



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

Thomas H. Gibson<sup>1</sup>, Lawrence Mitchell<sup>2</sup>, David A. Ham<sup>1</sup>, and Colin J. Cotter<sup>1</sup>

<sup>1</sup>Department of Mathematics, Imperial College London, London, SW7 2AZ, UK

<sup>2</sup>Department of Computer Science, Durham University, Durham, DH1 3LE, UK

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

Received: 1 April 2019 – Discussion started: 26 April 2019

Revised: 20 November 2019 – Accepted: 25 November 2019 – Published: 25 February 2020

**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 ongoing 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 a 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 the Portable, Extensible Toolkit for Scientific Computation (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., 2012a), 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 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 (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 PDEs, the hybridization technique provides a mechanism for enabling the static condensation of more complex linear systems. First introduced by Fraeijs de Veubeke (1965) and analyzed further by Brezzi and Fortin (1991), Cockburn et al. (2009a), and Boffi et al. (2013), the hybridization method 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, which exploit the approximation properties of the Lagrange multipliers. This enables the manufacturing of fields exhibiting superconvergent phenomena or enhanced conservation properties (Arnold and Brezzi, 1985; Brezzi et al., 1985; Bramble and Xu, 1989; Stenberg, 1991; Cockburn et al., 2009b, 2010b). 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 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., 2019). Our work is implemented in the Firedrake finite element library and the PETSc solver library (Balay et al., 1997,

2019), 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 Sect. 1.1. The embedded DSL, called ‘‘Slate’’, is introduced in Sect. 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 Sect. 3, including a discussion on post-processing. We then outline in Sect. 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 static condensation. We first demonstrate our implementation on a manufactured problem derived from a second-order elliptic equation, starting in Sect. 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 Sect. 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 Sect. 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  $\Omega$  is denoted by  $\mathcal{E}_h^\circ := \partial\mathcal{T}_h \setminus \partial\Omega$ . 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  $dx$  and  $ds$  denote appropriate integration measures. The operation  $\cdot$  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  $-$  denote 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., 2012a), 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 establishing our notation for a general finite element form following the convention of Alnæs et al. (2014). We define a real-valued *multi-linear form* as an operator which maps a list of *arguments*  $\mathbf{v} = (v_0, \dots, v_{\alpha-1}) \in V_0 \times \dots \times V_{\alpha-1}$  into  $\mathbb{R}$ :

$$a : V_0 \times \dots \times V_{\alpha-1} \rightarrow \mathbb{R}, \quad a \mapsto a(v_0, \dots, v_{\alpha-1}) = a(\mathbf{v}), \quad (4)$$

where  $a$  is linear in each argument  $v_k$ . The *arity* of a form is  $\alpha$ , an integer denoting the total number of form arguments. In traditional finite element nomenclature (for  $\alpha \leq 2$ ),  $V_0$  is referred to as the space of *test functions* and  $V_1$  as the space of *trial functions*. Each  $V_k$  are referred to as *argument spaces*. Forms with arity  $\alpha = 0, 1$ , or  $2$  are best interpreted as the more familiar mathematical objects: scalars (0-forms), linear forms or functionals (1-forms), and bilinear forms (2-forms), respectively.

If a given form  $a$  is parameterized by one or more *coefficients*, say  $\mathbf{c} = (c_0, \dots, c_q) \in C_0 \times \dots \times C_q$ , where  $\{C_k\}_{k=0}^q$  are *coefficient spaces*, then we write

$$a : C_0 \times \dots \times C_q \times V_0 \times \dots \times V_{\alpha-1} \rightarrow \mathbb{R}, \quad a \mapsto a(c_0, \dots, c_q; v_0, \dots, v_{\alpha-1}) = a(\mathbf{c}; \mathbf{v}). \quad (5)$$

From here on, we shall work exclusively with forms that are linear in  $\mathbf{v}$  and possibly nonlinear in the coefficients  $\mathbf{c}$ . This is reasonable since nonlinear methods based on Newton iterations produce linear problems via Gâteaux differentiation of

a nonlinear form corresponding to a PDE (also known as the *form Jacobian*). We refer the interested reader to Alnæs et al. (2014, Sect. 2.1.2) for more details. For clarity, we present examples of multi-linear forms of arity  $\alpha = 0, 1$ , and  $2$  that frequently appear in finite element discretizations:

$$a(\kappa; v, u) := (\nabla v, \kappa \nabla u)_{\mathcal{T}_h} \equiv \sum_{K \in \mathcal{T}_h} \int_K \nabla v \cdot (\kappa \nabla u) \, dx, \quad \kappa \in C_0, u \in V_1, v \in V_0, \alpha = 2, q = 1, \quad (6)$$

$$a(f; v) := (v, f)_{\mathcal{T}_h} \equiv \sum_{K \in \mathcal{T}_h} \int_K v f \, dx, \quad f \in C_0, v \in V_0, \alpha = 1, q = 1, \quad (7)$$

$$a(f, g; ) := (f - g, f - g)_{\mathcal{T}_h} \equiv \sum_{K \in \mathcal{T}_h} \int_K |f - g|^2 \, dx, \quad g \in C_1, f \in C_0, \alpha = 0, q = 2, \quad (8)$$

$$a(\gamma, \sigma) := \langle \gamma, [[\sigma]] \rangle_{\partial \mathcal{T}_h} \equiv \sum_{e \in \mathcal{E}_h^\circ} \int_e \gamma [[\sigma]] \, ds + \sum_{e \in \mathcal{E}_h^\partial} \int_e \gamma \sigma \cdot \mathbf{n} \, ds, \quad \sigma \in V_1, \gamma \in V_0, \alpha = 2, q = 0. \quad (9)$$

In general, a finite element form will consist of integrals over various geometric domains: integration over cells  $\mathcal{T}_h$ , interior facets  $\mathcal{E}_h^\circ$ , and exterior facets  $\mathcal{E}_h^\partial$ . Therefore, we express a general multi-linear form in terms of integrals over each set of geometric entities:

$$a(\mathbf{c}; \mathbf{v}) = \sum_{K \in \mathcal{T}_h} \int_K \mathcal{I}_K^T(\mathbf{c}; \mathbf{v}) \, dx + \sum_{e \in \mathcal{E}_h^\circ} \int_e \mathcal{I}_e^{\mathcal{E}, \circ}(\mathbf{c}; \mathbf{v}) \, ds + \sum_{e \in \mathcal{E}_h^\partial} \int_e \mathcal{I}_e^{\mathcal{E}, \partial}(\mathbf{c}; \mathbf{v}) \, ds, \quad (10)$$

where  $\mathcal{I}_K^T$  denotes a cell integrand on  $K \in \mathcal{T}_h$ ,  $\mathcal{I}_e^{\mathcal{E}, \circ}$  is an integrand on the interior facet  $e \in \mathcal{E}_h^\circ$ , and  $\mathcal{I}_e^{\mathcal{E}, \partial}$  is an integrand defined on the exterior facet  $e \in \mathcal{E}_h^\partial$ . 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 interior facet integrands  $\mathcal{I}_e^{\mathcal{E}, \circ}(\mathbf{c}; \mathbf{v})$  can be decomposed into two independent parts on each interior facet  $e$ : one for the positive restriction (+) and the negative restriction (-). That is, for each  $e \in \mathcal{E}_h^\circ$ , we may write  $\mathcal{I}_e^{\mathcal{E}, \circ}(\mathbf{c}; \mathbf{v}) = \mathcal{I}_{e^+}^{\mathcal{E}, \circ}(\mathbf{c}; \mathbf{v}) + \mathcal{I}_{e^-}^{\mathcal{E}, \circ}(\mathbf{c}; \mathbf{v})$ . This allows us to express the integral over an interior facet  $e$  connecting two adjacent elements, say  $K^+$  and  $K^-$ , as the sum

of integrals:

$$\begin{aligned} \int_{e \subset \partial K^+ \cup \partial K^-} \mathcal{I}_e^{\mathcal{E}, \circ}(\mathbf{c}; \mathbf{v}) \, ds &= \int_{e \subset \partial K^+} \mathcal{I}_{e^+}^{\mathcal{E}, \circ}(\mathbf{c}; \mathbf{v}) \, ds \\ &+ \int_{e \subset \partial K^-} \mathcal{I}_{e^-}^{\mathcal{E}, \circ}(\mathbf{c}; \mathbf{v}) \, ds. \end{aligned} \quad (11)$$

The local contribution of Eq. (10) in each cell  $K$ , along with its associated facets  $e \subset \partial K$ , is then

$$\begin{aligned} a_K(\mathbf{c}; \mathbf{v}) &= \int_K \mathcal{I}_K^{\mathcal{T}}(\mathbf{c}; \mathbf{v}) \, dx + \sum_{e \subset \partial K} \int_{\partial \Omega'_e} \mathcal{I}_e^{\mathcal{E}, \circ}(\mathbf{c}; \mathbf{v}) \, ds \\ &+ \sum_{e \subset \partial K \cap \partial \Omega'_e} \int \mathcal{I}_e^{\mathcal{E}, \partial}(\mathbf{c}; \mathbf{v}) \, ds. \end{aligned} \quad (12)$$

We call Eq. (12) the *cell-local* contribution of  $a(\mathbf{c}; \mathbf{v})$ , with

$$a(\mathbf{c}; \mathbf{v}) = \sum_{K \in \mathcal{T}_h} a_K(\mathbf{c}; \mathbf{v}). \quad (13)$$

To make matters concrete, let us suppose  $a(\mathbf{c}; \mathbf{v})$  is a bilinear form with arguments  $\mathbf{v} = (v_0, v_1) \in V_0 \times V_1$ . Now let  $\{\Phi_i\}_{i=1}^N$  and  $\{\Psi_j\}_{j=1}^M$  denote bases for  $V_0$  and  $V_1$ , respectively. Then the global  $N \times M$  matrix  $\mathbf{A}$  corresponding to  $a(\mathbf{c}; v_0, v_1)$  has its entries defined via

$$\begin{aligned} \mathbf{A}_{ij} &= a(\mathbf{c}; \Phi_i, \Psi_j) = \sum_{K \in \mathcal{T}_h} \mathbf{A}_{K,ij}, \mathbf{A}_{K,ij} \\ &= a_K(\mathbf{c}; \Phi_i, \Psi_j). \end{aligned} \quad (14)$$

By construction,  $\mathbf{A}_{K,ij} \neq \mathbf{0}$  if and only if  $\Phi_i$  and  $\Psi_j$  take non-zero values in  $K$ . Now we introduce the *cell-node map*  $i = e(K, \hat{i})$  as the mapping from the local node number  $\hat{i}$  in  $K$  to the global node number  $i$ . Suppose there are  $n$  and  $m$  nodes defining the degrees of freedom for  $V_0$  and  $V_1$ , respectively, in  $K$ . Then all non-zero entries of  $\mathbf{A}_{K,ij}$  arise from integrals involving basis functions with local indices corresponding to the global indices  $i, j$ :

$$\begin{aligned} \mathbf{A}_{\hat{i}\hat{j}}^K &:= a_K(\mathbf{c}; \Phi_{e(K, \hat{i})}, \Psi_{e(K, \hat{j})}), \hat{i} \in [1, \dots, n], \\ \hat{j} &\in [1, \dots, m]. \end{aligned} \quad (15)$$

These local contributions are collected in the  $n \times m$  dense matrix  $\mathbf{A}^K$ , which we call the *element tensor*. The global matrix  $\mathbf{A}$  is assembled from the collection of element tensors:  $\mathbf{A} \leftarrow \{\mathbf{A}^K\}_{K \in \mathcal{T}_h}$ . For details on the general evaluation of finite element basis functions and multi-linear forms, we refer the reader to Kirby (2004), Kirby and Logg (2006), Logg et al. (2012b), and Homolya et al. (2018). Further details on the global assembly of finite element operators, with a particular focus on code generation, are summarized in the work of Logg and Wells (2010) and Markall et al. (2013).

In standard finite element software packages, the element tensor is mapped entry-wise into a global sparse array using

the cell-node map  $e(K, \cdot)$ . Within Firedrake, this operation is handled by PyOP2 (Rathgeber et al., 2012) and serves as the main user-facing abstraction for global finite element assembly. For many applications, one may want to produce a new global operator by algebraically manipulating different element tensors. This is relatively invasive in numerical code, as it requires bypassing direct operator assembly to produce the new tensor. This is precisely the scope of Slate.

Like UFL, Slate relies on the grammar of the host-language: Python. The entire Slate language is implemented as a Python module which defines its types (classes) and operations on said types. Together, this forms a high-level language for expressing dense linear algebra on element tensors. 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 (for example, coefficient vectors of a finite element function) and
2. expressions consisting of algebraic operations on terminal tensors.

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 algebraic operations on locally assembled arrays. We summarize all currently supported Slate abstractions here.

### 2.1.1 Terminal tensors

In Slate, one associates a tensor with data on a cell either by using a multi-linear form, or assembled coefficient data:

- `Tensor( $a(\mathbf{c}; \mathbf{v})$ )`  
associates a form, expressed in UFL, with its local element tensor

$$\mathbf{A}^K \leftarrow a_K(\mathbf{c}; \mathbf{v}), \text{ for all } K \in \mathcal{T}_h. \quad (16)$$

The form arity  $\alpha$  of  $a_K(\mathbf{c}; \mathbf{v})$  determines the *rank* of the corresponding `Tensor`; i.e., scalars, vectors, and matrices are produced from scalars, linear forms, and bilinear forms, respectively.<sup>1</sup> The *shape* of the element tensor is determined by both the number of arguments and total number of degrees of freedom local to the cell.

- `AssembledVector( $f$ )`,  
where  $f$  is some finite element function. The function  $f \in V$  is expressed in terms of the finite element basis

<sup>1</sup>Similarly to UFL, Slate is capable of abstractly representing arbitrary rank tensors. However, only rank  $\leq 2$  tensors are typically used in most finite element applications, and therefore we currently only generate code for those ranks.

of  $V$ :  $f(x) = \sum_{i=1}^N f_i \Phi_i(x)$ . The result is the local coefficient vector of  $f$  on  $K$ :

$$\mathbf{F}^K \leftarrow \left\{ f_{e(K, \hat{i})} \right\}_{\hat{i}=1}^n, \quad (17)$$

where  $e(K, \hat{i})$  is the local node numbering and  $n$  is the number of nodes local to the cell  $K$ .

