



Citation for published version:

Betteridge, J, Gibson, T, Graham, I & Müller, E 2021, 'Multigrid preconditioners for the hybridized Discontinuous Galerkin discretisation of the shallow water equations', *Journal of Computational Physics*, vol. 426, 109948. <https://doi.org/10.1016/j.jcp.2020.109948>

DOI:

[10.1016/j.jcp.2020.109948](https://doi.org/10.1016/j.jcp.2020.109948)

Publication date:

2021

Document Version

Peer reviewed version

[Link to publication](#)

Publisher Rights

CC BY-NC-ND

University of Bath

Alternative formats

If you require this document in an alternative format, please contact:
openaccess@bath.ac.uk

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Multigrid preconditioners for the hybridised Discontinuous Galerkin discretisation of the shallow water equations

Jack Betteridge^{a,b}, Thomas H. Gibson^{b,c}, Ivan G. Graham^a, Eike H. Müller^{a,1,*}

^a*Department of Mathematical Sciences, University of Bath, Claverton Down, Bath BA2 7AY, United Kingdom*

^b*Department of Computing and Department of Mathematics, Imperial College, South Kensington Campus, London SW7 2AZ, United Kingdom*

^c*Naval Postgraduate School, 1 University Circle, Monterey, CA 93943, USA*

Abstract

Numerical climate- and weather-prediction models require the fast solution of the equations of fluid dynamics. Discontinuous Galerkin (DG) discretisations have several advantageous properties. They can be used for arbitrary domains and support a structured data layout, which is particularly important on modern chip architectures. For smooth solutions, higher order approximations can be particularly efficient since errors decrease exponentially in the polynomial degree. Due to the wide separation of timescales in atmospheric dynamics, semi-implicit time integrators are highly efficient, since the implicit treatment of fast waves avoids tight constraints on the time step size, and can therefore improve overall efficiency. However, if implicit-explicit (IMEX) integrators are used, a large linear system of equations has to be solved in every time step. A particular problem for DG discretisations of velocity-pressure systems is that the normal Schur-complement reduction to an elliptic system for the pressure is not possible since the numerical fluxes introduce artificial diffusion terms. For the shallow water equations, which form an important model system, hybridised DG methods have been shown to overcome this issue. However, no attention has been paid to the efficient solution of the resulting linear system of equations. In this paper we address this issue and show that the elliptic system for the flux unknowns can be solved efficiently by using a non-nested multigrid algorithm. The method is implemented in the Firedrake library and we demonstrate the excellent performance of the algorithm both for an idealised stationary flow problem in a flat domain and for non-stationary setups in spherical geometry from the well-known testsuite in [Williamson et al. (1992) JCP, 102(1), pp.211-224]. In the latter case the performance of our bespoke multigrid preconditioner (although itself not highly optimised) is comparable to that of a highly optimised direct solver.

Keywords: multigrid, elliptic PDE, Hybridised Discontinuous Galerkin, preconditioners, atmospheric modelling

1. Introduction

The dynamics of geophysical fluids is characterised by phenomena which occur at very different time scales. While large-scale flow in the atmosphere is described by the slow evolution of quasi-geostrophic modes and non-linearities are weak, fast acoustic waves play an important rôle in coupling phenomena at different scales and rapid adjustment to geostrophic balance. However, acoustic waves carry very little energy and hence do not have to be represented to high accuracy in computer simulations. Efficient numerical models for weather- and climate-prediction often exploit this scale-separation by using semi-implicit time integrators. Treating the fast waves implicitly (and thereby removing any numerical instabilities due to acoustic modes) avoids tight restrictions on the time step size. It is sufficient to chose the timestep such that it allows the stable and accurate representation of the slow modes, which may be treated explicitly. Examples of popular time-integrators are semi-implicit semi-Lagrangian methods used by the UK Met Office [1] and implicit-explicit (IMEX) methods [2, 3, 4, 5] (see also [6] for applications in atmospheric modelling).

*Corresponding author

¹email: e.mueller@bath.ac.uk

In a computer model the equations of fluid-dynamics have to be represented on a grid with finite resolution. While simple finite-difference- or finite-volumes schemes on structured meshes (such as latitude-longitude grids) work well for local area models, they suffer from the convergence of grid lines at the poles on global grids [7]. For a fully explicit time discretisation, this leads to artificially tight constraints on the timestep size due to the CFL condition because of the small grid cells at the poles. Although semi-implicit discretisations are not restricted by this stability constraint and can be run with larger timestep sizes, the linear system for the pressure correction, which has to be solved in every time step, becomes increasingly ill-conditioned and expensive to solve due to large Courant number at the poles. Furthermore, all-to-all communications limit scalability in massively parallel implementations. Finite element methods on quasi-uniform unstructured global grids overcome this problem, but care has to be taken when solving the resulting linear systems which arise in semi-implicit timestepping. The standard approach uses a Schur-complement reduction to a pressure system. Naively, this results in a dense Schur-complement operator in pressure space. For conforming finite element methods this can be traced back to the fact that the mass-matrices are non-diagonal. In this case the problem is easily overcome by under-integrating and co-locating the quadrature points with nodal basis points, or forming an approximate Schur-complement [8].

Discontinuous Galerkin (DG) discretisations have recently been explored as a promising alternative in atmospheric modelling; the review in [9] highlights their excellent parallel scalability and discusses current issues of the approach. DG methods have several advantages, especially on modern chip architectures: since the unknowns are associated with cells of the mesh, data access can be coalesced, which is particularly important for chips with wide vector units. For sufficiently smooth solutions, higher order methods are numerically efficient since errors decrease exponentially as the polynomial degree of the DG approximation increases. Matrix-free implementations are particularly efficient on chips with limited memory bandwidth since they have a very high computational intensity; sum factorisation techniques can further reduce the computational complexity [10]. Unfortunately, for DG methods the numerical flux introduces off-diagonal matrix elements which cause issues for semi-implicit discretisations. The flux generates artificial diffusion terms, which are numerical artifacts and can not be removed in an obvious way. This rules out the naive Schur-complement approach for solving the implicit linear system.

In [11] a solution to this problem is presented for the shallow water equations (SWEs), a popular two-dimensional model system for atmospheric dynamics. The authors use the hybridised Discontinuous Galerkin (HDG) method (see e.g. [12, 13, 14, 15]), which is adapted for the SWEs in [16, 17]. Introducing a set of flux-unknowns on the facets of the mesh allows the reduction to a sparse Schur-complement system in the flux space. In [11] this problem is assembled and solved with a direct method via LU factorisation. While this works for small problems, it is expected that the cost of this direct solver grows rapidly as the problem size increases, even if regularities in the sparsity pattern are exploited (see [18] for a general overview of direct methods). This approach is therefore clearly unfeasible for large scale applications, especially if they are eventually to be used in time critical operational forecasting systems.

The central aim of this paper is to present an alternative solution which overcomes those issues and to demonstrate the computational performance of our approach. We introduce a new non-nested multigrid method for solving the linear systems arising in IMEX time integrators for the SWEs, when discretised using a hybridised (high order) DG method. The method is strongly influenced by the work of Gopalakrishnan and Tan [19] and Cockburn, Dubois, Gopalakrishnan and Tan [20]. In particular [19] introduces a non-nested multigrid V-cycle for the hybridised mixed method for standard second order elliptic PDEs discretised using Raviart Thomas elements for the velocity approximation. This has been extended in [20] to the multigrid solution of the HDG discretisation for the same elliptic PDE. In this paper we adapt this approach to solve the linear and non-linear shallow water equations. The non-linear SWEs present additional challenges since they require the construction of a Lax-Friedrich flux, which results in a vector-valued flux space [17] for which the method in [20] is not directly applicable.

More specifically, the new contributions of this paper are:

1. The non-nested multigrid method of [20] is extended to the HDG solution of the linearised SWEs with IMEX methods.
2. To solve the non-linear SWEs by this approach we construct novel intergrid- and coarse-level operators for vector-valued flux spaces.
3. Our HDG solvers and multigrid algorithms are implemented in the Slate language [21], which forms

part of the open source Firedrake library [22] for the solution of finite element problems based on composable abstractions.

