incompressible Euler equations are used as a proxy for the full incompressible Navier-Stokes equations. The incompressible Euler equations model an inviscid incompressible flow with a stratified background density profile that is not dependent on time and only depends on the vertical direction; a perturbation density, $\rho'(\mathbf{x}, t)$, is overlaid on the background stratification and as noted does vary in time and space. The equations are given as

$$\frac{\partial \mathbf{u}}{\partial t} = \mathbf{u} \cdot \nabla \mathbf{u} - \frac{1}{\rho_0} \nabla p - g \frac{\rho'}{\rho_0} \mathbf{e}_z \quad (1)$$

$$\nabla \cdot \mathbf{u} = 0 \quad (2)$$

$$\frac{\partial \rho'}{\partial t} = -\nabla \cdot \mathbf{u}(\rho' + \bar{\rho}) \quad (3)$$

where $\rho(x, z, t) = \rho_0 + \bar{\rho}(z) + \rho'(x, z, t)$ is density stratified in the Boussinesq approximation [18] with $\bar{\rho}(z)$ the background stratification, $\mathbf{u}(x, z, t)$ the velocity, p the pressure, and $\mathbf{e}_z$ the unit vector in the vertical direction. It is assumed that $\rho_0 \gg \bar{\rho} \gg \rho'$. In the pressure-projection method Eqs. (1)-(2) are not solved directly. Instead, a projection operator $\mathbb{P}$ is used to solve

$$\frac{\partial \mathbf{u}}{\partial t} = \mathbb{P} \left(\mathbf{u} \cdot \nabla \mathbf{u} - g \frac{\rho'}{\rho_0} \mathbf{e}_z \right) \quad (4)$$

in which the projection operator $\mathbb{P}$ is defined as

$$\mathbb{P} \mathbf{u} := \mathbf{u} - \nabla \Delta^{-1} (\nabla \cdot \mathbf{u}). \quad (5)$$

This decouples the solution of pressure from velocity, and allows for a sequential solution algorithm in which the velocity is advected by the nonlinear term prior to being projected into a divergence-free space by $\mathbb{P}$.

While viscosity is neglected here, the difficulties encountered in stabilizing the pressure projection method for the Navier-Stokes equations are all encountered here as well. The presence of viscosity will only aid in damping the numerical instabilities driven by the nonlinear advection term, so stability in the inviscid case is more difficult to achieve due to the absence of physical viscous dissipation. In fact, the discussion of stability and under-resolution is primarily manifest in advection-dominated Navier-Stokes simulations in which a broad range of scales is present due to the lack of strong viscosity; in this sense then, the incompressible Euler simulations presented here capture the essence of the difficulty in simulating incompressible fully-viscous flows.

The numerical method used to model the density-stratified inviscid incompressible Euler equations is the spectral multi-domain penalty method (SMPM), a high-order discontinuous collocation-based variant of the spectral element method [13, 10] which has been previously shown to be effective in simulating high-Reynolds number environmental flows [7, 8] using the pressure-projection method. In particular, we will use as an example the inviscid propagation of an internal solitary wave in a density-stratified channel as a test bed for evaluating the efficacy of the various post-processing methods. In these simulations, the initial conditions propagate as waves in a non-dispersive non-dissipative fashion through the domain while retaining their form. Thus, the degree to which these solutions maintain their structure is a good heuristic for the efficacy of these post-processing methods. It should be mentioned that although the SMPM is chosen to demonstrate these methods, these ideas are readily extensible to discontinuous Galerkin (DG) discretizations, and especially to high-order DG discretizations.

The paper is organized as follows. In Section 2 the exact and weak null-space projection

methods, their motivation and the notation used are described. In Section 3 two simulations with each of the three methods are conducted and compared. Both simulations are of the same propagating internal solitary wave in tank, and the simulations differ in their mesh resolution. In Section 4 contains a discussion of the results presented in Section 3, along with a computational assessment of the spectrum and numerical conditioning properties of all of the methods compared in this paper. Finally, we conclude with a short discussion of applicability to other numerical methods as well as a discussion of future work related to the ideas outlined herein.

2. Methods

This section summarizes the the exact null-space projection (ENP) as outlined in Ref. [25], the weak null-space projection which is the contribution of this work, and the numerical method.

2.1. Numerical method and notation

In the 2D SMPM, each element is assumed to be smoothly and invertibly mapped from the unit square $[-1, 1] \times [-1, 1]$ and the element connectivity is logically cartesian (each element has a single neighbor in each of the North, South, East, and West directions). Within each element lies a two-dimensional Gauss-Lobatto-Legendre (GLL) grid; denote as n the number of GLL points per direction per element, and m_x and m_z the number of x and z elements in the grid². Thus, the total number of grid points is $r = n^2 m_x m_z$. On the GLL grid, a two-dimensional nodal Lagrange interpolant basis of polynomial order $n + 1$ is constructed such that each basis function has unit value on one of the n^2 GLL points and zero on all of the others. This nodal basis is used for approximating functions and their derivatives which are calculated by way of spectral differentiation matrices [4] which compute derivatives of nodally-represented functions by multiplying the nodal values by derivatives of the Lagrange

²Here z is the vertical direction as is convention in environmental fluid mechanics.

interpolants themselves. The SMPM is a discontinuous method and so C^0/C^1 inter-element continuity and boundary conditions are only weakly enforced, which yields operator matrices that are amenable to domain decomposition algorithms for obtaining fast solutions [15].