### 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 the following.

- $\mathbf{A} + \mathbf{B}$ , the addition of two equal shaped tensors:  $\mathbf{A}^K + \mathbf{B}^K$ .
- $\mathbf{A} * \mathbf{B}$ , a contraction over the last index of  $\mathbf{A}$  and the first index of  $\mathbf{B}$ . This is the usual multiplicative operation on matrices, vectors, and scalars:  $\mathbf{A}^K \mathbf{B}^K$ .
- $-\mathbf{A}$ , the additive inverse (negation) of a tensor:  $-\mathbf{A}^K$ .
- $\mathbf{A} . \mathbf{T}$ , the transpose of a tensor:  $(\mathbf{A}^K)^T$ .
- $\mathbf{A} . \text{inv}$ , the inverse of a square tensor:  $(\mathbf{A}^K)^{-1}$ .
- $\mathbf{A} . \text{solve}(\mathbf{B}, \text{decomposition}=\text{'...'})$ , the result,  $\mathbf{X}^K$ , of solving a local linear system  $\mathbf{A}^K \mathbf{X}^K = \mathbf{B}^K$ , optionally specifying a factorization strategy.
- $\mathbf{A} . \text{blocks}[\text{indices}]$ , where  $\mathbf{A}$  is a tensor from a mixed finite element space. This allows for the extraction of sub-blocks, which are 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 element tensors have the form

$$\mathbf{A}^K = \begin{bmatrix} \mathbf{A}_{00}^K & \mathbf{A}_{01}^K & \dots & \mathbf{A}_{0m}^K \\ \mathbf{A}_{10}^K & \mathbf{A}_{11}^K & \dots & \mathbf{A}_{1m}^K \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{A}_{n0}^K & \mathbf{A}_{n1}^K & \dots & \mathbf{A}_{nm}^K \end{bmatrix}. \quad (18)$$

The associated submatrix of Eq. (18) with indices  $\mathbf{i} = (\mathbf{p}, \mathbf{q})$ ,  $\mathbf{p} = \{p_1, \dots, p_r\}$ ,  $\mathbf{q} = \{q_1, \dots, q_c\}$ , is

$$\mathbf{A}_{\mathbf{p}\mathbf{q}}^K = \begin{bmatrix} \mathbf{A}_{p_1 q_1}^K & \dots & \mathbf{A}_{p_1 q_c}^K \\ \vdots & \ddots & \vdots \\ \mathbf{A}_{p_r q_1}^K & \dots & \mathbf{A}_{p_r q_c}^K \end{bmatrix} = \mathbf{A}^K . \text{blocks}[\mathbf{p}, \mathbf{q}], \quad (19)$$

where  $\mathbf{p} \subseteq \{0, \dots, n\}$ ,  $\mathbf{q} \subseteq \{0, \dots, m\}$ .

Each `Tensor` object knows all the information about the underlying UFL form that defines it, such as form arguments, coefficients, and the underlying finite element space(s) it operates on. This information is propagated through as unary or binary transformations are applied. The unary and binary operations shown here provide the necessary algebraic framework for a large class of problems, some of which we present in this paper.

In Firedrake, Slate expressions are transformed into low-level code by a *linear algebra compiler*. The compiler interprets Slate expressions as a *syntax tree*, where the tree is visited to identify what local arrays need to be assembled and the sequence of array operations. At the time of this work, our compiler generates C++ code, using the templated library Eigen (Guennebaud et al., 2015) for dense linear algebra. The translation from Slate to C++ is fairly straightforward, as all operations supported by Slate have a representation in Eigen.

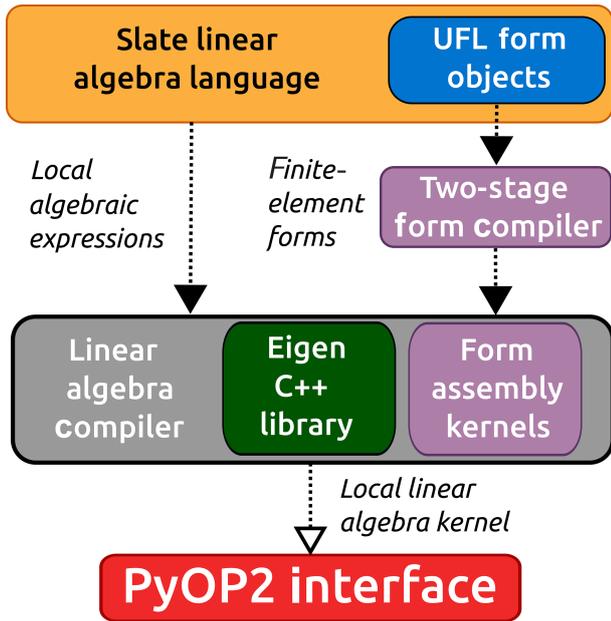
The compiler pass will generate a single “macro” kernel, which performs the dense linear algebra operations represented in Slate. The resulting code will also include (often multiple) function calls to local assembly kernels generated by TSFC (Two Stage Form Compiler) (Homolya et al., 2018) to assemble all necessary sub-blocks of an element tensor. All code generated by the linear algebra compiler conforms to the application programming interface (API) of the PyOP2 framework, as detailed by Rathgeber et al. (2012, Sect. 3). Figure 1 provides an illustration of the complete tool chain.

Most optimization of the resulting dense linear algebra code is handled directly by Eigen. In the case of unary and binary operations such as `A.inv` and `A.solve(B)`, stable default behaviors are applied by the linear algebra compiler. For example, `A.solve(B)` without a specified factorization strategy will default to using an in-place LU factorization with partial pivoting. For local matrices smaller than  $5 \times 5$ , the inverse is translated directly into Eigen’s `A.inverse()`, which employs stable analytic formulas. For larger matrices, the linear algebra replaces `A.inv` with an LU factorization.<sup>2</sup> Currently, we only support direct matrix factorizations for solving local linear systems. However, it would not be difficult to extend Slate to support more general solution techniques like iterative methods.

### 3 Examples

We now present 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 are chosen for clarity and to demonstrate key features of the Slate language. For our discussion, we use a model elliptic equation defined in a computational domain

<sup>2</sup>For more details on solving linear equations in Eigen, see [https://eigen.tuxfamily.org/dox/group\\_\\_TutorialLinearAlgebra.html](https://eigen.tuxfamily.org/dox/group__TutorialLinearAlgebra.html) (last access: 3 January 2020).



**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.

$\Omega$ . Consider the second-order PDE with both Dirichlet and Neumann boundary conditions:

$$-\nabla \cdot (\kappa \nabla p) + cp = f, \quad \text{in } \Omega, \quad (20)$$

$$p = p_0, \quad \text{on } \partial\Omega_D, \quad (21)$$

$$-\kappa \nabla p \cdot \mathbf{n} = g, \quad \text{on } \partial\Omega_N, \quad (22)$$

where  $\partial\Omega_D \cup \partial\Omega_N = \partial\Omega$  and  $\kappa, c : \Omega \rightarrow \mathbb{R}^+$  are positive-valued coefficients. To obtain a mixed formulation of Eqs. (20)–(22), we introduce the auxiliary velocity variable  $\mathbf{u} = -\kappa \nabla p$ . We then obtain the first-order system of PDEs:

$$\mu \mathbf{u} + \nabla p = 0, \quad \text{in } \Omega, \quad (23)$$

$$\nabla \cdot \mathbf{u} + cp = f, \quad \text{in } \Omega, \quad (24)$$

$$p = p_0, \quad \text{on } \partial\Omega_D, \quad (25)$$

$$\mathbf{u} \cdot \mathbf{n} = g, \quad \text{on } \partial\Omega_N, \quad (26)$$

where  $\mu = \kappa^{-1}$ .

### 3.1 Hybridization of mixed methods

To motivate our discussion in this section, we start by recalling the mixed method for Eqs. (23)–(26). Methods of this type seek approximations  $(\mathbf{u}_h, p_h)$  in finite-dimensional sub-

spaces  $U_h \times V_h \subset H(\text{div}; \Omega) \times L^2(\Omega)$ , defined by

$$U_h = \{\mathbf{w} \in H(\text{div}; \Omega) : \mathbf{w}|_K \in U(K), \forall K \in \mathcal{T}_h, \mathbf{w} \cdot \mathbf{n} = g \text{ on } \partial\Omega_N\}, \quad (27)$$

$$V_h = \{\phi \in L^2(\Omega) : \phi|_K \in V(K), \forall K \in \mathcal{T}_h\}. \quad (28)$$

The space  $U_h$  consists of  $H(\text{div})$ -conforming piecewise vector polynomials, where choices of  $U(K)$  typically include the Raviart–Thomas (RT), Brezzi–Douglas–Marini (BDM), or Brezzi–Douglas–Fortin–Marini (BDFM) elements (Raviart and Thomas, 1977; Nédélec, 1980; Brezzi et al., 1985, 1987). The space  $V_h$  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), and Shipton et al. (2018) for a discussion of mixed methods relevant for geophysical fluid dynamics. Two examples of such discretizations are presented in Sect. 5.2.

The mixed formulation of Eqs. (23)–(26) is arrived at by multiplying Eqs. (23)–(24) by test functions and integrating by parts. The resulting finite element problem reads as follows: find  $(\mathbf{u}_h, p_h) \in U_h \times V_h$  satisfying

$$(\mathbf{w}, \mu \mathbf{u}_h)_{\mathcal{T}_h} - (\nabla \cdot \mathbf{w}, p_h)_{\mathcal{T}_h} = -\langle \mathbf{w} \cdot \mathbf{n}, p_0 \rangle_{\partial\Omega_D}, \quad \forall \mathbf{w} \in U_{h,0}, \quad (29)$$

$$(\phi, \nabla \cdot \mathbf{u}_h)_{\mathcal{T}_h} + (\phi, cp_h)_{\mathcal{T}_h} = (\phi, f)_{\mathcal{T}_h}, \quad \forall \phi \in V_h, \quad (30)$$

where  $U_{h,0}$  is the subspace of  $U_h$  with functions 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:

$$\mathbf{u}_h = \sum_{i=1}^{N_u} U_i \Psi_i, \quad p_h = \sum_{i=1}^{N_p} P_i \xi_i, \quad (31)$$

where  $\{\Psi_i\}_{i=1}^{N_u}$  and  $\{\xi_i\}_{i=1}^{N_p}$  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 Eqs. (29)–(30) produces the following discrete saddle point system:

$$\begin{bmatrix} \mathbf{A} & -\mathbf{B}^T \\ \mathbf{B} & \mathbf{D} \end{bmatrix} \begin{Bmatrix} \mathbf{U} \\ \mathbf{P} \end{Bmatrix} = \begin{Bmatrix} \mathbf{F}_0 \\ \mathbf{F}_1 \end{Bmatrix}. \quad (32)$$

where  $\mathbf{U} = \{U_i\}_{i=1}^{N_u}$ ,  $\mathbf{P} = \{P_i\}_{i=1}^{N_p}$  are the coefficient vectors, and

$$\mathbf{A}_{ij} = (\Psi_i, \mu \Psi_j)_{\mathcal{T}_h}, \quad (33)$$

$$\mathbf{B}_{ij} = (\xi_i, \nabla \cdot \Psi_j)_{\mathcal{T}_h}, \quad (34)$$

$$\mathbf{D}_{ij} = (\xi_i, c \xi_j)_{\mathcal{T}_h}, \quad (35)$$

$$\mathbf{F}_{0,j} = -(\Psi_j \cdot \mathbf{n}, p_0)_{\partial\Omega_D}, \quad (36)$$

$$\mathbf{F}_{1,j} = (\xi_j, f)_{\mathcal{T}_h}. \quad (37)$$

Methods to efficiently invert such systems include  $H(\text{div})$ -multigrid (Arnold et al., 2000) (requiring complex overlapping Schwarz smoothers), global Schur complement factorizations (which require an approximation to the inverse of the *dense*<sup>3</sup> elliptic Schur complement  $\mathbf{D} + \mathbf{B}\mathbf{A}^{-1}\mathbf{B}^T$ ), 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, 1991; Boffi et al., 2013).

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

$$U_h^d = \left\{ \mathbf{w} \in [L^2(\Omega)]^n : \mathbf{w}|_K \in U(K), \forall K \in \mathcal{T}_h \right\}. \quad (38)$$

The vector finite element space  $U_h^d$  is a subspace of  $[L^2(\Omega)]^n$  consisting of local  $H(\text{div})$  functions, but normal components are no longer required to be 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:

$$M_h = \{ \gamma \in L^2(\partial\mathcal{T}_h) : \gamma|_e \in M(e), \forall e \in \partial\mathcal{T}_h \}, \quad (39)$$

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 or edges in three dimensions.

Deriving the hybridizable mixed system is accomplished through integration by parts over each element  $K$ . Testing with  $\mathbf{w} \in U_h^d(K)$  and integrating Eq. (23) over the cell  $K$  produces

$$\begin{aligned} & (\mathbf{w}, \mu \mathbf{u}_h^d)_K - (\nabla \cdot \mathbf{w}, p_h)_K + (\mathbf{w} \cdot \mathbf{n}, \lambda_h)_{\partial K} \\ & = -(\mathbf{w} \cdot \mathbf{n}, p_0)_{\partial K \cap \partial\Omega_D}. \end{aligned} \quad (40)$$

The trace function  $\lambda_h$  is introduced in the surface integral as an approximation to  $p|_{\partial K}$ . An additional constraint equation,

<sup>3</sup>The Schur complement, while elliptic, is globally dense due to the fact that  $\mathbf{A}$  has a dense inverse. This is a result of velocities in  $U_h$  having continuous normal components across cell interfaces.

called the *transmission condition*, is added to close the system. The resulting hybridizable formulation reads as follows: find  $(\mathbf{u}_h^d, p_h, \lambda_h) \in U_h^d \times V_h \times M_h$  such that

$$\begin{aligned} & (\mathbf{w}, \mu \mathbf{u}_h^d)_{\mathcal{T}_h} - (\nabla \cdot \mathbf{w}, p_h)_{\mathcal{T}_h} + \langle [[\mathbf{w}]], \lambda_h \rangle_{\partial\mathcal{T}_h \setminus \partial\Omega_D} \\ & = -(\mathbf{w} \cdot \mathbf{n}, p_0)_{\partial\Omega_D}, \forall \mathbf{w} \in U_h^d, \end{aligned} \quad (41)$$

$$(\phi, \nabla \cdot \mathbf{u}_h^d)_{\mathcal{T}_h} + (\phi, c p_h)_{\mathcal{T}_h} = (\phi, f)_{\mathcal{T}_h}, \forall \phi \in V_h, \quad (42)$$

$$\langle \gamma, [[\mathbf{u}_h^d]] \rangle_{\partial\mathcal{T}_h \setminus \partial\Omega_D} = \langle \gamma, g \rangle_{\partial\Omega_N}, \forall \gamma \in M_{h,0}, \quad (43)$$

where  $M_{h,0}$  denotes the space of traces vanishing on  $\partial\Omega_D$ . The transmission condition Eq. (43) enforces both the continuity of  $\mathbf{u}_h^d \cdot \mathbf{n}$  across element boundaries as well as the boundary condition  $\mathbf{u}_h^d \cdot \mathbf{n} = g$  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  $H(\text{div})$ -conforming counterpart. Specifically, the formulations in Eqs. (41)–(42) and (29)–(30) are solving equivalent problems if the normal components of  $\mathbf{w} \in U_h$  lie in the same polynomial space as the trace functions (Arnold and Brezzi, 1985).

The discrete matrix system arising from Eqs. (41)–(43) has the general form

$$\begin{bmatrix} \mathbf{A}_{00} & \mathbf{A}_{01} & \mathbf{A}_{02} \\ \mathbf{A}_{10} & \mathbf{A}_{11} & \mathbf{A}_{12} \\ \mathbf{A}_{20} & \mathbf{A}_{21} & \mathbf{A}_{22} \end{bmatrix} \begin{Bmatrix} \mathbf{U}^d \\ \mathbf{P} \\ \Lambda \end{Bmatrix} = \begin{Bmatrix} \mathbf{F}_0 \\ \mathbf{F}_1 \\ \mathbf{F}_2 \end{Bmatrix}, \quad (44)$$

where the discrete system is produced by expanding functions in terms of the finite element bases for  $U_h^d$ ,  $V_h$ , 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, Eq. (44) has a considerable advantage over Eq. (32) in the following ways.