4. Based on this, the superior performance of the non-nested multigrid approach (compared to a naive Schur-complement preconditioner) for both the linear- and non-linear SWEs is demonstrated for a simplified stationary test flow in a flat domain.
5. For more complex non-stationary flow problems from the standard Williamson et al. test suite [23] the non-nested multigrid preconditioner is compared to a highly-optimised direct solver from the MUMPS package [24, 25].

Our approach differs from the more algebraic method presented in [26, 27], which constructs a multigrid hierarchy directly on the trace-space defined on the skeleton of the mesh; this is achieved by using hp -refinement and statically condensing out unknowns on fine edges which are not represented on the coarser mesh.

We stress that the aim of this work is to quantify the advantages of replacing the direct solver (such as the one used in [11]) with our new multigrid-method. Our experiments on the stationary flow problem indicate substantial advantages of our new multigrid method. While for flows in a spherical geometry both considered solvers show similar performance on a single compute node, it is highly likely that an iterative solver with a multigrid preconditioner will be superior to a direct method when solving very large problems on massively parallel machines.

In the long run it would be important to properly compare the HDG method to other discretisation schemes such as finite volume. However a proper treatment of this would need a detailed assessment of both accuracy and efficiency of implementations and solvers. While we feel this is worth doing, it is well beyond the scope of the present paper.

Structure. This paper is organised as follows: In Section 2 we briefly review the physics of the shallow water equations. The hybridised DG discretisation is described in Section 3 and the non-nested multigrid preconditioner for the resulting linear systems in IMEX timestepping is presented in Section 4. The specific integrators considered in this paper are introduced in Section 5 with focus on the linear solvers used for implicit methods. The implementation in the Firedrake library is described in Section 6 and results are presented in Section 7. We finally conclude and outline directions for future work in Section 8. Some more technical issues such as the explicit expressions for key bilinear forms are relegated to the appendices.

2. Shallow water equations

The shallow water equations (SWEs) describe the evolution of a layer of fluid under gravitational and inertial forces. In a rotating frame a Coriolis term has to be included as well. The SWEs can be written in conservative form for the geopotential height $\phi(\mathbf{x}, t)$ and the column-integrated momentum $\mathbf{u}(\mathbf{x}, t)$, which are defined for all points $\mathbf{x} \in \Omega$ in a bounded domain $\Omega \subseteq \mathbb{R}^2$. The time evolution of ϕ and $\mathbf{u} = (u, v)$ is described by the following two equations:

$$\begin{aligned} \frac{\partial \phi}{\partial t} + c_g \nabla \cdot \mathbf{u} &= 0 && \text{(mass conservation)} \\ \frac{\partial \mathbf{u}}{\partial t} + c_g \nabla \cdot \left(\frac{\mathbf{u} \otimes \mathbf{u}}{\phi + \phi_B} + \frac{1}{2} (\phi_B + \phi)^2 \right) &= \mathbf{R} := c_g (\phi_B + \phi) \nabla \phi_B - f \mathbf{u}^\perp && \text{(momentum conservation)} \end{aligned} \tag{1}$$

Suitable boundary conditions have to be included at the boundary of Ω . The geopotential height is defined as $\phi = gH$ where H denotes the perturbation of the fluid relative to its state at rest and $g = 9.81 \text{ m s}^{-2}$ is the gravitational acceleration. Furthermore $\phi_B = gH_B$ where H_B is the bathymetry of the sea-floor, measured from the surface at rest (see Fig. 1). The momentum $\mathbf{u}(\mathbf{x}, t)$ is related to the column-averaged velocity of the fluid $\mathbf{U}(\mathbf{x}, t)$ at each point $\mathbf{x} \in \Omega$ by $\mathbf{u} = (\phi_B + \phi)\mathbf{U}$. $f = f(\mathbf{x}) \in \mathbb{R}$ is the latitude-dependent Coriolis parameter, and the vector $\mathbf{u}^\perp = (-v, u)$ is obtained by rotating \mathbf{u} by 90 degrees in the counter-clockwise direction.

Eq. (1) has been non-dimensionalised as follows: let \bar{H}_B be a reference height, and define $\bar{\phi}_B = g\bar{H}_B$. Further assume that any horizontal distances are measured in units of some reference length scale \bar{L} and

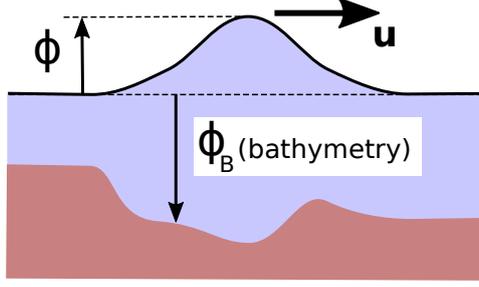


Figure 1: Fields used in the formulation of the shallow water equations

times in units of a reference time scale \bar{T} . ϕ , ϕ_B are given in units of $\bar{\phi}_B$ and the momentum \mathbf{u} in units of $\bar{\phi}_B^{3/2}$. The dimensionless reference gravity wave speed which multiplies the flux terms in Eq. (1) is then given by

$$c_g = \frac{\bar{T}}{\bar{L}} \bar{\phi}_B^{-1/2}.$$

For convenience, we introduce the variable $q = (\phi, \mathbf{u}) = (\phi, u, v)$, which allows us to write Eq. (1) in compact form as

$$\frac{\partial q}{\partial t} + \nabla \cdot \mathcal{F}(q) = s(q), \quad \text{or equivalently} \quad \frac{\partial q_i}{\partial t} + \partial_j \mathcal{F}_{ij}(q) = s_i(q). \quad (2)$$

Here and in the following summation over repeated indices is implicitly assumed. The exact form of $\mathcal{F} : \mathbb{R}^3 \rightarrow \mathbb{R}^{3 \times 2}$ and $s : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ can be deduced from Eq. (1) and is given by (see [11])

$$\mathcal{F}(q) = \begin{pmatrix} u & v \\ \frac{u^2}{\phi + \phi_B} + \phi_B \phi + \frac{1}{2} \phi^2 & \frac{uv}{\phi + \phi_B} \\ \frac{uv}{\phi + \phi_B} & \frac{v^2}{\phi + \phi_B} + \phi_B \phi + \frac{1}{2} \phi^2 \end{pmatrix}, \quad s(q) = \begin{pmatrix} 0 \\ c_g \phi \partial_1 \phi_B + f v \\ c_g \phi \partial_2 \phi_B - f u \end{pmatrix}.$$

The SWEs can be linearised about 0 (this is a good approximation if $\phi, |\mathbf{u}| \ll \phi_B$ at every point in the domain) to obtain

$$\frac{\partial \phi}{\partial t} + c_g \nabla \cdot \mathbf{u} = 0, \quad \frac{\partial \mathbf{u}}{\partial t} + c_g \phi_B \nabla \phi = -f \mathbf{u}^\perp. \quad (3)$$

It is important to note that in a non-rotating frame ($f = 0$) fast gravity waves are supported by Eq. (3). This can be seen by taking the divergence of the second equation and inserting it into the time derivative of the first, to obtain the wave equation

$$\frac{\partial^2 \phi}{\partial t^2} - c_g^2 \nabla \cdot (\phi_B \nabla \phi) = 0.$$

The (local) speed of the gravity waves is given by $c = c_g \sqrt{\phi_B}$. More generally, in a rotating frame the fundamental solutions of the linear SWEs in Eq. (3) are inertia-gravity (Poincaré) waves. In this case the frequency ω is related to the wave vector \mathbf{k} by the dispersion relation $\omega = \pm \sqrt{f^2 + c_g^2 \phi_B \mathbf{k}^2}$. In analogy to Eq. (2), the linearised SWEs in Eq. (3) can be written as

$$\frac{\partial q}{\partial t} + \nabla \cdot \mathcal{F}_L(q) = s(q) \quad \Leftrightarrow \quad \frac{\partial q_i}{\partial t} + \partial_j (\mathcal{F}_L)_{ij}(q) = s_i(q) \quad (4)$$

with $q = (\phi, u, v)$ and the linear operator $\mathcal{F}_L : \mathbb{R}^3 \rightarrow \mathbb{R}^{3 \times 2}$ given by

$$\mathcal{F}_L(q) = c_g \begin{pmatrix} u & v \\ \phi_B \phi & 0 \\ 0 & \phi_B \phi \end{pmatrix}.$$