On this grid, denote the differentiation operators $D_x, D_z \in \mathbb{R}^{r \times r}$. As unsymmetric matrices that operate element-by-element, each is permutation-equivalent to a block-diagonal matrix. The operations of gradient and divergence are defined and denoted as

$$G = \begin{bmatrix} D_x \\ D_z \end{bmatrix} \in \mathbb{R}^{2r \times r} \quad D = [D_x, D_z] \in \mathbb{R}^{r \times 2r} \quad (6)$$

and it is emphasized that unlike in symmetric discretizations, $G \neq D^T$. Finally we denote the Laplacian as $L \in \mathbb{R}^{r \times r}$. With this notation, the discrete pressure projection operator $\mathbb{P} \in \mathbb{R}^{2r \times 2r}$ is defined

$$\mathbb{P} = I - GL^{-1}D. \quad (7)$$

Finally we denote as $\mathcal{N}$ the nonlinear term in the momentum equations

$$\mathcal{N}(\mathbf{u}) = \mathbf{u} \cdot \nabla \mathbf{u} - g \frac{\rho'}{\rho_0} \mathbf{e}_z. \quad (8)$$

2.2. No post-processing

Using first-order forward Euler time-stepping, a single time-step of the pressure-projection algorithm is given in Algorithm 1, where Δt denotes the time-step. In practice an adaptive first-order forward Euler time-stepping method was employed, but for simplicity we do not include the adaptivity in the algorithmic summary. The choice of a low-accuracy time-stepping method was made primarily for convenience. This work is focused on studying spurious divergence and discontinuity due to the discontinuous spatial discretization; a high-order time-stepping method would perhaps delay the onset of these artifacts but would not eliminate them.

Algorithm 1 A single time-step of the pressure projection method with no post-processing.

Require: u, ρ at time t .

- 1: $u^* = u + \Delta t \mathcal{N}(u)$
 - 2: $\rho \leftarrow \rho - \Delta t Du(\rho' + \bar{\rho})$
 - 3: $u \leftarrow \mathbb{P}u^*$
 - 4: **return** u, ρ at time $t + \Delta t$.
-

2.2.1. The projection method in a discontinuous discretization

Recalling that $\mathbb{P} = I - GL^{-1}D$ where L is the Poisson operator, it is worth pausing to discuss why this projection method fails to yield an exactly divergence-free velocity field in the discrete case. Consider the definition of $L \in \mathbb{R}^{r \times r}$ as a matrix. L certainly contains the product of the divergence with the gradient, $DG \in \mathbb{R}^{r \times r}$, but must also contain terms related to inter-element continuity conditions and boundary conditions. Denoting these additional terms as $C \in \mathbb{R}^{r \times r}$, write L as

$$L = DG + C. \quad (9)$$

In the case of the spectral multi-domain penalty method C represents the penalty conditions that enforce inter-element continuity and boundary conditions [10]; in the case of discontinuous Galerkin C represents the inter-element fluxes in the DG residual [25]. Now consider the divergence of the updated velocity,

$$D\mathbb{P}u = D(I - GL^{-1}D)u = Du - DGL^{-1}Du. \quad (10)$$

If $L = DG$, then it is clear that $D\mathbb{P}u = 0$, but, as already stated, L contains other components encapsulated in C , and thus Eq. (10) imposes conditions on C to ensure that $D\mathbb{P}u = 0$. Denoting as $v = L^{-1}Du$ the solution of the Laplace problem, note that v satisfies $Lv = Du$ and so

$$DGv = Du - Cv. \quad (11)$$

Substituting this expression into Eq. (10) for $L^{-1}Du$, we obtain

$$D\mathbb{P}u = D(I - GL^{-1}D)u = Du - (Du - Cv) = Cv. \quad (12)$$

This shows that the divergence of the projected velocity is equal to the discontinuity (as measured by C) in the solution $v \in \mathbb{R}^r$ of the Laplace problem $Lv = Du$. As discussed above, the inter-element continuity conditions embedded within C may be of flux-type or C^0/C^1 type depending on the discretization method being employed. Nevertheless, assuming a discontinuous discretization, in general it is not true that $Cv = 0$, thus in general the projection method outlined in Algorithm 1 does not yield an exactly divergence-free velocity; the divergence of u is equal to the discontinuity of the Poisson solution v . Finally it should be noted that although the aliasing caused by the nonlinear advection term is often the culprit of non-solenoidal vector fields, and thus instability, this analysis (corroborating that in Ref. [25]) shows that there is another source of spurious divergence: the discontinuity embedded within the pressure projection operator for discontinuous element-based methods.

2.3. Exact null-space projection

To address the fact that the divergence-free condition is not satisfied (as described in Section 2.2.1), Steinmoeller et al. [25] construct a post-processing method that is now summarized in this section. Noting that as a matrix D is element-local, computing the null space basis of D is trivially parallelizable. A local SVD of each block of D can be computed in parallel and the right singular vectors corresponding to zero singular values can be retained. Assuming that D has a right nullity of dimension k , denote as $N \in \mathbb{R}^{2r \times k}$ the matrix whose columns are these null-value singular vectors. N as a matrix is orthonormal and $||DN|| = 0$. The columns of N are exactly all of the solenoidal vectors supported on the grid, and so, as in Ref. [25], these vectors are used to project the velocity u into the column space of N by

$$u = NN^T \mathbb{P} u^*. \quad (13)$$

Denoting the composition of the pressure projection with the exact null-space projection as $\mathbb{P}_e = NN^T \mathbb{P}$, observe that, by construction, $\mathbb{P}_e u$ is exactly divergence-free for any velocity $u \in \mathbb{R}^{2r}$ since $\|DNx\| = 0$ for all vectors $x \in \mathbb{R}^{2n}$. This algorithm, which we call *exact null-space projection*, is summarized in Algorithm 2.