1. Since both  $U_h^d$  and  $V_h$  are discontinuous spaces,  $\mathbf{U}^d$  and  $\mathbf{P}$  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,  $\Lambda$ :

$$\mathbf{S}\Lambda = \mathbf{E}, \quad (45)$$

where  $\mathbf{S} \leftarrow \{\mathbf{S}^K\}_{K \in \mathcal{T}_h}$  and  $\mathbf{E} \leftarrow \{\mathbf{E}^K\}_{K \in \mathcal{T}_h}$  are assembled via the local element tensors:

$$\mathbf{S}^K = \mathbf{A}_{22}^K - \begin{bmatrix} \mathbf{A}_{20}^K & \mathbf{A}_{21}^K \end{bmatrix} \begin{bmatrix} \mathbf{A}_{00}^K & \mathbf{A}_{01}^K \\ \mathbf{A}_{10}^K & \mathbf{A}_{11}^K \end{bmatrix}^{-1} \begin{bmatrix} \mathbf{A}_{02}^K \\ \mathbf{A}_{12}^K \end{bmatrix}, \quad (46)$$

$$\mathbf{E}^K = \mathbf{F}_2^K - \begin{bmatrix} \mathbf{A}_{20}^K & \mathbf{A}_{21}^K \end{bmatrix} \begin{bmatrix} \mathbf{A}_{00}^K & \mathbf{A}_{01}^K \\ \mathbf{A}_{10}^K & \mathbf{A}_{11}^K \end{bmatrix}^{-1} \begin{Bmatrix} \mathbf{F}_0^K \\ \mathbf{F}_1^K \end{Bmatrix}. \quad (47)$$

Note that the inverse of the block matrix in Eqs. (46) and (47) is *never* evaluated globally; the elimination can be performed locally by performing a sequence of Schur complement reductions within each cell.

2. The matrix  $\mathbf{S}$  is sparse, symmetric, positive-definite, and spectrally equivalent to the dense Schur complement  $\mathbf{D} + \mathbf{B}\mathbf{A}^{-1}\mathbf{B}^T$  from Eq. (32) of the original mixed formulation (Gopalakrishnan, 2003; Cockburn et al., 2009a).
3. Once  $\mathbf{\Lambda}$  is computed, both  $\mathbf{U}^d$  and  $\mathbf{P}$  can be recovered locally in each element. This can be accomplished in a number of ways. One way is to compute  $\mathbf{P}^K$  by solving

$$\begin{aligned} \left( \mathbf{A}_{11}^K - \mathbf{A}_{10}^K \left( \mathbf{A}_{00}^K \right)^{-1} \mathbf{A}_{01}^K \right) \mathbf{P}^K &= \mathbf{F}_1^K - \mathbf{A}_{10}^K \\ \left( \mathbf{A}_{00}^K \right)^{-1} \mathbf{F}_0^K - \left( \mathbf{A}_{12}^K - \mathbf{A}_{10}^K \left( \mathbf{A}_{00}^K \right)^{-1} \mathbf{A}_{02}^K \right) \mathbf{\Lambda}^K, \end{aligned} \quad (48)$$

followed by solving for  $(\mathbf{U}^d)^K$ :

$$\mathbf{A}_{00}^K (\mathbf{U}^d)^K = \mathbf{F}_0^K - \mathbf{A}_{01}^K \mathbf{P}^K - \mathbf{A}_{02}^K \mathbf{\Lambda}^K. \quad (49)$$

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  $\mathbf{U}^d$  and  $\mathbf{P}$ , respectively, in Sect. 3.3.

Figure 2 displays the corresponding Slate code for assembling the trace system, solving Eq. (45), and recovering the eliminated unknowns. For a complete reference on how to formulate the hybridized mixed system Eqs. (41)–(43) in UFL, we refer the reader to Alnæs et al. (2014). Complete Firedrake code using Slate to solve a hybridizable mixed system is also publicly available in Zenodo/Tabula-Rasa (2019, “Code verification”). We remark that, in the case of this hybridizable system, Eq. (44) contains zero-valued blocks which can simplify the resulting expressions in Eqs. (46)–(47) and (48)–(49). This is not true in general and therefore the expanded form using all sub-blocks of Eq. (44) is presented for completeness.

### 3.2 Hybridization of discontinuous Galerkin methods

The 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 into the definition of the local solver.

Deriving the LDG-H formulation follows exactly from standard DG methods. All prognostic variables are sought in the discontinuous spaces  $U_h \times V_h \subset [L^2(\Omega)]^n \times L^2(\Omega)$ . Within a cell  $K$ , integration by parts yields

$$\begin{aligned} (\mathbf{w}, \mu \mathbf{u}_h)_K - (\nabla \cdot \mathbf{w}, p_h)_K + \langle \mathbf{w} \cdot \mathbf{n}, \widehat{p} \rangle_{\partial K} &= 0, \\ \forall \mathbf{w} \in U(K), \end{aligned} \quad (50)$$

$$\begin{aligned} -(\nabla \phi, \mathbf{u}_h)_K + \langle \phi, \widehat{\mathbf{u}} \cdot \mathbf{n} \rangle_{\partial K} + (\phi, c p_h)_K &= (\phi, f)_K, \\ \forall \phi \in V(K), \end{aligned} \quad (51)$$

where  $U(K)$  and  $V(K)$  are vector and scalar polynomial spaces, respectively. Now, we define the numerical fluxes  $\widehat{p}$  and  $\widehat{\mathbf{u}}$  to be functions of the trial unknowns and a new independent unknown in the trace space  $M_h$ :

$$\widehat{\mathbf{u}}(\mathbf{u}_h, p_h, \lambda_h; \tau) = \mathbf{u}_h + \tau (p_h - \widehat{p}) \mathbf{n}, \quad (52)$$

$$\widehat{p}(\lambda_h) = \lambda_h, \quad (53)$$

where  $\lambda_h \in M_h$  is a function approximating  $p$  on  $\partial \mathcal{T}_h$  and  $\tau$  is a positive stabilization function that may vary on each facet  $e \in \partial \mathcal{T}_h$ . We further require that  $\lambda_h$  satisfies the Dirichlet condition for  $p$  on  $\partial \Omega_D$  in an  $L^2$ -projection sense. The full LDG-H formulation reads as follows. Find  $(\mathbf{u}_h, p_h, \lambda_h) \in U_h \times V_h \times M_h$  such that

$$\begin{aligned} (\mathbf{w}, \mu \mathbf{u}_h)_{\mathcal{T}_h} - (\nabla \cdot \mathbf{w}, p_h)_{\mathcal{T}_h} + \langle [[\mathbf{w}]], \lambda_h \rangle_{\partial \mathcal{T}_h} &= 0, \\ \forall \mathbf{w} \in U_h, \end{aligned} \quad (54)$$

$$\begin{aligned} -(\nabla \phi, \mathbf{u}_h)_{\mathcal{T}_h} + \langle \phi, [[\mathbf{u}_h + \tau (p_h - \lambda_h) \mathbf{n}]] \rangle_{\partial \mathcal{T}_h} \\ + (\phi, c p_h)_{\mathcal{T}_h} = (\phi, f)_{\mathcal{T}_h}, \forall \phi \in V_h, \end{aligned} \quad (55)$$

$$\begin{aligned} \langle \gamma, [[\mathbf{u}_h + \tau (p_h - \lambda_h) \mathbf{n}]] \rangle_{\partial \mathcal{T}_h \setminus \partial \Omega_D} = \langle \gamma, g \rangle_{\partial \Omega_N}, \\ \forall \gamma \in M_h, \end{aligned} \quad (56)$$

$$\langle \gamma, \lambda_h \rangle_{\partial \Omega_D} = \langle \gamma, p_0 \rangle_{\partial \Omega_D}, \forall \gamma \in M_h. \quad (57)$$

Equation (56) is the transmission condition, which enforces the continuity of  $\widehat{\mathbf{u}} \cdot \mathbf{n}$  on  $\partial \mathcal{T}_h$  and  $q$ . Equation (57) ensures  $\lambda_h$  satisfies the Dirichlet condition. This ensures that the numerical flux is single-valued on the facets. Hence, the LDG-H method defines a *conservative* DG method (Cockburn et al., 2010b). Note that the choice of  $\tau$  has a significant influence on the expected convergence rates of the computed solutions.

The LDG-H method retains the advantages of standard DG methods while also enabling the assembly of reduced linear systems through static condensation. The matrix system arising from Eqs. (54)–(57) has the same general form as the hybridized mixed method in Eq. (44), except all sub-blocks are

```

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  Sexp = A[2, 2] - A[2, :2] * A[:2, :2].inv * A[:2, 2] # Slate expression for  $S^K$ 
9  Eexp = F[2] - A[2, :2] * A[:2, :2].inv * F[:2] # Slate expression for  $E^K$ 
10 S = assemble(Sexp, bcs=[...]) # Assemble  $S$ 
11 E = assemble(Eexp) # Assemble  $E$ 
12 lambda_h = Function(M) # Function to store the result:  $\Lambda$ 
13
14 # Solve for the Lagrange multipliers:  $\Lambda$ 
15 solve(S, lambda_h, E, solver_parameters={"ksp_type": "preonly", "pc_type": "lu"})
16 p_h = Function(V) # Function to store the result:  $P$ 
17 u_h = Function(U) # Function to store the result:  $U^d$ 
18 Lambda = AssembledVector(lambda_h) # Local coefficient vector:  $\Lambda^K$ 
19 P = AssembledVector(p_h) # Local coefficient vector:  $P^K$ 
20
21 # Intermediate expressions
22 Sd = A[1, 1] - A[1, 0] * A[0, 0].inv * A[0, 1]
23 Sl = A[1, 2] - A[1, 0] * A[0, 0].inv * A[0, 2]
24
25 # Slate expressions for local recovery
26 p_sys = Sd.solve(F[1] - A[1, 0] * A[0, 0].inv * F[0] - Sl * Lambda,
27                 decomposition="PartialPivLu")
28 u_sys = A[0, 0].solve(F[0] - A[0, 1] * P - A[0, 2] * Lambda,
29                     decomposition="PartialPivLu")
30 assemble(p_sys, p_h) # Solve for  $P$ 
31 assemble(u_sys, u_h) # Solve for  $U^d$ 

```

**Figure 2.** Firedrake code for solving Eq. (44) via static condensation and local recovery, given UFL expressions  $a$ ,  $L$  for Eqs. (41)–(43). 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 Eqs. (46) and (47), respectively. Any vanishing conditions on the trace variables should be provided as boundary conditions during operator assembly (line 10). Lines 26 and 28 are expressions for Eqs. (48) and (49) (using LU). Code generation occurs in lines 10, 11, 30, and 31. 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.

now populated with non-zero entries due to the coupling of trace functions with both  $p_h$  and  $u_h$ . However, all previous properties of the discrete matrix system from Sect. 3.1 still apply. The Slate expressions for the local elimination and reconstruction operations will be identical to those illustrated in Fig. 2. 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) and Cockburn (2016).

### 3.3 Local post-processing

For both mixed (Arnold and Brezzi, 1985; Brezzi et al., 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 1 order of accuracy over the unprocessed solution). These methods can be expressed as local solves on each element and are straightforward to implement using Slate. In this section, 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 Sect. 3.1 and 3.2, using the standard set of operations on element tensors summarized in Sect. 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 hybridizable mixed and LDG-H methods. We proceed by posing the finite element systems cell-wise.

Let  $\mathcal{P}_k(K)$  denote a polynomial space of degree  $\leq k$  on a cell  $K \in \mathcal{T}_h$ . Then for a given pair of computed solutions  $u_h, p_h$  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:

$$(\nabla w, \nabla p_h^*)_K = -(\nabla w, \kappa^{-1} \mathbf{u}_h)_K, \forall w \in \mathcal{P}_{k+1}^{\perp, l}(K), \quad (58)$$

$$(v, p_h^*)_K = (v, p_h)_K, \forall v \in \mathcal{P}_l(K), \quad (59)$$

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  $\mathbf{u}_h$ , the approximation of  $-\kappa \nabla p$ . In practice, the space  $\mathcal{P}_{k+1}^{\perp, l}(K)$  may be constructed using an orthogonal hierarchical basis, and solving Eqs. (58)–(59) amounts to inverting a symmetric positive definite system in each cell of the mesh.

At the time of this work, Firedrake does not support the construction of such a finite element basis. However, we can 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

$$\begin{aligned} (\nabla w, \nabla p_h^*)_K + (w, \psi)_K &= -(\nabla w, \kappa^{-1} \mathbf{u}_h)_K, \\ \forall w \in \mathcal{P}_{k+1}(K), \end{aligned} \quad (60)$$

$$(\phi, p_h^*)_K = (\phi, p_h)_K, \forall \phi \in \mathcal{P}_l(K), \quad (61)$$

where  $0 \leq l \leq k$ . The local problems Eqs. (60)–(61) and Eqs. (58)–(59) are equivalent, with the Lagrange multiplier  $\psi$  enforcing orthogonality 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 (Stenberg, 1991; Cockburn et al., 2010a). 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., 2009b, 2010b). We demonstrate the increased accuracy in computed solutions in Sect. 5.1. An abridged example using Firedrake and Slate to solve the local linear systems is provided in Fig. 3.

### 3.3.2 Post-processing of the flux

Our second example illustrates a procedure that uses the numerical flux of an HDG discretization for Eqs. (23)–(26). Within the context of the LDG-H method, we can use the numerical trace in Eq. (52) to produce a vector field that is  $H(\text{div})$ -conforming. The technique we outline here follows that of Cockburn et al. (2009b).

Let  $\mathcal{T}_h$  be a mesh consisting of simplices. On each cell  $K \in \mathcal{T}_h$ , 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

$$(\mathbf{r}, \mathbf{u}_h^*)_K = (\mathbf{r}, \mathbf{u}_h)_K, \forall \mathbf{r} \in [\mathcal{P}_{k-1}(K)]^n, \quad (62)$$

$$\langle \mu, \mathbf{u}_h^* \cdot \mathbf{n} \rangle_e = \langle \mu, \widehat{\mathbf{u}} \cdot \mathbf{n} \rangle_e, \forall \mu \in \mathcal{P}_k(e) \quad (63)$$

for all facets  $e$  on  $\partial K$ , where  $\widehat{\mathbf{u}}$  is the numerical flux defined in Eq. (52). 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  $\tau$  producing a solvable system for Eqs. (54)–(57). However,
2.  $\mathbf{u}_h^* \in H(\text{div}; \Omega)$ . That is,  $[[[\mathbf{u}_h^*]]]_e = 0, \forall e \in \mathcal{E}_h^\circ$ .
3. Additionally, the divergence of  $\mathbf{u}_h^*$  converges at a rate of  $k+1$ .

The Firedrake implementation using Slate is similar to the scalar post-processing example (see Fig. 3); the cell-wise linear systems Eqs. (62)–(63) 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 Sect. 2.1. We use the very sensitive parameter dependency in the post-processing methods to validate our software implementation in Zenodo/Tabula-Rasa (2019, “Code-verification”).

## 4 Static condensation as a preconditioner

Slate enables static condensation approaches to be expressed very concisely. Nonetheless, the 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: PETSc. Firedrake uses PETSc as its main solver abstraction framework and can provide *operator-based preconditioners* for solving linear systems as PC objects expressed in Python via petsc4py (Dalcin et al., 2011). For a comprehensive overview on solving linear systems using PETSc, we refer the interested reader to Balay et al. (2019, Sect. 4).