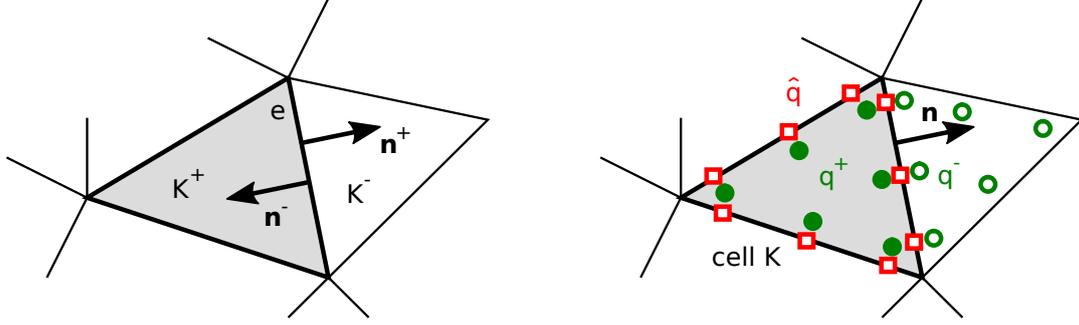


Figure 2: Two cells K^+ , K^- and their respective outer normal vectors \mathbf{n}^+ , \mathbf{n}^- on the joint edge $e = \partial K^+ \cap \partial K^-$ (left). Degrees of freedom of the hybridised DG discretisation (right).

3. IMEX HDG methods for the shallow water equations

3.1. Spatial discretisation

To discretise the SWEs with the discontinuous Galerkin (DG) method, we first construct a mesh Ω_h of non-overlapping cells K that partition the domain such that $\Omega = \cup_{K \in \Omega_h} K$. The mesh spacing is defined as $h = \max_{K \in \Omega_h} \text{diam}(K)$. Let $\mathcal{E}_h = \cup_{K, K' \in \Omega_h} (K \cap K')$ be the set of facets of the mesh; \mathcal{E}_h is often called the “*skeleton of Ω_h* ”. We further assume that the mesh consists of triangular or quadrilateral cells and has no hanging nodes. More specifically, each facet $e \in \mathcal{E}_h$ is a straight line which forms the side of exactly two triangles or quadrilaterals. For simplicity we use periodic boundary conditions in this paper, which implies that all facets are internal.

By construction, functions defined in the DG space introduced below are not continuous across cell-boundaries. As a result, special care has to be taken when defining quantities on the mesh skeleton \mathcal{E}_h . For two neighbouring cells K^+ and K^- that share a facet $e = K^+ \cap K^-$, let \mathbf{n}^- be the outward unit normal on the boundary ∂K^- of element K^- , and \mathbf{n}^+ the corresponding outward normal of cell K^+ (see Fig. 2), such that $\mathbf{n}^+ = -\mathbf{n}^-$ on the shared facet e . On each facet $e \in \mathcal{E}_h$, the following average $\{\!\{ \cdot \}\!\}$, difference $\langle\!\langle \cdot \rangle\!\rangle$ and jump $[\![\cdot]\!]$ operators can be defined for scalar (s) and vector valued (\mathbf{a}) quantities:

$$\begin{aligned} \{\!\{ s \}\!\} &= \frac{1}{2} (s^+ + s^-), & \{\!\{ \mathbf{a} \}\!\} &= \frac{1}{2} (\mathbf{a}^+ + \mathbf{a}^-), \\ \langle\!\langle s \rangle\!\rangle &= s^+ - s^-, & \langle\!\langle \mathbf{a} \rangle\!\rangle &= \mathbf{a}^+ - \mathbf{a}^-, \\ [s] &= s^+ \mathbf{n}^+ + s^- \mathbf{n}^-, & [\![\mathbf{a}]\!] &= \mathbf{a}^+ \cdot \mathbf{n}^+ + \mathbf{a}^- \cdot \mathbf{n}^-. \end{aligned}$$

At any point $\mathbf{x} \in e$ on the facet e and for any scalar- or vector-valued piecewise polynomial functions $s : \Omega \rightarrow \mathbb{R}$, $\mathbf{a} : \Omega \rightarrow \mathbb{R}^2$ we define $s^\pm(\mathbf{x}) = \lim_{\epsilon \rightarrow 0^+} s(\mathbf{x} \mp \epsilon \mathbf{n}^\pm)$ and $\mathbf{a}^\pm(\mathbf{x}) = \lim_{\epsilon \rightarrow 0^+} \mathbf{a}(\mathbf{x} \mp \epsilon \mathbf{n}^\pm)$. Note that the jump $[s]$ of a scalar quantity s is vector-valued and vice versa.

We write $\mathcal{P}^p(D)$ for the space of polynomials of degree at most p on an arbitrary domain D . The (scalar-valued) DG space of order p on the mesh Ω_h is defined as

$$V_h := \{ \psi \in L^2(\Omega_h) : \psi|_K \in \mathcal{P}^p(K), \forall K \in \Omega_h \}. \quad (5)$$

Case 1: Linear SWEs. It is instructive to first discuss the DG discretisation of the linear SWEs in Eq. (4). This can be found by assuming that² $q = (\phi, u, v) \in V_h \times V_h \times V_h := W_h$, then multiplying by a test function $v \in W_h$ and integrating by parts in each cell K to obtain

$$\left(\frac{\partial q_i}{\partial t}, v_i \right)_K - ((\mathcal{F}_L)_{ij}(q), \partial_j v_i)_K + ((\mathcal{F}_L^*)_{ij}(q^+, q^-) n_j, v_i)_{\partial K} = (s_i(q), v_i)_K \quad \text{for all } v \in W_h. \quad (6)$$

Here \mathbf{n} is the outward normal on each facet of cell K , with $q^\pm(\mathbf{x}) = \lim_{\epsilon \rightarrow 0^+} q(\mathbf{x} \mp \epsilon \mathbf{n})$. Round and angular brackets denote L_2 scalar products in function spaces defined on the cell K and its boundary

²In a slight abuse of notation, we use the same variable names for the continuous fields q, u, v, ϕ etc. in Section 2 and for the DG fields discussed from this section onwards.

∂K respectively:

$$(f, g)_K = (f, g)_{L_2(K)} = \int_K f(\mathbf{x})g(\mathbf{x}) dx, \quad \langle f, g \rangle_{\partial K} = (f, g)_{L_2(\partial K)} = \int_{\partial K} f(\mathbf{x})g(\mathbf{x}) ds.$$

The numerical flux $\mathcal{F}_L^*(q^+, q^-)$ appearing in Eq. (6) (still to be defined) will depend on the values q^+, q^- of the field on both sides of the facets comprising ∂K and will couple the unknowns between neighbouring cells. For the linear problem both the local Lax-Friedrichs (Rusanov) flux $\mathcal{F}_L^{*(\text{LF})}$ [28] and the upwind flux $\mathcal{F}_L^{*(\text{up})}$ (see e.g. [29]) are candidates for the numerical flux \mathcal{F}_L^* . As shown in [16, 17], those are defined as

$$(\mathcal{F}_L^{*(\text{LF})})_{ij}(q^+, q^-)n_j = \{(\mathcal{F}_L)_{ij}(q)\}n_j + \frac{c_g}{2}\sqrt{\phi_B}\langle\langle q_i \rangle\rangle \quad (7a)$$

$$(\mathcal{F}_L^{*(\text{up})})_{ij}(q^+, q^-)n_j = \{(\mathcal{F}_L)_{ij}(q)\}n_j + \frac{c_g}{2}\sqrt{\phi_B}B_{ik}\langle\langle q_k \rangle\rangle \quad (7b)$$

where the 3×3 matrix B defined on ∂K is given by

$$B = \begin{pmatrix} 1 & 0 & 0 \\ 0 & n_x^2 & n_x n_y \\ 0 & n_x n_y & n_y^2 \end{pmatrix}.$$

If Eq. (6) is integrated in time with a semi-implicit method, a coupled sparse linear system for the fields (ϕ, \mathbf{u}) has to be solved in each timestep. The standard approach in many atmospheric models is to eliminate the momentum unknowns \mathbf{u} from the *discrete* equations and solve the resulting elliptic Schur-complement system for the height perturbation ϕ , followed by reconstruction of the momentum field (see e.g. [1] for the equivalent procedure applied to the full Navier Stokes equations). Unfortunately, for the DG discretisation considered here, the numerical flux in Eq. (6) introduces artificial diffusion terms in the momentum equation. Although those terms disappear in the limit $h \rightarrow 0$, they result in a dense Schur-complement operator which renders this naive Schur-complement approach unfeasible here. It is possible to address this issue by forming an approximate Schur-complement and using the resulting approximate solver as a preconditioner for a Krylov subspace method. A similar approach is pursued to deal with the non-diagonal velocity mass matrix in conforming mixed finite element discretisations in [8, 30]. However, this requires additional iterations over the large (ϕ, \mathbf{u}) -system. As our numerical results in Section 7 demonstrate, this approximate Schur-complement method (which essentially drops the artificial diffusion terms in the preconditioner) is not competitive with the hybridised DG method pursued in this paper since the latter allows the formulation of an exact Schur-complement in a suitably defined finite element space on the facets.

