



Slate: extending Firedrake's domain-specific abstraction to hybridized solvers for geoscience and beyond

Thomas H. Gibson¹, Lawrence Mitchell², David A. Ham¹, and Colin J. Cotter¹

¹Department of Mathematics, Imperial College London, London, SW7 2AZ, UK

²Department of Computer Science, Durham University, Durham, DH1 3LE, UK

Correspondence: Thomas H. Gibson (t.gibson15@imperial.ac.uk)

Abstract. Within the finite element community, discontinuous Galerkin (DG) and mixed finite element methods have become increasingly popular in simulating geophysical flows. However, robust and efficient solvers for the resulting saddle-point and elliptic systems arising from these discretizations continue to be an on-going challenge. One possible approach for addressing this issue is to employ a method known as hybridization, where the discrete equations are transformed such that classic static condensation and local post-processing methods can be employed. However, it is challenging to implement hybridization as performant parallel code within complex models, whilst maintaining separation of concerns between applications scientists and software experts. In this paper, we introduce a domain-specific abstraction within the Firedrake finite element library that permits the rapid execution of these hybridization techniques within a code-generating framework. The resulting framework composes naturally with Firedrake's solver environment, allowing for the implementation of hybridization and static condensation as runtime-configurable preconditioners via the Python interface to PETSc, petsc4py. We provide examples derived from second order elliptic problems and geophysical fluid dynamics. In addition, we demonstrate that hybridization shows great promise for improving the performance of solvers for mixed finite element discretizations of equations related to large-scale geophysical flows.

1 Introduction

The development of simulation software is an increasingly important aspect of modern scientific computing, in the geosciences in particular. Such software requires a vast range of knowledge spanning several disciplines, ranging from applications expertise to mathematical analysis to high-performance computing and low-level code optimization. Software projects developing automatic code generation systems have become quite popular in recent years, as such systems help create a separation of concerns which focuses on a particular complexity independent from the rest. This allows for agile collaboration between computer scientists with hardware and software expertise, computational scientists with numerical algorithm expertise, and domain scientists such as meteorologists, oceanographers and climate scientists. Examples of such projects in the domain of finite element methods include FreeFEM++ (Hecht, 2012), Sundance (Long et al., 2010), the FEniCS Project (Logg et al., 2012), Feel++ (Prud'Homme et al., 2012), and Firedrake (Rathgeber et al., 2016).



The finite element method (FEM) is a mathematically robust framework for computing numerical solutions of partial differential equations (PDEs) that has become increasingly popular in fluids and solids models across the geosciences, with a formulation that is highly amenable to code-generation techniques. A description of the weak formulation of the PDEs, together with appropriate discrete function spaces, is enough to characterize the finite element problem. Both the FEniCS and
5 Firedrake projects employ the *Unified Form Language* (UFL) (Alnæs et al., 2014) to specify the finite element integral forms and discrete spaces necessary to properly define the finite element problem. UFL is a highly expressive domain-specific language (DSL) embedded in Python, which provides the necessary abstractions for code generation systems.

There are classes of finite element discretizations resulting in discrete systems that can be solved more efficiently by directly manipulating local tensors. For example, the static condensation technique for the reduction of global finite element systems
10 (Guyan, 1965; Irons, 1965) produces smaller globally-coupled linear systems by eliminating interior unknowns to arrive at an equation for the degrees of freedom defined on cell-interfaces only. This procedure is analogous to the point-wise elimination of variables used in staggered finite difference codes, such as the ENDGame dynamical core (Melvin et al., 2010; Wood et al., 2014) of the UK Meteorological Office (Met Office), but requires the local inversion of finite element systems. For finite element discretizations of coupled equations relevant to geophysical flows, the hybridization technique (Arnold and Brezzi,
15 1985; Brezzi and Fortin, 2012; Cockburn et al., 2009a) introduces Lagrange multipliers enforcing certain continuity constraints. Local static condensation can then be applied to the augmented system to produce a reduced equation for the multipliers. Methods of this type are often accompanied by local post-processing techniques that produce superconvergent approximations, or enhanced conservation properties (Bramble and Xu, 1989; Cockburn et al., 2010b, 2009b). These procedures require invasive manual intervention during the equation assembly process in intricate numerical code.

In this paper, we provide a simple yet effective high-level abstraction for localized dense linear algebra on systems derived from finite element problems. Using embedded DSL technology, we provide a means to enable the rapid development of hybridization and static condensation techniques within an automatic code-generation framework. In other words, the main contribution of this paper is in solving the problem of automatically translating from the mathematics of static condensation and hybridization to compiled code. This automated translation facilitates the separation of concerns between applications
20 scientists and computational/computer scientists, and facilitates the automated optimization of compiled code. This framework provides an environment for the development and testing of numerics relevant to the Gung-Ho Project, an initiative by the UK Met Office in designing the next-generation atmospheric dynamical core using mixed finite element methods (Melvin et al., 2018). Our work is implemented in the Firedrake finite element library and the PETSc (Balay et al., 1997, 2016) solver library, accessed via the Python interface `petsc4py` (Dalcin et al., 2011).

The rest of the paper is organized as follows. We introduce common notation used throughout the paper in Section 1.1. The embedded DSL, called “Slate”, is introduced in Section 2, which allows concise expression of localized linear algebra operations on finite element tensors. We provide some contextual examples for static condensation and hybridization in Section 3, including a discussion on post-processing. We then outline in Section 4 how, by interpreting static condensation techniques as a preconditioner, we can go further, and automate many of the symbolic manipulations necessary for hybridization and
35 static condensation. We first demonstrate our implementation on a manufactured problem derived from a second-order elliptic



equation, starting in Section 5. The first example compares a hybridizable discontinuous Galerkin (HDG) method with an optimized continuous Galerkin method. Section 5.2 illustrates the composability and relative performance of hybridization for compatible mixed methods applied to a semi-implicit discretization of the nonlinear rotating shallow water equations. Our final example in Section 5.3 demonstrates time-step robustness of a hybridizable solver for a compatible finite element discretization of a rotating linear Boussinesq model. Conclusions follow in Section 6.

1.1 Notation

We begin by establishing notation used throughout this paper. Let \mathcal{T}_h denote a tessellation of $\Omega \subset \mathbb{R}^n$, the computational domain, consisting of polygonal elements K associated with a mesh size parameter h , and $\partial\mathcal{T}_h = \{e \in \partial K : K \in \mathcal{T}_h\}$ the set of facets of \mathcal{T}_h . The set of facets *interior* to the domain Ω is denoted by $\mathcal{E}_h^\circ \equiv \partial\mathcal{T}_h \setminus \partial\Omega$. Similarly, we denote the set of *exterior* facets as $\mathcal{E}_h^\partial = \partial\mathcal{T}_h \cap \partial\Omega$. For brevity, we denote the finite element integral forms over \mathcal{T}_h and any facet set $\Gamma \subset \partial\mathcal{T}_h$ by

$$(u, v)_K = \int_K u \cdot v \, dx, \quad \langle u, v \rangle_e = \int_e u \cdot v \, ds, \quad (1)$$

$$(u, v)_{\mathcal{T}_h} = \sum_{K \in \mathcal{T}_h} (u, v)_K, \quad \langle u, v \rangle_\Gamma = \sum_{e \in \Gamma} \langle u, v \rangle_e, \quad (2)$$

where \cdot should be interpreted as standard multiplication for scalar functions or a dot product for vector functions.

For any double-valued vector field \mathbf{w} on a facet $e \in \partial\mathcal{T}_h$, we define the jump of its normal component across e by

$$[[\mathbf{w}]]_e = \begin{cases} \mathbf{w}|_{e^+} \cdot \mathbf{n}_{e^+} + \mathbf{w}|_{e^-} \cdot \mathbf{n}_{e^-}, & e \in \mathcal{E}_h^\circ \\ \mathbf{w}|_e \cdot \mathbf{n}_e, & e \in \mathcal{E}_h^\partial \end{cases} \quad (3)$$

where $+$ and $-$ denotes arbitrarily but globally defined sides of the facet. Here, \mathbf{n}_{e^+} and \mathbf{n}_{e^-} are the unit normal vectors with respect to the positive and negative sides of the facet e . Whenever the facet domain is clear by the context, we omit the subscripts for brevity and simply write $[[\cdot]]$.

2 A system for localized algebra on finite element tensors

We present an expressive language for dense linear algebra on the elemental matrix systems arising from finite element problems. The language, which we call *Slate*, provides typical mathematical operations performed on matrices and vectors, hence the input syntax is comparable to high-level linear algebra software such as MATLAB. The Slate language provides basic abstract building blocks which can be used by a specialized compiler for linear algebra to generate low-level code implementations.

Slate is heavily influenced by the Unified Form Language (UFL) (Alnæs et al., 2014; Logg et al., 2012), a DSL embedded in Python which provides symbolic representations of finite element forms. The expressions can be compiled by a *form compiler*, which translates UFL into low level code for the local assembly of a form over the cells and facets of a mesh. In a similar manner, Slate expressions are compiled to low level code that performs the requested linear algebra element-wise on a mesh.



2.1 An overview of Slate

To clarify conventions and the scope of Slate, we start by considering a general form. Suppose we have a finite element form:

$$a(\mathbf{c}; \mathbf{v}) = \sum_{K \in \mathcal{T}_h} \int_K \mathcal{I}^c(\mathbf{c}; \mathbf{v}) dx + \sum_{e \in \mathcal{E}_h^\circ} \int_e \mathcal{I}^{f,\circ}(\mathbf{c}; \mathbf{v}) ds + \sum_{e \in \mathcal{E}_h^\partial} \int_e \mathcal{I}^{f,\partial}(\mathbf{c}; \mathbf{v}) ds, \quad (4)$$

where dx and ds denote appropriate integration measures. The integral form in (4) is uniquely determined by its lists (possibly of 0-length) of arbitrary coefficient functions $\mathbf{c} = (c_0, \dots, c_p)$ in the associated finite element spaces, arguments $\mathbf{v} = (v_0, \dots, v_q)$ describing any test or trial functions, and its integrand expressions for each integral type: \mathcal{I}^c , $\mathcal{I}^{f,\circ}$, $\mathcal{I}^{f,\partial}$. The form $a(\mathbf{c}; \mathbf{v})$ describes a finite element form *globally* over the entire problem domain. Here, we will consider the case where the integrand $\mathcal{I}^{f,\circ}(\mathbf{c}; \mathbf{v})$ can be decomposed into two independent parts for each facet e : one for the positive restriction (+) and the negative restriction (-).

10 The contribution of (4) in each cell K of the mesh \mathcal{T}_h is simply

$$a(\mathbf{c}; \mathbf{v})|_K = \int_K \mathcal{I}^c(\mathbf{c}; \mathbf{v})|_K dx + \sum_{e \in \partial K \setminus \partial \Omega} \int_e \mathcal{I}^{f,\circ}(\mathbf{c}; \mathbf{v})|_e ds + \sum_{e \in \partial K \cap \partial \Omega} \int_e \mathcal{I}^{f,\partial}(\mathbf{c}; \mathbf{v})|_e ds. \quad (5)$$

We call (5) the *cell-local* contribution of $a(\mathbf{c}; \mathbf{v})$, with

$$a(\mathbf{c}; \mathbf{v}) = \sum_{K \in \mathcal{T}_h} a(\mathbf{c}; \mathbf{v})|_K. \quad (6)$$

Equation (5) produces an element tensor which is mapped into a global data structure. However, before doing so, one may want to produce a new local tensor by algebraically manipulating different element tensors. This is precisely the job of Slate.

The Slate language consists of two primary abstractions for linear algebra:

1. terminal element tensors corresponding to multi-linear integral forms (matrices, vectors, and scalars), or assembled data (coefficient vectors); and
2. expressions consisting of operations on terminal tensors.

20 The composition of binary and unary operations on terminal tensors produces a *Slate expression*. Such expressions can be composed with other Slate objects in arbitrary ways, resulting in concise representations of complex operations on locally assembled arrays.

2.1.1 Terminal tensors:

In Slate, one associates a tensor with data on an element either by using a form, or assembled coefficient data:

25 – $\text{Tensor}(a(\mathbf{c}; \mathbf{v}))$
 associates a form, expressed in UFL, with its local element tensor:

$$A^K \leftarrow a(\mathbf{c}; \mathbf{v})|_K, \text{ for all } K \in \mathcal{T}_h. \quad (7)$$



The number of arguments v determine the *rank* of `Tensor`, i.e. scalars, vectors, and matrices are produced from 0-forms, 1-forms, and 2-forms¹ respectively.

– `AssembledVector(f)`

where f is some finite element function. The result associates a function with its local coefficient vectors.

5 2.1.2 Symbolic linear algebra:

Slate supports typical binary and unary operations in linear algebra, with a high-level syntax close to mathematics. At the time of this paper, these include:

– `A + B`, the addition of two equal shaped tensors.

– `A * B`, a contraction over the last index of `A` and the first index of `B`. This is the usual multiplicative operation on matrices, vectors, and scalars.

– `-A`, the additive inverse (negation) of a tensor.

– `A.T`, the transpose of a tensor.

– `A.inv`, the inverse of a square tensor.

– `A.solve(B, decomposition="...")`, the result of solving a local linear system $AX = B$ for X , optionally specifying a direct factorization strategy.

– `A.blocks[indices]`, where `A` is a tensor from a mixed finite element space, allows extraction of subblocks of the tensor indexed by field (slices are allowed). For example, if a matrix A corresponds to the bilinear form $a : V \times W \rightarrow \mathbb{R}$, where $V = V_0 \times \dots \times V_n$ and $W = W_0 \times \dots \times W_m$ are product spaces consisting of finite element spaces $\{V_i\}_{i=0}^n$, $\{W_i\}_{i=0}^m$, then the cell-local tensors have the form:

$$A^K = \begin{bmatrix} A_{00}^K & A_{01}^K & \dots & A_{0m}^K \\ A_{10}^K & A_{11}^K & \dots & A_{1m}^K \\ \vdots & \vdots & \ddots & \vdots \\ A_{n0}^K & A_{n1}^K & \dots & A_{nm}^K \end{bmatrix}. \quad (8)$$

The associated submatrix of (8) with indices $\mathbf{i} = (\mathbf{p}, \mathbf{q})$, $\mathbf{p} = \{p_1, \dots, p_r\}$, $\mathbf{q} = \{q_1, \dots, q_c\}$, is

$$A_{\mathbf{pq}}^K = \begin{bmatrix} A_{p_1 q_1}^K & \dots & A_{p_1 q_c}^K \\ \vdots & \ddots & \vdots \\ A_{p_r q_1}^K & \dots & A_{p_r q_c}^K \end{bmatrix} = A^K.\text{blocks}[\mathbf{p}, \mathbf{q}], \quad (9)$$

where $\mathbf{p} \subset \{1, \dots, n\}$, $\mathbf{q} \subset \{1, \dots, m\}$.

¹As with UFL, Slate is capable of abstractly representing arbitrary rank tensors. However, only rank ≤ 2 tensors are typically used in most finite element applications and therefore we currently only generate code for those ranks.



These building blocks may be arbitrarily composed, giving a symbolic expression for the linear algebra we wish to perform on each cell during assembly. They provide the necessary algebraic framework for a large class of problems, some of which we present in this paper.

In Firedrake, these *Slate expressions* are transformed into low-level code by a *linear algebra compiler*. This uses the form compiler, TSFC (Homolya et al., 2018), to compile kernels for the assembly of terminal tensors and generates a dense linear algebra kernel to be iterated cell-wise. At the time of this work, our compiler generates C++ code, using the templated library Eigen (Guennebaud et al., 2015) for dense linear algebra. During execution, the local computations in each cell are mapped into global data objects via appropriate indirection mappings using the PyOP2 framework (Rathgeber et al., 2012). Figure 1 provides an illustration of the complete tool-chain.