Suppose we wish to solve a linear system:  $\mathbf{Ax} = \mathbf{b}$ . We can think of (left) preconditioning the system in residual form:

$$\mathbf{r} = \mathbf{r}(\mathbf{A}, \mathbf{b}) \equiv \mathbf{b} - \mathbf{Ax} = \mathbf{0} \quad (64)$$

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

$$\mathcal{P}\mathbf{r} = \mathcal{P}\mathbf{b} - \mathcal{P}\mathbf{Ax} = \mathbf{0}. \quad (65)$$

```

1 # Define spaces for the higher-order pressure approximation and Lagrange multipliers
2 DGk1 = FunctionSpace(mesh, "DG", degree + 1)
3 DGO = FunctionSpace(mesh, "DG", 0)
4 W = DGk1 * DGO
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 the computed pressure p_h and flux u_h in the 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 p_h^*
15 p_star = Function(DGk1, name="Post-processed scalar")
16 assemble(E.blocks[0], p_star) # Assemble only the first field (pressure)

```

**Figure 3.** Example of local post-processing using Firedrake and Slate. Here, we locally solve the mixed system defined in Eqs. (58)–(59). 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.

Given a current iterate  $\mathbf{x}_i$  the residual at the  $i$ th iteration is simply  $\mathbf{r}_i \equiv \mathbf{b} - \mathbf{A}\mathbf{x}_i$ , and  $\mathcal{P}$  acts on the residual to produce an approximation to the error  $\boldsymbol{\epsilon}_i \equiv \mathbf{x} - \mathbf{x}_i$ . If  $\mathcal{P}$  is an application of an exact inverse, the residual is converted into an exact (up to numerical roundoff) error.

We will denote the application of a particular Krylov subspace method (KSP) for the linear system Eq. (64) as  $\mathcal{K}_x(\mathbf{r}(\mathbf{A}, \mathbf{b}))$ . Upon preconditioning the system via  $\mathcal{P}$  as in Eq. (65), we write

$$\mathcal{K}_x(\mathcal{P}\mathbf{r}(\mathbf{A}, \mathbf{b})). \quad (66)$$

If Eq. (66) is solved directly via  $\mathcal{P} = \mathbf{A}^{-1}$ , then  $\mathcal{P}\mathbf{r}(\mathbf{A}, \mathbf{b}) = \mathbf{A}^{-1}\mathbf{b} - \mathbf{x}$ . So Eq. (66) then becomes  $\mathcal{K}_x(\mathbf{r}(\mathbf{I}, \mathbf{A}^{-1}\mathbf{b}))$ , producing the exact solution of Eq. (64) 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,  $\mathcal{P}$ .

#### 4.1 Interfacing with PETSc via custom preconditioners

The implementation of preconditioners for the systems considered in this paper requires the 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  $H(\text{div}) \times L^2$  mixed problems and a generic interface for statically condensing finite element 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 or nonlinear solver (Newton-based methods) at *runtime*.

##### 4.1.1 A static condensation interface for hybridization

As discussed in Sect. 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 cell-wise static condensation. To facilitate this, we provide a PETSc PC static condensation interface: `firedrake.SPC`. This preconditioner takes the discretized system as in Eq. (44) 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

$$\begin{bmatrix} \mathbf{A}_{e,e} & \mathbf{A}_{e,c} \\ \mathbf{A}_{c,e} & \mathbf{A}_{c,c} \end{bmatrix} \begin{Bmatrix} \mathbf{X}_e \\ \mathbf{X}_c \end{Bmatrix} = \begin{Bmatrix} \mathbf{R}_e \\ \mathbf{R}_c \end{Bmatrix}, \quad (67)$$

where  $\mathbf{X}_e$  is the vector of unknowns to be eliminated,  $\mathbf{X}_c$  is the vector of unknowns for the condensed field, and  $\mathbf{R}_e$  and  $\mathbf{R}_c$  are the incoming right-hand sides. The partitioning in Eq. (67) is determined by the solver 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 `firedrake.SPC` to cell-wise eliminate fields 0 and 1; the resulting condensed system is associated with field 2.

The `firedrake.SPC` preconditioner can be interpreted as a Schur complement method for Eq. (67) of the

form

$$\mathcal{P} = \begin{bmatrix} \mathbf{I} & -\mathbf{A}_{e,e}^{-1}\mathbf{A}_{e,c} \\ \mathbf{0} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{A}_{e,e}^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{S}^{-1} \end{bmatrix} \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ -\mathbf{A}_{c,e}\mathbf{A}_{e,e}^{-1} & \mathbf{I} \end{bmatrix}, \quad (68)$$

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

In practice, the only globally coupled system requiring iterative inversion is  $\mathbf{S}$ :

$$\mathcal{K}_{\mathbf{X}_c}(\mathcal{P}_1 r(\mathbf{S}, \mathbf{R}_s)), \quad (69)$$

where  $\mathbf{R}_s = \mathbf{R}_c - \mathbf{A}_{c,e}\mathbf{A}_{e,e}^{-1}\mathbf{R}_e$  is the condensed right-hand side and  $\mathcal{P}_1$  is a preconditioner for  $\mathbf{S}$ . Once  $\mathbf{X}_c$  is computed,  $\mathbf{X}_e$  is reconstructed by inverting the system  $\mathbf{X}_e = \mathbf{A}_{e,e}^{-1}(\mathbf{R}_c - \mathbf{A}_{e,c}\mathbf{X}_c)$  cell-wise.

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  $\mathbf{S}$  via the appropriate options prefix (`condensed_field`). The resulting KSP for Eq. (69) 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.

#### 4.1.2 Preconditioning mixed methods via hybridization

The preconditioner `firedrake.HybridizationPC` expands on the previous one, this time taking an  $H(\text{div}) \times L^2$  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 Sect. 4.1.1.

More precisely, let  $\mathbf{A}\mathbf{X} = \mathbf{R}$  be the incoming mixed saddle point problem, where  $\mathbf{R} = \mathbf{R}_U \ \mathbf{R}_P^T$ ,  $\mathbf{X} = \mathbf{U} \ \mathbf{P}^T$ , and  $\mathbf{U}$  and  $\mathbf{P}$  are the velocity and scalar unknowns, respectively. Then this preconditioner replaces  $\mathbf{A}\mathbf{X} = \mathbf{R}$  with the extended problem:

$$\begin{bmatrix} \widehat{\mathbf{A}} & \mathbf{C}^T \\ \mathbf{C} & \mathbf{0} \end{bmatrix} \begin{Bmatrix} \widehat{\mathbf{X}} \\ \boldsymbol{\Lambda} \end{Bmatrix} = \begin{Bmatrix} \widehat{\mathbf{R}} \\ \mathbf{R}_g \end{Bmatrix} \quad (70)$$

where  $\boldsymbol{\Lambda}$  are the Lagrange multipliers,  $\widehat{\mathbf{R}} = \widehat{\mathbf{R}}_U \ \mathbf{R}_P^T$ ,  $\widehat{\mathbf{R}}_U$  and  $\mathbf{R}_P$  are the right-hand sides for the flux and scalar equations, respectively, and  $\widehat{\cdot}$  indicates modified matrices and co-vectors with discontinuous functions. Here,  $\widehat{\mathbf{X}} = \mathbf{U}^d \ \mathbf{P}^T$

are the hybridizable (discontinuous) unknowns to be determined and  $\mathbf{C}\widehat{\mathbf{X}} = \mathbf{R}_g$  is the matrix representation of the transmission condition for the hybridizable mixed method (see Eq. 43).

The application of `firedrake.HybridizationPC` can be interpreted as the Schur complement reduction of Eq. (70):

$$\widehat{\mathcal{P}} = \begin{bmatrix} \mathbf{I} & -\widehat{\mathbf{A}}^{-1}\mathbf{C}^T \\ \mathbf{0} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \widehat{\mathbf{A}}^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{S}^{-1} \end{bmatrix} \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ -\mathbf{C}\widehat{\mathbf{A}}^{-1} & \mathbf{I} \end{bmatrix}, \quad (71)$$

where  $\mathbf{S}$  is the Schur complement matrix  $\mathbf{S} = -\mathbf{C}\widehat{\mathbf{A}}^{-1}\mathbf{C}^T$ . As before, a single global system for  $\mathbf{A}$  can be assembled cell-wise using Slate-generated kernels. Configuring the solver for inverting  $\mathbf{S}$  is done via the PETSc options prefix: `-hybridization`. The recovery of  $\mathbf{U}^d$  and  $\mathbf{P}$  happens in the same manner as `firedrake.SPCPC`.

Since the hybridizable flux solution is constructed in the broken  $H(\text{div})$  space  $U_h^d$ , we must project the computed solution into  $U_h \subset H(\text{div})$ . This can be done cheaply via local facet averaging. The resulting solution is then updated via  $\mathbf{U} \leftarrow \Pi_{\text{div}}\mathbf{U}^d$ , where  $\Pi_{\text{div}} : U_h^d \rightarrow U_h$  is a projection operator. This ensures that the residual for the original mixed problem is properly evaluated to test for solver convergence. With  $\widehat{\mathcal{P}}$  as in Eq. (71), the preconditioning operator for the original system  $\mathbf{A}\mathbf{X} = \mathbf{R}$  then has the form

$$\mathcal{P} = \boldsymbol{\Pi}\widehat{\mathcal{P}}\boldsymbol{\Pi}^T, \quad \boldsymbol{\Pi} = \begin{bmatrix} \Pi_{\text{div}} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} & \mathbf{0} \end{bmatrix}. \quad (72)$$

We note here that assembly of the right-hand side for the  $\mathbf{A}$  system requires special attention. Firstly, when Neumann conditions are present, then  $\mathbf{R}_g$  is not necessarily  $\mathbf{0}$ . 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 into the generated Slate expressions. A more subtle issue that requires extra care is the incoming right-hand side tested in the  $H(\text{div})$  space  $U_h$ .

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

$$\widehat{\mathbf{R}}_U(\mathbf{w}) = \mathbf{R}_U(\mathbf{w}). \quad (73)$$

We can construct such an  $\widehat{\mathbf{R}}_U$  satisfying Eq. (73) in the following way. By construction, we have for each basis function  $\boldsymbol{\Psi}_i \in U_h$

$$\boldsymbol{\Psi}_i = \begin{cases} \boldsymbol{\Psi}_i^d & \boldsymbol{\Psi}_i \text{ associated with an exterior,} \\ & \text{facet node} \\ \boldsymbol{\Psi}_i^{d,+} + \boldsymbol{\Psi}_i^{d,-} & \boldsymbol{\Psi}_i \text{ associated with an interior,} \\ & \text{facet node} \\ \boldsymbol{\Psi}_i^d & \boldsymbol{\Psi}_i \text{ associated with a cell,} \\ & \text{interior node} \end{cases},$$

(74) mixed problem:

where  $\Psi_i^d, \Psi_i^{d,\pm} \in U_h^d$ , and  $\Psi_i^{d,\pm}$  are basis functions corresponding to the positive and negative restrictions associated with the  $i$ th facet node.<sup>4</sup> We then define our broken right-hand side via the local definition:

$$\widehat{\mathbf{R}}_U(\Psi_i^d) = \frac{\mathbf{R}_U(\Psi_i)}{N_i}, \quad (75)$$

where  $N_i$  is the number of cells that the degree of freedom corresponding to  $\Psi_i \in U_h$  is topologically associated with. Using Eq. (74), Eq. (75), and the fact that  $\mathbf{R}_U$  is linear in its argument, we can verify that our construction of  $\widehat{\mathbf{R}}_U$  satisfies Eq. (73).

### 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 Sect. 4, Slate is accessed indirectly and requires no manually written solver code for hybridization or static condensation or local recovery. All parallel results were obtained on a single fully loaded compute node of dual-socket Intel E5-2630v4 (Xeon) processors with  $2 \times 10$  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.

The 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 Sect. 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.

#### 5.1 HDG 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):

$$-\nabla \cdot \nabla p + p = f, \quad \text{in } \Omega = [0, 1]^3, \quad (76)$$

$$p = g, \quad \text{on } \partial\Omega, \quad (77)$$

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 Eqs. (76)–(77) as the

<sup>4</sup>These are the two broken parts of  $\Psi_i$  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.

$$\mathbf{u} + \nabla p = 0, \quad (78)$$

$$\nabla \cdot \mathbf{u} + p = f, \quad (79)$$

$$p = g, \quad \text{on } \partial\Omega. \quad (80)$$

We start with linear polynomial approximations, up to cubic, for the LDG-H discretization of Eqs. (78)–(80). 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).

A continuous Galerkin (CG) discretization of the primal problem Eqs. (76)–(77) 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 orders 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.

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.

To invert the CG system, we use a conjugate gradient solver with Hypre’s BoomerAMG implementation of algebraic multigrid (AMG) as a preconditioner (Falgout et al., 2006). For the HDG method, we use the preconditioner described in Sect. 4.1.1 and the same solver setup as the CG method for the trace system. While the trace operator is indeed symmetric and positive-definite, one should keep in mind that conclusions regarding the performance of off-the-shelf AMG packages on the HDG trace system are still relatively unclear. As a result, efforts on developing more efficient multigrid strategies are a topic of ongoing interest (Cockburn et al., 2014; Kronbichler and Wall, 2018).

To avoid over-solving, we iterate to a relative tolerance such that the *discretization error* is minimal for a given mesh. In other words, the solvers are configured to terminate when there is no further reduction in the  $L^2$  error of the computed solution compared with the analytic solution. This means we are not iterating to a fixed solver tolerance across all mesh resolutions. Therefore, we can expect the total number of Krylov iterations (for both the CG and HDG methods) to increase as the mesh resolution becomes finer. The rationale behind this approach is to directly compare the execution time to solve for the best possible approximation to the solution given a fixed resolution.

### 5.1.1 Error versus execution time

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 all setup costs, 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 and Krylov iterations to reach discretization error for each mesh are summarized in Fig. 4.

The HDG method of order  $k - 1$  ( $\text{HDG}_{k-1}$ ) with post-processing, as expected, produces a solution which is as accurate as the CG method of order  $k$  ( $\text{CG}_k$ ). While the full HDG system is never explicitly assembled, the larger execution time is a result of several factors. The primary factor is that the total number of trace unknowns for the  $\text{HDG}_1$ ,  $\text{HDG}_2$ , and  $\text{HDG}_3$  discretizations is roughly 4, 3, and 2 times larger, respectively, than the corresponding number of CG unknowns. Therefore, each iteration is more expensive. We also observe that the trace system requires more Krylov iterations to reach discretization error, which appears to improve relative to the CG method as the approximation order increases. Further analysis on a multigrid methods for HDG systems is required to draw further conclusions. The main computational bottleneck in HDG methods is the global linear solver. We therefore expect our implementation to be dominated by the cost associated with inverting the trace operator. If one considers just the time-to-solution, the CG method is clearly ahead of the HDG method. However, the superior scaling, local conservation, and stabilization properties of the HDG method make it a particularly appealing choice for fluid dynamics applications (Yakovlev et al., 2016; Kronbichler and Wall, 2018). Therefore, the development of good preconditioning strategies for the HDG method is critical for its competitive use.

### 5.1.2 Breakdown of solver time

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 significantly 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.

Figure 5 displays a breakdown of total 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 more expensive than the CG method. This is expected given the larger degree-of-freedom count and expensive global solve. The raw numerical breakdown of the HDG and CG solvers are shown in Table 1. We isolate each component of the HDG method contributing to the total execution time. Local op-

erations 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 scalar solution. For all  $k$ , our HDG implementation is solver-dominated as expected.