3.2. Hybridised DG discretisation

As described in [11, section 2], IMEX time integrators separate the operator which describes the time evolution of the shallow water equations into a linear part, which is treated implicitly, and an explicit non-linear part. By using a *hybridised* DG discretisation for the linear part it is possible to use the efficient non-nested multigrid techniques which are discussed in this paper to solve the implicit linear problem. The explicit part can be discretised with a standard DG method.

The key idea of the hybridised DG method described in [16, 17] is to introduce a new variable \hat{q} , which is associated with the skeleton \mathcal{E}_h of the mesh, see Fig. 2 (right). As we will show below, it is possible to form an exact Schur-complement in \hat{q} . For this, recall that the facets $e \in \mathcal{E}_h$ are straight lines bounding the cells. On these we define the piecewise polynomial space

$$\Lambda_h = \{\lambda \in L^2(\mathcal{E}_h) : \lambda|_e \in \mathcal{P}^p(e), \forall e \in \mathcal{E}_h\} \quad \text{and set} \quad \hat{\Lambda}_h = \Lambda_h \times \Lambda_h \times \Lambda_h. \quad (8)$$

This allows the introduction of a flux $(\hat{\mathcal{F}}_L)_{ij}(q, \hat{q})$ which only depends on the value q inside the cell K and the field \hat{q} on its boundary. To see this, recall that for hyperbolic conservation laws the numerical upwind flux can be constructed by solving a Riemann problem on the interface between two cells (see e.g. section 2.4 of [31]). The solution contains both the field q inside the cell and an intermediate state q^* , which is usually eliminated with a jump condition. Since this jump condition involves the value of fields on both sides of the interface, the resulting flux in Eq. (7b) depends on q^+ and q^- . As argued

in [16, 17], the HDG method represents the intermediate state as $\hat{q} = q^*$ in the space $\hat{\Lambda}_h$ instead of eliminating q^* ; the jump condition is enforced weakly.

$\hat{\mathcal{F}}_L$ replaces the numerical flux \mathcal{F}_L^* in Eq. (6) and is given by either $\hat{\mathcal{F}}_L^{(\text{LF})}$ (Lax-Friedrichs) or $\hat{\mathcal{F}}_L^{(\text{up})}$ (upwind) defined on the boundary ∂K of each cell as

$$(\hat{\mathcal{F}}_L^{(\text{LF})})_{ij}(q, \hat{q})n_j = (\mathcal{F}_L)_{ij}(q)n_j + c_g \sqrt{\phi_B}(q_i - \hat{q}_i) \quad (9a)$$

$$(\hat{\mathcal{F}}_L^{(\text{up})})_{ij}(q, \hat{q})n_j = (\mathcal{F}_L)_{ij}(q)n_j + c_g \sqrt{\phi_B} B_{ik}(q_k - \hat{q}_k). \quad (9b)$$

To close the system of equations, continuity of the flux is enforced weakly by requiring that on each facet $e \in \mathcal{E}_h$

$$\langle \{(\hat{\mathcal{F}}_L)_{ij}(q, \hat{q})n_j\}, \hat{v}_i \rangle_e = 0 \quad \text{for all test functions } \hat{v} \in \hat{\Lambda}_h. \quad (10)$$

To simplify notation, introduce the following abbreviations:

$$\begin{aligned} \mathcal{L}(q, v) &= \sum_{K \in \Omega_h} \{((\mathcal{F}_L)_{ij}(q), \partial_j v_i)_K - \langle (\mathcal{F}_L^*)_{ij}(q^+, q^-)n_j, v_i \rangle_{\partial K}\} \\ \hat{\mathcal{L}}(q, \hat{q}, v) &= \sum_{K \in \Omega_h} \{((\mathcal{F}_L)_{ij}(q), \partial_j v_i)_K - \langle (\hat{\mathcal{F}}_L)_{ij}(q, \hat{q})n_j, v_i \rangle_{\partial K}\} \\ \mathcal{N}_0(q, v) &= \sum_{K \in \Omega_h} (s_i(q), v_i)_K \quad (\text{see Eq. (6)}) \\ \mathcal{M}(q, v) &= \sum_{K \in \Omega_h} (q_i, v_i)_K \\ \Xi(q, \hat{q}, \hat{v}) &= \sum_{e \in \mathcal{E}_h} \langle \{(\hat{\mathcal{F}}_L)_{ij}(q, \hat{q})n_j\}, \hat{v}_i \rangle_e \end{aligned} \quad (11)$$

With the definitions in Eq. (11), after summing over all cells $K \in \Omega_h$ the weak form of the linear SWEs discretised with DG in Eq. (6) can be written in condensed form as

$$\frac{\partial \mathcal{M}(q, v)}{\partial t} = \mathcal{N}_0(q, v) + \mathcal{L}(q, v) \quad \text{for all } v \in W_h. \quad (12)$$

The equivalent HDG discretisation of Eq. (6) is obtained by replacing \mathcal{F}_L^* by $\hat{\mathcal{F}}_L$ in Eq. (12) and enforcing the continuity condition in Eq. (10). This results in

$$\frac{\partial \mathcal{M}(q, v)}{\partial t} = \mathcal{N}_0(q, v) + \hat{\mathcal{L}}(q, \hat{q}, v), \quad \text{subject to} \quad \Xi(q, \hat{q}, \hat{v}) = 0 \quad \text{for all } v \in W_h, \hat{v} \in \hat{\Lambda}_h. \quad (13)$$

As has been shown in [17], the solution q of the original, non-hybridized system in Eq. (12) is identical to the solution q of the HDG problem given in Eq. (13). In general, the additional field $\hat{q} = (\hat{\phi}, \hat{\mathbf{u}}) = (\hat{\phi}, \hat{u}, \hat{v}) \in \hat{\Lambda}_h$ introduced in Eq. (13) on the facets has three components. As will be discussed in Section 4.1 below, $\hat{\mathbf{u}}$ can be eliminated for the upwind flux, whereas $\hat{\phi}$ does not enter the equations if the Lax-Friedrichs method is used.

Case 2: non-linear SWEs. Similarly, the weak form of the non-linear SWEs in Eq. (2) can be written for each cell $K \in \Omega_h$ as

$$\left(\frac{\partial q_i}{\partial t}, v_i \right)_K - (\mathcal{F}_{ij}(q), \partial_j v_i)_K + \langle \mathcal{F}_{ij}^*(q^+, q^-)n_j, v_i \rangle_{\partial K} = (s_i(q), v_i)_K \quad \text{for all } v \in W_h. \quad (14)$$

where $\mathcal{F}^* = \mathcal{F}^{(\text{LF})}$ is the non-linear Lax-Friedrichs flux

$$\mathcal{F}_{ij}^{(\text{LF})}(q^+, q^-)n_j = \{\{\mathcal{F}_{ij}(q)\}\}n_j + \frac{c_g}{2} \tau^* \langle \{q_i\} \rangle \quad (15)$$

with $\tau^* = \max \left\{ |\mathbf{n} \cdot \mathbf{u}^+| + \sqrt{\phi_B + \phi^+}, |\mathbf{n} \cdot \mathbf{u}^-| + \sqrt{\phi_B + \phi^-} \right\}$. Further define

$$\mathcal{N}(q, v) = \sum_{K \in \Omega_h} \{((\mathcal{F} - \mathcal{F}_L)_{ij}(q), \partial_j v_i)_K - \langle (\mathcal{F}^* - \mathcal{F}_L^*)_{ij}(q^+, q^-)n_j, v_i \rangle_{\partial K} + (s_i(q), v_i)_K\} \quad (16)$$