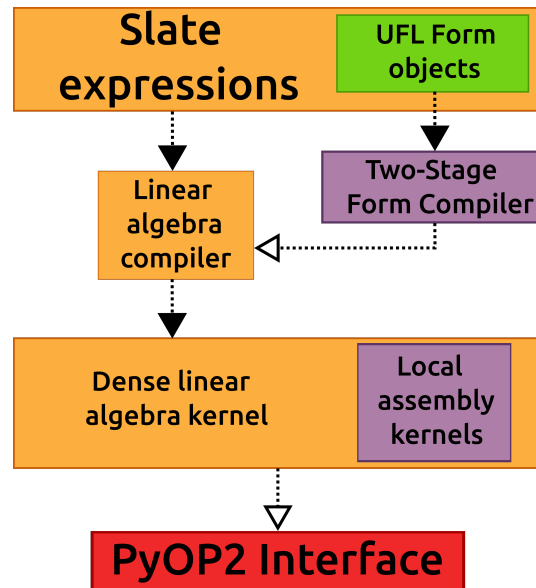


Figure 1. The Slate language wraps UFL objects describing the finite element system. The resulting Slate expressions are passed to a specialized linear algebra compiler, which produces a single “macro” kernel assembling the local contributions and executes the dense linear algebra represented in Slate. The kernels are passed to the Firedrake’s PyOP2 interface, which wraps the Slate kernel in a mesh-iteration kernel. Parallel scheduling, code generation, and compilation occurs after the PyOP2 layer.

10 3 Examples

We now present a few examples and discuss solution methods which require element-wise manipulations of finite element systems and their specification in Slate. We stress here that Slate is not limited to these model problems; rather these examples were chosen for clarity and to demonstrate key features of the Slate language. In Sections 4 and 5, we discuss more intricate ways Slate is used in custom preconditioners.



For the hybridization of mixed and discontinuous Galerkin methods, we use a model elliptic equation. Consider the second-order PDE with both Dirichlet and Neumann boundary conditions:

$$-\nabla \cdot (\kappa \nabla p) + cp = f \text{ in } \Omega, \quad (10)$$

$$p = p_0 \text{ on } \partial\Omega_D, \quad (11)$$

$$5 \quad -\kappa \nabla p \cdot \mathbf{n} = g \text{ on } \partial\Omega_N, \quad (12)$$

where $\partial\Omega_D \cup \partial\Omega_N = \partial\Omega$ and $\kappa, c : \Omega \rightarrow \mathbb{R}^+$ are positive-valued coefficients. Rewriting as a first-order system, we obtain the mixed problem:

$$\mu \mathbf{u} + \nabla p = 0 \text{ in } \Omega, \quad (13)$$

$$\nabla \cdot \mathbf{u} + cp = f \text{ in } \Omega, \quad (14)$$

$$10 \quad p = p_0 \text{ on } \partial\Omega_D, \quad (15)$$

$$\mathbf{u} \cdot \mathbf{n} = g \text{ on } \partial\Omega_N, \quad (16)$$

where $\mu = \kappa^{-1}$ and $\mathbf{u} = -\kappa \nabla p$ is the velocity variable.

3.1 Hybridization of mixed methods

To motivate our discussion in this section, we start by recalling the mixed method for (13)–(16). Methods of this type seek approximations (\mathbf{u}_h, p_h) in finite-dimensional subspaces $\mathbf{U}_h \times V_h \subset \mathbf{H}(\text{div}) \times L^2$, defined by:

$$\mathbf{U}_h = \{\mathbf{w} \in \mathbf{H}(\text{div}; \Omega) : \mathbf{w}|_K \in \mathbf{U}(K), \forall K \in \mathcal{T}_h, \mathbf{w} \cdot \mathbf{n} = g \text{ on } \partial\Omega_N\}, \quad (17)$$

$$V_h = \{\phi \in L^2(\Omega) : \phi|_K \in V(K), \forall K \in \mathcal{T}_h\}. \quad (18)$$

The space \mathbf{U}_h consists of $\mathbf{H}(\text{div})$ -conforming piecewise vector polynomials, where choices of $\mathbf{U}(K)$ typically include the Raviart-Thomas (RT), Brezzi-Douglas-Marini (BDM), or Brezzi-Douglas-Fortin-Marini (BDFM) elements (Brezzi et al., 1987, 1985; Nédélec, 1980; Raviart and Thomas, 1977). V_h is the Lagrange family of discontinuous polynomials. These spaces are of particular interest when simulating geophysical flows, since choosing the right pairing results in stable discretizations with desirable conservation properties and avoids spurious computational modes. We refer the reader to Cotter and Shipton (2012); Cotter and Thuburn (2014); Natale et al. (2016); Shipton et al. (2018) for a discussion of mixed methods relevant for geophysical fluid dynamics. Two examples of such a discretization is presented in Section 5.2.

The mixed finite element formulation of (13)–(16) reads as follows: find $(\mathbf{u}_h, p_h) \in \mathbf{U}_h \times V_h$ satisfying

$$(\mathbf{w}, \mu \mathbf{u}_h)_{\mathcal{T}_h} - (\nabla \cdot \mathbf{w}, p_h)_{\mathcal{T}_h} = -\langle \mathbf{w} \cdot \mathbf{n}, p_0 \rangle_{\partial\Omega_D}, \quad \forall \mathbf{w} \in \mathbf{U}_{h,0}, \quad (19)$$

$$(\phi, \nabla \cdot \mathbf{u}_h)_{\mathcal{T}_h} + (\phi, cp_h)_{\mathcal{T}_h} = (\phi, f)_{\mathcal{T}_h}, \quad \forall \phi \in V_h, \quad (20)$$



where $U_{h,0}$ is the space of functions in U_h whose normal components vanish on $\partial\Omega_N$. The discrete system is obtained by first expanding the solutions in terms of the finite element bases:

$$\mathbf{u}_h = \sum_{i=1}^{N_u} U_i \Psi_i, \quad p_h = \sum_{i=1}^{N_p} P_i \xi_i, \quad (21)$$

where $\{\Psi_i\}_i$ and $\{\xi_i\}_i$ are bases for U_h and V_h respectively. Here, U_i and P_i are the coefficients to be determined. As per standard Galerkin-based finite element methods, taking $\mathbf{w} = \Psi_j$, $j \in \{1, \dots, N_u\}$ and $\phi = \xi_j$, $j \in \{1, \dots, N_p\}$ in (19)–(20) produces the discrete saddle point system:

$$\begin{bmatrix} A & -B^T \\ B & C \end{bmatrix} \begin{Bmatrix} U \\ P \end{Bmatrix} = \begin{Bmatrix} F_0 \\ F_1 \end{Bmatrix}. \quad (22)$$

where $U = \{U_i\}$, $P = \{P_i\}$ are the coefficient vectors, and

$$A_{ij} = (\Psi_i, \mu \Psi_j)_{\mathcal{T}_h}, \quad (23)$$

$$10 \quad B_{ij} = (\xi_i, \nabla \cdot \Psi_j)_{\mathcal{T}_h}, \quad (24)$$

$$C_{ij} = (\xi_i, c \xi_j)_{\mathcal{T}_h}, \quad (25)$$

$$F_{0,j} = -\langle \Psi_j \cdot \mathbf{n}, p_0 \rangle_{\partial\Omega_D}, \quad (26)$$

$$F_{1,j} = (\xi_j, f)_{\mathcal{T}_h}. \quad (27)$$

Methods to efficiently invert such systems include $\mathbf{H}(\text{div})$ -multigrid (Arnold et al., 2000) (requiring complex overlapping-Schwarz smoothers), global Schur-complement factorizations (which require an approximation to the inverse of the *dense*² elliptic Schur-complement $C + BA^{-1}B^T$), or auxiliary space multigrid (Hiptmair and Xu, 2007). Here, we focus on a solution approach using a hybridized mixed method (Arnold and Brezzi, 1985; Brezzi and Fortin, 2012).

The hybridization technique replaces the original system with a discontinuous variant, decoupling the velocity degrees of freedom between cells. This is done by replacing the discrete solution space for \mathbf{u}_h with the “broken” space U_h^d , defined as:

$$20 \quad U_h^d = \{\mathbf{w} \in [L^2(\Omega)]^n : \mathbf{w}|_K \in U(K), \forall K \in \mathcal{T}_h\}. \quad (28)$$

The vector finite element space U_h^d is a subspace of $[L^2(\Omega)]^n$ consisting of local $\mathbf{H}(\text{div})$ functions, but normal components are no longer continuous on $\partial\mathcal{T}_h$. The approximation space for p_h remains unchanged.

Next, Lagrange multipliers are introduced as an auxiliary variable in the space M_h , defined only on cell-interfaces:

$$M_h = \{\gamma \in L^2(\partial\mathcal{T}_h) : \gamma|_e \in M(e), \forall e \in \partial\mathcal{T}_h\}, \quad (29)$$

25 where $M(e)$ denotes a polynomial space defined on each facet. We call M_h the space of approximate traces. Functions in M_h are discontinuous across vertices in two-dimensions, and vertices/edges in three-dimensions.

²The Schur-complement, while elliptic, is globally dense due to the fact that A has a dense inverse. This is a result of velocities in U_h having continuous normal components across cell-interfaces.



Deriving the hybridizable mixed system is accomplished through integration by parts over each element K . Testing with $\mathbf{w} \in \mathbf{U}_h^d(K)$ and integrating (13) over K produces:

$$(\mathbf{w}, \mu \mathbf{u}_h^d)_K - (\nabla \cdot \mathbf{w}, p_h)_K + \langle \mathbf{w} \cdot \mathbf{n}, \lambda_h \rangle_{\partial K} = -\langle \mathbf{w} \cdot \mathbf{n} p_0 \rangle_{\partial K \cap \partial \Omega_D}. \quad (30)$$

The trace function λ_h is introduced in surface integrals and approximates p_h on elemental boundaries. An additional constraint equation is added to close the system. The resulting hybridizable formulation reads: find $(\mathbf{u}_h^d, p_h, \lambda_h) \in \mathbf{U}_h^d \times V_h \times M_h$ such that

$$(\mathbf{w}, \mu \mathbf{u}_h^d)_{\mathcal{T}_h} - (\nabla \cdot \mathbf{w}, p_h)_{\mathcal{T}_h} + \langle \llbracket \mathbf{w} \rrbracket, \lambda_h \rangle_{\partial \mathcal{T}_h \setminus \partial \Omega_D} = -\langle \mathbf{w} \cdot \mathbf{n}, p_0 \rangle_{\partial \Omega_D}, \quad \forall \mathbf{w} \in \mathbf{U}_h^d, \quad (31)$$

$$(\phi, \nabla \cdot \mathbf{u}_h^d)_{\mathcal{T}_h} + (\phi, cp_h)_{\mathcal{T}_h} = (\phi, f)_{\mathcal{T}_h}, \quad \forall \phi \in V_h, \quad (32)$$

$$\langle \gamma, \llbracket \mathbf{u}_h^d \rrbracket \rangle_{\partial \mathcal{T}_h \setminus \partial \Omega_D} = \langle \gamma, g \rangle_{\partial \Omega_N}, \quad \forall \gamma \in M_{h,0}, \quad (33)$$

where $M_{h,0}$ denotes the space of traces vanishing on $\partial \Omega_D$. The constraint in (33) enforces both continuity of the normal components of \mathbf{u}_h^d across elemental boundaries, as well as the boundary condition on $\partial \Omega_N$. If the space of Lagrange multipliers M_h is chosen appropriately, then the “broken” velocity \mathbf{u}_h^d , albeit sought a priori in a discontinuous space, will coincide with its $\mathbf{H}(\text{div})$ -conforming counterpart. Specifically, the formulations in (31)–(32) and (19)–(20) are solving equivalent problems if the normal components of $\mathbf{w} \in \mathbf{U}_h$ lie in the same polynomial space as the trace functions (Arnold and Brezzi, 1985).

The discrete matrix system arising from (31)–(33) has the general form:

$$\begin{bmatrix} A_{00} & A_{01} & A_{02} \\ A_{10} & A_{11} & A_{12} \\ A_{20} & A_{21} & A_{22} \end{bmatrix} \begin{Bmatrix} U^d \\ P \\ \Lambda \end{Bmatrix} = \begin{Bmatrix} F_0 \\ F_1 \\ F_2 \end{Bmatrix}, \quad (34)$$

where the discrete system is produced by expanding functions in terms of the finite element bases for \mathbf{U}_h^d , V_h , and M_h like before. Upon initial inspection, it may not appear to be advantageous to replace our original formulation with this augmented equation-set; the hybridizable system has substantially more total degrees of freedom. However, (34) has a considerable advantage over (22) in the following ways:

1. Since both \mathbf{U}_h^d and V_h are discontinuous spaces, U^d and P are coupled only within the cell. This allows us to simultaneously eliminate both unknowns via *local* static condensation to produce a significantly smaller global (hybridized) problem for the trace unknowns, Λ :

$$S\Lambda = E, \quad (35)$$

where $S = \{S^K\}_{K \in \Omega_h}$ and $E = \{E^K\}_{K \in \Omega_h}$ are assembled by gathering the local contributions:

$$S^K = A_{22}^K - \begin{bmatrix} A_{20}^K & A_{21}^K \end{bmatrix} \begin{bmatrix} A_{00}^K & A_{01}^K \\ A_{10}^K & A_{11}^K \end{bmatrix}^{-1} \begin{bmatrix} A_{02}^K \\ A_{12}^K \end{bmatrix}, \quad (36)$$

$$E^K = F_2^K - \begin{bmatrix} A_{20}^K & A_{21}^K \end{bmatrix} \begin{bmatrix} A_{00}^K & A_{01}^K \\ A_{10}^K & A_{11}^K \end{bmatrix}^{-1} \begin{Bmatrix} F_0^K \\ F_1^K \end{Bmatrix}. \quad (37)$$



Note that the inverse of the block matrix in (36) and (37) is *never* evaluated globally; the elimination can be performed locally by performing a sequence of Schur-complement reductions within each cell.

2. The matrix S is sparse, symmetric, positive-definite, and spectrally similar to the dense Schur-complement $C+BA^{-1}B^T$ from (22) of the original mixed formulation (Cockburn et al., 2009a).
- 5 3. Once Λ is computed, both U^d and P can be recovered locally in each element. This can be accomplished in a number ways. One way is to compute P^K by solving:

$$\left(A_{11}^K - A_{10}^K (A_{00}^K)^{-1} A_{01}^K \right) P^K = F_1^K - A_{10}^K (A_{00}^K)^{-1} F_0^K - \left(A_{12}^K - A_{10}^K (A_{00}^K)^{-1} A_{02}^K \right) \Lambda^K, \quad (38)$$

followed by solving for $(U^d)^K$:

$$A_{00}^K (U^d)^K = F_0^K - A_{01}^K P^K - A_{02}^K \Lambda^K. \quad (39)$$

- 10 Similarly, one could rearrange the order in which each variable is reconstructed.

4. If desired, the solutions can be improved further through local post-processing. We highlight two such procedures, for U^d and P respectively, in Section 3.3.

- 15 Listing 1 displays the corresponding Slate code for assembling the trace system, solving (35), and recovering the eliminated unknowns. For a complete reference on how to formulate the hybridized mixed system (31)–(33) in UFL, we refer the reader to Alnæs et al. (2014). We remark that, in the case of this hybridizable system, (34) contains zero-valued blocks which can simplify the resulting expressions in (36)–(37) and (38)–(39). This is not true in general and therefore the expanded form using all sub-blocks of (34) is presented for completeness.