Algorithm 2 A single time-step of the pressure projection method with exact null-space projection.

Require: u, ρ at time t .

- 1: $u^* = u + \Delta t \mathcal{N}(u)$
 - 2: $\rho \leftarrow \rho - \Delta t Du(\rho' + \bar{\rho})$
 - 3: $u \leftarrow \mathbb{P}_e u^*$
 - 4: **return** u, ρ at time $t + \Delta t$.
-

Finally, it is worth noting that $\mathbb{P}_e$ is a projection in its own right since it is a composition of two projections.

2.4. Weak null-space projection

While on one hand, the element-local nature of the divergence operator D allows for the easy computation of its null-space basis N , this element-local nature also means that NN^T as a projection may introduce discontinuities at the element interfaces. Generally speaking, such discontinuities lead to steep gradients at the element interfaces and consequently instability. Motivated by this intuition, we now describe a second method of post processing the pressure-projection operator $\mathbb{P}$ which we call *weak null-space projection* (WNP).

First, consider the following minimization problem in which k is the column dimension of N , and $u \in \mathbb{R}^{2r}$ is a velocity vector:

$$\lambda' = \arg \min_{\lambda \in \mathbb{R}^k} \|N\lambda - u\|^2. \quad (14)$$

This least squares problem seeks a divergence-free velocity vector $N\lambda$ that minimizes the distance to u . Of course the solution to this problem is $\lambda' = N^T u$ and thus the divergence-free velocity is $u' = NN^T u$, which is the ENP method outlined in the previous section.

Treating the ENP as a least squares problem, we seek to penalize the minimization of the residual for discontinuity in the velocity. Suppose there is an inter-element discontinuity operator $E \in \mathbb{R}^{2r \times 2r}$ that computes the inter-element discontinuity in a velocity $u \in \mathbb{R}^{2r}$. For example, in the discontinuous Galerkin approximation, E might be the numerical fluxes between elements, or, in the SMPM, it may be the inter-element penalty conditions. In any case, to penalize both divergence and discontinuity in the residual, we write the following least squares problem:

$$u' = \arg \min_{u' \in \mathbb{R}^{2r}} \|u' - u\|^2 + \alpha_1 \|Du'\|^2 + \alpha_2 \|Eu'\|^2. \quad (15)$$

The solution of this least squares problem is given by the solution u' to its normal equations,

$$(I + \alpha_1 D^T D + \alpha_2 E^T E)u' = u, \quad (16)$$

and a unique solution exists for all $\alpha_1, \alpha_2 > 0$, since the normal equations are then a symmetric, positive definite linear system and thus invertible. Composing the solution to the normal equations with the pressure-projection, we denote the modified projection operator as

$$\mathbb{P}_w = (I + \alpha_1 D^T D + \alpha_2 E^T E)^{-1} \mathbb{P} \quad (17)$$

and note the algorithm for one time-step of this method is summarized in Algorithm 3.

Algorithm 3 A single time-step of the pressure projection method with weak null-space projection.

Require: u, ρ at time t .

- 1: $u^* = u + \Delta t \mathcal{N}(u)$
 - 2: $\rho \leftarrow \rho - \Delta t Du(\rho' + \bar{\rho})$
 - 3: $u \leftarrow \mathbb{P}_w u^*$
 - 4: **return** u, ρ at time $t + \Delta t$.
-

Notice a few things about this method. First, if $\alpha_1 = \alpha_2 = 0$ then $\mathbb{P}_w = \mathbb{P}$ in which case no post-processing is being done and Algorithm 3 is identical to Algorithm 1. Second, notice that $\mathbb{P}_w$ is not a projection in the linear algebraic sense since $\mathbb{P}_w^2 \neq \mathbb{P}_w$. However, notice that all eigenvalues of $(I + \alpha_1 D^T D + \alpha_2 E^T E)^{-1}$ lie within $[0, 1]$; since projections also satisfy this property (although they *only* have zero and one as eigenvalues, a fact that follows from idempotence) we claim that $\mathbb{P}_w$ is similar to a projection in the sense that it cannot inflate the norm of any vector it is applied to. In other words, $\|\mathbb{P}_w\|_2 < 1$. Finally, note that the solution of the regularized least squares problem defined in Eq. (15) will *not* be exactly divergence-free, but the degree to which the divergence-free condition is satisfied can be controlled by the magnitude of α_1 , with a similar argument holding for continuity and α_2 .

One may well ask why an exact projection onto the intersection of the spaces of divergence-free and continuous functions was not constructed. This is primarily because the construction of an orthogonal basis for the space of continuous functions is impractical, as it requires the factorization of matrices of dimension r that are not element-local (e.g. obtaining the right null space basis of E). Moreover, it is not clear that such a projection would even be desired, given the suggested numerical benefits of allowing a discontinuous discretization [6, 22, 13].

Notice also that the weak null-space projection method is a generalization, in a sense, of both the pressure-projection and exact null-space projection. If $\alpha_1 = \alpha_2 = 0$ then there is no post-processing being done and this method is identical to Algorithm 1. If $\alpha_2 = 0$ then this method emphasizes minimizing divergence at the cost of continuity, and its solution approaches that of Algorithm 2 in the limit that $\alpha_1 \gg 1^3$. Thus, weak null-space projection is a generalization of both of the other two methods which appear as special cases of it.