With the definitions in Eqs. (11) and (16), after summing over all cells $K \in \Omega_h$ the weak form of the DG discretisation in Eq. (14) can be written in condensed form as

$$\frac{\partial \mathcal{M}(q, v)}{\partial t} = \mathcal{N}(q, v) + \mathcal{L}(q, v) \quad \text{for all } v \in W_h. \quad (17)$$

As above, the equivalent HDG discretisation of Eq. (14) is obtained by replacing \mathcal{F}_L^* by $\widehat{\mathcal{F}}_L$ in the linear term \mathcal{L} in Eq. (17) and enforcing the continuity condition in Eq. (10). This results in

$$\frac{\partial \mathcal{M}(q, v)}{\partial t} = \mathcal{N}(q, v) + \widehat{\mathcal{L}}(q, \widehat{q}, v), \quad \text{subject to } \Xi(q, \widehat{q}, \widehat{v}) = 0 \quad \text{for all } v \in W_h, \widehat{v} \in \widehat{\Lambda}_h. \quad (18)$$

By splitting the right hand side into two terms, we have isolated the fast dynamics due to gravity waves in $\widehat{\mathcal{L}}(q, \widehat{q}, v)$, whereas any slower modes are described by $\mathcal{N}(q, v)$. This observation is crucial for the construction of suitable semi-implicit time integrators, which avoid numerical instabilities due to the fast gravity waves.

3.3. Time discretisation

To integrate Eq. (18) (Eq. (13) can be dealt with in exactly the same way) in time, the term $\mathcal{N}(q, v)$ is integrated explicitly, whereas $\widehat{\mathcal{L}}(q, \widehat{q}, v)$ is treated implicitly. Denoting the solution at time $t = n\Delta t$ by $q^{(n)}$, a simple scheme would for example be the following “ θ -method”:

$$\frac{\mathcal{M}(q^{(n+1)}, v) - \mathcal{M}(q^{(n)}, v)}{\Delta t} = \mathcal{N}(q^{(n)}, v) + \theta \widehat{\mathcal{L}}(q^{(n+1)}, \widehat{q}^{(n+1)}, v) + (1 - \theta) \mathcal{L}(q^{(n)}, v) \quad \text{for all } v \in W_h \quad (19)$$

with $\theta \in [0, 1]$ and subject to the flux condition Eq. (10). The equivalent expression for the linear SWEs in Eq. (13) can be obtained by replacing $\mathcal{N} \rightarrow \mathcal{N}_0$ in Eq. (19). At each time step, calculating the field $q^{(n+1)}$ requires the solution of the *linear* system

$$\begin{aligned} \mathcal{M}(q^{(n+1)}, v) - \theta \Delta t \widehat{\mathcal{L}}(q^{(n+1)}, \widehat{q}^{(n+1)}, v) &= \mathcal{R}(q^{(n)}, \widehat{q}^{(n)}, v) \\ &:= \mathcal{M}(q^{(n)}, v) + \Delta t \left(\mathcal{N}(q^{(n)}, v) + (1 - \theta) \mathcal{L}(q^{(n)}, v) \right) \quad (20) \end{aligned}$$

subject to $\Xi(q^{(n+1)}, \widehat{q}^{(n+1)}, \widehat{v}) = 0 \quad \text{for all } v \in W_h, \widehat{v} \in \widehat{\Lambda}_h.$

Note that for a purely explicit method ($\theta = 0$), the system matrix arising from the left hand side of the system in Eq. (20) is simply the mass matrix (obviously, in this case Eq. (19) reduces to the explicit Euler integrator). On the other hand, s -stage IMEX methods [2, 3, 4, 5, 6] are a generalisation of the method in Eq. (19). They require the construction of s intermediate states $\{(Q^{(i)}, \widehat{Q}^{(i)})\}_{i=1}^s$ which are obtained by solving the system of linear equations

$$\mathcal{M}(Q^{(i)}, v) = \mathcal{M}(q^{(n)}, v) + \Delta t \sum_{j=1}^{i-1} a_{ij} \mathcal{N}(Q^{(j)}, v) + \Delta t \sum_{j=1}^i \tilde{a}_{ij} \widehat{\mathcal{L}}(Q^{(j)}, \widehat{Q}^{(j)}, v), \quad (21)$$

$$\text{with } \Xi(Q^{(i)}, \widehat{Q}^{(i)}, \widehat{v}) = 0 \quad \text{for each stage } i = 1, \dots, s \text{ and all } v \in W_h, \widehat{v} \in \widehat{\Lambda}_h.$$

The state $q^{(n+1)}$ at the next timestep satisfies

$$\mathcal{M}(q^{(n+1)}, v) = \mathcal{M}(q^{(n)}, v) + \Delta t \sum_{i=1}^s b_i \mathcal{N}(Q^{(i)}, v) + \Delta t \sum_{i=1}^s \tilde{b}_i \widehat{\mathcal{L}}(Q^{(i)}, \widehat{Q}^{(i)}, v). \quad (22)$$

and is obtained by an additional mass-solve. Each particular IMEX method is determined by a choice of coefficients a_{ij} , \tilde{a}_{ij} , b_i , \tilde{b}_i which are commonly known as the Butcher tableau coefficients for the explicit and implicit terms. The equations in Eq. (21) can be solved stage-by-stage. At every stage a *linear* system of the abstract general form

$$\widehat{\mathcal{A}}(q, \widehat{q}, v, \widehat{v}) = \mathcal{M}(q, v) - \alpha \Delta t \widehat{\mathcal{L}}(q, \widehat{q}, v) + \Xi(q, \widehat{q}, \widehat{v}) = \mathcal{R}(v) \quad \text{for all } v \in W_h, \widehat{v} \in \widehat{\Lambda}_h. \quad (23)$$

with some positive constant $\alpha \in \mathbb{R}^+$ has to be solved³. Typically the value of α is different in each stage of the IMEX method and the right hand side \mathcal{R} depends on previously calculated fields. To see that $\hat{\mathcal{A}}$ in Eq. (23) is indeed linear in the combined field $(q, \hat{q}) \in W_h \times \hat{\Lambda}_h$, use the definitions of \mathcal{M} , $\hat{\mathcal{L}}$ and Ξ in Eqs. (9a), (9b), (11) and observe that

$$\hat{\mathcal{A}}(c_1 q_1 + c_2 q_2, c_1 \hat{q}_1 + c_2 \hat{q}_2, v, \hat{v}) = c_1 \hat{\mathcal{A}}(q_1, \hat{q}_1, v, \hat{v}) + c_2 \hat{\mathcal{A}}(q_2, \hat{q}_2, v, \hat{v}) \quad \text{for all } v \in W_h, \hat{v} \in \hat{\Lambda}_h$$

and for all $c_1, c_2 \in \mathbb{R}$, $(q_1, \hat{q}_1), (q_2, \hat{q}_2) \in W_h \times \hat{\Lambda}_h$.

4. Schur-complement multigrid solver

Constructing efficient solvers for Eq. (23) is the topic of this paper. In the following section we give explicit expressions for $\hat{\mathcal{A}}$ for the Lax-Friedrichs and upwind flux. We then describe a Schur complement approach which reduces the problem to solving an elliptic system in the flux variable \hat{q} on the facets. This flux system is solved with a new non-nested multigrid preconditioner based on ideas in [20].