Listing 1. Firedrake code for solving (34) via static condensation and local recovery, given UFL expressions a , L for (31)–(33). Arguments of the mixed space $U_h^d \times V_h \times M_h$ are indexed by 0, 1, and 2 respectively. Lines 8 and 9 are symbolic expressions for (36) and (37) respectively. Any vanishing conditions on the trace variables should be provided as boundary conditions during operator assembly (line 12). Lines 27 and 29 are expressions for (38) and (39) (using LU). Code generation occurs in lines 12, 13, 31, and 32. A global linear solver for the reduced system is created and used in line 15. Configuring the linear solver is done by providing an appropriate Python dictionary of solver options for the PETSc library.

```

1  # Element tensors defining the local 3-by-3 block system
2  _A = Tensor(a)
3  _F = Tensor(L)
4
5  # Extracting blocks for Slate expression of the reduced system
6  A = _A.blocks
7  F = _F.blocks
8  S = A[2, 2] - A[2, :2] * A[:2, :2].inv * A[:2, 2]
9  E = F[2] - A[2, :2] * A[:2, :2].inv * F[:2]
10
11 # Assemble and solve: SΛ = E
12 Smat = assemble(S, bcs=[...])
13 Evec = assemble(E)
14 lambda_h = Function(M)
15 solve(Smat, lambda_h, Evec, solver_parameters={"ksp_type": "preonly",
16                                               "pc_type": "lu"})
17 p_h = Function(V) # Function to store the result: P
18 u_h = Function(U) # Function to store the result: U^d
19
20 # Intermediate expressions
21 Sd = A[1, 1] - A[1, 0] * A[0, 0].inv * A[0, 1]
22 Sl = A[1, 2] - A[1, 0] * A[0, 0].inv * A[0, 2]
23 Lambda = AssembledVector(lambda_h) # Local coefficient vector for Λ
24 P = AssembledVector(p_h) # Local coefficient vector for P
25
26 # Local solve expressions for P and U^d
27 p_sys = Sd.solve(F[1] - A[1, 0] * A[0, 0].inv * F[0] - Sl * Lambda,
28                 decomposition="PartialPivLu")
29 u_sys = A[0, 0].solve(F[0] - A[0, 1] * P - A[0, 2] * Lambda,
30                      decomposition="PartialPivLu")
31 assemble(p_sys, p_h)
32 assemble(u_sys, u_h)
    
```



3.2 Hybridization of discontinuous Galerkin methods

The hybridized discontinuous Galerkin (HDG) method is a natural extension of discontinuous Galerkin (DG) discretizations. Here, we consider a specific HDG discretization, namely the LDG-H method (Cockburn et al., 2010b). Other forms of HDG that involve local lifting operators can also be implemented in this software framework by the introduction of additional local (i.e., discontinuous) variables in the definition of the local solver.

To construct the LDG-H discretization, we define the DG numerical fluxes \widehat{p} and $\widehat{\mathbf{u}}$ to be functions of the trial unknowns and a new independent unknown in the trace space M_h :

$$\widehat{\mathbf{u}}(\mathbf{u}_h, p_h, \lambda_h; \tau) = \mathbf{u}_h + \tau(p_h - \widehat{p})\mathbf{n}, \quad (40)$$

$$\widehat{p}(\lambda_h) = \lambda_h, \quad (41)$$

where $\lambda_h \in M_h$ is a function approximating the trace of p on $\partial\mathcal{T}_h$ and τ is a positive function that may vary on each facet $e \in \partial\mathcal{T}_h$. The full LDG-H formulation reads as follows. Find $(\mathbf{u}_h, p_h, \lambda_h) \in \mathbf{U}_h \times V_h \times M_h$ such that

$$(\mathbf{w}, \mu\mathbf{u}_h)_{\mathcal{T}_h} - (\nabla \cdot \mathbf{w}, p_h)_{\mathcal{T}_h} + \langle \llbracket \mathbf{w} \rrbracket, \lambda_h \rangle_{\partial\mathcal{T}_h} = 0, \quad \forall \mathbf{w} \in \mathbf{U}_h, \quad (42)$$

$$-(\nabla\phi, \mathbf{u}_h)_{\mathcal{T}_h} + \langle \phi, \llbracket \mathbf{u}_h + \tau(p_h - \lambda_h)\mathbf{n} \rrbracket \rangle_{\partial\mathcal{T}_h} + (\phi, cp_h)_{\mathcal{T}_h} = (\phi, f)_{\mathcal{T}_h}, \quad \forall \phi \in V_h, \quad (43)$$

$$\langle \gamma, \llbracket \mathbf{u}_h + \tau(p_h - \lambda_h)\mathbf{n} \rrbracket \rangle_{\partial\mathcal{T}_h \setminus \partial\Omega_D} + \langle \gamma, \lambda_h \rangle_{\partial\Omega_D} = \langle \gamma, g \rangle_{\partial\Omega_N} + \langle \gamma, p_0 \rangle_{\partial\Omega_D}, \quad \forall \gamma \in M_h, \quad (44)$$

Equation (44) enforces continuity of the numerical flux $\widehat{\mathbf{u}}$ on $\partial\mathcal{T}_h$, which in turn produces a flux that is single-valued on the facets. Note that the choice of τ has a significant influence on the expected convergence rates of the computed solutions.

The matrix system arising from (42)–(44) has the same general form as that of the hybridized mixed method in (34), except all sub-blocks are now populated with non-zero entries due to the coupling of trace functions with both p_h and \mathbf{u}_h . However, all previous properties of the discrete matrix system from Section 3.1 still apply. The Slate expressions for the local elimination and reconstruction operations will be identical to those of Listing 1. For the interested reader, a unified analysis of hybridization methods (both mixed and DG) for second-order elliptic equations is presented in Cockburn et al. (2009a); Cockburn (2016).

3.3 Local post-processing

For both mixed (Arnold and Brezzi, 1985; Bramble and Xu, 1989; Stenberg, 1991) and discontinuous Galerkin methods (Cockburn et al., 2010b, 2009b), it is possible to locally post-process solutions to obtain superconvergent approximations (gaining one order of accuracy over the unprocessed solution). These methods can be expressed as local solves on each element, and so, in addition to static condensation, the Slate language also provides access to code generation for local post-processing of computed solutions.

Here, we present two post-processing techniques: one for scalar fields, and another for the vector unknown. The Slate code follows naturally from previous discussions in Sections 3.1 and 3.2, using the standard set of operations on local tensors summarized in Section 2.1.



3.3.1 Post-processing of the scalar solution

Our first example is a modified version of the procedure presented by Stenberg (1991) for enhancing the accuracy of the scalar solution. This was also highlighted within the context of hybridizing eigenproblems by Cockburn et al. (2010a). This post-processing technique can be used for both the hybridized mixed and LDG-H methods.

- 5 Let $\mathcal{P}_k(K)$ denote a polynomial space of degree $\leq k$ on an element $K \in \mathcal{T}_h$. Then for a given pair of computed solutions \mathbf{u}_h, p_h of the hybridized methods, we define the post-processed scalar $p_h^* \in \mathcal{P}_{k+1}(K)$ as the unique solution of the local problem:

$$(\nabla w, \nabla p_h^*)_K = -(\nabla w, \kappa^{-1} \mathbf{u}_h)_K, \quad \forall w \in \mathcal{P}_{k+1}^{\perp, l}(K), \quad (45)$$

$$(v, p_h^*)_K = (v, p_h)_K, \quad \forall v \in \mathcal{P}_l(K), \quad (46)$$

- 10 where $0 \leq l \leq k$. Here, the space $\mathcal{P}_{k+1}^{\perp, l}(K)$ denotes the L^2 -orthogonal complement of $\mathcal{P}_l(K)$. This post-processing method directly uses the definition of the flux $\mathbf{u}_h = -\kappa \nabla p_h$ to construct the local problem. In practice, the space $\mathcal{P}_{k+1}^{\perp, l}(K)$ may be constructed using an orthogonal hierarchical basis, and solving (45)–(46) amounts to inverting a local symmetric positive definite system.

- At the time of this work, Firedrake does not support the construction of such a finite element basis. However, we can
 15 introduce Lagrange multipliers to enforce the orthogonality constraint. The resulting local problem then becomes the following mixed system: find $(p_h^*, \psi) \in \mathcal{P}_{k+1}(K) \times \mathcal{P}_l(K)$ such that

$$(\nabla w, \nabla p_h^*)_K + (w, \psi)_K = -(\nabla w, \kappa^{-1} \mathbf{u}_h)_K, \quad \forall w \in \mathcal{P}_{k+1}(K), \quad (47)$$

$$(\phi, p_h^*)_K = (\phi, p_h)_K, \quad \forall \phi \in \mathcal{P}_l(K), \quad (48)$$

- where $0 \leq l \leq k$. The local problems (47)–(48) and (45)–(46) are equivalent, with the Lagrange multiplier ψ enforcing orthog-
 20 onality of test functions in $\mathcal{P}_{k+1}(K)$ with functions in $\mathcal{P}_l(K)$.

- This post-processing method produces a new approximation which superconverges at a rate of $k+2$ for hybridized mixed methods (Arnold and Brezzi, 1985; Cockburn et al., 2010a; Stenberg, 1991). For the LDG-H method, $k+2$ superconvergence is achieved when $\tau = \mathcal{O}(1)$ and $\tau = \mathcal{O}(h)$, but only $k+1$ convergence is achieved when $\tau = \mathcal{O}(1/h)$ (Cockburn et al., 2010b, 2009b). We demonstrate the increased accuracy in computed solutions in Section 5.1. An abridged example using Firedrake
 25 and Slate is provided in Listing 2.



Listing 2. Example of local post-processing using Firedrake and Slate. Here, we locally solve the mixed system defined in (45)–(46). The corresponding symbolic local tensors are defined in lines 9 and 11. The Slate expression for directly inverting the local system is written in line 12. In line 16, a Slate-generated kernel is produced which solves the resulting linear system in each cell. Since we are not interested in the multiplier, we only return the block corresponding to the new pressure field.

```

1 # Define spaces for the higher-order pressure approximation and Lagrange multipliers
2 DGk1 = FunctionSpace(mesh, "DG", degree + 1)
3 DG0 = FunctionSpace(mesh, "DG", 0)
4 W = DGk1 * DG0
5 p, psi = TrialFunctions(W)
6 w, phi = TestFunctions(W)
7
8 # Create local Slate tensors for the post-processing system
9 K = Tensor((inner(grad(p), grad(w)) + inner(psi, w) + inner(p, phi))*dx)
10 # Use computed pressure p_h and flux u_h in right-hand side
11 F = Tensor((-inner(u_h, grad(w)) + inner(p_h, phi))*dx)
12 E = K.inv * F
13
14 # Function for the post-processed scalar field p_h^*
15 p_star = Function(DGk1, name="Post-processed scalar")
16 assemble(E.blocks[0], p_star) # Only want the first field (pressure)
    
```

3.3.2 Post-processing of the flux

Our second example illustrates a procedure that uses the numerical flux of an HDG discretization for (13)–(16). Within the context of the LDG-H method, we can use the numerical trace in (40) to produce a vector field that is $\mathbf{H}(\text{div})$ -conforming. The technique we outline here follows that of Cockburn et al. (2009b).

Let \mathcal{T}_h be a mesh consisting of simplices³. On each element $K \in \mathcal{T}_h$, we define a new function \mathbf{u}_h^* to be the unique element of the local Raviart-Thomas space $[\mathcal{P}_k(K)]^n + \mathbf{x}\mathcal{P}_k(K)$ satisfying

$$(\mathbf{r}, \mathbf{u}_h^*)_K = (\mathbf{r}, \mathbf{u}_h)_K, \quad \forall \mathbf{r} \in [\mathcal{P}_{k-1}(K)]^n, \quad (49)$$

$$\langle \mu, \mathbf{u}_h^* \cdot \mathbf{n} \rangle_e = \langle \mu, \hat{\mathbf{u}}_h \cdot \mathbf{n} \rangle_e, \quad \forall \mu \in \mathcal{P}_k(e), \forall e \in \partial K. \quad (50)$$

10 This local problem produces a new velocity \mathbf{u}_h^* with the following properties:

1. \mathbf{u}_h^* converges at the *same* rate as \mathbf{u}_h for all choices of τ producing a solvable system for (42)–(44). However,
2. $\mathbf{u}_h^* \in \mathbf{H}(\text{div}; \Omega)$. That is,

$$[[\mathbf{u}_h^*]]_e = 0, \quad \forall e \in \mathcal{E}_h^\circ. \quad (51)$$

³This particular post-processing strategy only works on triangles and tetrahedra.



3. Additionally, the divergence of \mathbf{u}_h^* convergences at a rate of $k + 1$.

The Firedrake implementation using Slate is similar to the scalar post-processing example (see Listing 2); the element-wise linear systems (49)–(50) can be expressed in UFL, and therefore the necessary Slate expressions to invert the local systems follows naturally from the set of operations presented in Section 2.1. We use the very sensitive parameter dependency in the post-processing methods to validate our software implementation in Zenodo/Tabula-Rasa (2019).

4 Static condensation as a preconditioner

Slate enables static condensation approaches to be expressed very concisely. Nonetheless, application of a particular approach to different variational problems using Slate still requires a certain amount of code repetition. By formulating each form of static condensation as a preconditioner, code can be written once and then applied to any mathematically suitable problem. Rather than writing the static condensation by hand, in many cases, it is sufficient to just select the appropriate, Slate-based, preconditioner.

For context, it is helpful to frame the problem in the particular context of the solver library. Firedrake uses PETSc to provide linear solvers, and we implement our preconditioners as PETSc PC objects. These are defined to act on the problem residual, and return a correction to the solution. Specifically, we can think of (left) preconditioning the matrix equation in residual form:

$$r = r(A, b) \equiv b - Ax = 0 \quad (52)$$

by an operator \mathbf{P} (which may not necessarily be linear) as a transformation into an *equivalent* system of the form

$$\mathbf{P}r = \mathbf{P}(b - Ax) = 0. \quad (53)$$

Given a current iterate x_i the residual at the i -th iteration is simply $r_i \equiv b - Ax_i$, and \mathbf{P} acts on the residual to produce an approximation to the error $\epsilon_i \equiv x - x_i$. If \mathbf{P} is an application of an exact inverse, the residual is converted into an exact (up to numerical round-off) error.

We will denote the application of particular Krylov subspace method (KSP) for the linear system (52) as $\mathcal{K}_x(r(A, b))$. Upon preconditioning the system via \mathbf{P} as in (53), we write

$$\mathcal{K}_x(\mathbf{P}r(A, b)). \quad (54)$$

If (54) is solved directly via the application of A^{-1} , then $\mathbf{P}r(A, b) = A^{-1}b - x$. So, we have that $\mathcal{K}_x(\mathbf{P}r(A, b)) = \mathcal{K}_x(r(I, A^{-1}b))$ produces the exact solution of (52) in a single iteration of \mathcal{K} . Having established notation, we now present our implementation of static condensation via Slate by defining the appropriate operator, \mathbf{P} .

4.1 Interfacing with PETSc via custom preconditioners

The implementation of preconditioners for these systems requires manipulation not of assembled matrices, but rather their symbolic representation. To do this, we use the preconditioning infrastructure developed by Kirby and Mitchell (2018), which



gives preconditioners written in Python access to the symbolic problem description. In Firedrake, this means all derived preconditioners have direct access to the UFL representation of the PDE system. From this mathematical specification, we manipulate this appropriately via Slate and provide operators assembled from Slate expressions to PETSc for further algebraic preconditioning. Using this approach, we have developed a static condensation interface for the hybridization of $\mathbf{H}(\text{div}) \times L^2$ mixed problems, and a generic interface for statically condensing hybridized systems. The advantage of writing even the latter as a preconditioner is the ability to switch out the solution scheme for the system, even when nested inside a larger set of coupled equations at *runtime*.

4.1.1 A static condensation interface for hybridization

As discussed in sections 3.1 and 3.2, one of the main advantages of using a hybridizable variant of a DG or mixed method is that such systems permit the use of element-wise condensation and recovery. To facilitate this, we provide a PETSc PC static condensation interface, SCPC. This preconditioner takes the discretized system as in (34), and performs the local elimination and recovery procedures. Slate expressions are generated from the underlying UFL problem description.

More precisely, the incoming system has the form:

$$\begin{bmatrix} A_{e,e} & A_{e,c} \\ A_{c,e} & A_{c,c} \end{bmatrix} \begin{Bmatrix} X_e \\ X_c \end{Bmatrix} = \begin{Bmatrix} R_e \\ R_c \end{Bmatrix}, \quad (55)$$

where X_e is the vector of unknowns to be eliminated, X_c is the vector of unknowns for the condensed field, and R_e, R_c are the incoming right-hand sides. The partitioning in (55) is determined by the SCPC option: `pc_sc_eliminate_fields`. Field indices are provided in the same way one configures solver options to PETSc. These indices determine which field(s) to statically condense into. For example, on a three-field problem (with indices 0, 1, and 2), setting `-pc_sc_eliminate_fields 0, 1` will configure SCPC to cell-wise eliminate field 0 and 1; the resulting condensed system is associated with field 2.