Both trace operator and right-hand-side assembly are dominated by the costs of inverting a local square mixed matrix coupling the scalar and velocity unknowns, which is performed directly via an LU factorization. This is also the case for backwards substitution. They should all therefore be of the same magnitude in time spent. We observe that this is the case across all degrees, with times ranging between approximately 6 % and 11 % of total execution time for assembling the condensed system. Back-substitution takes roughly the same time as the static condensation and forward elimination stages (approximately 12 % of execution time on average). Finally, the additional cost of post-processing accrues negligible time (roughly 2 % of execution time across all degrees). This is a small cost for an increase in order of accuracy.

We note that the 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 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 Hybridizable mixed methods for the shallow water equations

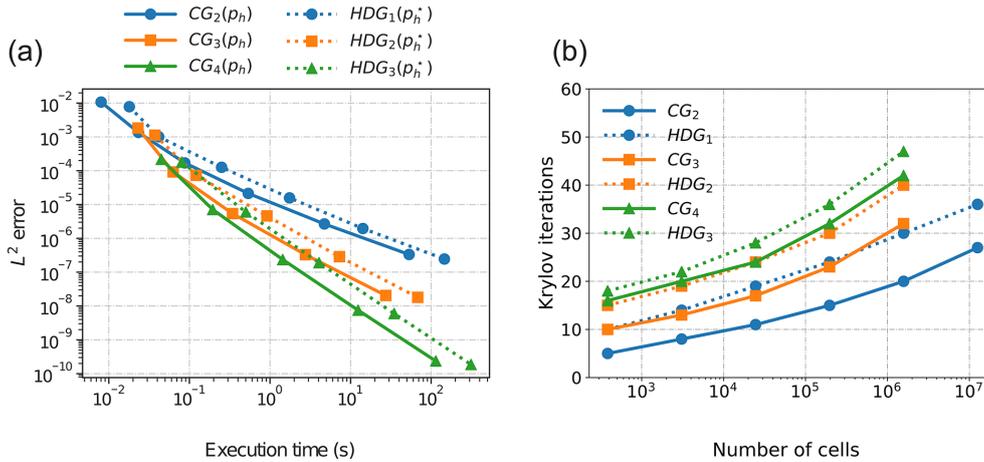
A primary motivator for our interest in hybridizable methods revolves around developing efficient solvers for problems in geophysical flows. In this section, we present some results integrating the nonlinear, rotating shallow water equations on the sphere using test case 5 (flow past an isolated mountain)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

The model equations we consider are the vector-invariant rotating nonlinear shallow water system defined on a two-dimensional spherical surface  $\Omega$  embedded in  $\mathbb{R}^3$ :

$$\frac{\partial \mathbf{u}}{\partial t} + \left( \nabla^\perp \cdot \mathbf{u} + f \right) \mathbf{u}^\perp + \nabla \left( g(D + b) + \frac{1}{2} |\mathbf{u}|^2 \right) = 0, \quad (81)$$

$$\frac{\partial D}{\partial t} + \nabla \cdot (\mathbf{u}D) = 0, \quad (82)$$

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  $\Omega$ . After discretizing in time and space



**Figure 4.** Comparison of continuous Galerkin and LDG-H solvers for the model three-dimensional positive-definite Helmholtz equation. Panel (a): a log–log plot showing the error against execution time for the CG and HDG with post-processing ( $\tau = 1$ ) methods. Panel (b): a log–linear plot showing Krylov iterations of the AMG-preconditioned conjugate gradient algorithm (to reach discretization error) against number of cells.

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

Stage	HDG <sub>1</sub>		HDG <sub>2</sub>		HDG <sub>3</sub>	
	$t_{\text{stage}}$ (s)	% $t_{\text{total}}$	$t_{\text{stage}}$ (s)	% $t_{\text{total}}$	$t_{\text{stage}}$ (s)	% $t_{\text{total}}$
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 <sub>2</sub>		CG <sub>3</sub>		CG <sub>4</sub>	
	$t_{\text{stage}}$ (s)	% $t_{\text{total}}$	$t_{\text{stage}}$ (s)	% $t_{\text{total}}$	$t_{\text{stage}}$ (s)	% $t_{\text{total}}$
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	

using a semi-implicit scheme and Picard linearization, following Natale et al. (2016), we must solve a sequence of the saddle point system at each time step of the form

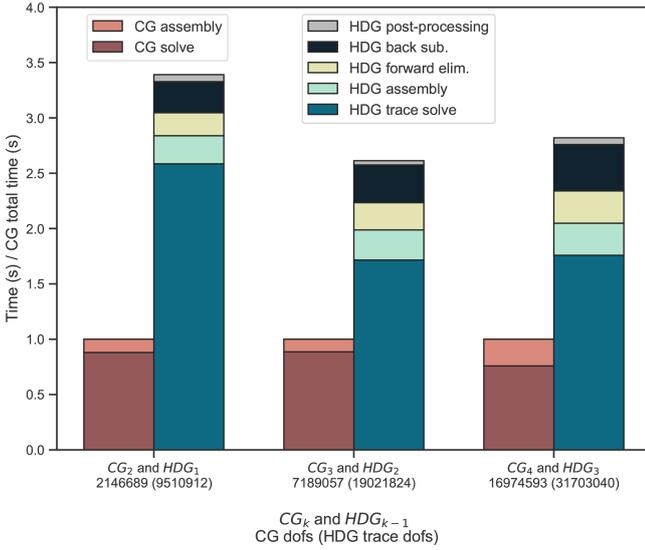
$$\begin{bmatrix} \mathbf{A} & -g \frac{\Delta t}{2} \mathbf{B}^T \\ H \frac{\Delta t}{2} \mathbf{B} & \mathbf{M} \end{bmatrix} \begin{Bmatrix} \Delta \mathbf{U} \\ \Delta \mathbf{D} \end{Bmatrix} = \begin{Bmatrix} \mathbf{R}_u \\ \mathbf{R}_D \end{Bmatrix}. \quad (83)$$

See Appendix A for a complete description of the entire discretization strategy. The system Eq. (83) is the matrix equation corresponding to the linearized equations in Eqs. (A3)–(A4).

The Picard updates  $\Delta \mathbf{U}$  and  $\Delta \mathbf{D}$  are sought in the mixed finite element spaces  $U_h \subset H(\text{div})$  and  $V_h \subset L^2$ , respectively. Stable mixed finite element pairings correspond to the well-known RT and BDM mixed methods, such as  $\text{RT}_k \times \text{DG}_{k-1}$

or  $\text{BDM}_k \times \text{DG}_{k-1}$ . These also fall within the set of compatible mixed spaces ideal for geophysical fluid dynamics (Cotter and Shipton, 2012; Natale et al., 2016; Melvin et al., 2019). 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 by many operational dynamical cores) is analogous to the Arakawa C-grid finite difference discretization.

In staggered finite difference models, the standard approach for solving Eq. (83) is to neglect the Coriolis term and eliminate the velocity unknown  $\Delta \mathbf{U}$  to obtain a discrete elliptic equation for  $\Delta \mathbf{D}$ , where smoothers like Richardson iterations or relaxation methods are convergent. This is more problematic in the compatible finite element frame-



**Figure 5.** Breakdown of the  $CG_k$  and  $HDG_{k-1}$  execution times on a  $6 \cdot 64^3$  simplicial mesh.

work, since  $\mathbf{A}$  has a dense inverse. Instead, we use the preconditioner described in Sect. 4.1.2 to form the equivalent hybridizable formulation, where both  $\Delta \mathbf{U}$  and  $\Delta \mathbf{D}$  are eliminated locally to produce a sparse elliptic equation for the Lagrange multipliers.

### 5.2.1 Atmospheric flow over a mountain

As a test problem, we solve test case 5 of Williamson et al. (1992), on the surface of a sphere with radius  $R = 6371$  km. We refer the reader to Cotter and Shipton (2012) and 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  $(RT_1, DG_0)$  (lowest-order RT method) and  $(BDM_2, DG_1)$  (next-to-lowest-order BDM method) for the velocity and depth spaces. A mesh of the sphere is generated from seven refinements of an icosahedron, resulting in a triangulation  $\mathcal{T}_h$  consisting of 327 680 elements in total. The grid information for both mixed methods is summarized in Table 2.

We run for a total of 25 time steps, with a fixed number of four 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 the generalized minimal residual method (GMRES)<sup>5</sup> acting on the system Eq. (83) with an approximate Schur complement preconditioner:

$$\mathcal{P}_{SC} = \begin{bmatrix} \mathbf{I} & g \frac{\Delta t}{2} \mathbf{A}^{-1} \mathbf{B}^T \\ \mathbf{0} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{A}^{-1} & \mathbf{0} \\ \mathbf{0} & \tilde{\mathbf{S}}^{-1} \end{bmatrix} \quad (84)$$

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

$$\begin{bmatrix} \mathbf{I} & \mathbf{0} \\ -H \frac{\Delta t}{2} \mathbf{B} \mathbf{A}^{-1} & \mathbf{I} \end{bmatrix},$$

where  $\tilde{\mathbf{S}} = \mathbf{M} + gH \frac{\Delta t^2}{4} \mathbf{B} \text{diag}(\mathbf{A})^{-1} \mathbf{B}^T$  and  $\text{diag}(\mathbf{A})$  is a diagonal approximation to the velocity mass matrix (plus the addition of a Coriolis matrix). The Schur complement system is inverted via GMRES due to the asymmetry from the Coriolis term, with the inverse of  $\tilde{\mathbf{S}}$  as the preconditioning operator. The sparse approximation  $\tilde{\mathbf{S}}$  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  $10^8$ .  $\mathbf{A}^{-1}$  is computed approximately using a single application of incomplete LU (ILU) with zero fill-in.

Next, we use only the application of our hybridization preconditioner (no outer Krylov method), which replaces the original linearized mixed system with its hybridizable equivalent. After hybridization, we have the following extended problem for the Picard updates: find  $(\Delta \mathbf{u}_h^d, \Delta D_h, \lambda_h) \in U_h^d \times V_h \times M_h$  satisfying

$$\begin{aligned} & (\mathbf{w}, \Delta \mathbf{u}_h^d)_{\mathcal{T}_h} + \frac{\Delta t}{2} (\mathbf{w}, f(\Delta \mathbf{u}_h^d)^\perp)_{\mathcal{T}_h} \\ & - \frac{\Delta t}{2} (\nabla \cdot \mathbf{w}, g \Delta D_h)_{\mathcal{T}_h} + \langle [[\mathbf{w}]], \lambda_h \rangle_{\partial \mathcal{T}_h} \\ & = \hat{\mathbf{R}}_{\mathbf{u}}, \forall \mathbf{w} \in U_h^d, \end{aligned} \quad (85)$$

$$(\phi, \Delta D_h)_{\mathcal{T}_h} + \frac{\Delta t}{2} (\phi, H \nabla \cdot \Delta \mathbf{u}_h^d)_{\mathcal{T}_h} = R_D, \quad \forall \phi \in V_h, \quad (86)$$

$$\langle \gamma, [[\Delta \mathbf{u}_h^d]] \rangle_{\partial \mathcal{T}_h} = 0, \quad \forall \gamma \in M_h. \quad (87)$$

Note that the space  $M_h$  is chosen such that the trace functions, when restricted to a facet  $e \in \partial \mathcal{T}_h$ , are in the same polynomial space as  $\Delta \mathbf{u}_h \cdot \mathbf{n}|_e$ . Moreover, it can be shown that the Lagrange multiplier  $\lambda_h$  is an approximation to the depth unknown  $\Delta t g \Delta D / 2$  restricted to  $\partial \mathcal{T}_h$ .

The resulting three-field problem in Eqs. (85)–(87) produces the following matrix equation:

$$\begin{bmatrix} \hat{\mathbf{A}} & \mathbf{C}^T \\ \mathbf{C} & \mathbf{0} \end{bmatrix} \begin{Bmatrix} \Delta \mathbf{X} \\ \boldsymbol{\Lambda} \end{Bmatrix} = \begin{Bmatrix} \hat{\mathbf{R}}_{\Delta X} \\ \mathbf{0} \end{Bmatrix}, \quad (88)$$

where  $\hat{\mathbf{A}}$  is the discontinuous operator coupling  $\Delta \mathbf{X} = \Delta \mathbf{U}^d \quad \Delta \mathbf{D}^T$  and  $\mathbf{R}_{\Delta X} = \hat{\mathbf{R}}_{\mathbf{u}} \quad \mathbf{R}_D^T$  are the problem residuals. An exact Schur complement factorization is performed on Eq. (88), using Slate to generate the local elimination kernels. We use the same set of solver options for the inversion of  $\tilde{\mathbf{S}}$  in Eq. (84) to invert the Lagrange multiplier system. The increments  $\Delta \mathbf{U}^d$  and  $\Delta \mathbf{D}$  are recovered locally, using Slate-generated kernels. Once recovery is complete,  $\Delta \mathbf{U}^d$  is projected back into the conforming  $H(\text{div})$  finite element space via  $\Delta \mathbf{U} \leftarrow \Pi_{\text{div}} \Delta \mathbf{U}^d$ . Based on the discussion in

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

Mixed method	No. of cells	$\Delta x$	Discretization properties		
			Velocity unknowns	Depth unknowns	Total (millions)
$RT_1 \times DG_0$	327 680	$\approx 43$ km	491 520	327 680	0.8 M
$BDM_2 \times DG_1$			2 457 600	983 040	3.4 M

Sect. 4.1.2, we apply

$$\mathcal{P}_{\text{hybrid}} = \Pi \left( \begin{bmatrix} \mathbf{I} & -\hat{\mathcal{A}}^{-1} \mathbf{C}^T \\ \mathbf{0} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \hat{\mathcal{A}}^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{S}^{-1} \end{bmatrix} \right) \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ -\mathbf{C} \hat{\mathcal{A}}^{-1} & \mathbf{I} \end{bmatrix} \Pi^T. \quad (89)$$

Table 3 displays a summary of our findings. The advantages of a hybridizable method versus a mixed method are more clearly realized in this experiment. When using hybridization, we observe a significant reduction in time spent in the implicit solver compared to the approximate Schur complement approach. This is primarily because we have reduced the number of “outer” iterations to zero; the hybridization preconditioner is performing an *exact* factorization of the global hybridizable system. This is empirically supported when considering per-application solve times. The values reported in Table 4 show the average cost of a single outer GMRES iteration (which includes the application of  $\mathcal{P}_{\text{SC}}$ ) and a single application of  $\mathcal{P}_{\text{hybrid}}$ . Hybridization and the approximate Schur complement preconditioner are comparable in terms of average execution time, with hybridization being slightly faster. This further demonstrates that the primary cause for the longer execution time of the latter is directly related to the additional outer iterations induced from using an approximate factorization. In terms of over all time-to-solution, the hybridizable methods are clearly ahead of the original mixed methods.

We also measure the relative reductions in the problem residual of the linear system Eq. (83). Our hybridization preconditioner 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. In other words, the reduction in the residual for the trace system translates into an overall reduction in the residual for the mixed system by the same factor.

The test case was run up to day 15 on a coarser resolution (20 480 simplicial cells with  $\Delta x \approx 210$  km) and a time-step size  $\Delta t = 500$  s. Snapshots of the entire simulation are provided in Fig. 6 using the semi-implicit scheme described in Appendix A. The results we have obtained for days 5, 10, and 15 are comparable to the corresponding results of Nair et al. (2005), Ullrich et al. (2010), and Kang et al. (2020). We refer the reader to Shipton et al. (2018) for further demonstrations

of shallow water test cases featuring the use of the hybridization preconditioner described in Sect. 4.1.2.