4.1. Linear equation

Recall that $q = (\phi, \mathbf{u}) \in W_h$ in the domain, $\hat{q} = (\hat{\phi}, \hat{\mathbf{u}}) \in \hat{\Lambda}_h$ on the facets, and write the corresponding test-functions as $v = (\psi, \mathbf{w}) \in W_h$, $\hat{v} = (\hat{\psi}, \hat{\mathbf{w}}) \in \hat{\Lambda}_h$. Using the definitions in Eqs. (7a) and (7b), the bilinear form $\hat{\mathcal{A}}(q, \hat{q}, v, \hat{v})$ on the left hand side of Eq. (23) can be written down explicitly for the upwind- and Lax-Friedrichs flux. As shown in [17], the resulting bilinear form $\mathcal{A}^{(\text{up})}$ for the upwind flux does not depend on $\hat{\mathbf{u}}$. For the Lax-Friedrichs flux the form $\mathcal{A}^{(\text{LF})}$ does not depend on $\hat{\phi}$. Explicit expressions for those forms are given as follows:

Upwind flux:

$$\begin{aligned} \hat{\mathcal{A}}^{(\text{up})}(q, \hat{q}, v, \hat{v}) &= \hat{\mathcal{A}}^{(\text{up})}(\phi, \mathbf{u}, \hat{\phi}, \psi, \mathbf{w}, \hat{\psi}) := (\phi\psi + \mathbf{u} \cdot \mathbf{w})_{\Omega_h} - c_g \alpha \Delta t \left[(\mathbf{u} \cdot \nabla \psi + \phi_B \phi \nabla \cdot \mathbf{w})_{\Omega_h} \right. \\ &\quad \left. - \left([\mathbf{u}\psi] + 2\sqrt{\phi_B} \left(\{\phi\psi\} - \hat{\phi}\{\psi\} \right) + \phi_B \hat{\phi} [\mathbf{w}] \right)_{\mathcal{E}_h} \right] \\ &\quad + \left(\hat{\psi} \left[[\mathbf{u}] + 2\sqrt{\phi_B} \left(\{\phi\} - \hat{\phi} \right) \right] \right)_{\mathcal{E}_h} \end{aligned} \quad (24)$$

Lax-Friedrichs flux:

$$\begin{aligned} \hat{\mathcal{A}}^{(\text{LF})}(q, \hat{q}, v, \hat{v}) &= \hat{\mathcal{A}}^{(\text{LF})}(\phi, \mathbf{u}, \hat{\mathbf{u}}, \psi, \mathbf{w}, \hat{\mathbf{w}}) := (\phi\psi + \mathbf{u} \cdot \mathbf{w})_{\Omega_h} - c_g \alpha \Delta t \left[(\mathbf{u} \cdot \nabla \psi + \phi_B \phi \nabla \cdot \mathbf{w})_{\Omega_h} \right. \\ &\quad \left. - \left(\hat{\mathbf{u}}[\psi] + 2\sqrt{\phi_B} (\{\mathbf{u} \cdot \mathbf{w}\} - \hat{\mathbf{u}} \cdot \{\mathbf{w}\}) + \phi_B [\phi \mathbf{w}] \right)_{\mathcal{E}_h} \right] \\ &\quad + \left(\hat{\mathbf{w}} \cdot \left[\phi_B [\phi] + 2\sqrt{\phi_B} (\{\mathbf{u}\} - \hat{\mathbf{u}}) \right] \right)_{\mathcal{E}_h} \end{aligned} \quad (25)$$

The system in Eq. (23) can be solved with a Schur-complement approach by eliminating the field q defined on the elements and leaving a problem for \hat{q} defined on the facets. For this write

$$\hat{\mathcal{A}}(q, \hat{q}, v, \hat{v}) = (\mathfrak{A}q, v)_{\Omega_h} + (\hat{q}, \mathfrak{B}^\top v)_{\mathcal{E}_h} + (\mathfrak{C}q, \hat{v})_{\mathcal{E}_h} - (\mathfrak{M}\hat{q}, \hat{v})_{\mathcal{E}_h} = \mathcal{R}(v) = (f, v)_{\Omega_h} \quad (26)$$

where $f \in W_h$ is the Riesz representer of $\mathcal{R}(v)$. The scalar products in Eq. (26) are defined as

$$\begin{aligned} (q, v)_{\Omega_h} &= (\phi\psi + \mathbf{u} \cdot \mathbf{w})_{\Omega_h} := \int_{\Omega} (\phi(\mathbf{x})\psi(\mathbf{x}) + \mathbf{u}(\mathbf{x}) \cdot \mathbf{w}(\mathbf{x})) \, dx, \\ (\hat{q}, \hat{v})_{\mathcal{E}_h} &= \left(\hat{\phi}\hat{\psi} + \hat{\mathbf{u}} \cdot \hat{\mathbf{w}} \right)_{\mathcal{E}_h} := \int_{\mathcal{E}_h} \left(\hat{\phi}(\mathbf{x})\hat{\psi}(\mathbf{x}) + \hat{\mathbf{u}}(\mathbf{x}) \cdot \hat{\mathbf{w}}(\mathbf{x}) \right) \, ds \end{aligned}$$

³Note that the “ θ -method” in Eq. (19) could be written as a 2-stage IMEX method with

$$a = \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix} \quad \tilde{a} = \begin{pmatrix} 0 & 0 \\ 1 - \theta & \theta \end{pmatrix} \quad b = \begin{pmatrix} 1 \\ 0 \end{pmatrix} \quad \tilde{b} = \begin{pmatrix} 1 - \theta \\ \theta \end{pmatrix}$$

but this would be inefficient since it requires an additional mass-solve in the first stage.

and the linear operators

$$\mathfrak{A} : W_h \rightarrow W_h, \quad \mathfrak{B} : \widehat{\Lambda}_h \rightarrow W_h, \quad \mathfrak{C} : W_h \rightarrow \widehat{\Lambda}_h, \quad \mathfrak{M} : \widehat{\Lambda}_h \rightarrow \widehat{\Lambda}_h. \quad (27)$$

Explicit expressions for the operators in Eq. (27) can be deduced from Eqs. (25) and (24) and are given both for the Lax-Friedrichs- and for the upwind flux in Appendix B.

4.2. Schur complement solve

Let the inverse $\mathfrak{F} : W_h \rightarrow W_h$ of \mathfrak{A} be defined in the obvious way as

$$(\mathfrak{F}\mathfrak{A}q, v)_{\Omega_h} = (q, v)_{\Omega_h} \quad \text{for all } q, v \in W_h. \quad (28)$$

It can now be shown that Eq. (23) can be solved in three steps.

Step 1. Given f in Eq. (26), calculate the modified right hand side

$$\widehat{\mathcal{R}}(\widehat{v}) = (\mathfrak{C}\mathfrak{F}f, \widehat{v})_{\varepsilon_h} \quad \text{for all } \widehat{v} \in \widehat{\Lambda}_h.$$

Step 2. Solve the Schur-complement system

$$\widehat{\mathcal{S}}(\widehat{q}, \widehat{v}) = (\mathfrak{M}\widehat{q}, \widehat{v})_{\varepsilon_h} + (\mathfrak{C}\mathfrak{F}\mathfrak{B}\widehat{q}, \widehat{v})_{\varepsilon_h} = \widehat{\mathcal{R}}(\widehat{v}) \quad \text{for all } \widehat{v} \in \widehat{\Lambda}_h \quad (29)$$

for the flux unknown \widehat{q} .

Step 3. Reconstruct the solution q as

$$q = \mathfrak{F}(f - \mathfrak{B}\widehat{q}). \quad (30)$$

Note that the operator $\widehat{\mathcal{A}}$ defined in Eq. (23) contains the term $\Xi(q, \widehat{q}, \widehat{v})$ and therefore the Schur-complement solution procedure implicitly enforces the continuity constraint in the second line of Eq. (20). The crucial observation is that both \mathfrak{A} and \mathfrak{F} only couple unknowns within one grid cell and therefore can both be represented by a set of cell-local matrices. Since further \mathfrak{B} and \mathfrak{C} couple the unknowns in one grid cell to the flux unknowns on its boundary, the matrix representation of the operator $\mathfrak{C}\mathfrak{F}\mathfrak{B} : \widehat{\Lambda}_h \rightarrow \widehat{\Lambda}_h$ in Eq. (29) is sparse. Recall further that for the Lax-Friedrichs flux only $\widehat{\mathbf{u}}$ and $\widehat{\mathbf{w}}$ appear in Eq. (23) whereas for the upwind flux, the equation only contains $\widehat{\phi}$ and $\widehat{\psi}$. Therefore Eq. (29) reduces to

$$\widehat{\mathcal{S}}^{(\text{up})}(\widehat{\phi}, \widehat{\psi}) = \widehat{\mathcal{R}}^{(\text{up})}(\widehat{\psi}) \quad \text{for all } \widehat{\psi} \in \Lambda_h \quad (\text{upwind}) \quad (31a)$$

$$\widehat{\mathcal{S}}^{(\text{LF})}(\widehat{\mathbf{u}}, \widehat{\mathbf{w}}) = \widehat{\mathcal{R}}^{(\text{LF})}(\widehat{\mathbf{w}}) \quad \text{for all } \widehat{\mathbf{w}} \in \Lambda_h \times \Lambda_h \quad (\text{Lax-Friedrichs}) \quad (31b)$$

in those cases. Before discussing the solution of Eq. (29) in Section 4.4, we briefly discuss the bilinear form in Eq. (24) in the context of previous work.