In exact arithmetic, the SCPC preconditioner applies the inverse of the Schur-complement factorization of (55):

$$\mathbf{P} = \begin{bmatrix} I & A_{e,e}^{-1}A_{e,c} \\ 0 & I \end{bmatrix} \begin{bmatrix} A_{e,e}^{-1} & 0 \\ 0 & S^{-1} \end{bmatrix} \begin{bmatrix} I & 0 \\ A_{c,e}A_{e,e}^{-1} & I \end{bmatrix}, \quad (56)$$

where $S = A_{c,c} - A_{c,e}A_{e,e}^{-1}A_{e,c}$ is the Schur-complement operator for the X_c system. The distinction here from block preconditioners via `fieldsplit` (Brown et al., 2012), for example, is that \mathbf{P} does not require global actions; by design $A_{e,e}^{-1}$ can be inverted locally and S is sparse. As a result, S can be assembled or applied exactly via Slate-generated kernels.

In practice, the only globally coupled system requiring iterative inversion is S :

$$\mathcal{K}_{X_c}(\mathbf{P}_1 r(S, R_E)), \quad (57)$$

where R_E is the condensed right-hand side and \mathbf{P}_1 is another possible choice of preconditioner for S . Once X_c is computed, X_e is reconstructed element-wise via inverting the local systems.

By construction, this preconditioner is suitable for both hybridized mixed and HDG discretizations. It can also be used within other contexts, such as the static condensation of continuous Galerkin discretizations (Guyan, 1965; Irons, 1965) or



primal-hybrid methods (Devloo et al., 2018). As with any PETSc preconditioner, solver options can be specified for inverting S via the appropriate options prefix (`condensed_field`). The resulting KSP for (57) is compatible with existing solvers and external packages provided through the PETSc library. This allows users to experiment with a direct method and then switch to a more parallel-efficient iterative solver without changing the core application code.

5 4.1.2 Preconditioning mixed methods via hybridization

The preconditioner `HybridizationPC` expands on the previous one, this time taking an $\mathbf{H}(\text{div}) \times L^2$ system and automatically forming the hybridizable problem. This is accomplished through manipulating the UFL objects representing the discretized PDE. This includes replacing argument spaces with their discontinuous counterparts, introducing test functions on an appropriate trace space, and providing operators assembled from Slate expressions in a similar manner as described in section 4.1.1.

More precisely, let $AX = R$ be the incoming mixed saddle point problem, where $R = \begin{Bmatrix} R_U & R_P \end{Bmatrix}^T$, $X = \begin{Bmatrix} U & P \end{Bmatrix}^T$, and U and P are the velocity and scalar unknowns respectively. Then this preconditioner replaces $AX = R$ with the augmented system:

$$\begin{bmatrix} \hat{A} & K^T \\ K & 0 \end{bmatrix} \begin{Bmatrix} \hat{X} \\ \Lambda \end{Bmatrix} = \begin{Bmatrix} \hat{R} \\ R_\Lambda \end{Bmatrix} \quad (58)$$

where $\hat{R} = \begin{Bmatrix} \hat{R}_U & R_P \end{Bmatrix}^T$, \hat{R}_U , R_P are the right-hand sides for the flux and scalar equations respectively, and $\hat{\cdot}$ indicates modified matrices and co-vectors with discontinuous functions. Here, $\hat{X} = \begin{Bmatrix} U^d & P \end{Bmatrix}^T$ are the hybridized (discontinuous) unknowns to be determined.

The preconditioning operator for the hybrid-mixed system (58) has the form:

$$\hat{P} = \begin{bmatrix} I & \hat{A}^{-1}K^T \\ 0 & I \end{bmatrix} \begin{bmatrix} \hat{A}^{-1} & 0 \\ 0 & S^{-1} \end{bmatrix} \begin{bmatrix} I & 0 \\ K\hat{A}^{-1} & I \end{bmatrix}, \quad (59)$$

where S is the Schur-complement matrix $S = -K\hat{A}^{-1}K^T$. As before, a single globally coupled system for Λ is required. The recovery of U^d and P happens in the same manner as SCPC.

Since the flux is constructed in a discontinuous space \mathbf{U}_h^d , we must project the computed solution into $\mathbf{U}_h \subset \mathbf{H}(\text{div})$. This can be done cheaply via local facet averaging. The resulting solution is then updated via $U \leftarrow \Pi_{\text{div}} U^d$, where $\Pi_{\text{div}} : \mathbf{U}_h^d \rightarrow \mathbf{U}_h$ is the projection mapping. This ensures the residual for the original mixed problem is properly evaluated to test for convergence.

With \hat{P} as in (59), the preconditioning operator for the original $AX = R$ system is:

$$P = \Pi \hat{P} \Pi^T, \quad \Pi = \begin{bmatrix} \Pi_{\text{div}} & 0 & 0 \\ 0 & I & 0 \end{bmatrix}. \quad (60)$$

We note here that assembly of the right-hand for the Λ system requires special attention. Firstly, when Neumann conditions are present, then R_Λ is not necessarily 0. Since the hybridization preconditioner has access to the entire Python context (which



includes a list of boundary conditions and the spaces in which they are applied), surface integrals on the exterior boundary are added where appropriate and incorporated in the generated Slate expressions. A more subtle issue that requires extra care is the incoming right-hand side tested in U_h .

The situation we are given is that we have $R_U = R_U(\mathbf{w})$ for $\mathbf{w} \in U_h$, but require $\widehat{R}_U(\mathbf{w}^d)$ for $\mathbf{w}^d \in U_h^d$. For consistency, we also require for any $\mathbf{w} \in U_h$ that

$$\widehat{R}_U(\mathbf{w}) = R_U(\mathbf{w}). \quad (61)$$

We can construct such a \widehat{R}_U satisfying (61) in the following way. By construction of the space U_h^d , we have for $\Psi_i \in U_h$:

$$\Psi_i = \begin{cases} \Psi_i^d & \Psi_i \text{ associated with an exterior facet node,} \\ \Psi_i^{d,+} + \Psi_i^{d,-} & \Psi_i \text{ associated with an interior facet node,} \\ \Psi_i^d & \Psi_i \text{ associated with a cell interior node,} \end{cases} \quad (62)$$

where $\Psi_i^d, \Psi_i^{d,\pm} \in U_h^d$, and $\Psi_i^{d,\pm}$ are functions corresponding to the positive and negative restrictions associated with the i -th facet node⁴. We then define our “broken” right-hand side via the local definition:

$$\widehat{R}_U(\Psi_i^d) = \frac{R_U(\Psi_i)}{N_i}, \quad (63)$$

where N_i is the number of cells that the degree of freedom corresponding to the basis function $\Psi_i \in U_h$ touches. Using (62), (63), and the fact that R_U is linear in its argument, we can verify that our construction of \widehat{R}_U satisfies (61).

5 Numerical studies

We now present results utilizing the Slate DSL and our static condensation preconditioners for a set of test problems. Since we are using the interfaces outlined in Section 4, Slate is accessed indirectly and requires no manually-written solver code for hybridization or static condensation/local recovery. All parallel results were obtained on a single fully-loaded compute node of dual-socket Intel E5-2630v4 (Xeon) processors with 2×10 cores (2 threads per core) running at 2.2GHz. In order to avoid potential memory effects due to the operating system migrating processes between sockets, we pin MPI processes to cores.

Verification of the generated code is performed using parameter-sensitive convergence tests. The study consists of running a variety of discretizations spanning the methods outlined in Section 3. Details and numerical results are made public and can be viewed in Zenodo/Tabula-Rasa (2019) (see “Code and data availability”). All results are in full agreement with the theory.

⁴These are the two “broken” parts of Ψ_i on a particular facet connecting two elements. That is, for two adjacent cells, a basis function in U_h for a particular facet node can be decomposed into two basis functions in U_h^d defined on their respective sides of the facet.



5.1 LDG-H method for a three-dimensional elliptic equation

In this section, we take a closer look at the LDG-H method for the model elliptic equation (sign-definite Helmholtz):

$$-\nabla \cdot \nabla p + p = f \text{ in } \Omega = [0, 1]^3, \quad (64)$$

$$p = g \text{ on } \partial\Omega, \quad (65)$$

- 5 where f and g are chosen such that the analytic solution is $p = \exp\{\sin(\pi x)\sin(\pi y)\sin(\pi z)\}$. We use a regular mesh consisting of $6 \cdot N^3$ tetrahedral elements ($N \in \{4, 8, 16, 32, 64\}$). First, we reformulate (64)–(65) as the mixed problem:

$$\mathbf{u} + \nabla p = 0, \quad (66)$$

$$\nabla \cdot \mathbf{u} + p = f, \quad (67)$$

$$p = g \text{ on } \partial\Omega. \quad (68)$$

- 10 We start with linear polynomial approximations, up to cubic, for the LDG-H discretization of (66)–(68). Additionally, we compute a post-processed scalar approximation p_h^* of the HDG solution. This raises the approximation order of the computed solution by an additional degree. In all numerical studies here, we set the HDG parameter $\tau = 1$. All results were computed in parallel, utilizing a single compute node (described previously).


15 A continuous Galerkin (CG) discretization of the primal problem (64)–(65) serves as a reference for this experiment. Due to the superconvergence in the post-processed solution for the HDG method, we use CG discretizations of polynomial order 2, 3, and 4. This takes into account the enhanced accuracy of the HDG solution, despite being initially computed as a lower-order approximation. We therefore expect both methods to produce equally accurate solutions to the model problem.

20 Our aim here is not to compare the performance of HDG and CG, which has been investigated elsewhere (for example, see Kirby et al. (2012); Yakovlev et al. (2016)). Instead, we provide a reference that the reader might be more familiar with in order to evaluate whether our software framework produces a sufficiently performant HDG implementation relative to what might be expected.


25 For the CG discretization, we use a matrix-explicit iterative solver consisting of the conjugate gradient method preconditioned with hypre’s boomerAMG implementation of algebraic multigrid (AMG) Falgout et al. (2006). We use the preconditioner described in Section 4.1.1 to statically condense the LDG-H system, using the same iterative solver as the CG method for the Lagrange multipliers. This is sensible, as the global trace operator defines a symmetric positive-definite operator. To avoid over-solving, we iterate to a relative tolerance such that the discretization error is minimal for a given mesh.

5.1.1 Error versus execution time

30 The total execution time is recorded for the CG and HDG solvers, which includes the setup time for the AMG preconditioner, matrix-assembly, and the time-to-solution for the Krylov method. In the HDG case, we include the time spent building the Schur-complement for the traces, local recovery of the scalar and flux approximations, and post-processing. The L^2 -error against execution time is summarized in Figure 2a.



(a) Error against execution time for the CG and HDG with post-processing ($\tau = 1$) methods.



(b) Krylov iterations of the AMG-preconditioned conjugate gradient algorithm (to reach discretization error) against number of cells.

Figure 2. Comparison of continuous Galerkin and LDG-H solvers for the model three-dimensional positive-definite Helmholtz equation.

The HDG method of order $k - 1$ (HDG_{k-1}) with post-processing, as expected, produces a solution which is as accurate as the CG method of order k (CG_k). While the full HDG system is never explicitly assembled, the larger execution time is a result of several factors. The total number of trace unknowns for the HDG_1 , HDG_2 , and HDG_3 discretizations is roughly four, three, and two times larger (resp.) than the corresponding number of CG unknowns. Therefore, each iteration is more expensive. Moreover, we also observe that the trace system requires more Krylov iterations to reach discretization error. The gap in total number of iterations starts to close as the approximation degree increases (see Figure 2b). The extra cost of HDG due to the larger degree-of-freedom count and the need to perform local tensor inversion is offset by the local conservation and stabilization properties which are useful for fluid dynamics applications.

5.1.2 Break down of solver time

- 10 The HDG method requires many more degrees of freedom than CG or primal DG methods. This is largely due to the fact that the HDG method simultaneously approximates the primal solution and its velocity. The global matrix for the traces is larger than the one for the CG system at low polynomial order. The execution time for HDG is then compounded by a more expensive global solve. We remind the reader that our goal here is to verify that the solve time for HDG is as might be expected given the problem size.
- 15 Figure 3 displays the execution times on a simplicial mesh consisting of 1.5 million elements. The execution times have been normalized by the CG total time, showing that the HDG method is roughly 3 times the execution time of the CG method. This is expected given the larger degree-of-freedom count. The raw numerical breakdown of the HDG and CG solvers are shown

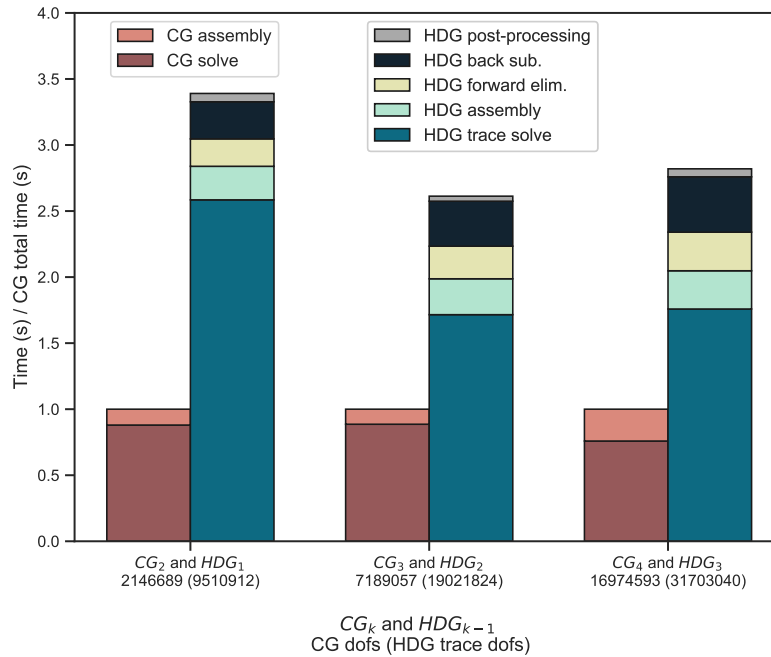


Figure 3. Break down of the CG_k and HDG_{k-1} execution times on a $6 \cdot 64^3$ simplicial mesh.

Table 1. Breakdown of the raw timings for the HDG_{k-1} ($\tau = 1$) and CG_k methods, $k = 2, 3$, and 4. Each method corresponds to a mesh size $N = 64$ on a fully-loaded compute node.

Stage	HDG_1		HDG_2		HDG_3	
	t_{stage} (s)	% t_{total}	t_{stage} (s)	% t_{total}	t_{stage} (s)	% t_{total}
Matrix assembly (static cond.)	1.05	7.49 %	6.95	10.40 %	31.66	10.27 %
Forward elimination	0.86	6.13 %	6.32	9.45 %	31.98	10.37 %
Trace solve	10.66	76.24 %	43.89	65.66 %	192.31	62.36 %
Back substitution	1.16	8.28 %	8.71	13.03 %	45.81	14.85 %
Post processing	0.26	1.86 %	0.98	1.46 %	6.62	2.15 %
HDG Total	13.98		66.85		308.37	
	CG_2		CG_3		CG_4	
	t_{stage} (s)	% t_{total}	t_{stage} (s)	% t_{total}	t_{stage} (s)	% t_{total}
Matrix assembly (monolithic)	0.50	12.01 %	2.91	11.39 %	26.37	24.11 %
Solve	3.63	87.99 %	22.67	88.61 %	82.99	75.89 %
CG Total	4.12		25.59		109.36	



in Table 1. We isolate each component of the HDG method contributing to the total execution time. Local operations include static condensation (trace operator assembly), forward elimination (right-hand side assembly for the trace system), backwards substitution to recover the scalar and velocity unknowns, and local post-processing of the primal solution. For all k , our HDG implementation is solver-dominated.