Finally note that the idea of solving a least squares problem as a means of enforcing the inter-element continuity and boundary conditions within a Navier-Stokes simulation has

³This can be proven by considering the eigendecomposition of the normal equations and taking the limit as $\alpha_2 \rightarrow \infty$

been explored previously in the case of least-squares spectral element methods [21], but these methods suffer from the ill-conditioning of the normal equations of already ill-conditioned element matrices. By using these least-squares ideas as a post-processing method alone, we can somewhat control the ill-conditioning of the normal equations, as will be discussed in Section 4.1.

3. Results

In this section, the three algorithms for post-processing are compared in two simulations of a propagating internal solitary wave in a density-stratified tank of water. The three methods compared are henceforth called *pressure projection*, *exact null-space projection*, and *weak null-space projection methods* and correspond respectively to Algorithms 1, 2, and 3. The simulations model a propagating internal solitary wave and are initialized by an eigenfunction solution of the Dubreil-Jacotin-Long (DJL) equation [9]. These waves are exact solutions of the Euler-Boussinesq equations meaning that through a balance of nonlinearity and dispersion, they preserve their shape as they propagate through the domain. Thus, a simulation ought to explicitly preserve the waveform as it propagates, and so this initial condition provides a convenient test of the accuracy of a numerical method.

The computational domain represents a fluid-filled tank and is a rectangle of width 12 m and height 0.15m. The wave has a phase speed of $c = 0.104$ m/s, and a wavelength of about 1.5 meters. The wave is allowed to propagate for approximately 60 seconds during which time it travels about four wavelengths. Stability of these methods is assessed by examining the duration for which the norm of the velocity remains bounded. Two simulations of this physical problem are considered; in Section 3.1 the mesh is chosen to be of high-enough resolution to clearly resolve the solitary wave. In Section 3.2 the same simulation is repeated but with a lower-resolution mesh that only marginally resolves the wave scale. All simulations use a simple adaptive first-order Euler time-stepping in which is time-step is adjusted to

maintain a CFL number of less than 0.1. This restrictive bound on the CFL condition is due to the low-order time-integration method used, but the focus of this study is divergence-free enforcement in the projection method and not the time-stepping method, so the time-stepping method is kept simple for ease of implementation. In all examples of the WNP method, values of $\alpha_1 = \alpha_2 = 10^4$ are used; these values are obtained empirically via experimentation, although it was found that a broad range of values of α_1 and α_2 produced very similar results (see the discussion in Section 4.3).

3.1. Well-resolved simulation

Figure 1: The initial conditions of the well-resolved simulation in which $n = 10, m_x = 12, m_z = 6$.

The well resolved simulation uses a mesh with $m_x = 16$ elements in the horizontal direction and $m_z = 6$ elements in the vertical direction. Each element has $n = 10$ grid points per direction, resulting in 100 Gauss-Lobatto-Legendre points per element and a total of $r = 9600$ grid points. The initial condition and a snapshot of the first 5 meters of the mesh are depicted in Fig. 1. The wave is propagating to the right.

The results of this simulation are summarized in Fig. 2 in four different ways. Fig. 2(a) depicts the norm divergence of the flow at each time-step. This is calculated by taking the discrete 2-norm of the divergence of the discrete velocity vector and normalizing by the norm of the velocity, and is defined as

$$\frac{\|Du\|}{\|u\|}(t) = \frac{\|D_x u_x(t) + D_z u_z(t)\|}{\sqrt{u_x(t)^2 + u_z(t)^2}} \quad (18)$$

Figure 2: The divergence norm, discontinuity norm, and total mass of the well-resolved simulation as a function of time. Algorithm 1 is shown in red, Algorithm 2 in blue, and Algorithm 3 in green. All of the data in the above plots are shown beginning from the second time-step onwards, as the first time-step has zero divergence and zero discontinuity for all methods.

and as such is a scalar function of time. The results in Fig. 2(a) demonstrate that the rapid growth of divergence in the pressure-projection simulation (red) drives instability, leading to blowup at about time $t = 5$ seconds. As postulated in Section 2.2.1, the divergence grows quickly at the inter-element interfaces and eventually spreads into the interior of the elements. By contrast, both the post-processed simulations (green and blue) remain stable throughout the simulation window, though it is clear that they reduce divergence to differing degrees – the exact null-space projection reduces divergence near machine precision (blue curve), while the weak null-space projection (green curve) reduces the divergence to about $\mathcal{O}(10^{-10})$; tolerably small but not machine precision. Nevertheless, since both post-processing algorithms are capable of maintaining stability by reducing the divergence of the flow, it is clear that divergence is the root of the instability. Namely, discontinuities in the solution of the Poisson problem (c.f. Section 2.2.1) lead to spurious divergence at the element interfaces which in turn feeds into the nonlinear terms in the momentum equation which amplifies these numerical artifacts until the simulation becomes unstable.

Fig. 2(b) depicts the degree to which the velocity field is continuous across element

boundaries. Discontinuity in this context is defined as

$$\frac{\|Eu\|}{\|u\|}(t) = \frac{\|E[u_x; u_z]\|}{\sqrt{u_x(t)^2 + u_z(t)^2}} \quad (19)$$

where $E \in \mathbb{R}^{2r \times 2r}$ is the inter-element continuity misfit functional as defined in the SMPM, which includes both C^0 and C^1 continuity conditions [10]. The rapid growth of discontinuity in the pressure-projection simulation (red) is just due to the instability; as $\|u\|$ grows, naturally $\|Eu\| / \|u\|$ grows at a rate proportional to the operator norm of E . The exact null-space projected velocity (blue) is discontinuous at about three orders of magnitude below the norm of the velocity, and appears to grow slowly over time. This is still relatively continuous, but the growth in time suggests that over time the accumulated discontinuity may need to be addressed. Finally, the weak null-space projected velocity (green) is the most continuous; the discontinuity norm is about five orders of magnitude smaller than the norm of the velocity. Here too the continuity appears to worsen over time.