### 5.3 Hybridizable methods for a 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:

$$\frac{\partial \mathbf{u}}{\partial t} + 2\boldsymbol{\Omega} \times \mathbf{u} = -\nabla p + b\hat{\mathbf{z}}, \quad (90)$$

$$\frac{\partial p}{\partial t} = -c^2 \nabla \cdot \mathbf{u}, \quad (91)$$

$$\frac{\partial b}{\partial t} = -N^2 \mathbf{u} \cdot \hat{\mathbf{z}}, \quad (92)$$

where  $\mathbf{u}$  is the fluid velocity,  $p$  the pressure,  $b$  the buoyancy,  $\boldsymbol{\Omega}$  the planetary angular rotation vector,  $c$  the speed of sound ( $\approx 343 \text{ms}^{-1}$ ), and  $N$  the buoyancy frequency ( $\approx 0.01 \text{s}^{-1}$ ). Equations (90)–(92) permit fast-moving acoustic waves driven by perturbations in  $b$ . This is the 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  $H_\Omega$ . 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 finite element de Rham complexes:  $V_h^0 \xrightarrow{\partial_z} V_h^1$  and  $U_h^0 \xrightarrow{\nabla^\perp} U_h^1 \xrightarrow{\nabla} U_h^2$ . We can then construct the three-dimensional complex:  $W_h^0 \xrightarrow{\nabla} W_h^1 \xrightarrow{\nabla \times} W_h^2 \xrightarrow{\nabla \cdot}$

**Table 3.** Preconditioner solve times for a 25-step run with  $\Delta t = 100s$ . 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 outermost solver.

Preconditioner and solver details					
Mixed method	Preconditioner	$t_{\text{total}}$ (s)	Avg. outer its.	Avg. inner its.	$\frac{t_{\text{total}}^{\text{SC}}}{t_{\text{total}}^{\text{hybrid}}}$
$RT_1 \times DG_0$	approx. Schur. ( $\mathcal{P}_{\text{SC}}$ )	15.137	2	8	3.413
	hybridization ( $\mathcal{P}_{\text{hybrid}}$ )	4.434	none	2	
$BDM_2 \times DG_1$	approx. Schur. ( $\mathcal{P}_{\text{SC}}$ )	300.101	4	9	5.556
	hybridization ( $\mathcal{P}_{\text{hybrid}}$ )	54.013	none	6	

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

Preconditioner	Stage	$RT_1 \times DG_0$		$BDM_2 \times DG_1$	
		$t_{\text{stage}}$ (s)	% $t_{\text{total}}$	$t_{\text{stage}}$ (s)	% $t_{\text{total}}$
approx. Schur ( $\mathcal{P}_{\text{SC}}$ )	Schur solve	0.07592	91.28 %	0.78405	93.53 %
	invert velocity operator: $\mathbf{A}$	0.00032	0.39 %	0.00678	0.81 %
	apply inverse: $\mathbf{A}^{-1}$	0.00041	0.49 %	0.00703	0.84 %
	gmres other	0.00652	7.84 %	0.04041	4.82 %
	Total	0.08317		0.83827	
hybridization ( $\mathcal{P}_{\text{hybrid}}$ )	Transfer: $\mathbf{R}_{\Delta X} \rightarrow \widehat{\mathbf{R}}_{\Delta X}$	0.00322	7.26 %	0.00597	1.10 %
	Forward elim.: $-\mathbf{C}\widehat{\mathbf{A}}^{-1}\widehat{\mathbf{R}}_{\Delta X}$	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: $\Pi_{\text{div}}\Delta\mathbf{U}^d$	0.00264	5.96 %	0.00516	0.96 %
	Total	0.04434		0.54013	

$W_h^3$ , where

$$W_h^0 = U_h^0 \otimes V_h^0, \quad (93)$$

$$W_h^1 = \text{HCurl}(U_h^1 \otimes V_h^0) \oplus \text{HCurl}(U_h^0 \otimes V_h^1) =: W_h^{1,h} \oplus W_h^{1,v}, \quad (94)$$

$$W_h^2 = \text{HDiv}(U_h^1 \otimes V_h^1) \oplus \text{HDiv}(U_h^2 \otimes V_h^0) =: W_h^{2,h} \oplus W_h^{2,v}, \quad (95)$$

$$W_h^3 = U_h^2 \otimes V_h^1. \quad (96)$$

Here,  $\text{HCurl}$  and  $\text{HDiv}$  denote operators which ensure that 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 in Firedrake. For the analysis of compatible finite element discretizations and their relation to the complex Eqs. (93)–(96), we refer the reader to Natale et al. (2016). Each discretization used in this section is constructed from more familiar finite element families, shown in Table 5.

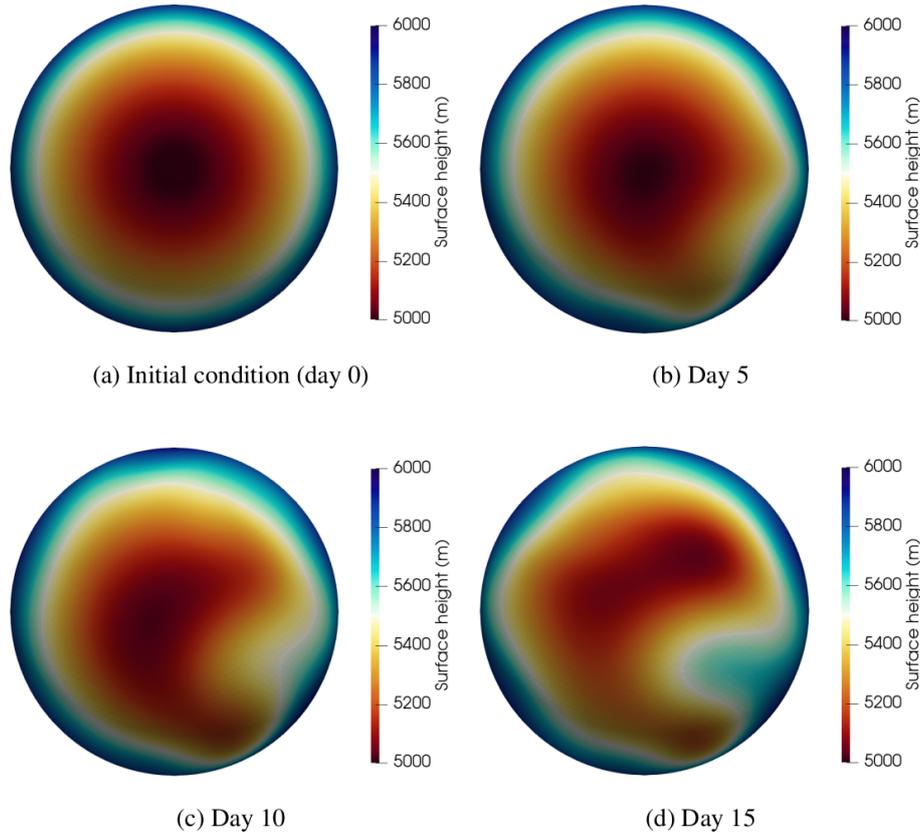
### 5.3.1 Compatible finite element discretization

A compatible finite element discretization of Eqs. (90)–(92) constructs solutions in the following finite element spaces:

$$\mathbf{u}_h \in \mathring{W}_h^2, \quad p_h \in W_h^3, \quad b_h \in W_h^b, \quad (97)$$

where  $\mathring{W}_h^2$  is the subspace of  $W_h^2 \subset H(\text{div})$  whose functions  $\mathbf{w}$  satisfy  $\mathbf{w} \cdot \mathbf{n} = 0$  on  $\partial\Omega$ ,  $W_h^3 \subset L^2$ , and  $W_h^b \equiv U_h^2 \otimes V_h^0$ . Note that  $W_h^b$  is just the scalar version of the vertical velocity space.<sup>6</sup> That is,  $W_h^b$  and  $W_h^{2,v}$  have the same number of

<sup>6</sup>The choice of  $W_h^b$  in Eq. (97) 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



**Figure 6.** 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  $\Delta t = 500$  s. The linear system during each Picard cycle was solved using the hybridization preconditioner.

**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.

Mixed method	Compatible finite element spaces				
	$V_h^0$	$V_h^1$	$U_h^0$	$U_h^1$	$U_h^2$
$RT_k$	$CG_k([0, H_\Omega])$	$DG_{k-1}([0, H_\Omega])$	$CG_k(\Delta)$	$RT_k(\Delta)$	$DG_{k-1}(\Delta)$
$BDFM_{k+1}$	$CG_{k+1}([0, H_\Omega])$	$DG_k([0, H_\Omega])$	$CG_{k+1}(\Delta)$	$BDFM_{k+1}(\Delta)$	$DG_k(\Delta)$
$RTCF_k$	$CG_k([0, H_\Omega])$	$DG_{k-1}([0, H_\Omega])$	$Q_k(\square)$	$RTCF_k(\square)$	$DQ_{k-1}(\square)$

degrees of freedom but differ in how they are pulled back to the reference element.

To obtain the discrete system, we simply multiply Eqs. (90)–(92) by test functions  $\mathbf{w} \in \mathring{W}_h^2$ ,  $\phi \in W_h^3$ , and  $\eta \in W_h^b$  and integrate by parts. We introduce the increments  $\delta \mathbf{u}_h \equiv \mathbf{u}_h^{n+1} - \mathbf{u}_h^n$ , and set  $\mathbf{u}_0 \equiv \mathbf{u}_h^n$  (similarly for  $\delta p_h$ ,  $p_0$ ,  $\delta b_h$ , and  $b_0$ ). Using an implicit midpoint rule discretization, we need to solve the following mixed problem at each time step:

et al., 2010). One could also collocate  $b_h$  with  $p_h$  ( $b_h \in W_h^3$ ), 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.

find  $\delta \mathbf{u}_h \in \mathring{W}_h^2$ ,  $\delta p_h \in W_h^3$  and  $\delta b_h \in W_h^b$  such that

$$\begin{aligned}
 (\mathbf{w}, \delta \mathbf{u}_h)_{\mathcal{T}_h} + \frac{\Delta t}{2} (\mathbf{w}, 2\boldsymbol{\Omega} \times \delta \mathbf{u}_h)_{\mathcal{T}_h} - \frac{\Delta t}{2} \\
 (\nabla \cdot \mathbf{w}, \delta p_h)_{\mathcal{T}_h} - \frac{\Delta t}{2} (\mathbf{w}, \delta b_h \hat{\mathbf{z}})_{\mathcal{T}_h} &= r_u, \quad \forall \mathbf{w} \in \mathring{W}_h^2 \tag{98}
 \end{aligned}$$

$$(\phi, \delta p_h)_{\mathcal{T}_h} + \frac{\Delta t}{2} c^2 (\phi, \nabla \cdot \delta \mathbf{u}_h)_{\mathcal{T}_h} = r_p, \quad \forall \phi \in W_h^3, \tag{99}$$

$$(\eta, \delta b_h)_{\mathcal{T}_h} + \frac{\Delta t}{2} N^2 (\eta, \delta \mathbf{u}_h \cdot \hat{\mathbf{z}})_{\mathcal{T}_h} = r_b, \quad \forall \eta \in W_h^b, \tag{100}$$

where the residuals are  $r_u = -\Delta t(\mathbf{w}, 2\boldsymbol{\Omega} \times \mathbf{u}_0)_{\mathcal{T}_h}$ ,  $r_p = -c^2 \Delta t(\phi, \nabla \cdot \mathbf{u}_0)_{\mathcal{T}_h}$ , and  $r_b = -N^2 \Delta t(\eta, \mathbf{u}_0 \cdot \hat{\mathbf{z}})_{\mathcal{T}_h}$ .

The resulting matrix equations have the form

$$\begin{bmatrix} \mathbf{A}_u & -\frac{\Delta t}{2} \mathbf{D}^T & -\frac{\Delta t}{2} \mathbf{Q}^T \\ \frac{\Delta t}{2} c^2 \mathbf{D} & \mathbf{M}_p & \mathbf{0} \\ \frac{\Delta t}{2} N^2 \mathbf{Q} & \mathbf{0} & \mathbf{M}_b \end{bmatrix} \begin{Bmatrix} \mathbf{U} \\ \mathbf{P} \\ \mathbf{B} \end{Bmatrix} = \begin{Bmatrix} \mathbf{R}_u \\ \mathbf{R}_p \\ \mathbf{R}_b \end{Bmatrix}, \quad (101)$$

where  $\mathbf{A}_u = \mathbf{M}_u + \frac{\Delta t}{2} \mathbf{C}_\Omega$ ,  $\mathbf{C}_\Omega$  is the asymmetric matrix associated with the Coriolis term,  $\mathbf{M}_u$ ,  $\mathbf{M}_p$ , and  $\mathbf{M}_b$  are mass matrices,  $\mathbf{D}$  is the weak divergence term, and  $\mathbf{Q}$  is an operator containing the vertical components of  $\delta \mathbf{u}_h$ . In the absence of orography, we can use the point-wise expression for the buoyancy,

$$\delta b_h = r_b - \frac{\Delta t}{2} N^2 \delta \mathbf{u}_h \cdot \hat{\mathbf{z}}, \quad (102)$$

and eliminate  $\delta b_h$  from Eq. (101) by substituting Eq. (102) into Eq. (98). This produces the following mixed velocity–pressure system:

$$\mathcal{A} \begin{Bmatrix} \mathbf{U} \\ \mathbf{P} \end{Bmatrix} = \begin{bmatrix} \tilde{\mathbf{A}}_u & -\frac{\Delta t}{2} \mathbf{D}^T \\ c^2 \frac{\Delta t}{2} \mathbf{D} & \mathbf{M}_p \end{bmatrix} \begin{Bmatrix} \mathbf{U} \\ \mathbf{P} \end{Bmatrix} = \begin{Bmatrix} \tilde{\mathbf{R}}_u \\ \mathbf{R}_p \end{Bmatrix}, \quad (103)$$

where  $\tilde{\mathbf{A}}_u = \mathbf{A}_u + \frac{\Delta t^2}{4} N^2 \mathbf{Q}^T \mathbf{M}_b^{-1} \mathbf{Q}$  and  $\tilde{\mathbf{R}}_u = \mathbf{R}_u + \frac{\Delta t}{2} \mathbf{Q}^T \mathbf{M}_b^{-1} \mathbf{R}_b$  are the modified velocity operator and right-hand side, respectively. Note that in our elimination strategy,  $\tilde{\mathbf{A}}_u$  corresponds to the bilinear form obtained after eliminating the buoyancy at the equation level:

$$\tilde{\mathbf{A}}_u \leftarrow (\mathbf{w}, \delta \mathbf{u})_{\mathcal{T}_h} + \frac{\Delta t}{2} (\mathbf{w}, 2\boldsymbol{\Omega} \times \delta \mathbf{u})_{\mathcal{T}_h} + \frac{\Delta t^2}{4} N^2 (\mathbf{w} \cdot \hat{\mathbf{z}}, \delta \mathbf{u} \cdot \hat{\mathbf{z}})_{\mathcal{T}_h}. \quad (104)$$

A similar construction holds for  $\tilde{\mathbf{R}}_u$ . Once Eq. (103) is solved,  $\delta b_h$  is reconstructed by solving

$$\mathbf{M}_b \mathbf{B} = \mathbf{R}_b - \frac{\Delta t}{2} N^2 \mathbf{Q} \mathbf{U}. \quad (105)$$

Equation (105) can be efficiently inverted using the conjugate gradient method.