5 Both operator and right-hand side assembly are dominated by the costs of inverting a local square mixed matrix coupling the primal and dual variables, which is performed directly via an LU factorization. They should therefore be of the same magnitude in time spent. We observe that this is the case across all degrees (ranging between approximately 6—11% of total execution time for both local operations). Back-substitution takes roughly the same time as the static condensation and forward elimination stages (between 8—15% of execution time across all k). This is expected, as these operations are all dominated
 10 by the cost of inverting the local matrix coupling the scalar and velocity degrees of freedom. The slight increase in time is due splitting the local equations into two local solvers: one for p_h and another for the velocity. Finally, the additional cost of post-processing accrues negligible time (roughly 2% of execution time across all degrees).

We note that caching of local tensors does not occur. Each pass to perform the local eliminations and backwards reconstructions rebuilds the local element tensors. It is not clear at this time whether the performance gained from avoiding rebuilding
 15 the local operators will offset the memory costs of storing the local matrices. Moreover, in time-dependent problems where the operators may contain state-dependent variables, rebuilding local matrices will be necessary in each time-step regardless.

5.2 Hybridized-mixed solver for the shallow water equations

A primary motivator for our interest in hybridized methods revolves around developing efficient solvers for problems in geophysical flows. In section, we present some results integrating the nonlinear shallow water equations on the sphere using test
 20 case 5 (flow past a mountain) from Williamson et al. (1992). We use the framework of compatible finite elements (Cotter and Shipton, 2012; Cotter and Thuburn, 2014).

5.2.1 Semi-implicit finite element discretization

We start with the vector-invariant rotating nonlinear shallow water system defined on a two-dimensional spherical surface Ω embedded in \mathbb{R}^3 :

$$25 \quad \frac{\partial \mathbf{u}}{\partial t} + (\nabla^\perp \cdot \mathbf{u} + f) \mathbf{u}^\perp + \nabla \left(g(D + b) + \frac{1}{2} |\mathbf{u}|^2 \right) = 0, \quad (69)$$

$$\frac{\partial D}{\partial t} + \nabla \cdot (\mathbf{u}D) = 0, \quad (70)$$

where \mathbf{u} is the fluid velocity, D is the depth field, f is the Coriolis parameter, g is the acceleration due to gravity, b is the bottom topography, and $(\cdot)^\perp \equiv \hat{\mathbf{z}} \times \cdot$, with $\hat{\mathbf{z}}$ being the unit normal to the surface Ω .

After discretizing in space and time using a semi-implicit scheme and Picard linearization, following Natale and Cotter
 30 (2017) and based on Melvin et al. (2018) (see Appendix A for a summary of the entire solution strategy), we must solve an



indefinite saddle point system at each step:

$$\begin{bmatrix} A & B \\ C & D \end{bmatrix} \begin{Bmatrix} \Delta U \\ \Delta D \end{Bmatrix} = \begin{Bmatrix} R_u \\ R_D \end{Bmatrix}. \quad (71)$$

The approximations ΔU and ΔD are sought in the mixed finite element spaces $U_h \subset \mathbf{H}(\text{div})$ and $V_h \subset L^2$ respectively. Pairings including the RT or BDM spaces such as $\text{RT}_k \times \text{DG}_{k-1}$ or $\text{BDM}_k \times \text{DG}_{k-1}$ fall within the set of compatible mixed spaces ideal for geophysical fluids (Cotter and Shipton, 2012; Natale et al., 2016; Melvin et al., 2018). In particular, the lowest-order RT method ($\text{RT}_1 \times \text{DG}_0$) on a structured quadrilateral grid (such as the latitude-longitude grid used many operational dynamical cores) corresponds to the Arakawa C-grid finite difference discretization.

In staggered finite difference models, the standard approach for solving (71) is to neglect the Coriolis term and eliminate the velocity unknowns ΔU to obtain a discrete elliptic problem for which smoothers like Richardson iterations or relaxation methods are convergent. This is more problematic in the compatible finite element framework, since A has a dense inverse. Instead, we use the preconditioner described in Section 4.1.2 to form the hybridized problem and eliminate both ΔU and ΔD locally.

5.2.2 Atmospheric flow over a mountain

As a test problem, we solve test case 5 of Williamson et al. (1992), on the surface of an Earth-sized sphere. We refer the reader to Cotter and Shipton (2012); Shipton et al. (2018) for a more comprehensive study on mixed finite elements for shallow water systems of this type. We use the mixed finite element pairs $(\text{RT}_1, \text{DG}_0)$ (lowest-order RT method) and $(\text{BDM}_2, \text{DG}_1)$ (next-to-lowest order BDM method) for the velocity and depth spaces. The sphere mesh is generated from 7 refinements of an icosahedron, resulting in a triangulation consisting of 327,680 elements in total. The discretization information is summarized in Table 2.

Table 2. The number of unknowns to be determined are summarized for each compatible finite element method. Resolution is the same for both methods.

Discretization properties					
Mixed method	# cells	Δx	Velocity unknowns	Depth unknowns	Total (millions)
$\text{RT}_1 \times \text{DG}_0$	327,680	≈ 43 km	491,520	327,680	0.8 M
$\text{BDM}_2 \times \text{DG}_1$			2,457,600	983,040	3.4 M

We run for a total of 25 time-steps, with a fixed number of 4 Picard iterations in each time-step. We compare the overall simulation time using two different solver configurations for the implicit linear system. First, we use a flexible variant of



GMRES⁵ acting on the outer system with an approximate Schur complement preconditioner:

$$P_{sc} = \begin{bmatrix} I & 0 \\ CA^{-1} & I \end{bmatrix} \begin{bmatrix} A^{-1} & 0 \\ 0 & \tilde{S}^{-1} \end{bmatrix} \begin{bmatrix} I & -A^{-1}B \\ 0 & I \end{bmatrix}, \quad (72)$$

where $\tilde{S} = D - C\text{diag}(A)^{-1}B$, and $\text{diag}(A)$ is a diagonal approximation to the velocity mass matrix. The Schur-complement system is inverted via GMRES due to the asymmetry from the Coriolis term, with the inverse of \tilde{S} as the preconditioning operator. The sparse approximation \tilde{S} is inverted using PETSc's smoothed aggregation multigrid (GAMG). The Krylov method is set to terminate once the preconditioned residual norm is reduced by a factor of 10^8 . A^{-1} is computed approximately using a single application of incomplete LU (zero fill-in).

Next, we use only the application of our hybridization preconditioner (no outer iterations), which replaces the original linearized mixed system with its hybrid-mixed equivalent. After hybridization, we have the problem: find $(\Delta \mathbf{u}_h^d, \Delta D_h, \lambda_h) \in \mathbf{U}_h^d \times V_h \times M_h$ such that

$$(\mathbf{w}, \Delta \mathbf{u}_h^d)_{\mathcal{T}_h} + \frac{\Delta t}{2} (\mathbf{w}, f(\Delta \mathbf{u}_h^d)^\perp)_{\mathcal{T}_h} - \frac{\Delta t}{2} (\nabla \cdot \mathbf{w}, g \Delta D_h)_{\mathcal{T}_h} + \langle \llbracket \mathbf{w} \rrbracket, \lambda_h \rangle_{\partial \mathcal{T}_h} = \hat{R}_u, \quad \forall \mathbf{w} \in \mathbf{U}_h^d, \quad (73)$$

$$(\phi, \Delta D_h)_{\mathcal{T}_h} + \frac{\Delta t}{2} (\phi, H \nabla \cdot \Delta \mathbf{u}_h^d)_{\mathcal{T}_h} = R_D, \quad \forall \phi \in V_h, \quad (74)$$

$$\langle \gamma, \llbracket \Delta \mathbf{u}_h^d \rrbracket \rangle_{\partial \mathcal{T}_h} = 0, \quad \forall \gamma \in M_h. \quad (75)$$

Note that the space M_h is chosen such that these trace functions when restricted to a facet $e \in \partial \mathcal{T}_h$ are from the same polynomial space as $\Delta \mathbf{u}_h \cdot \mathbf{n}$ restricted to that same facet. Additionally, it can be shown that λ_h is an approximation to $\Delta t g \Delta D / 2$.

The resulting three-field problem for (73)–(75) has the matrix form:

$$\begin{bmatrix} \hat{A} & K^T \\ K & 0 \end{bmatrix} \begin{Bmatrix} \Delta X \\ \Lambda \end{Bmatrix} = \begin{Bmatrix} \hat{R}_{\Delta X} \\ 0 \end{Bmatrix} \quad (76)$$

where \hat{A} is the discontinuous operator coupling $\Delta X = \left\{ \Delta U^d \quad \Delta D \right\}^T$, and $R_{\Delta X} = \left\{ \hat{R}_u \quad R_D \right\}^T$ are the problem residuals. An exact Schur-complement factorization is performed on (76), using Slate to generate the local elimination kernels. We use the same set of solver options for the inversion of \tilde{S} in (72) to invert the Lagrange multiplier system. The increments ΔU^d and ΔD are recovered locally, using Slate-generated kernels. Once recovery is complete, ΔU^d is injected back into the conforming $\mathbf{H}(\text{div})$ finite element space via $\Delta U \leftarrow \Pi_{\text{div}} \Delta U^d$. Based on the discussion in section 4.1.2, we apply:

$$P_{\text{hybrid}} = \Pi \left(\begin{bmatrix} I & \hat{A}^{-1} K^T \\ 0 & I \end{bmatrix} \begin{bmatrix} \hat{A}^{-1} & 0 \\ 0 & S^{-1} \end{bmatrix} \begin{bmatrix} I & 0 \\ K \hat{A}^{-1} & I \end{bmatrix} \right) \Pi^T. \quad (77)$$

Table 3 displays a summary of our findings. When using hybridization, we observe a significant reduction in time spent during the implicit solve stage compared to the approximate Schur-complement approach. This is primarily because we reduce

⁵We use a flexible version of GMRES on the outer system since we use an additional Krylov solver to iteratively invert the Schur-complement.



Table 3. Preconditioner solve times for a 25-step run with $\Delta t = 100$ s. These are cumulative times in each stage of the two preconditioners throughout the entire profile run. We display the average iteration count (rounded to the nearest integer) for both the outer and the inner Krylov solvers. The significant speedup when using hybridization is a direct result of eliminating the outer-most solver.

Preconditioner and solver details					
Mixed method	Preconditioner	t_{total} (s)	Avg. outer its.	Avg. inner its.	$\frac{t_{\text{total}}^{\text{SC}}}{t_{\text{total}}^{\text{hybrid}}}$
$\text{RT}_1 \times \text{DG}_0$	approx. Schur. (\mathbf{P}_{SC})	15.137	2	8	3.413
	hybridization ($\mathbf{P}_{\text{hybrid}}$)	4.434	None	2	
$\text{BDM}_2 \times \text{DG}_1$	approx. Schur. (\mathbf{P}_{SC})	300.101	4	9	5.556
	hybridization ($\mathbf{P}_{\text{hybrid}}$)	54.013	None	6	

Table 4. Breakdown of the cost (average) of a single application of the preconditioned flexible GMRES algorithm and hybridization. Hybridization takes approximately the same time per iteration.

Preconditioner	Stage	$\text{RT}_1 \times \text{DG}_0$		$\text{BDM}_2 \times \text{DG}_1$	
		t_{stage} (s)	% t_{total}	t_{stage} (s)	% t_{total}
approx. Schur (\mathbf{P}_{SC})	Schur solve	0.07592	91.28 %	0.78405	93.53 %
	invert velocity mass: A	0.00032	0.39 %	0.00678	0.81 %
	apply inverse: A^{-1}	0.00041	0.49 %	0.00703	0.84 %
	gmres other	0.00652	7.84 %	0.04041	4.82 %
	Total	0.08317		0.83827	
hybridization ($\mathbf{P}_{\text{hybrid}}$)	Transfer: $R_{\Delta X} \rightarrow \widehat{R}_{\Delta X}$	0.00322	7.26 %	0.00597	1.10 %
	Forward elim.: $-K \widehat{A}^{-1} \widehat{R}_{\Delta X}$	0.00561	12.64 %	0.12308	22.79 %
	Trace solve	0.02289	51.63 %	0.28336	52.46 %
	Back sub.	0.00986	22.23 %	0.12220	22.62 %
	Projection: $\Pi_{\text{div}} \Delta \widehat{U}$	0.00264	5.96 %	0.00516	0.96 %
	Total	0.04434		0.54013	



the number of required “outer” iterations to zero; the hybridization preconditioner is performing an *exact* factorization of the global hybridized operator. This is empirically supported when considering the per-iteration solve times, summarized in Table 4. Hybridization and the approximate Schur complement preconditioner are comparable in terms of average execution time, with hybridization being slightly faster per application. This further demonstrates that the primary cause for the longer execution time of the latter is a direct result of the additional outer iterations induced from using an approximate factorization to achieve the same reduction in the overall residual.

We also measure the reductions in the true-residual of the linear system (71). Our hybridized method reduces the residual by a factor of 10^8 on average, which coincides with the specified relative tolerance for the Krylov method on the trace system. Snapshots of a (coarser) 15 day simulation are provided in Figure 4 using the semi-implicit scheme described in this paper. We refer the reader to Shipton et al. (2018) for an exposition of shallow water test cases featuring the use of a hybridized implicit solver (as described in Section 4.1.2).

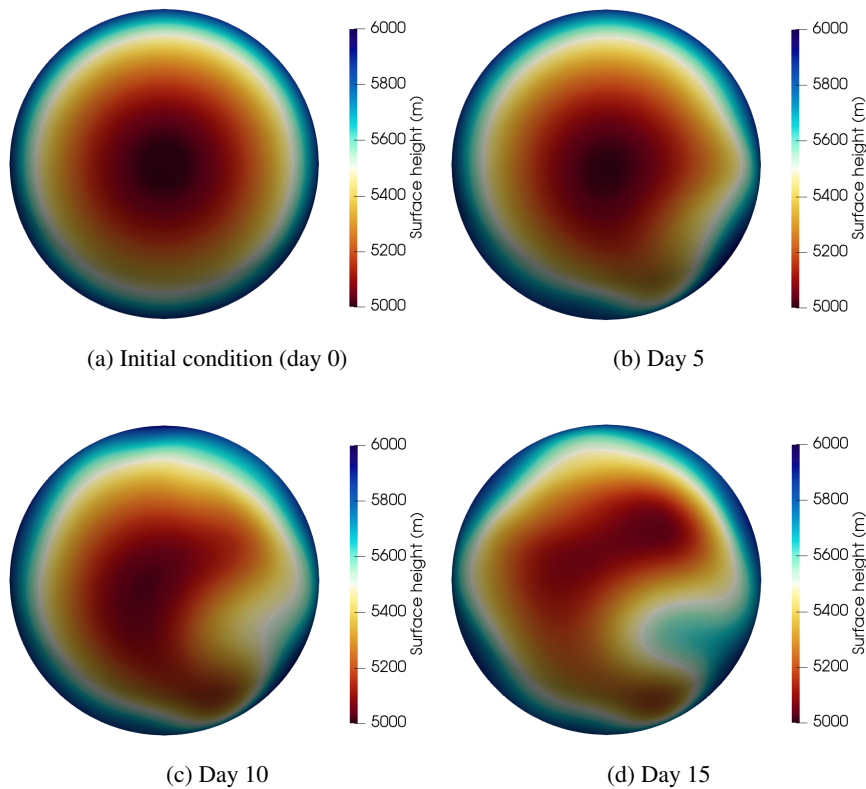


Figure 4. Snapshots (view from the northern pole) from the isolated mountain test case. The surface height (m) at days 5, 10, and 15. The snapshots were generated on a mesh with 20,480 simplicial cells, a $BDM_2 \times DG_1$ discretization, and $\Delta t = 500$ seconds. The linear system during each Picard cycle was solved using the hybridization preconditioner.