Finally we measure the conservation properties of these methods by examining the percent change in mass (Fig. 2(c)) over time. Assuming a unit transverse direction, the mass is computed with Gauss-Lobatto-Legendre quadrature-discretized version of the following integral:

$$\text{total perturbed mass}(t) = \sum_k \int_{\Omega_k} \rho'(x, z, t) d\Omega = \sum_k \sum_{ij} w_i w_j \rho_k(x_i, z_j, t) \quad (20)$$

where Ω_k is an element in the mesh and w_j are the Gauss-Lobatto-Legendre quadrature weights.

The change in mass is shown as a percentage in Figs. 2(c). Both the exact (blue) and weak (red) null-space projections conserve mass relatively well. The exact null-space projection is slightly less conservative losing about 0.25% of total mass; the weak projection conserves mass to 5 decimal places. Of note is that both post-processing methods evolve the wave at the correct phase speed as determined by the DJL computation. Finally, the wave at

Figure 3: The exact null-space projection method at time $t = 60$ s in the well-resolved simulation in which $n = 10, m_x = 12, m_z = 6$.

the final time of $t = 60$ seconds as computed by the exact null-space projection and weak null-space projection methods is depicted in Figures 3 and 4. The two solutions are visually indistinguishable and differ less than 1% numerically, a discrepancy likely due to the slightly different final physical time that results from the adaptive time-stepping.

Figure 4: The weak null-space projection method at time $t = 60$ s in the well-resolved simulation in which $n = 10, m_x = 12, m_z = 6$.

3.2. Marginally-resolved simulation

The marginally-resolved simulation uses a mesh with one-quarter of the resolution of the previous mesh. This coarse mesh has $m_x = 8$ elements in the horizontal direction and $m_z = 3$ elements in the vertical direction. Each element has $n = 10$ grid points per direction, resulting in 100 Gauss-Lobatto-Legendre points per element and a total of $r = 2400$ grid points. The initial condition and a snapshot of the first 5 meters of this mesh are depicted in

Figure 5: The initial conditions of the marginally-resolved simulation in which $n = 10, m_x = 6, m_z = 3$.

Fig. 5; again the wave propagates to the right. The results of this simulation are summarized in Fig. 6.

Figure 6: The divergence norm, discontinuity norm, and the total mass of the marginally-resolved simulation as a function of time. Algorithm 1 is shown in red, Algorithm 2 in blue, and Algorithm 3 in green. All of the data in the above plots are shown beginning from the second time-step onwards, as the first time-step has zero divergence and zero discontinuity for all methods.

Notice first in Fig. 6(a) that, although the ENP method (blue) is divergence-free throughout the time that it is stable, it is only stable until time $t = 16$ seconds. Conversely, the weak null-space projection (WNP) method remains stable throughout the simulation window of 60 seconds, even though it is several orders of magnitude more divergent. A possible explanation of this is found in Fig. 6(b), which shows that the ENP method is three orders of magnitude more discontinuous than the WNP method, and the discontinuity grows rapidly just before the simulation becomes unstable. Examining Fig. 6(c), it is also clear that the ENP does not

conserve mass (as the WNP does) prior to becoming unstable. Finally note that the velocity fields exhibit high-frequency oscillations at element interfaces that are likely due to a lack of filtering or de-aliasing.

Figure 7: The weak null-space projection method at time $t = 60\text{s}$ in the marginally-resolved simulation in which $n = 10, m_x = 12, m_z = 6$.

4. Discussion

4.1. Conditioning of the weak projection operator as a function of α_1 and α_2

The 2-norm condition number of a symmetric positive definite matrix A is defined as the ratio of its maximum to its minimum eigenvalue,

$$\kappa_2(A) = \frac{\lambda_{\max}(A)}{\lambda_{\min}(A)}. \quad (21)$$

For $\alpha_1 = \alpha_2 = \alpha$, the normal equations have the form $I + \alpha(D^T D + E^T E)$, which means that their condition number as matrix is

$$\kappa_2(I + \alpha D^T D + \alpha E^T E) = \frac{1 + \alpha \lambda_{\max}(D^T D + E^T E)}{1 + \alpha \lambda_{\min}(D^T D + E^T E)}. \quad (22)$$

Since there exist velocity vectors that are both divergence-free and continuous (for example, any constant velocity), we know that there is a vector $v \in \mathbb{R}^{2r}$ such that $(D^T D + E^T E)v = 0$, which means that $\lambda_{\min}(D^T D + E^T E) = 0$. Thus the condition number of the normal equations is

$$\kappa_2(I + \alpha D^T D + \alpha E^T E) = 1 + \alpha \lambda_{\max}(D^T D + E^T E). \quad (23)$$

This means that the conditioning of the normal equations scales linearly with α , and so the greater the degree to which we penalize discontinuity and divergence, the more difficult and less accurate the solution of the linear system becomes.