### 5.3.2 Preconditioning the mixed velocity pressure system

The primary difficulty is finding efficient solvers for Eq. (103). This was studied by Mitchell and Müller (2016) within the context of developing a preconditioner which is robust against parameters, like  $\Delta t$  and mesh resolution. However, the implicit treatment of the Coriolis term was not taken into account. We consider two preconditioning strategies.

One strategy proposed by Mitchell and Müller (2016) is to build a preconditioner based on the Schur complement factorization of  $\mathcal{A}$  in Eq. (103):

$$\mathcal{A}^{-1} = \begin{bmatrix} \mathbf{I} & \frac{\Delta t}{2} \tilde{\mathbf{A}}_u^{-1} \mathbf{D}^T \\ \mathbf{0} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \tilde{\mathbf{A}}_u^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{H}^{-1} \end{bmatrix} \quad (106)$$

$$\begin{bmatrix} \mathbf{I} & \mathbf{0} \\ -c^2 \frac{\Delta t}{2} \mathbf{D} \tilde{\mathbf{A}}_u^{-1} & \mathbf{I} \end{bmatrix},$$

where  $\mathbf{H} = \mathbf{M}_p + c^2 \frac{\Delta t^2}{4} \mathbf{D} \tilde{\mathbf{A}}_u^{-1} \mathbf{D}^T$  is the dense pressure Helmholtz operator. Because we have chosen to include the Coriolis term, the operator  $\mathbf{H}$  is nonsymmetric and has the form

$$\mathbf{H} = \mathbf{M}_p + c^2 \frac{\Delta t^2}{4} \mathbf{D} \left( \tilde{\mathbf{M}}_u + \frac{\Delta t}{2} \mathbf{C}_\Omega \right)^{-1} \mathbf{D}^T, \quad (107)$$

where  $\tilde{\mathbf{M}}_u$  is a modified mass matrix. As  $\Delta t$  increases, the contribution of  $\mathbf{C}_\Omega$  becomes more prominent in  $\mathbf{H}$ , making sparse approximations of  $\mathbf{H}$  more challenging. We shall elaborate on this further below when we present the results of our second solver strategy.

Our preferred strategy solves the hybridizable formulation of the system Eq. (103). Let  $W_h^{2,d}$  denote the broken version of  $W_h^2$  and  $M_h$  the space of Lagrange multipliers. Then the hybridizable formulation for the velocity–pressure system reads as follows: find  $\delta \mathbf{u}_h^d \in W_h^{2,d}$ ,  $\delta p_h \in W_h^3$ , and  $\lambda_h \in M_h$  such that

$$\begin{aligned} & (\mathbf{w}, \delta \mathbf{u}_h^d)_{\mathcal{T}_h} + \frac{\Delta t}{2} (\mathbf{w}, 2\boldsymbol{\Omega} \times \delta \mathbf{u}_h^d)_{\mathcal{T}_h} \\ & + \frac{\Delta t^2}{4} N^2 (\mathbf{w} \cdot \hat{\mathbf{z}}, \delta \mathbf{u}_h^d \cdot \hat{\mathbf{z}})_{\mathcal{T}_h} - \frac{\Delta t}{2} (\nabla \cdot \mathbf{w}, \delta p_h)_{\mathcal{T}_h} \\ & + \langle [[\mathbf{w}]], \lambda_h \rangle_{\partial \mathcal{T}_h} = \tilde{r}_u, \quad \forall \mathbf{w} \in W_h^{2,d} \end{aligned} \quad (108)$$

$$(\phi, \delta p_h)_{\mathcal{T}_h} + \frac{\Delta t}{2} c^2 (\phi, \nabla \cdot \delta \mathbf{u}_h)_{\mathcal{T}_h} = r_p, \quad \forall \phi \in W_h^3, \quad (109)$$

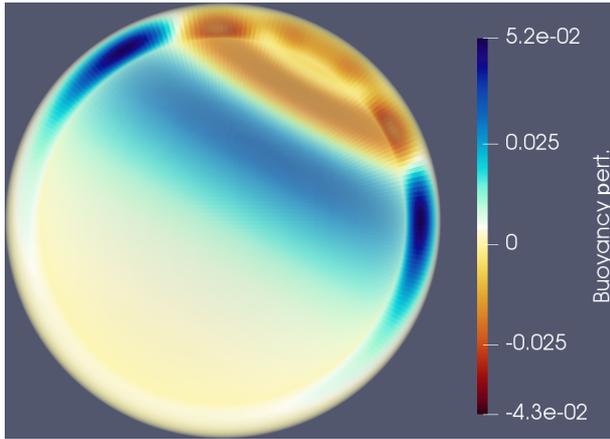
$$\langle \gamma, [[[\delta \mathbf{u}_h^d]]] \rangle_{\partial \mathcal{T}_h} = 0, \quad \forall \gamma \in M_h. \quad (110)$$

The system Eqs. (108)–(110) is automatically formed by the Firedrake preconditioner: `firedrake.HybridizationPC`. We then locally eliminate the velocity and pressure after hybridization, producing the following condensed problem:

$$\mathbf{H}_\theta \boldsymbol{\Lambda} = \mathbf{E}, \quad \mathbf{H}_\theta = \begin{bmatrix} \mathbf{C} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \hat{\tilde{\mathbf{A}}}_u & -\frac{\Delta t}{2} \hat{\mathbf{D}}^T \\ c^2 \frac{\Delta t}{2} \hat{\mathbf{D}} & \mathbf{M}_p \end{bmatrix}^{-1} \quad (111)$$

$$\begin{bmatrix} \mathbf{C}^T \\ \mathbf{0} \end{bmatrix}, \quad \mathbf{E} = \begin{bmatrix} \mathbf{C} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \hat{\tilde{\mathbf{A}}}_u & -\frac{\Delta t}{2} \hat{\mathbf{D}}^T \\ c^2 \frac{\Delta t}{2} \hat{\mathbf{D}} & \mathbf{M}_p \end{bmatrix}^{-1} \begin{Bmatrix} \hat{\mathbf{R}}_u \\ \mathbf{R}_p \end{Bmatrix},$$

where  $\hat{\cdot}$  denotes matrices or vectors with discontinuous test and trial functions. The nonsymmetric operator  $\mathbf{H}_\theta$  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



**Figure 7.** 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 are a simple solid-body rotation:  $\mathbf{u} = 20\mathbf{e}_\lambda$ , where  $\mathbf{e}_\lambda$  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,  $R = 6371\text{km}/125$  is the planet radius, and  $\lambda_\phi = 2/3$ . The equations are discretized using the lowest-order method RTCF<sub>1</sub>, with 24 576 quadrilateral cells in the horizontal and 64 extrusion levels. The velocity–pressure system is solved using hybridization.

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. 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 nonsymmetric problems can be found in Bramble et al. (1994, 1988) and Mandel (1986). A gravity wave test using our solution strategy and hybridization preconditioner is illustrated in Fig. 7 for a problem on a condensed Earth (radius scaled down by a factor of 125) and 10 km lid.

### 5.3.3 Robustness against acoustic Courant number with implicit Coriolis

In this final experiment, we repeat a similar study to that presented in Mitchell and Müller (2016). Our setup closely resembles the gravity wave test of Skamarock and Klemp (1994) extended to a spherical annulus. We initialize the velocity in a simple solid-body rotation and introduce a localized buoyancy anomaly. A Coriolis term is introduced as a function of the Cartesian coordinate  $z$  and is constant along lines of latitude ( $f$  plane):  $2\boldsymbol{\Omega} = 2\Omega_r \frac{z}{R} \hat{\mathbf{z}}$ , with angular rotation rate  $\Omega_r = 7.292 \times 10^{-5} \text{s}^{-1}$ . We fix the resolution of

the problem and run the solver over a range of  $\Delta t$ . 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  $\Delta x$  is the horizontal resolution.

Note 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 setup mirrors that of actual global weather models; we extrude a spherical mesh of the Earth upwards to a height of 85 km. The setup for the different discretizations (including degrees of freedom for the velocity–pressure and hybridized systems) is presented in Table 6.

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

$$\tilde{\mathbf{H}} = \mathbf{M}_p + c^2 \frac{\Delta t^2}{4} \mathbf{D}(\text{Diag}(\tilde{\mathbf{A}}_{\mathbf{u}}))^{-1} \mathbf{D}^T \quad (112)$$

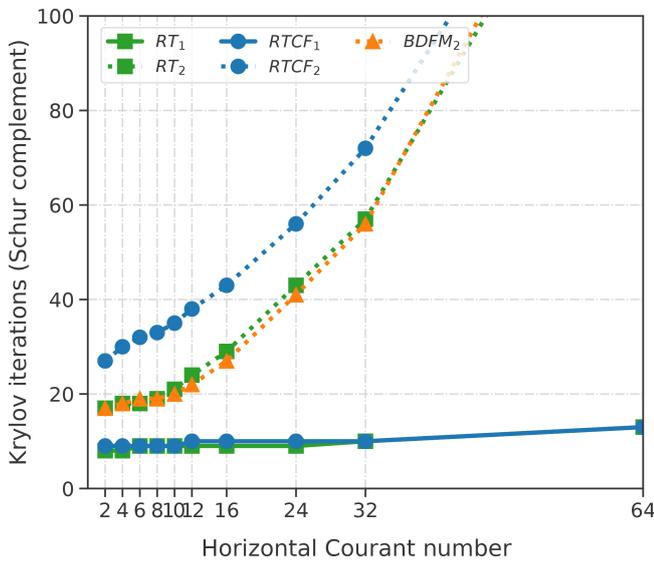
served as a good preconditioner, leading to a system that was amenable to multigrid methods and resulted in a Courant-number-independent solver. However, when the Coriolis term is included, this is no longer the case: the diagonal approximation  $\text{Diag}(\tilde{\mathbf{A}}_{\mathbf{u}})$  becomes worse with increasing  $\lambda_C$ . To demonstrate this, we solve the gravity wave system on a low-resolution grid (10 km lid, 10 vertical levels, maintaining the same cell aspect ratio as in Table 6) using the Schur complement factorization Eq. (106). LU factorizations are applied to invert both  $\tilde{\mathbf{A}}_{\mathbf{u}}^{-1}$  and  $\tilde{\mathbf{H}}^{-1}$ . Inverting the Schur complement  $\tilde{\mathbf{H}}^{-1}$  is done using preconditioned GMRES iterations, and a flexible-GMRES algorithm is used on the full velocity–pressure system. If  $\tilde{\mathbf{H}}^{-1}$  is a good approximation to  $\mathbf{H}^{-1}$ , then we should see low iteration counts in the Schur complement solve. Figure 8 shows the results of this study for a range of Courant numbers.

For the lower-order methods, the number of iterations to invert  $\mathbf{H}$  grows slowly but remains under control. Increasing the approximation degree by 1 results in degraded performance. As  $\Delta t$  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 canceled out due to asymmetry.

Hybridization avoids this problem entirely: we always construct an exact Schur complement and only have to worry about solving the trace system Eq. (111). We now show that this approach (described in Sect. 5.3.2) is much more robust to changes in  $\Delta t$ . We use the same workstation as for the three-dimensional CG/HDG problem in Sect. 5.1 (executed with a total of 40 MPI processes). Figure 9 shows the parameter test for all the discretizations described in Table 6. We see that, in terms of total number of GCR iterations needed

**Table 6.** Grid setup and discretizations for the acoustic Courant number study. The total number of degrees of freedom (dofs) for the mixed (velocity and pressure) and hybridizable (velocity, pressure, and trace) discretizations are shown in the last two columns (millions). The vertical resolution is fixed across all discretizations.

Discretizations and grid information						
Mixed method	No. of horiz. cells	No. of vert. layers	$\Delta x$	$\Delta z$	$U$ - $P$ dofs	Hybrid. dofs
RT <sub>1</sub>	81 920	85	86 km	1 000 m	24.5 M	59.3 M
RT <sub>2</sub>	5 120	85	346 km	1 000 m	9.6 M	17.4 M
BDFM <sub>2</sub>	5 120	85	346 km	1 000 m	10.5 M	18.3 M
RTCF <sub>1</sub>	98 304	85	78 km	1 000 m	33.5 M	83.7 M
RTCF <sub>2</sub>	6 144	85	312 km	1 000 m	16.7 M	29.3 M



**Figure 8.** 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 1 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.

to invert the trace system, hybridization is far more controlled as Courant number increases. 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 Fig. 9).

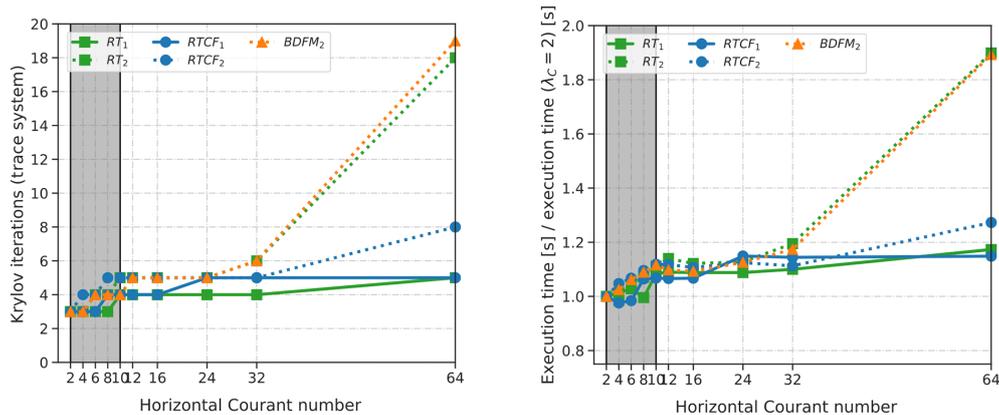
In Fig. 9b, we display the ratio of execution time and the time-to-solution at the lowest Courant number of 2. We perform this normalization to better compare the lower

and higher-order methods (and discretizations on triangular prisms vs. extruded quadrilaterals). The calculation of the ratios includes the time needed to eliminate and reconstruct the hybridized velocity–pressure variables. The fact that the hybridization solver remains close to 1 demonstrates that the entire solution procedure is largely  $\lambda_C$ -independent until around  $\lambda_C = 32$ . The overall trend is largely the same as the number of Krylov iterations to reach solver convergence. This is due to our hybridization approach being solver dominated, with local operations like forward elimination together with local recovery taking approximately one-third of the total execution time for each method. The percentage breakdown of the hybridization solver is similar to what is already presented in Sect. 5.1.2.

Implicitly treating the Coriolis term has been discussed for semi-implicit discretizations of large-scale geophysical flows (Temperton, 1997; Cullen, 2001; Nechaev and Yaremchuk, 2004). The addition of the Coriolis term presents a particular challenge to the solution finite element discretizations of these equations since it increases the difficulty of finding a good sparse approximation of the nonsymmetric elliptic operator. Hybridization shows promise here, as it allows for the assembly of the elliptic equation that both captures the effects of rotation and results in a sparse linear system.

## 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 are 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 code 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



**Figure 9.** Courant number parameter test run on a fully loaded compute node. Both figures display the hybridized solver for each discretization, described in Table 6. Panel (a) displays total iteration count (preconditioned GCR) to solve the trace system to a relative tolerance of  $10^{-5}$ . Panel (b) displays the relative work of each solver, which takes into account the time required to forward eliminate and locally recover the velocity and pressure.