4.3. Relationship to previous work in [20]

For the HDG upwind flux, the Schur-complement system in Eq. (29) is very similar to the elliptic problem written down in Theorem 2.1 of [20]. To see this, observe that for constant c , τ_K Eq. (2.5) in [20] can be written in our notation as

$$\widehat{\mathcal{A}}^{(\text{CGDT})}(\phi, \mathbf{u}, \widehat{\phi}, \psi, \mathbf{w}, \widehat{\psi}) = b(\mathbf{w}, \psi)$$

for some right hand side b . The bilinear form $\widehat{\mathcal{A}}^{(\text{CGDT})}$ is given by

$$\begin{aligned} \widehat{\mathcal{A}}^{(\text{CGDT})}(\phi, \mathbf{u}, \widehat{\phi}, \psi, \mathbf{w}, \widehat{\psi}) := & (\mathbf{u} \cdot \mathbf{w})_{\Omega_h} - c^{-1} \left[(\kappa \mathbf{u} \cdot \nabla \psi + \phi \nabla \cdot \mathbf{w})_{\Omega_h} \right. \\ & \left. - \left(\kappa [\mathbf{u}\psi] + 2\kappa\tau_K \left(\{\!\!\{ \phi\psi \}\!\!\} - \widehat{\phi} \{\!\!\{ \psi \}\!\!\} \right) + \widehat{\phi} [\mathbf{w}] \right)_{\varepsilon_h} \right] \\ & + \left(\widehat{\psi} \left[[\mathbf{u}] + 2\tau_K \left(\{\!\!\{ \phi \}\!\!\} - \widehat{\phi} \right) \right] \right)_{\varepsilon_h} \end{aligned} \quad (32)$$

where κ is a constant. If we assume spatially non-varying bathymetry ($\phi_B = \text{constant}$) for the shallow water problem and set $\kappa = \phi_B^{-1}$, $c^{-1} = \phi_B c_g \alpha \Delta t$, $\tau_K = \sqrt{\phi_B}$ in Eq. (32), the resulting bilinear form differs from the one given in our Eq. (24) only by a zero order term, namely:

$$\widehat{\mathcal{A}}^{(\text{up})}(\phi, \mathbf{u}, \widehat{\phi}, \psi, \mathbf{w}, \widehat{\psi}) = \widehat{\mathcal{A}}^{(\text{CGDT})}(\phi, \mathbf{u}, \widehat{\phi}, \psi, \mathbf{w}, \widehat{\psi}) + (\phi, \psi)_{\Omega_h}.$$

It can then be shown [32] that in addition to slightly modifying the lifting operators $\vec{\mathcal{Q}}$ and \mathcal{U} defined on page 5 of [20], this gives rise to an additional, zero order term in the elliptic problem given in Theorem 2.1 of [20]. This zero order term can be identified with the $(\mathfrak{M}\widehat{q}, \widehat{v})_{\mathcal{E}_h}$ contribution to $\widehat{S}(\widehat{q}, \widehat{v})$ written down in Eq. (29) above. The additional term has a positive sign and does not affect the positive definiteness of \widehat{S} . In fact, since it adds a positive shift to all eigenvalues, it reduces the condition number of the corresponding matrix and allows the use of a shallow multigrid hierarchy as discussed in [33, 34]. The above observations motivated the construction of the non-nested multigrid algorithm described in Section 4.4. Further mathematical analysis and an extension of the proof in [20] to the HDG discretisation of the shallow water equations with the Lax-Friedrichs flux is beyond the scope of this paper and could be pursued in a future publication.

4.4. Non-nested multigrid algorithm

In [11] a direct method is used to solve the linear system in Eq. (29) for the flux unknowns. This is inefficient for large problems since the cost grows rapidly as the number of unknowns increases. To overcome this issue we now describe how Eq. (29) can be solved with an iterative Krylov-subspace method, which is preconditioned with a bespoke multigrid algorithm. For this let V_c be a coarse function space, which is not necessarily a subspace of $\widehat{\Lambda}_h$. Assume that there is prolongation operator $\pi : V_c \rightarrow \widehat{\Lambda}_h$, which naturally induces a corresponding restriction $\rho : \widehat{\Lambda}_h^* \rightarrow V_c^*$ on the dual spaces (marked by an asterisk *). For any $\mathcal{R} \in \widehat{\Lambda}_h^*$ the restriction is given as

$$\rho(\mathcal{R}) = \mathcal{R}_c = \mathcal{R} \circ \pi \in V_c^*, \quad \text{in other words } \mathcal{R}_c(q_c) = \mathcal{R}(\pi q_c) \text{ for all } q_c \in V_c.$$

Further assume the existence of a coarse space bilinear form⁴ $\mathcal{S}_c : V_c \times V_c \mapsto \mathbb{R}$. Based on this, a two-level V-cycle for approximately solving Eq. (29) is written down in Alg. 1 and shown schematically in Fig. 3. The left-hand side represents the hybridised DG space $\widehat{\Lambda}_h$ on the facets while the non-nested coarse space V_c on the triangular mesh induced by the facets is shown in the centre of the figure; in this example we assume that V_c is the lowest order piecewise linear conforming space (usually known as P1). Functions are transported between the two spaces $\widehat{\Lambda}_h$ and V_c with the p -restriction ρ and p -prolongation π described above. The right-hand side of the figure represents one level of coarsening of the P1 space; the latter is applied recursively in a GMG iteration. More generally, the gray box in Fig. 3 symbolises the solution of the coarse space problem with a geometric- or algebraic multigrid algorithm. For the upwind flux, the application of geometric multigrid is straightforward since the coarse level elliptic problem in Eq. (33) below is discretised in a conforming piecewise linear function space.

Algorithm 1 Non-nested two-level V-cycle for approximately solving $\widehat{S}(\widehat{q}, \widehat{v}) = \mathcal{R}(\widehat{v})$ in Eq. (29)

- 1: Presmooth $\widehat{q} \mapsto \text{Smooth}(\widehat{\mathcal{R}}, \widehat{q}; n_{\text{pre}})$
 - 2: Calculate residual $\widehat{\mathcal{B}}(v) = \widehat{\mathcal{R}}(\widehat{v}) - \widehat{S}(\widehat{q}, \widehat{v})$ for all $\widehat{v} \in \widehat{\Lambda}_h$
 - 3: Restrict residual $\mathcal{R}_c(v_c) = \widehat{\mathcal{B}}(\pi v_c)$ for all $v_c \in V_c$
 - 4: Solve $\mathcal{S}_c(q_c, v_c) = \mathcal{R}_c(v_c)$ for all $v_c \in V_c$
 - 5: Update $\widehat{q} \mapsto \widehat{q} + \pi q_c$
 - 6: Postsmooth $\widehat{q} \mapsto \text{Smooth}(\widehat{\mathcal{R}}, \widehat{q}; n_{\text{post}})$
-

In Algorithm 1 “Smooth($\widehat{\mathcal{R}}, \widehat{q}; n$)” stands for n applications of a stationary smoother for the equation $\widehat{S}(\widehat{q}, \widehat{v}) = \widehat{\mathcal{R}}(\widehat{v})$, such as a Jacobi iteration or a more sophisticated Chebyshev smoother. We use the following coarse level spaces, prolongation operators and bilinear forms for the upwind- and Lax-Friedrichs flux:

⁴One possible way (not pursued here) of defining this would be to set $\mathcal{S}_c(q_c, v_c) = \widehat{S}(\pi q_c, \pi v_c)$ for all $q_c, v_c \in V_c$.