4.2. Conditioning of the weak projection operator as a function of n

The conditioning of the normal equations as a function of n can be estimated theoretically by noting that the conditioning of the operator $D^T D + E^T E$ scales approximately as $h^{-2}n^4$, since $\kappa(D) \propto h^{-1}n^2$ [14]. For the purpose of this discussion, let's note this as

$$\kappa(D^T D + E^T E) \approx \kappa(D^T D) = Cn^4 \quad (24)$$

where $C = C(h)$ is a constant dependent on the element width h . Then, noting from the previous section that the conditioning of the normal equations is given as

$$\kappa(I + \alpha(D^T D + E^T E)) = 1 + \alpha\lambda_{max}(D^T D + E^T E), \quad (25)$$

and assuming that $\kappa(D^T D) \propto \lambda_{max}(D^T D)$, we get that

$$\kappa(I + \alpha(D^T D + E^T E)) \propto 1 + \alpha Cn^4. \quad (26)$$

Thus, it is shown that the conditioning of the weak projection method worsens rapidly with increasing n , behavior which is expected when using spectral differentiation [14]. To confirm this, a numerical experiment was carried out whose results are summarized in Fig. 8, confirming that the conditioning worsens far more quickly in n than in α .

4.3. Study of various values of the regularization coefficients

Finding an optimal choice of the regularization coefficients α_1 and α_2 is in general a difficult problem, especially in this context in which optimality is really determined by stability. To obtain an insight into the general performance of the weak projection method as a function of $\alpha = \alpha_1 = \alpha_2$, the low-resolution problem (Section 3.2) was solved for values of $\alpha = [10^{-5}, 10^{-4}, \dots, 10^5]$. The results of this set of simulations are summarized in Fig. 9,

Figure 8: The condition number of $I + \alpha(D^T D + E^T E)$ as a function of α and n , the number of GLL points per direction per element, as computed by a series of numerical experiments varying n and α .

Figure 9: The divergence norm, discontinuity norm, maximum stable time, and condition number of the marginally-resolved simulation as a function of $\alpha = \alpha_1 = \alpha_2$.

in which the divergence norm, the discontinuity norm, the maximum stable time, and the condition number are all plotted as a function of α at the final time the simulation remained stable. As α grows, generally the divergence (Fig. 9(a)) and discontinuity (Fig. 9(b)) in the simulation decrease, but the condition number (Fig. 9(d)) grows linearly (see discussion in Section 4.1). This results in increased difficulty in solving the normal equations and eventually a loss of accuracy that manifests as an increase in divergence and discontinuity (as observed in the $\alpha > 10^4$ regime in Figs. 9(a) and (b)). For all $\alpha > 10^0$ the simulation is stable throughout the simulation window of 60 seconds. While this analysis is not conclusive, it suggests that an optimal choice of α exists, one that minimizes the condition number while still ensuring a sufficiently divergence-free and continuous velocity field that remains stable.

4.4. Spectral analysis of pressure update operators

Figure 10: The divergence (a) and discontinuity (b) of the left singular vectors that span the range space of each pressure update operator as a function of α . In blue is the exact null space projection and in red is pressure projection with no post-processing. Darker colors mean larger values of α with $\alpha \in [10^{-5}, 10^5]$. The singular vectors were sorted by their divergence norm (left) and discontinuity norm (right), respectively.

While it is clear that the weak null-space formulation is effective at reducing discontinuity and divergence, it is instructive to observe exactly how the range space of the pressure update operator is modified as $\alpha = \alpha_1 = \alpha_2$ is changed. Recall that given a matrix $A \in \mathbb{R}^{n \times n}$ the singular value decomposition (SVD) is given as

$$A = USV^T \quad (27)$$

where $U, V \in \mathbb{R}^{n \times n}$ are unitary and $S \in \mathbb{R}^{n \times n}$ is a diagonal matrix. The range space of A is given by the singular vectors (e.g. columns of U) that correspond to non-zero singular values of A (e.g. the diagonal entries of S). It is the space of all possible vectors $Ax \in \mathbb{R}^n$ where $x \in \mathbb{R}^n$, and so in the context of the velocity update operators in the projection method the range space represents all possible updated velocities. To study the divergence and the continuity properties of the range space of the projection methods, then, we compute the SVD of $\mathbb{P}$, $\mathbb{P}_e$, and $\mathbb{P}_w(\alpha)$ for various values of α , and for a grid with parameters $(m_x, m_z, n) = (8, 3, 4)$. Then for each left singular vector u_j and corresponding singular value σ_j we can compute its

divergence

$$||\sigma_j Du_j|| \tag{28}$$

and discontinuity

$$||\sigma_j Eu_j||. \tag{29}$$

The magnitude of these values characterize the degree to which each singular vector is divergent and discontinuous, and the set of these values help characterize the divergence and discontinuity of the spectrum and range space of each of our pressure update operators $\mathbb{P}$, $\mathbb{P}_e$, and $\mathbb{P}_w(\alpha)$.

This computation is shown in Fig. 10, with the divergence shown on the left and the discontinuity shown on the right. Each curve depicts a descending sorted set of values $\{\sigma_j ||Du_j||\}_{j=1}^{2r}$ or $\{\sigma_j ||Eu_j||\}_{j=1}^{2r}$ plotted against j . The values for the pressure update operator with no post-projection are shown in red, and the values for the exact null-space projection are shown in blue. In shades of gray are the values for the weak null-space projection, ranging for $\alpha = [10^{-5}, 10^{-4}, \dots, 10^5]$ with darker shades representing larger values of α .