5.3 Rotating linear Boussinesq model

As a final example, we consider the simplified atmospheric model obtained from a linearization of the compressible Boussinesq equations in a rotating domain:

$$\frac{\partial \mathbf{u}}{\partial t} + 2\boldsymbol{\Omega} \times \mathbf{u} = \nabla p + b\hat{\mathbf{z}}, \quad (78)$$

$$5 \quad \frac{\partial p}{\partial t} = -c^2 \nabla \cdot \mathbf{u}, \quad (79)$$

$$\frac{\partial b}{\partial t} = -N^2 \mathbf{u} \cdot \hat{\mathbf{z}}, \quad (80)$$

where \mathbf{u} is the fluid velocity, p the pressure, b is the buoyancy, $\boldsymbol{\Omega}$ the planetary angular rotation vector, c is the speed of sound ($\approx 343\text{ms}^{-1}$), and N is the buoyancy frequency ($\approx 0.01\text{s}^{-1}$). Equations (78)–(80) permit fast-moving acoustic waves driven by perturbations in b . This is the same model presented in Skamarock and Klemp (1994), which uses a quadratic equation of state to avoid some of the complications of the full compressible Euler equations (the hybridization of which we shall address in future work). We solve these equations subject to the rigid-lid condition $\mathbf{u} \cdot \mathbf{n} = 0$ on all boundaries.

Our domain consists of a spherical annulus, with the mesh constructed from a horizontal “base” mesh of the surface of a sphere of radius R , extruded upwards by a height H_Ω . The vertical discretization is a structured one-dimensional grid, which facilitates the staggering of thermodynamic variables, such as b . We consider two kinds of meshes: one obtained by extruding an icosahedral sphere mesh, and another from a cubed sphere.

Since our mesh has a natural tensor-product structure, we construct suitable finite element spaces constructed by taking the tensor product of a horizontal space with a vertical space. To ensure our discretization is “compatible,” we use the one- and two-dimensional de-Rham complexes: $V_0 \xrightarrow{\partial_z} V_1$ and $U_0 \xrightarrow{\nabla^\perp} U_1 \xrightarrow{\nabla} U_2$. We can then construct the three-dimensional complex: $W_0 \xrightarrow{\nabla} W_1 \xrightarrow{\nabla \times} W_2 \xrightarrow{\nabla} W_3$, where

$$20 \quad W_0 = U_0 \otimes V_0, \quad (81)$$

$$W_1 = \text{HCurl}(U_1 \otimes V_0) \oplus \text{HCurl}(U_0 \otimes V_1) = W_1^v \oplus W_1^h, \quad (82)$$

$$W_2 = \text{HDiv}(U_2 \otimes V_0) \oplus \text{HDiv}(U_1 \otimes V_1) = W_2^v \oplus W_2^h, \quad (83)$$

$$W_3 = U_2 \otimes V_1. \quad (84)$$

Here, HCurl and HDiv denote operators which ensure the correct Piola transformations are applied when mapping from physical to reference element. We refer the reader to (McRae et al., 2016) for an overview of constructing tensor-product finite element spaces. For the analysis of compatible finite element discretizations and their relation to the complex (81)–(84), we refer the reader to Natale et al. (2016); Cotter and Thuburn (2014); Cotter and Shipton (2012). Each discretization used in this section is constructed from more familiar finite element families, shown in Table 5.



Table 5. Vertical and horizontal spaces for the three-dimensional compatible finite element discretization of the linear Boussinesq model. The RT_k and $BDFM_{k+1}$ methods are constructed on triangular prism elements, while the $RTCF_k$ method is defined on extruded quadrilateral elements.

Compatible finite element spaces					
Mixed method	V_0	V_1	U_0	U_1	U_2
RT_k	$CG_k([0, H_\Omega])$	$DG_{k-1}([0, H_\Omega])$	$CG_k(\Delta)$	$RT_k(\Delta)$	$DG_{k-1}(\Delta)$
$BDFM_{k+1}$	$CG_{k+1}([0, H_\Omega])$	$DG_k([0, H_\Omega])$	$CG_{k+1}(\Delta)$	$BDFM_{k+1}(\Delta)$	$DG_k(\Delta)$
$RTCF_k$	$CG_k([0, H_\Omega])$	$DG_{k-1}([0, H_\Omega])$	$Q_k(\square)$	$RTCF_k(\square)$	$DQ_{k-1}(\square)$

5.3.1 Compatible finite element discretization

We seek the unknown fields from (78)–(80) in the following finite element spaces:

$$\mathbf{u} \in W_2^0, \quad p \in W_3, \quad b \in W_b, \quad (85)$$

where W_2^0 is the subspace of W_2 satisfying the no-slip condition, and $W_b \equiv U_2 \otimes V_0$. Note that W_b is just the vertical part of the velocity space ⁶. That is, W_b and W_2^v have the same number of degrees of freedom, but differ in how they are pulled back to the reference element. For ease of notation, we write W_2 in place of W_2^0 .

To obtain the discrete system, we simply multiply equations (78)–(80) by test functions $\mathbf{w} \in W_2$, $\phi \in W_3$ and $\eta \in W_b$ and integrate by parts. We introduce the increments $\delta \mathbf{u} \equiv \mathbf{u}^{(n+1)} - \mathbf{u}^{(n)}$, and set $\mathbf{u}_0 \equiv \mathbf{u}^{(n)}$ (similarly for δp , p_0 , δb , and b_0). Using an implicit midpoint rule discretization, we need to solve the following linear variational problem at each time-step: find

10 $\delta \mathbf{u} \in W_2$, $\delta p \in W_3$ and $\delta b \in W_b$ such that

$$(\mathbf{w}, \delta \mathbf{u})_{T_h} + \Delta t (\mathbf{w}, \boldsymbol{\Omega} \times \delta \mathbf{u})_{T_h} - \frac{\Delta t}{2} (\nabla \cdot \mathbf{w}, \delta p)_{T_h} - \frac{\Delta t}{2} (\mathbf{w}, \delta b \hat{\mathbf{z}})_{T_h} = r_u, \quad (86)$$

$$(\phi, \delta p)_{T_h} + \frac{\Delta t}{2} c^2 (\phi, \nabla \cdot \delta \mathbf{u})_{T_h} = r_p, \quad (87)$$

$$(\eta, \delta b)_{T_h} + \frac{\Delta t}{2} N^2 (\eta, \delta \mathbf{u} \cdot \hat{\mathbf{z}})_{T_h} = r_b. \quad (88)$$

for all $\mathbf{w} \in W_2$, $\phi \in W_3$ and $\eta \in W_b$. The resulting matrix equations have the form:

$$15 \begin{bmatrix} A_u & -\frac{\Delta t}{2} D^T & -\frac{\Delta t}{2} Q^T \\ \frac{\Delta t}{2} c^2 D & M_p & 0 \\ \frac{\Delta t}{2} N^2 Q & 0 & M_b \end{bmatrix} \begin{Bmatrix} U \\ P \\ B \end{Bmatrix} = \begin{Bmatrix} R_u \\ R_p \\ R_b \end{Bmatrix}, \quad (89)$$

where $A_u = M_u + \Delta t C_\Omega$, C_Ω is the matrix associated with the Coriolis term, M_u , M_p , M_b are mass matrices, D is the weak divergence term, and Q is the operator coupling $\delta \mathbf{u}$ and δb .

⁶The choice for W_b in (85) corresponds to a Charney-Phillips vertical staggering of the buoyancy variable, which is the desired approach for the UK Met Office's Unified Model (Melvin et al., 2010). One could also collocate b with p ($b \in W_3$), which corresponds to a Lorenz staggering. This however supports a computational mode which is exacerbated by fast-moving waves. We restrict our discussion to the former case.



To solve (89), we can use the buoyancy equation to arrive at an elimination procedure by substituting:

$$\delta b = r_b - \frac{\Delta t}{2} N^2 \delta \mathbf{u} \cdot \hat{\mathbf{z}} \quad (90)$$

into (86). Note that in a planar geometry, (90) is satisfied point-wise. However, this is not true when orography is present or on the surface of a sphere. Hence this is a further approximation of the equations. After solving for \mathbf{u} and p , we reconstruct b from Equation (88) which requires solving the mass matrix M_b . Fortunately, M_b is well-conditioned (independent of the grid resolution/time-step) and decoupled columnwise; it can be inverted with a small number of conjugate gradient iterations.

The full solution procedure is outlined as follows. First, we approximately eliminate the buoyancy to obtain a mixed system for the velocity and pressure:

$$\mathcal{A} \begin{Bmatrix} U \\ P \end{Bmatrix} = \begin{bmatrix} \tilde{A}_u & -\frac{\Delta t}{2} D^T \\ c^2 \frac{\Delta t}{2} D & M_p \end{bmatrix} \begin{Bmatrix} U \\ P \end{Bmatrix} = \begin{Bmatrix} \tilde{R}_u \\ R_p \end{Bmatrix}, \quad (91)$$

where $\tilde{A}_u = A_u + \frac{\Delta t^2}{4} N^2 Q^T M_b^{-1} Q$ and $\tilde{R}_u = R_u + \frac{\Delta t}{2} Q^T M_b^{-1} R_b$ is the modified velocity operator and right-hand side respectively. Note that in our elimination strategy, the expression for \tilde{A}_u corresponds to the bilinear form obtained after substituting the point-wise expression for δb :

$$\tilde{A}_u \leftarrow (\mathbf{w}, \delta \mathbf{u})_{T_h} + \Delta t (\mathbf{w}, \boldsymbol{\Omega} \times \delta \mathbf{u})_{T_h} + \frac{\Delta t^2}{4} N^2 (\mathbf{w} \cdot \hat{\mathbf{z}}, \delta \mathbf{u} \cdot \hat{\mathbf{z}})_{T_h}. \quad (92)$$

A similar construction holds for \tilde{R}_u . Once (91) is solved, δb is reconstructed by solving:

$$M_b B + \frac{\Delta t}{2} N^2 Q U = R_b \quad \implies \quad B = M_b^{-1} R_b - \frac{\Delta t}{2} N^2 M_b^{-1} Q U. \quad (93)$$

Equation (93) can be efficiently inverted using a preconditioned conjugate gradient method (preconditioned by block ILU with zero in-fill).

5.3.2 Preconditioning the mixed velocity pressure system

The primary difficulty is find robust solvers for (91). This was studied by Mitchell and Müller (2016) within the context of developing a robust preconditioner to withstand fast-moving acoustic waves. This is critical, as (78)–(80) support fast waves driven by perturbations to the pressure field. However, the implicit treatment of the Coriolis term was not taken into account. We consider two preconditioning strategies.

The first strategy follows from Mitchell and Müller (2016). As with the rotating shallow water system in Section 5.2, we can construct a preconditioner based on the Schur-complement factorization of \mathcal{A} in (91):

$$\mathcal{A}^{-1} = \begin{bmatrix} I & \frac{\Delta t}{2} \tilde{A}_u^{-1} D^T \\ 0 & I \end{bmatrix} \begin{bmatrix} \tilde{A}_u^{-1} & 0 \\ 0 & H^{-1} \end{bmatrix} \begin{bmatrix} I & 0 \\ -c^2 \frac{\Delta t}{2} D \tilde{A}_u^{-1} & I \end{bmatrix}, \quad (94)$$

where $H = M_p + c^2 \frac{\Delta t^2}{4} D \tilde{A}_u^{-1} D^T$ is the dense pressure Helmholtz operator. Because we have chosen to include the Coriolis term, the operator H is non-symmetric, and we require a sparse approximation to

$$H = M_p + c^2 \frac{\Delta t^2}{4} D \left(\tilde{M}_u + \Delta t C \boldsymbol{\Omega} \right)^{-1} D^T \quad (95)$$



where \widetilde{M}_u is the modified velocity mass matrix. As Δt increases, the contribution of C_Ω becomes more prominent in H , making sparse approximations of H more challenging. We elaborate on this further below when we present the results of our second solver.

Our second strategy revolves around the hybridization of the system defined in (91), using Firedrake's `HybridizationPC`.

- 5 The space of traces is generated on the faces of triangular-prisms and extruded quadrilaterals, with appropriate polynomial degrees such that every trace function lies in the same polynomial space as $\delta \mathbf{u} \cdot \mathbf{n}|_f$ for all faces f .

We locally eliminate U and P after hybridization, and produce the resulting system for the traces:

$$H_\partial \Lambda = E, \quad H_\partial = K \widehat{\mathcal{A}}^{-1} K^T, \quad E = K \widehat{\mathcal{A}}^{-1} \begin{Bmatrix} \widehat{R}_u \\ R_p \end{Bmatrix}, \quad (96)$$

- where $\widehat{\mathcal{A}}$ is the result of hybridizing \mathcal{A} and H_∂ is the statically condensed trace operator to be inverted. The non-symmetric operator H_∂ is inverted using a preconditioned generalized conjugate residual (GCR) Krylov method, as suggested in Thomas et al. (2003). For our choice of preconditioner, we follow strategies outlined in Elman et al. (2001) and employ an algebraic multigrid method (V-cycle) with GMRES (five iterations) smoothers on the coarse levels. The GMRES smoothers are preconditioned with block ILU on each level. For the finest level, block ILU produces a line smoother (necessary for efficient solution on thin domains) when the trace variable nodes are numbered in vertical lines, as is the case in our Firedrake implementation.
- 15 On the coarser levels, less is known about the properties of ILU under the AMG coarsening strategies, but as we shall see, we observe performance that suggests ILU is still behaving as a line smoother. More discussion on multigrid for non-symmetric problems can be found in Bramble et al. (1994, 1988); Mandel (1986). A gravity wave test using our solution strategy and hybridization preconditioner is illustrated in Figure 5 for a problem on a condensed Earth (radius scaled down by a factor of 125) and 10km lid.

20 5.3.3 Robustness against acoustic Courant number with implicit Coriolis

A desired property of the implicit linear solver is robustness with respect to the time-step size. In particular, it is desirable for the execution time to remain constant across a wide range of Δt . As Δt increases, the conditioning of the elliptic operator becomes worse. Therefore, iterative solvers need to work much harder to reduce the residual down to a specified tolerance.

- In this final experiment, we repeat a similar study to that presented in Mitchell and Müller (2016). We fix the resolution of the problem and run our solver for a range of Δt . We measure this by adjusting the horizontal acoustic Courant number $\lambda_C = c \frac{\Delta t}{\Delta x}$, where c is the speed of sound and Δx is the horizontal resolution. We remark that the range of Courant numbers used in this paper exceeds what is typical in operational forecast settings (typically between $\mathcal{O}(2)$ – $\mathcal{O}(10)$). The grid set up mirrors that of actual global weather models; we extrude a spherical mesh of the Earth upwards to a height of 85km. The set up for the different discretizations (including degrees of freedom for the velocity-pressure and hybridized systems) is presented in
- 30 Table 6.




Figure 5. Buoyancy perturbation ($y - z$ cross section) at $t = 3600$ s from a simple gravity wave test ($\Delta t = 100$ s). The initial conditions (in lat.-long. coordinates) for the velocity is a simple solid-body rotation: $\mathbf{u} = 20\mathbf{e}_\lambda$, where \mathbf{e}_λ is the unit vector pointing in the direction of decreasing longitude. A buoyancy anomaly is defined via $b = \frac{d^2}{d^2+q^2} \sin(\pi z/10000)$, where $q = R \cos^{-1}(\cos(\phi) \cos(\lambda - \lambda_\phi))$, $d = 5000$ m, $R = 6371$ km/125 is the planet radius, and $\lambda_\phi = 2/3$. The equations are discretized using the lowest-order method RTCF_1 , with 24,576 quadrilateral cells in the horizontal and 64 extrusion levels. The velocity-pressure system is solved using hybridization.

Table 6. Grid set up and discretizations for the acoustic Courant number study. The total unknowns (velocity and pressure) and hybridized unknowns (broken velocity, pressure, and trace) are shown in the last two columns (millions). The vertical resolution is fixed across all discretizations.