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) and 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 require detailed analysis of the resulting operator systems arising from hybridization, are a necessary next step to determine whether hybridization provides a scalable solution strategy for compatible finite elements in operational settings.

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

For some tessellation,  $\mathcal{T}_h$ , our semi-discrete mixed method for Eqs. (81)–(82) seeks approximations  $(\mathbf{u}_h, D_h) \in U_h \times V_h \subset H(\text{div}) \times L^2$  satisfying

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

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

where  $\tilde{\cdot}$  indicates that the value of the function should be taken from the upwind side of each facet. The discretization 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

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

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

where  $H$  is the mean layer depth and  $R_{\mathbf{u}}$  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 Eqs. (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 Eqs. (A1)–(A2) is

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

$$\begin{aligned} & \left( \phi, D_h^{n+1} - D_h^n \right)_{\mathcal{T}_h} - \Delta t (\nabla \phi, \mathbf{u}_h^* D_h^*)_{\mathcal{T}_h} \\ & + \Delta t \left\langle \left[ \left[ \phi \mathbf{u}_h^* \right] \right], \tilde{D}_h^* \right\rangle_{\partial \mathcal{T}_h} = 0, \forall \phi \in V_h, \end{aligned} \quad (\text{A6})$$

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  $R_{\mathbf{u}}$  and  $R_D$  would be to simply define these from Eqs. (A5)–(A6). However, this leads to a small critical time step for the stability of the scheme. To make the numerical scheme more stable, we define residuals as follows. For  $R_{\mathbf{u}}$ , we first solve for  $\mathbf{v}_h \in \mathbf{U}_h$  such that

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

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

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

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

$$\begin{aligned} & (\phi, E_h - D_h^n)_{\mathcal{T}_h} - \Delta t (\nabla \phi, \mathbf{u}_h^* E_h^*)_{\mathcal{T}_h} \\ & + \Delta t \left\langle \left[ \left[ \phi \mathbf{u}_h^* \right] \right], \tilde{E}_h^* \right\rangle_{\partial \mathcal{T}_h} = 0, \forall \phi \in V_h, \end{aligned} \quad (\text{A9})$$

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

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

This process can be thought of as iteratively solving for the average velocity and depth that satisfies the implicit midpoint rule discretization. Both Eqs. (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 Eqs. (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  $(\mathbf{u}_h^{n+1}, D_h^{n+1}) \leftarrow (\mathbf{u}_h^{n+1} + \Delta \mathbf{u}_h, D_h^{n+1} + \Delta D_h)$ , having initially chosen  $(\mathbf{u}_h^{n+1}, D_h^{n+1}) = (\mathbf{u}_h^n, D_h^n)$ .

*Code and data availability.* The contribution in this paper is available through open-source software provided by the Firedrake Project: <https://www.firedrakeproject.org/> (last access: 30 January 2020). 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).

*Author contributions.* THG is the principal author and developer of the software presented in this paper and main author of the text. Authors LM and DAH assisted and guided the software abstraction as a domain-specific language and edited text. CJC contributed to the formulation of the geophysical fluid dynamics and the design of the numerical experiments and edited text.

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

*Acknowledgements.* The authors would like to acknowledge funding from the Engineering and Physical Sciences Research Council (EPSRC) and the Natural Environment Research Council (NERC). The authors also wish to thank Andrew T. T. McRae for providing thoughtful comments on early drafts of this paper.

*Financial support.* This research has been supported by the Engineering and Physical Sciences Research Council (grant nos. EP/M011054/1, EP/L000407/1, and EP/L016613/1) and the Natural Environment Research Council (grant no. NE/K008951/1).

*Review statement.* This paper was edited by Simone Marras and reviewed by two anonymous referees.

## 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 Trans. Mathe. Softw. (TOMS)*, 40, 1–37, 2014.
- Arnold, D. N. and Brezzi, F.: Mixed and nonconforming finite element methods: implementation, postprocessing and error estimates, *ESAIM: Mathe. Modell. Num. Anal.*, 19, 7–32, 1985.
- Arnold, D. N., Falk, R. S., and Winther, R.: Multigrid in  $H(\text{div})$  and  $H(\text{curl})$ , *Num. Mathe.*, 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*, 163–202, Springer, 1997.
- Balay, S., Abhyankar, S., Adams, M. F., Brown, J., Brune, P., Buschelman, K., Dalcin, L., Eijkhout, V., Gropp, W. D., Karpeyev, D., Kaushik, D., Knepley, M. G., May, D. A., McInnes, L. C., Mills, R. T., Munson, T., Rupp, K., Sanan, P., Smith, B. F., Zampini, S., Zhang, H., and Zhang, H.: *PETSc Users Manual*, Tech. Rep. ANL-95/11 – Revision 3.11, Argonne National Laboratory, 2019.
- Bauer, W. and Cotter, C.: Energy-entropy conserving compatible finite element schemes for the rotating shallow water equations with slip boundary conditions, *J. Comput. Phys.*, 373, 171–187, <https://doi.org/10.1016/j.jcp.2018.06.071>, 2018.
- Boffi, D., Brezzi, F., and Fortin, M.: *Mixed finite element methods and applications*, vol. 44 of *Springer Series in Computational Mathematics*, Springer-Verlag New York, 2013.
- Bramble, J. H. and Xu, J.: A local post-processing technique for improving the accuracy in mixed finite-element approximations, *SIAM J. Num. Anal.*, 26, 1267–1275, 1989.
- Bramble, J. H., Pasciak, J. E., and Xu, J.: The analysis of multigrid algorithms for nonsymmetric and indefinite elliptic problems, *Math. Comput.*, 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 J. Num. Anal.*, 31, 1746–1763, 1994.
- Brezzi, F. and Fortin, M.: *Mixed and hybrid finite element methods*, vol. 15 of *Springer Series in Computational Mathematics*, Springer-Verlag New York, 1991.
- Brezzi, F., Douglas, J., and Marini, L. D.: Two families of mixed finite elements for second order elliptic problems, *Num. Mathe.*, 47, 217–235, 1985.
- Brezzi, F., Douglas, J., Durán, R., and Fortin, M.: Mixed finite elements for second order elliptic problems in three variables, *Num. Mathe.*, 51, 237–250, 1987.
- Brown, J., Knepley, M. G., May, D. A., McInnes, L. C., and Smith, B.: Composable linear solvers for multiphysics, in: *Parallel and Distributed Computing (ISPD)*, 2012 11th International Symposium on, 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*, 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 J. Num. Anal.*, 47, 1319–1365, 2009a.
- Cockburn, B., Guzmán, J., and Wang, H.: Superconvergent discontinuous Galerkin methods for second-order elliptic problems, *Mathe. Comput.*, 78, 1–24, 2009b.
- Cockburn, B., Gopalakrishnan, J., Li, F., Nguyen, N.-C., and Peraire, J.: Hybridization and postprocessing techniques for mixed eigenfunctions, *SIAM J. Num. Anal.*, 48, 857–881, 2010a.
- Cockburn, B., Gopalakrishnan, J., and Sayas, F.-J.: A projection-based error analysis of HDG methods, *Mathe. Comput.*, 79, 1351–1367, 2010b.
- Cockburn, B., Dubois, O., Gopalakrishnan, J., and Tan, S.: Multigrid for an HDG method, *IMA J. Num. Anal.*, 34, 1386–1425, 2014.
- Cotter, C. J. and Shipton, J.: Mixed finite elements for numerical weather prediction, *J. Comput. Phys.*, 231, 7076–7091, 2012.

- Cotter, C. J. and Thuburn, J.: A finite element exterior calculus framework for the rotating shallow-water equations, *J. Comput. Phys.*, 257, 1506–1526, 2014.
- Cullen, M.: Alternative implementations of the semi-Lagrangian semi-implicit schemes in the ECMWF model, *Q. J. Roy. Meteorol. Soc.*, 127, 2787–2802, 2001.
- Dalcin, L. D., Paz, R. R., Kler, P. A., and Cosimo, A.: Parallel distributed computing using python, *Adv. Water Resour.*, 34, 1124–1139, 2011.
- 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, *Int. J. Num. Method. Eng.*, 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 J. Sci. Comput.*, 23, 1291–1315, 2001.
- 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*, 267–294, Springer, 2006.
- Fraeijs de Veubeke, B.: Displacement and equilibrium models in the finite element method, in: *Stress Analysis*, edited by: Zienkiewicz, O. and Holister, G. S., John Wiley & Sons, reprinted in *Internat. J. Numer. Methods Engrg.*, 52, 287–342, 1965.
- Gopalakrishnan, J.: A Schwarz preconditioner for a hybridized mixed method, *Comput. Methods Appl. Math.*, 3, 116–134, 2003.
- Guennebaud, G., Jacob, B., Avery, P., Bachrach, A., and Barthelmy, S.: Eigen v3, 2010, available at: <http://eigen.tuxfamily.org> (last access: 7 February 2020), 2015.
- Guyan, R. J.: Reduction of stiffness and mass matrices, *AIAA J.*, 3, 380, 1965.
- Hecht, F.: New development in FreeFem++, *J. Num. Mathe.*, 20, 251–266, 2012.
- Hiptmair, R. and Xu, J.: Nodal auxiliary space preconditioning in  $H(\text{curl})$  and  $H(\text{div})$  spaces, *SIAM J. Num. Anal.*, 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 J. Sci. Comput.*, 40, C401–C428, <https://doi.org/10.1137/17M1130642>, 2018.
- Irons, B.: Structural eigenvalue problems-elimination of unwanted variables, *AIAA J.*, 3, 961–962, 1965.
- Kang, S., Giraldo, F. X., and Bui-Thanh, T.: IMEX HDG-DG: A coupled implicit hybridized discontinuous Galerkin and explicit discontinuous Galerkin approach for shallow water systems, *J. Comput. Phys.*, 401, 109010, 2020.
- Kirby, R. C.: Algorithm 839: FIAT, a new paradigm for computing finite element basis functions, *ACM Trans. Mathe. Softw. (TOMS)*, 30, 502–516, 2004.
- Kirby, R. C. and Logg, A.: A compiler for variational forms, *ACM Trans. Mathe. Softw. (TOMS)*, 32, 417–444, 2006.
- Kirby, R. C. and Mitchell, L.: Solver composition across the PDE/linear algebra barrier, *SIAM J. Sci. Comput.*, 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, *J. Sci. Comput.*, 51, 183–212, 2012.
- Kronbichler, M. and Wall, W. A.: A performance comparison of continuous and discontinuous Galerkin methods with fast multigrid solvers, *SIAM J. Sci. Comput.*, 40, A3423–A3448, 2018.
- Logg, A. and Wells, G. N.: DOLFIN: Automated finite element computing, *ACM Trans. Mathe. Softw. (TOMS)*, 37, 20, 2010.
- Logg, A., Mardal, K.-A., and Wells, G.: Automated solution of differential equations by the finite element method: The FEniCS book, Vol. 84, Springer Science & Business Media, 2012a.
- Logg, A., Ølgaard, K. B., Rognes, M. E., and Wells, G. N.: FFC: the FEniCS form compiler, *Automated Solution of Differential Equations by the Finite Element Method*, 227–238, 2012b.
- Long, K., Kirby, R., and van Bloemen Waanders, B.: Unified embedded parallel finite element computations via software-based Fréchet differentiation, *SIAM J. Sci. Comput.*, 32, 3323–3351, 2010.
- Mandel, J.: Multigrid convergence for nonsymmetric, indefinite variational problems and one smoothing step, *Appl. Mathe. Comput.*, 19, 201–216, 1986.
- Markall, G., Slemmer, A., Ham, D., Kelly, P., Cantwell, C., and Sherwin, S.: Finite element assembly strategies on multi-core and many-core architectures, *International J. Num. Method. Fluids*, 71, 80–97, 2013.
- 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 J. Sci. Comput.*, 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, *Q. J. Roy. Meteorol. Soc. A*, 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, *Q. J. Roy. Meteorol. Soc.*, 145, 2835–2853, <https://doi.org/10.1002/qj.3501>, 2019.
- Mitchell, L. and Müller, E. H.: High level implementation of geometric multigrid solvers for finite element problems: Applications in atmospheric modelling, *J. Comput. Phys.*, 327, 1–18, 2016.
- Nair, R. D., Thomas, S. J., and Loft, R. D.: A discontinuous Galerkin global shallow water model, *Mon. Weather Rev.*, 133, 876–888, 2005.
- Natale, A. and Cotter, C. J.: A variational  $H(\text{div})$  finite-element discretization approach for perfect incompressible fluids, *IMA J. Num. Anal.*, 38, p. drx033, <https://doi.org/10.1093/imanum/drx033>, 2017.
- Natale, A., Shipton, J., and Cotter, C. J.: Compatible finite element spaces for geophysical fluid dynamics, *Dynam. Stat. Clim. Syst.*, 1, dzw005, <https://doi.org/10.1093/climsys/dzw005>, 2016.
- Nechaev, D. and Yaremchuk, M.: On the approximation of the Coriolis terms in C-grid models, *Mon. Weather Rev.*, 132, 2283–2289, 2004.
- Nédélec, J.-C.: Mixed finite elements in  $\mathbb{R}^3$ , *Num. Mathe.*, 35, 315–341, 1980.
- 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*, 38, 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., 1116–1123, IEEE, 2012.
- 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 Trans. Mathe. Softw. (TOMS)*, 43, 1–27, 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*, 292–315, Springer, 1977.
- Shipton, J., Gibson, T. H., and Cotter, C. J.: Higher-order compatible finite element schemes for the nonlinear rotating shallow water equations on the sphere, *J. Comput. Phys.*, 375, 1121–1137, 2018.
- Skamarock, W. C. and Klemp, J. B.: Efficiency and accuracy of the Klemp-Wilhelmson time-splitting technique, *Mon. Weather Rev.*, 122, 2623–2630, 1994.
- Stenberg, R.: Postprocessing schemes for some mixed finite elements, *ESAIM: Mathe. Modell. Num. Anal.*, 25, 151–167, 1991.
- Temperton, C.: Treatment of the Coriolis terms in semi-Lagrangian spectral models, *Atmos.-Ocean*, 35, 293–302, 1997.
- Thomas, S. J., Hacker, J. P., Smolarkiewicz, P. K., and Stull, R. B.: Spectral preconditioners for nonhydrostatic atmospheric models, *Mon. Weather Rev.*, 131, 2464–2478, 2003.
- Ullrich, P. A., Jablonowski, C., and Van Leer, B.: High-order finite-volume methods for the shallow-water equations on the sphere, *J. Comput. Phys.*, 229, 6104–6134, 2010.
- Williamson, D. L., Drake, J. B., Hack, J. J., Jakob, R., and Swartztrauber, P. N.: A standard test set for numerical approximations to the shallow water equations in spherical geometry, *J. Comput. Phys.*, 102, 211–224, 1992.
- 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, *Q. J. Roy. Meteorol. Soc.*, 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, *J. Sci. Comput.*, 67, 192–220, 2016.
- Zenodo/Firedrake: Software used in 'Slate: extending Firedrake's domain-specific abstraction to hybridized solvers for geoscience and beyond', Zenodo, <https://doi.org/10.5281/zenodo.2587072>, 2019.
- Zenodo/Tabula-Rasa: Tabula Rasa: experimentation framework for hybridization and static condensation, Zenodo, <https://doi.org/10.5281/zenodo.2616031>, 2019.