First notice in the left panel of Fig. 10 that as α grows from 10^{-5} to 10^5 more and more modes of $\mathbb{P}_w(\alpha)$ become divergence-free, and the rate of decay of lower (higher index j) modes decay faster. This indicates, naturally, that as α grows, the range space of $\mathbb{P}_w(\alpha)$ becomes more and more solenoidal. Second, notice that for all values of α the vast majority of modes of $\mathbb{P}_w(\alpha)$ are less divergent than those of $\mathbb{P}$; this is a good indication that the post-processing will reduce the divergence in the velocity field. Finally, notice that all modes of $\mathbb{P}_e$ are less divergent than any of the other modes of any of the other operators, as is expected.

The right panel of Fig. 10 depicts the discontinuity of each of the singular vectors of the three operators $\mathbb{P}$, $\mathbb{P}_e$, and $\mathbb{P}_w(\alpha)$. Notice first that first few modes of $\mathbb{P}_e$, shown in blue, are

very discontinuous. These modes are the dominant modes of the range space of $\mathbb{P}_e$, and so contribute significantly to its amplification of discontinuity. Secondly, notice that the shape of $\mathbb{P}_w(\alpha)$ changes dramatically from $\alpha = 10^{-1}$ to $\alpha = 10^0$, exhibiting a steep drop in the discontinuity of the spectrum around the $j = 400$ mode. This transition is also the same place the stability properties reach a plateau as is evidenced in Fig. 9(c), indicating that discontinuity of the pressure projection operator and stability are at least well-correlated. With regards to choosing a value of α , it is observed that for $\alpha > 10^0$, the shape of the $\mathbb{P}_w(\alpha)$ curve doesn't change, but its overall value continues to decrease. Along with the evidence shown in Fig. 9(c), this suggests that a fundamental property of the weak projection operator is unchanged after $\alpha > 1$ in which regime the method is stable. And finally, as observed in practice, $\mathbb{P}_w$ exhibits significantly better continuity properties than either $\mathbb{P}$ and $\mathbb{P}_e$; this is not unexpected as the latter two projection methods do not consider continuity at all in their post-processing.

5. Summary and Conclusions

5.1. Summary

In this paper, a method for post-processing the projection operation in the pressure projection method was developed that reduces inter-element discontinuity along with divergence as part of a time-evolving incompressible Euler simulation. This method was compared with an exactly divergence-free projection method as developed in Ref.[25] in the framework of a simulation of a propagating internal solitary wave. While the exact null-space projection method (e.g. as in Ref. [25]) was better at mitigating divergence than the weak null-space projection, when the mesh resolution was coarsened so that the flow was only marginally resolved the weak null-space projection maintained stability for at least four times as long. The continuity and conservation properties of the flow field were also observed to be better in the weak null-space projection method, and both weak and exact null-space projection methods outperformed the simulation without any post-processing.

The drawback of the weak null-space projection method is that it requires the solution of a large, sparse, positive-definite linear system in solving the normal equations. The difficulty of computing such a solution at scale was not addressed here, and is likely a significant computational challenge for an iterative method. However, the benefits of this method, particularly the ability to remain stable in an under-resolved context (or conversely to reduce resolution while maintaining stability), suggest that this considerable computational effort is worth undertaking, and will be a focus of future work. It should be noted that there is a large body of literature [23] and software [1] for obtaining the efficient parallel solution of symmetric positive definite matrices, and although it is beyond the scope of this paper, it is anticipated that it will be possible to draw on this body of work in implementing a scalable solver for the normal equations.

5.2. Extension

The results described here are limited in the sense that they are only demonstrated on a single test case, that of a propagating internal solitary wave, and for a single discretization, the spectral multi-domain penalty method. However, the ideas described are rather general in that they apply to any discontinuous element-based discretization of the pressure projection method. A case can be made that although the efficacy of the weak projection method needs to be demonstrated on a broader class of problems, the spectral analysis presented in Section 4.4 is independent of the problem being solved and demonstrates that the WNP method preserves continuity and reduces divergence simultaneously.

5.3. Future Work

The natural extension is to ask whether α_1 need be equal to α_2 , and whether either should vary in time. In particular, if $\alpha_1(t)$ and $\alpha_2(t)$ are time-dependant, then it might be possible

to choose them such that

$$\alpha_1(t) \propto ||Du|| (t) \quad (30)$$

$$\alpha_2(t) \propto ||Eu|| (t) \quad (31)$$

as a way to dynamically mitigate the effects of divergence and discontinuity. So, if a flow were to begin to exhibit troublesome discontinuities or divergence, then the coefficients could be dynamically adjusted to compensate and presumably improve the divergence and continuity properties of the flow. It may even be possible to let α vary in space as well as time and become a weighting function for the residual that would locally adjust to values of the divergence and discontinuity. In this way it may be possible to minimize the penalty paid in ill-conditioning for large values of α while still reaping the benefits of a smoother more divergence-free solution in places where it is critical.

Finally, and as mentioned previously in Section 3.2, there are significant numerical artifacts that are likely due to aliasing effects due to the nonlinear term. If this is indeed due to the nonlinear term, then some sort of filtering or dealiasing is necessary to damp spurious high-frequency aliased components. Worth investigating, then, is how these post-processing methods would interact with filtering or dealiasing.

Acknowledgements

The authors would also like to thank the United States Department of Defense High-Performance Computing Modernization Office for the NDSEG fellowship and the National Science Foundation for CAREER award #0845558 for support. Marek Stastna is supported by the Natural Sciences and Engineering Research Council of Canada.