Discretizations and grid information						
Mixed method	# horiz. cells	# vert. layers	Δx	Δz	U - P dofs	Hybrid. dofs
RT_1	81,920	85	86 km	1,000 m	24.5 M	59.3 M
RT_2	5,120	85	346 km	1,000 m	9.6 M	17.4 M
BDFM_2	5,120	85	346 km	1,000 m	10.5 M	18.3 M
RTCF_1	98,304	85	78 km	1,000 m	33.5 M	83.7 M
RTCF_2	6,144	85	312 km	1,000 m	16.7 M	29.3 M

It was shown in Mitchell and Müller (2016) that using a sparse approximation of the pressure Schur-complement of the form:

$$\tilde{H} = M_p + c^2 \frac{\Delta t^2}{4} D \left(\text{Diag}(\tilde{M}_u) \right)^{-1} D^T \quad (97)$$

served as a good preconditioner, leading to a system that was amenable to multigrid methods and resulted in a Courant number independent solver. However, when the Coriolis term is included, this is no longer the case: the diagonal approximation to \tilde{M}_u becomes worse with increasing λ_C . To demonstrate this, we solve the mixed problem on a low-resolution grid (10km lid, 10 vertical levels, maintaining the same cell aspect ratio as in Table 6) using the Schur-complement factorization (94), and use LU factorizations to apply both \tilde{A}_u^{-1} and \tilde{H}^{-1} . H^{-1} is computed using preconditioned GMRES iterations, and a flexible-GMRES



algorithm is used on the full velocity-pressure system. If \tilde{H}^{-1} is a good approximation to H^{-1} we should see low iteration counts when inverting H . Figure 6 shows the results of this study for a range of Courant numbers.




Figure 6. Number of Krylov iterations to invert the Helmholtz system using \tilde{H}^{-1} as a preconditioner. The preconditioner is applied using a direct LU factorization within a GMRES method on the entire pressure equation. While the lowest-order methods grow slowly over the Courant number range, the higher-order (by only one approximation order) methods quickly degrade and diverge after the critical range $\lambda_C = \mathcal{O}(2) - \mathcal{O}(10)$. At $\lambda_C > 32$, the solvers take over 150 iterations.

For the lower-order methods, the number of iterations to invert H grow slowly but remain under control. Increasing the approximation degree by one results in degraded performance. As Δt increases, the number of Krylov iterations needed to invert the system to a relative tolerance of 10^{-5} grows rapidly. It is clear that this sparse approximation is not robust against Courant number. This can be explained by the fact that diagonalizing the velocity operator fails to take into account the effects of the Coriolis term (which appear in off-diagonal positions in the operator). Even if one were to use traditional mass-lumping (row-wise sums), the Coriolis effects are effectively cancelled out due to asymmetry.

Hybridization avoids this problem: we always construct an exact Schur complement, and only have to worry about solving the trace system. We now show that this approach (described in Section 5.3.2) is much more robust to increases in the Courant number. We use the same workstation as for the three-dimensional CG/HDG problem in Section 5.1 (executed with a total of 40 MPI processes). Figure 7 shows the parameter test for all the discretizations described in Table 6. We see that, in terms of total number of GCR iterations needed to invert the trace system, hybridization is far more robust against Courant number than the approximate Schur complement approach. They largely remain constant throughout the entire parameter range, only varying by an iteration or two. It is not until after $\lambda_C > 32$ that we begin to see a larger jump in the number of GCR iterations. This is expected, since the Coriolis operator causes the problem to become singular for very large Courant numbers. However, unlike



with the approximate Schur-complement solver, iteration counts are still under control. In particular, each method (lowest and higher order) remains constant throughout the critical range (shaded in gray in Figures 7a and 7b).

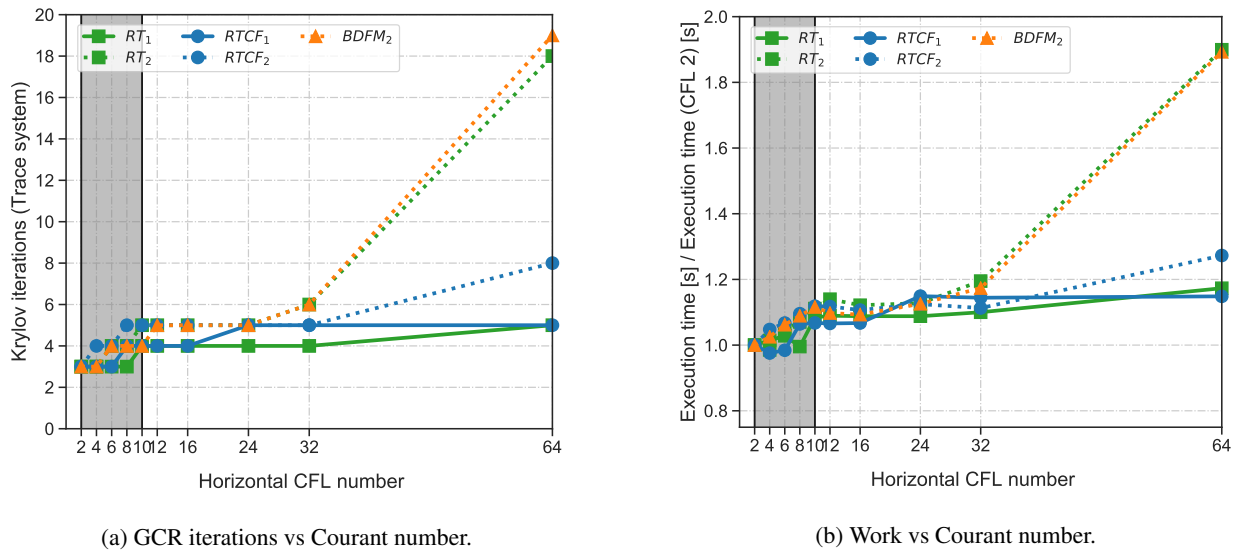


Figure 7. Courant number parameter test run on a fully-loaded compute node. Both figures display the hybridized solver for each discretization, described in Table 6. The left figure (a) displays the total iteration count (preconditioned GCR) to solve the trace system to a relative tolerance of 10^{-5} . The right figure (b) displays the relative work of each solver. Figure (b) takes into account not just the time-to-solution of the trace solver, but also the time required to forward eliminate and locally recover the velocity and pressure.

In Figure 7b, we display the ratio of execution time and the time-to-solution at the lowest Courant number of two. We perform this normalization to better compare the lower and higher order methods (and discretizations on triangular prisms vs extruded quadrilaterals). The calculation of the ratios include the time needed to eliminate and reconstruct the hybridized velocity/pressure variables. The fact that the hybridization solver remains close to one demonstrates that the entire solution procedure is largely λ_C -independent until around $\lambda_C = 32$. The overall trend is largely the same as what is observed in Figure 7a. This is due to our hybridization approach being solver dominated, with local operations like forward elimination together with local recovery taking approximately 1/3 of the total execution time for each method. The percentage breakdown of the hybridization solver is similar to what is already presented in Section 5.1.2.

Implicitly treating the Coriolis term has been discussed for semi-implicit discretizations of large-scale geophysical flows (Temperton, 1997; Côté and Staniforth, 1988; Cullen, 2001; Nechaev and Yaremchuk, 2004). Incorporating the Coriolis term in finite element discretizations is difficult, as this makes the challenge of finding a robust sparse approximation of the resulting elliptic operator even more difficult. Hybridization shows promise here, as we no longer require the inversion of a dense elliptic system. Instead, hybridization allows for the assembly of an elliptic equation that both captures the effects of rotation and results in a sparse linear system. In fact, the hybridization process mimics standard staggered finite difference elimination



procedures used in many global circulation models. In other words, hybridization is the finite element analogue of point-wise elimination strategies in finite difference codes.

6 Conclusions

We have presented Slate, and shown how this language can be used to create concise mathematical representations of localized linear algebra on the tensors corresponding to finite element forms. We have shown how this DSL can be used in tandem with UFL in Firedrake to implement solution approaches making use of automated code generation for static condensation, hybridization, and localized post-processing. Setup and configuration is done at runtime, allowing one to switch in different discretizations at will. In particular, this framework alleviates much of the difficulty in implementing such methods within intricate numerical, and paves the way for future low-level optimizations. In this way, the framework in this paper can be used to help enable the rapid development and exploration of new hybridization and static condensation techniques for a wide class of problems. We remark here that the reduction of global matrices via element-wise algebraic static condensation, as described in Guyan (1965); Irons (1965) is also possible using Slate, including other more general static condensation procedures outside the context of hybridization.

Our approach to preconditioner design revolves around its composable nature, in that these Slate-based implementations can be seamlessly incorporated into complicated solution schemes. In particular, there is current research in the design of dynamical cores for numerical weather prediction using implementations of hybridization and static condensation with Slate (Bauer and Cotter, 2018; Shipton et al., 2018). The performance of such methods for geophysical flows are a subject of ongoing investigation.

In this paper, we have provided some examples of hybridization procedures for compatible finite element discretizations of geophysical flows. These approaches avoid the difficulty in constructing sparse approximations of dense elliptic operators. Static condensation arising from hybridizable formulations can best be interpreted as producing an *exact* Schur-complement factorization on the global hybridizable system. This eliminates the need for outer iterations from a suitable Krylov method to solve the full mixed system, and replaces the original global mixed equations with a condensed elliptic system. More extensive performance benchmarks, which requires detailed analysis of the resulting operator systems arising from hybridization, is a necessary next-step to determine whether hybridization provides a scalable solution strategy for compatible finite elements in operational settings.

Code and data availability. The contribution in this paper is available through open-source software provided by the Firedrake Project: <https://www.firedrakeproject.org/>. We cite archives of the exact software versions used to produce the results in this paper. For all components of the Firedrake project used in this paper, see Zenodo/Firedrake (2019). The numerical experiments, full solver configurations, code-verification (including local-processing), and raw data are available in Zenodo/Tabula-Rasa (2019).



Appendix A: Semi-implicit method for the shallow water system

For some tessellation, \mathcal{T}_h , our semi-discrete mixed method for (69)–(70) seeks approximations $(\mathbf{u}_h, D_h) \in \mathbf{U}_h \times V_h \subset \mathbf{H}(\text{div}) \times L^2$ satisfying:

$$\begin{aligned} \left(\mathbf{w}, \frac{\partial \mathbf{u}_h}{\partial t} \right)_{\mathcal{T}_h} - (\nabla^\perp (\mathbf{w} \cdot \mathbf{u}_h^\perp), \mathbf{u}_h^\perp)_{\mathcal{T}_h} + (\mathbf{w}, f \mathbf{u}_h^\perp)_{\mathcal{T}_h} + \langle \llbracket \mathbf{n}^\perp \mathbf{w} \cdot \mathbf{u}_h^\perp \rrbracket, \tilde{\mathbf{u}}_h^\perp \rangle_{\partial \mathcal{T}_h} \\ - \left(\nabla \cdot \mathbf{w}, g(D_h + b) + \frac{1}{2} |\mathbf{u}_h|^2 \right)_{\mathcal{T}_h} = 0, \quad \forall \mathbf{w} \in \mathbf{U}_h, \end{aligned} \quad (\text{A1})$$

$$\left(\phi, \frac{\partial D_h}{\partial t} \right)_{\mathcal{T}_h} - (\nabla \phi, \mathbf{u}_h D_h)_{\mathcal{T}_h} + \langle \llbracket \phi \mathbf{u}_h \rrbracket, \tilde{D}_h \rangle_{\partial \mathcal{T}_h} = 0, \quad \forall \phi \in V_h, \quad (\text{A2})$$

where $\tilde{\cdot}$ indicates that the value of the function should be taken from the upwind side of each facet. The discretisation of the velocity advection operator is an extension of the energy-conserving scheme of Natale and Cotter (2017) to the shallow-water equations.

The time-stepping scheme follows a Picard iteration semi-implicit approach, where predictive values of the relevant fields are determined via an explicit step of the advection equations, and corrective updates are generated by solving an implicit linear system (linearized about a state of rest) for $(\Delta \mathbf{u}_h, \Delta D_h) \in \mathbf{U}_h \times V_h$, given by

$$(\mathbf{w}, \Delta \mathbf{u}_h)_{\mathcal{T}_h} + \frac{\Delta t}{2} (\mathbf{w}, f \Delta \mathbf{u}_h^\perp)_{\mathcal{T}_h} - \frac{\Delta t}{2} (\nabla \cdot \mathbf{w}, g \Delta D_h)_{\mathcal{T}_h} = -R_u[\mathbf{u}_h^{n+1}, D_h^{n+1}; \mathbf{w}], \quad \forall \mathbf{w} \in \mathbf{U}_h, \quad (\text{A3})$$

$$(\phi, \Delta D_h)_{\mathcal{T}_h} + \frac{H \Delta t}{2} (\phi, \nabla \cdot \Delta \mathbf{u}_h)_{\mathcal{T}_h} = -R_D[\mathbf{u}_h^{n+1}, D_h^{n+1}; \phi], \quad \forall \phi \in V_h, \quad (\text{A4})$$

where H is the mean layer depth, and R_u and R_D are residual linear forms that vanish when \mathbf{u}_h^{n+1} and D_h^{n+1} are solutions to the implicit midpoint rule time discretization of (A1)–(A2). The residuals are evaluated using the predictive values of \mathbf{u}_h^{n+1} and D_h^{n+1} .

The implicit midpoint rule time discretization of the non-linear rotating shallow water equations (A1)–(A2) is:

$$\begin{aligned} (\mathbf{w}, \mathbf{u}_h^{n+1} - \mathbf{u}_h^n)_{\mathcal{T}_h} - \Delta t \left(\nabla^\perp (\mathbf{w} \cdot \mathbf{u}_h^{*\perp}), \mathbf{u}_h^{*\perp} \right)_{\mathcal{T}_h} + \Delta t (\mathbf{w}, f \mathbf{u}_h^{*\perp})_{\mathcal{T}_h} \\ + \Delta t \langle \llbracket \mathbf{n}^\perp \mathbf{w} \cdot \mathbf{u}_h^{*\perp} \rrbracket, \tilde{\mathbf{u}}_h^{*\perp} \rangle_{\partial \mathcal{T}_h} \\ - \Delta t \left(\nabla \cdot \mathbf{w}, g(D_h^* + b) + \frac{1}{2} |\mathbf{u}_h^{*\perp}|^2 \right)_{\mathcal{T}_h} = 0, \quad \forall \mathbf{w} \in \mathbf{U}_h, \end{aligned} \quad (\text{A5})$$

$$(\phi, D_h^{n+1} - D_h^n)_{\mathcal{T}_h} - \Delta t (\nabla \phi, \mathbf{u}_h^* D_h^*)_{\mathcal{T}_h} + \Delta t \langle \llbracket \phi \mathbf{u}_h^* \rrbracket, \tilde{D}_h^* \rangle_{\partial \mathcal{T}_h} = 0, \quad \forall \phi \in V_h, \quad (\text{A6})$$

where $\mathbf{u}_h^* = (\mathbf{u}_h^{n+1} + \mathbf{u}_h^n)/2$ and $D_h^* = (D_h^{n+1} + D_h^n)/2$.