References

- [1] Balay, S., Abhyankar, S., Adams, M. F., Brown, J., Brune, P., Buschelman, K., Dalcin, L., Eijkhout, V., Gropp, W. D., Kaushik, D., Knepley, M. G., McInnes, L. C., Rupp, K.,

- Smith, B. F., Zampini, S., Zhang, H., 2015. PETSc Users Manual. Tech. Rep. ANL-95/11 - Revision 3.6, Argonne National Laboratory.
- URL <http://www.mcs.anl.gov/petsc>
- [2] Brown, D., Cortez, R., Minion, M., 2001. Accurate Projection Methods for the Incompressible NavierStokes Equations. *Journal of Computational Physics* 168, 464–499.
 - [3] Chorin, A. J., 1968. Numerical solutions of the Navier-Stokes equations.
 - [4] Costa, B., Don, W. S., 2000. On the computation of high order pseudospectral derivatives. *Applied Numerical Mathematics* 33, 151–159.
 - [5] Deville, M. O., Fischer, P. F., Mund, E. H., 2002. High-Order Methods for Incompressible Fluid Flow. Cambridge Monographs on Applied and Computational Mathematics.
 - [6] Diamessis, P. J., Domaradzki, J. A., Hesthaven, J. S., 2005. A spectral multidomain penalty method model for the simulation of high Reynolds number localized incompressible stratified turbulence. *Journal of Computational Physics* 202, 298–322.
 - [7] Diamessis, P. J., Redekopp, L. G., 2006. Numerical Investigation of Solitary Internal Wave-Induced Global Instability in Shallow Water Benthic Boundary Layers. *Journal of Physical Oceanography* 36, 784–812.
 - [8] Diamessis, P. J., Spedding, G. R., Domaradzki, J. A., 2011. Similarity scaling and vorticity structure in high-Reynolds-number stably stratified turbulent wakes. *Journal of Fluid Mechanics* 671, 52–95.
 - [9] Dunphy, M., Subich, C., Stastna, M., 2011. Spectral methods for internal waves: Indistinguishable density profiles and double-humped solitary waves. *Nonlinear Processes in Geophysics* 18, 351–358.

- [10] Escobar-Vargas, J., Diamessis, P., Sakai, T., 2014. A spectral quadrilateral multido-
main penalty method model for high Reynolds number incompressible stratified flows.
International Journal for Numerical Methods in Fluids 75 (March), 403–425.
- [11] Guermond, J., Mineev, P., Shen, J., 2006. An overview of projection methods for in-
compressible flows. Computer Methods in Applied Mechanics and Engineering 195,
6011–6045.
- [12] Guermond, J.-L., 1996. Some implementations of projection methods for navier-stokes
equations. Mathematical Modelling and Numerical Analysis 30, 637–667.
- [13] Hesthaven, J. S., 1998. A Stable Penalty Method for the Compressible Navier-Stokes
Equations: III. Multidimensional Domain Decomposition Schemes. SIAM Journal on
Scientific Computing 20 (1), 62–93.
- [14] Hesthaven, J. S., Gottlieb, S., Gottlieb, D., 2007. Spectral Methods for Time-Dependent
Problems. Cambridge University Press.
- [15] Joshi, S. M., Thomsen, G. N., Diamessis, P. J., 2016. Deflation-accelerated precon-
ditioning of the Poisson-Neumann Schur problem on long domains with a high-order
discontinuous element-based collocation method. Journal of Computational Physics 313,
209–232.
URL <http://arxiv.org/abs/1512.01756>
- [16] Karniadakis, G. E., Israeli, M., Orszag, S. a., 1991. High-order splitting methods for the
incompressible Navier-Stokes equations. Journal of Computational Physics 97, 414–443.
- [17] Kim, J., Moin, P., 1985. Application of a fractional-step method to incompressible
Navier-Stokes equations. Journal of Computational Physics 59, 308–323.
- [18] Kundu, P., Dowling, D., 2011. Fluid Mechanics, 5th Edition. Academic Press.

- [19] Liu, J.-G., Liu, J., Pego, R. L., 2010. Stable and accurate pressure approximation for unsteady incompressible viscous flow. *Journal of Computational Physics* 229, 3428–3453.
- [20] Marcus, P. S., 1984. Simulation of Taylor-Couette flow. Part 1. Numerical methods and comparison with experiment. *Journal of Fluid Mechanics* 146, 45–64.
- [21] Proot, M. M., Gerritsma, M. I., 2005. Application of the Least-Squares Spectral Element Method Using Chebyshev Polynomials to Solve the Incompressible Navier-Stokes Equations. *Numerical Algorithms* 38, 155–172.
- [22] Restelli, M., Giraldo, F., 2009. A Conservative Discontinuous Galerkin semi-implicit formulation for the Navier-Stokes equations in Nonhydrostatic Mesoscale Modeling. *SIAM Journal Scientific Computing* 31 (October), 2231–2257.
- [23] Saad, Y., 2003. *Iterative Methods for Sparse Linear Systems*. Vol. 3. Society for Industrial and Applied Mathematics.
- [24] Sani, R. L., Shen, J., Pironneau, O., Gresho, P. M., 2006. Pressure boundary condition for the time-dependent incompressible Navier-Stokes equations. *International Journal for Numerical Methods in Fluids* 50 (August 2004), 673–682.
- [25] Steinmoeller, D., Stastna, M., Lamb, K., 2013. A short note on the discontinuous Galerkin discretization of the pressure projection operator in incompressible flow. *Journal of Computational Physics* 251, 480–486.
- [26] Tomboulides, A., Israeli, M., Karniadakis, G., 1989. Efficient removal of boundary-divergence errors in time-splitting methods. *Journal of scientific computing* 4 (3), 291–308.