One approach to construct the residual functionals R_u and R_D would be to simply define these from (A5)–(A6). However, this leads to a small critical time-step for stability of the scheme. To make the numerical scheme more stable, we define



residuals as follows. For R_u , we first solve for $\mathbf{v}_h \in \mathbf{U}_h$ such that

$$\begin{aligned} (\mathbf{w}, \mathbf{v}_h - \mathbf{u}_h^n)_{\mathcal{T}_h} - \Delta t \left(\nabla^\perp \left(\mathbf{w} \cdot \mathbf{u}_h^{*\perp} \right), \mathbf{v}_h^{\#\perp} \right)_{\mathcal{T}_h} + \Delta t \left(\mathbf{w}, f \mathbf{v}_h^{\#\perp} \right)_{\mathcal{T}_h} \\ + \Delta t \langle \llbracket \mathbf{n}^\perp \mathbf{w} \cdot \mathbf{u}_h^{*\perp} \rrbracket, \tilde{\mathbf{v}}_h^{\#\perp} \rangle_{\partial \mathcal{T}_h} \\ - \Delta t \left(\nabla \cdot \mathbf{w}, g(D_h^* + b) + \frac{1}{2} |\mathbf{u}_h^*|^2 \right)_{\mathcal{T}_h} = 0, \quad \forall \mathbf{w} \in \mathbf{U}_h, \end{aligned} \quad (\text{A7})$$

5 where $\mathbf{v}_h^\# = (\mathbf{v}_h + \mathbf{u}_h^n)/2$. This is a linear variational problem. Then,

$$R_u[\mathbf{u}_h^{n+1}, D_h^{n+1}; \mathbf{w}] = (\mathbf{w}, \mathbf{v}_h - \mathbf{u}_h^{n+1})_{\mathcal{T}_h}. \quad (\text{A8})$$

Similarly, for R_D we first solve for $E_h \in V_h$ such that

$$(\phi, E_h - D_h^n)_{\mathcal{T}_h} - \Delta t (\nabla \phi, \mathbf{u}_h^* E_h^\#)_{\mathcal{T}_h} + \Delta t \langle \llbracket \phi \mathbf{u}_h^* \rrbracket, \tilde{E}_h^\# \rangle_{\partial \mathcal{T}_h} = 0, \quad \forall \phi \in V_h, \quad (\text{A9})$$

where $E_h^\# = (E_h + D_h^n)/2$. This is also a linear problem. Then,

$$10 \quad R_D[\mathbf{u}_h^{n+1}, D_h^{n+1}; \phi] = (\phi, E_h - D_h^{n+1})_{\mathcal{T}_h}. \quad (\text{A10})$$

This process can be thought of as iteratively solving for the average velocity and depth that satisfies the implicit midpoint rule discretisation. Both (A7) and (A9) can be solved separately, since there is no coupling between them. The fields \mathbf{v}_h and E_h are then used to construct the right-hand side for the implicit linearized system in (A3)–(A4). Once the system is solved, the solution $(\Delta \mathbf{u}_h, \Delta D_h)$ is then used to update the iterative values of \mathbf{u}_h^{n+1} and D_h^{n+1} according to $(\mathbf{u}_h^{n+1}, D_h^{n+1}) \leftarrow$

$$15 \quad (\mathbf{u}_h^{n+1} + \Delta \mathbf{u}_h, D_h^{n+1} + \Delta D_h), \text{ having initially chosen } (\mathbf{u}_h^{n+1}, D_h^{n+1}) = (\mathbf{u}_h^n, D_h^n).$$

Author contributions. T. H. Gibson is the principal author and developer of the software presented in this paper and main author of the text. Authors L. Mitchell and D. A. Ham assisted and guided the software abstraction as a domain-specific language, and edited text. C. J. Cotter contributed to the formulation of the geophysical fluid dynamics and the design of the numerical experiments, and edited text.

Competing interests. D. A. Ham is an executive editor of the journal. The other authors declare they have no other competing interests.

20 *Acknowledgements.* This work was supported by the Engineering and Physical Sciences Research Council (grant numbers: EP/M011054/1, EP/L000407/1, and EP/L016613/1), and the Natural Environment Research Council (grant number: NE/K008951/1).



References

- Alnæs, M. S., Logg, A., Ølgaard, K. B., Rognes, M. E., and Wells, G. N.: Unified form language: A domain-specific language for weak formulations of partial differential equations, *ACM Transactions on Mathematical Software (TOMS)*, 40, 9, 2014.
- Arnold, D. N. and Brezzi, F.: Mixed and nonconforming finite element methods: implementation, postprocessing and error estimates, *ESAIM: Mathematical Modelling and Numerical Analysis*, 19, 7–32, 1985.
- 5 Arnold, D. N., Falk, R. S., and Winther, R.: Multigrid in $H(\text{div})$ and $H(\text{curl})$, *Numerische Mathematik*, 85, 197–217, <https://doi.org/10.1007/s002110000137>, 2000.
- Balay, S., Gropp, W. D., McInnes, L. C., and Smith, B. F.: Efficient management of parallelism in object-oriented numerical software libraries, in: *Modern software tools for scientific computing*, pp. 163–202, Springer, 1997.
- 10 Balay, S., Abhyankar, S., Adams, M., Brune, P., Buschelman, K., Dalcin, L., Gropp, W., Smith, B., Karpeyev, D., Kaushik, D., et al.: *Petsc users manual revision 3.7*, Tech. rep., Argonne National Lab.(ANL), Argonne, IL (United States), 2016.
- Bauer, W. and Cotter, C.: Energy-entropy conserving compatible finite element schemes for the rotating shallow water equations with slip boundary conditions, *Journal of Computational Physics*, <https://doi.org/10.1016/j.jcp.2018.06.071>, 2018.
- Bramble, J. H. and Xu, J.: A local post-processing technique for improving the accuracy in mixed finite-element approximations, *SIAM Journal on Numerical Analysis*, 26, 1267–1275, 1989.
- 15 Bramble, J. H., Pasciak, J. E., and Xu, J.: The analysis of multigrid algorithms for nonsymmetric and indefinite elliptic problems, *Mathematics of Computation*, 51, 389–414, 1988.
- Bramble, J. H., Kwak, D. Y., and Pasciak, J. E.: Uniform convergence of multigrid V-cycle iterations for indefinite and nonsymmetric problems, *SIAM journal on numerical analysis*, 31, 1746–1763, 1994.
- 20 Brezzi, F. and Fortin, M.: *Mixed and hybrid finite element methods*, vol. 15, Springer Science & Business Media, 2012.
- Brezzi, F., Douglas, J., and Marini, L. D.: Two families of mixed finite elements for second order elliptic problems, *Numerische Mathematik*, 47, 217–235, 1985.
- Brezzi, F., Douglas, J., Durán, R., and Fortin, M.: Mixed finite elements for second order elliptic problems in three variables, *Numerische Mathematik*, 51, 237–250, 1987.
- 25 Brown, J., Knepley, M. G., May, D. A., McInnes, L. C., and Smith, B.: Composable linear solvers for multiphysics, in: *Parallel and Distributed Computing (ISPDC), 2012 11th International Symposium on*, pp. 55–62, IEEE, 2012.
- Cockburn, B.: Static condensation, hybridization, and the devising of the HDG methods, in: *Building Bridges: Connections and Challenges in Modern Approaches to Numerical Partial Differential Equations*, pp. 129–177, Springer, 2016.
- Cockburn, B., Gopalakrishnan, J., and Lazarov, R.: Unified hybridization of discontinuous Galerkin, mixed, and continuous Galerkin methods for second order elliptic problems, *SIAM Journal on Numerical Analysis*, 47, 1319–1365, 2009a.
- 30 Cockburn, B., Guzmán, J., and Wang, H.: Superconvergent discontinuous Galerkin methods for second-order elliptic problems, *Mathematics of Computation*, 78, 1–24, 2009b.
- Cockburn, B., Gopalakrishnan, J., Li, F., Nguyen, N.-C., and Peraire, J.: Hybridization and postprocessing techniques for mixed eigenfunctions, *SIAM Journal on Numerical Analysis*, 48, 857–881, 2010a.
- 35 Cockburn, B., Gopalakrishnan, J., and Sayas, F.-J.: A projection-based error analysis of HDG methods, *Mathematics of Computation*, 79, 1351–1367, 2010b.



- Côté, J. and Staniforth, A.: A two-time-level semi-Lagrangian semi-implicit scheme for spectral models, *Monthly weather review*, 116, 2003–2012, 1988.
- Cotter, C. J. and Shipton, J.: Mixed finite elements for numerical weather prediction, *Journal of Computational Physics*, 231, 7076–7091, 2012.
- 5 Cotter, C. J. and Thuburn, J.: A finite element exterior calculus framework for the rotating shallow-water equations, *Journal of Computational Physics*, 257, 1506–1526, 2014.
- Cullen, M.: Alternative implementations of the semi-Lagrangian semi-implicit schemes in the ECMWF model, *Quarterly Journal of the Royal Meteorological Society*, 127, 2787–2802, 2001.
- Dalcin, L. D., Paz, R. R., Kler, P. A., and Cosimo, A.: Parallel distributed computing using python, *Advances in Water Resources*, 34, 1124–1139, 2011.
- 10 Devloo, P., Faria, C., Farias, A., Gomes, S., Loula, A., and Malta, S.: On continuous, discontinuous, mixed, and primal hybrid finite element methods for second-order elliptic problems, *International Journal for Numerical Methods in Engineering*, 115, 1083–1107, <https://doi.org/10.1002/nme.5836>, 2018.
- Elman, H. C., Ernst, O. G., and O’leary, D. P.: A multigrid method enhanced by Krylov subspace iteration for discrete Helmholtz equations, *SIAM Journal on scientific computing*, 23, 1291–1315, 2001.
- 15 Falgout, R. D., Jones, J. E., and Yang, U. M.: The design and implementation of hypre, a library of parallel high performance preconditioners, in: *Numerical solution of partial differential equations on parallel computers*, pp. 267–294, Springer, 2006.
- Guennebaud, G., Jacob, B., Lenz, M., et al.: Eigen v3, 2010, URL <http://eigen.tuxfamily.org>, 2015.
- Guyan, R. J.: Reduction of stiffness and mass matrices, *AIAA journal*, 3, 380, 1965.
- 20 Hecht, F.: New development in FreeFem++, *Journal of numerical mathematics*, 20, 251–266, 2012.
- Hiptmair, R. and Xu, J.: Nodal auxiliary space preconditioning in $H(\text{curl})$ and $H(\text{div})$ spaces, *SIAM Journal on Numerical Analysis*, 45, 2483–2509, <https://doi.org/10.1137/060660588>, 2007.
- Homolya, M., Mitchell, L., Luporini, F., and Ham, D. A.: TSFC: a structure-preserving form compiler, *SIAM Journal on Scientific Computing*, 40, C401–C428, 2018.
- 25 Irons, B.: Structural eigenvalue problems-elimination of unwanted variables, *AIAA journal*, 3, 961–962, 1965.
- Kirby, R. C. and Mitchell, L.: Solver composition across the PDE/linear algebra barrier, *SIAM Journal on Scientific Computing*, 40, C76–C98, <https://doi.org/10.1137/17M1133208>, 2018.
- Kirby, R. M., Sherwin, S. J., and Cockburn, B.: To CG or to HDG: a comparative study, *Journal of Scientific Computing*, 51, 183–212, 2012.
- Logg, A., Mardal, K.-A., and Wells, G.: Automated solution of differential equations by the finite element method: The FEniCS book, vol. 84, 30 Springer Science & Business Media, 2012.
- Long, K., Kirby, R., and van Bloemen Waanders, B.: Unified embedded parallel finite element computations via software-based Fréchet differentiation, *SIAM Journal on Scientific Computing*, 32, 3323–3351, 2010.
- Mandel, J.: Multigrid convergence for nonsymmetric, indefinite variational problems and one smoothing step, *Applied Mathematics and Computation*, 19, 201–216, 1986.
- 35 McRae, A. T. T., Bercea, G.-T., Mitchell, L., Ham, D. A., and Cotter, C. J.: Automated generation and symbolic manipulation of tensor product finite elements, *SIAM Journal on Scientific Computing*, 38, S25–S47, 2016.



- Melvin, T., Dubal, M., Wood, N., Staniforth, A., and Zerroukat, M.: An inherently mass-conserving iterative semi-implicit semi-Lagrangian discretization of the non-hydrostatic vertical-slice equations, *Quarterly Journal of the Royal Meteorological Society: A journal of the atmospheric sciences, applied meteorology and physical oceanography*, 136, 799–814, 2010.
- Melvin, T., Benacchio, T., Shipway, B., Wood, N., Thuburn, J., and Cotter, C.: A mixed finite-element, finite-volume, semi-implicit discretization for atmospheric dynamics: Cartesian geometry, 2018.
- Mitchell, L. and Müller, E. H.: High level implementation of geometric multigrid solvers for finite element problems: Applications in atmospheric modelling, *Journal of Computational Physics*, 327, 1–18, 2016.
- Natale, A. and Cotter, C. J.: A variational H (div) finite-element discretization approach for perfect incompressible fluids, *IMA J. Numer. Anal.*, p. drx033, <https://doi.org/10.1093/imanum/drx033>, 2017.
- 10 Natale, A., Shipton, J., and Cotter, C. J.: Compatible finite element spaces for geophysical fluid dynamics, *Dynamics and Statistics of the Climate System*, 1, dzw005, 2016.
- Nechaev, D. and Yaremchuk, M.: On the approximation of the Coriolis terms in C-grid models, *Monthly weather review*, 132, 2283–2289, 2004.
- Nédélec, J.-C.: Mixed finite elements in \mathbb{R}^3 , *Numerische Mathematik*, 35, 315–341, 1980.
- 15 Prud’Homme, C., Chabannes, V., Doyeux, V., Ismail, M., Samake, A., and Pena, G.: Feel++: A computational framework for galerkin methods and advanced numerical methods, in: *ESAIM: Proceedings*, vol. 38, pp. 429–455, EDP Sciences, 2012.
- Rathgeber, F., Markall, G. R., Mitchell, L., Lorient, N., Ham, D. A., Bertolli, C., and Kelly, P. H. J.: PyOP2: A high-level framework for performance-portable simulations on unstructured meshes, in: *High Performance Computing, Networking, Storage and Analysis (SCC), 2012 SC Companion*., pp. 1116–1123, IEEE, 2012.
- 20 Rathgeber, F., Ham, D. A., Mitchell, L., Lange, M., Luporini, F., McRae, A. T. T., Bercea, G.-T., Markall, G. R., and Kelly, P. H. J.: Firedrake: automating the finite element method by composing abstractions, *ACM Transactions on Mathematical Software (TOMS)*, 43, 24, 2016.
- Raviart, P.-A. and Thomas, J.-M.: A mixed finite element method for 2-nd order elliptic problems, in: *Mathematical aspects of finite element methods*, pp. 292–315, Springer, 1977.
- Shipton, J., Gibson, T., and Cotter, C.: Higher-order compatible finite element schemes for the nonlinear rotating shallow water equations on the sphere, *Journal of Computational Physics*, 375, 1121–1137, 2018.
- 25 Skamarock, W. C. and Klemp, J. B.: Efficiency and accuracy of the Klemp-Wilhelmson time-splitting technique, *Monthly Weather Review*, 122, 2623–2630, 1994.
- Stenberg, R.: Postprocessing schemes for some mixed finite elements, *ESAIM: Mathematical Modelling and Numerical Analysis*, 25, 151–167, 1991.
- 30 Temperton, C.: Treatment of the Coriolis terms in semi-Lagrangian spectral models, *Atmosphere-Ocean*, 35, 293–302, 1997.
- Thomas, S. J., Hacker, J. P., Smolarkiewicz, P. K., and Stull, R. B.: Spectral preconditioners for nonhydrostatic atmospheric models, *Monthly Weather Review*, 131, 2464–2478, 2003.
- Williamson, D. L., Drake, J. B., Hack, J. J., Jakob, R., and Swarztrauber, P. N.: A standard test set for numerical approximations to the shallow water equations in spherical geometry, *Journal of Computational Physics*, 102, 211–224, 1992.
- 35 Wood, N., Staniforth, A., White, A., Allen, T., Diamantakis, M., Gross, M., Melvin, T., Smith, C., Vosper, S., Zerroukat, M., et al.: An inherently mass-conserving semi-implicit semi-Lagrangian discretization of the deep-atmosphere global non-hydrostatic equations, *Quarterly Journal of the Royal Meteorological Society*, 140, 1505–1520, 2014.



Yakovlev, S., Moxey, D., Kirby, R. M., and Sherwin, S. J.: To CG or to HDG: a comparative study in 3D, *Journal of Scientific Computing*, 67, 192–220, 2016.

Zenodo/Firedrake: Software used in 'Slate: extending Firedrake's domain-specific abstraction to hybridized solvers for geoscience and beyond', <https://doi.org/10.5281/zenodo.2587072>, 2019.

5 Zenodo/Tabula-Rasa: Tabula Rasa: experimentation framework for hybridization and static condensation, <https://doi.org/10.5281/zenodo.2616031>, 2019.