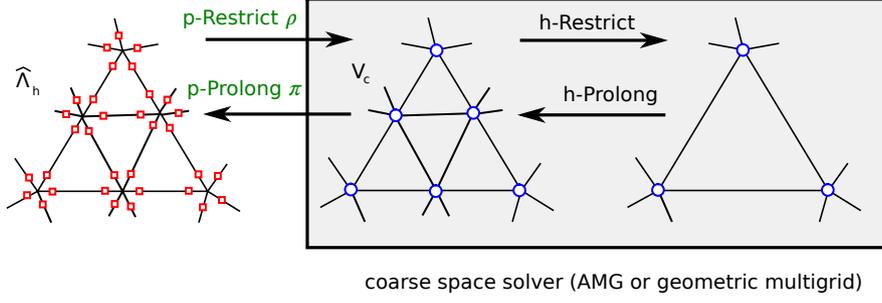


Figure 3: Multigrid algorithm. The coarse space problem in the space V_c is solved with AMG or geometric multigrid.

Upwind flux: The coarse space is chosen to be the scalar valued piecewise linear conforming finite element space on Ω with respect to the mesh Ω_h (usually called P1), defined by

$$V_c^{(\text{up})} = \{ \phi_c \in H^1(\Omega) : \phi_c|_K \in \mathcal{P}^1(K), \forall K \in \Omega_h \}.$$

In the following we assume that the polynomial degree of the DG space is $p \geq 1$. Since the restriction of $V_c^{(\text{up})}$ to the skeleton \mathcal{E}_h is a subspace of Λ_h , the prolongation $\pi^{(\text{up})} : V_c^{(\text{up})} \rightarrow \Lambda_h$ is the natural injection onto \mathcal{E}_h :

$$\pi^{(\text{up})} \phi_c = \widehat{\phi}, \quad \text{such that } \widehat{\phi} = \phi_c|_{\mathcal{E}_h} \text{ for all } \phi_c \in V_c^{(\text{up})}.$$

The coarse space bilinear form is the weak form of the second order positive definite div-grad-problem $\phi - (c_g \alpha \Delta t)^2 \nabla \cdot (\nabla \phi)$, which would arise formally from the time-implicit discretisation of the continuum equations in Eq. (3) after eliminating the field \mathbf{u} :

$$\mathcal{S}_c^{(\text{up})}(\phi_c, \psi_c) = (\phi_c, \psi_c)_{\Omega_h} + (c_g \alpha \Delta t)^2 (\nabla \phi_c \cdot \nabla \psi_c)_{\Omega_h} \quad \text{for all } \phi_c, \psi_c \in V_c^{(\text{up})} \quad (33)$$

Note that, as discussed in Section 4.3, the coarse level problem in [20] is very similar, but does not contain a zero order term.

Lax-Friedrichs flux: The coarse space is the lowest order Raviart-Thomas-space [35]

$$\mathbf{V}_c^{(\text{LF})} = \{ \mathbf{u}_c \in H(\text{div}, \Omega) : \mathbf{u}_c|_K = \boldsymbol{\alpha}_K + \beta_K \mathbf{x}, \forall K \in \Omega_h \}. \quad (34)$$

Here $H(\text{div}, \Omega)$ is the space of all vector fields for which both the components and the weak divergence are square-integrable. $\boldsymbol{\alpha}_K$ and β_K are constant on each cell K and are chosen such that \mathbf{u}_c has continuous normal component across facets. The prolongation $\pi^{(\text{LF})} : \mathbf{V}_c^{(\text{LF})} \rightarrow \widehat{\Lambda}_h$ is given by the L_2 projection on the skeleton \mathcal{E}_h :

$$\pi^{(\text{LF})} \mathbf{u}_c = \widehat{\mathbf{u}}, \quad \text{such that } \langle \widehat{\mathbf{u}}, \widehat{\mathbf{w}} \rangle_{\mathcal{E}_h} = \langle \llbracket \mathbf{u}_c \rrbracket, \widehat{\mathbf{w}} \rangle_{\mathcal{E}_h} \text{ for all } \widehat{\mathbf{w}} \in \widehat{\Lambda}_h.$$

In contrast to the upwind case, the coarse space bilinear-form $\mathcal{S}_c^{(\text{LF})}$ is now the weak form of the second order grad-div-problem $\mathbf{u} - (\alpha \Delta t)^2 \nabla (\nabla \cdot \mathbf{u})$, which would arise from the time-implicit discretisation of the continuum equations in Eq. (3) after eliminating the field ϕ :

$$\mathcal{S}_c^{(\text{LF})}(\mathbf{u}_c, \mathbf{w}_c) = (\mathbf{u}_c, \mathbf{w}_c)_{\Omega_h} + (c_g \alpha \Delta t)^2 ((\nabla \cdot \mathbf{u}_c), (\nabla \cdot \mathbf{w}_c))_{\Omega_h} \quad \text{for all } \mathbf{u}_c, \mathbf{w}_c \in \mathbf{V}_c^{(\text{LF})} \quad (35)$$

5. Timestepping methods and linear solver configuration

To study the performance of our solvers in a wider context we use a range of different timestepping methods.

5.1. IMEX methods

The following semi-implicit time stepping methods are used:

Theta(θ): The method in Eq. (19) with a particular value of the off-centering parameter θ . For $\theta = \frac{1}{2}$ this corresponds to treating the linear part of the PDE with the Crank-Nicholson method [36], which is second order in time.

ARS2(2,3,2): The ARS2 three-stage IMEX method which is also used in [11]. In general, ARS IMEX methods are developed in [2]; the matrices are written down explicitly in [6] and are given for ARS2(2,3,2) as

$$a = \begin{pmatrix} 0 & 0 & 0 \\ \gamma & 0 & 0 \\ \delta & 1 - \delta & 0 \end{pmatrix}, \quad \tilde{a} = \begin{pmatrix} 0 & 0 & 0 \\ 0 & \gamma & 0 \\ 0 & 1 - \gamma & \gamma \end{pmatrix}, \quad b = \tilde{b} = \begin{pmatrix} 0 \\ 1 - \gamma \\ \gamma \end{pmatrix}$$

with $\gamma = 1 - \frac{1}{\sqrt{2}}$ and $\delta = -\frac{2}{3}\sqrt{2}$.

SSP2(3,2,2): The strong stability preserving method described in [37] can be written as an IMEX method which is defined by the matrices

$$a = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix}, \quad \tilde{a} = \begin{pmatrix} 1/2 & 0 & 0 \\ -1/2 & 1/2 & 0 \\ 0 & 1/2 & 1/2 \end{pmatrix}, \quad b = \tilde{b} = \begin{pmatrix} 0 \\ 1/2 \\ 1/2 \end{pmatrix}.$$

ARS3(4,4,3): The ARS3 five-stage IMEX method with

$$a = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 \\ 1/2 & 0 & 0 & 0 & 0 \\ 11/18 & 1/18 & 0 & 0 & 0 \\ 5/6 & -5/6 & 1/2 & 0 & 0 \\ 1/4 & 7/4 & 3/4 & -7/4 & 0 \end{pmatrix}, \quad b = \begin{pmatrix} 1/4 \\ 7/4 \\ 3/4 \\ -7/4 \\ 0 \end{pmatrix},$$

$$\tilde{a} = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 1/2 & 0 & 0 & 0 \\ 0 & 1/6 & 1/2 & 0 & 0 \\ 0 & -1/2 & 1/2 & 1/2 & 0 \\ 0 & 3/2 & -3/2 & 1/2 & 1/2 \end{pmatrix}, \quad \tilde{b} = \begin{pmatrix} 0 \\ 3/2 \\ -3/2 \\ 1/2 \\ 1/2 \end{pmatrix}.$$

As discussed in Section 3.3, all semi-implicit methods require the solution of the ill-conditioned linear system in Eq. (23). In addition, mass-solvers are required to obtain the state $q^{(n+1)}$ at the next time step in Eq. (22). However, the corresponding mass matrices are block-diagonal. For each cell, a small, dense matrix has to be inverted, which is readily achieved with LU factorisation. Tab. 1 lists the order as well as the number of linear solves and function evaluations for all discussed methods, including the explicit time integrators which will be discussed in Section 5.3 below.

To assess the performance of the naive, approximate Schur-complement preconditioner, semi-implicit methods have also been implemented for the standard (non-hybridised) DG discretisation. The non-hybridised variant of the Theta(θ) method for Eq. (12) is

$$\frac{\mathcal{M}(q^{(n+1)}, v) - \mathcal{M}(q^{(n)}, v)}{\Delta t} = \mathcal{N}(q^{(n)}, v) + \theta \mathcal{L}(q^{(n+1)}, v) + (1 - \theta) \mathcal{L}(q^{(n)}, v) \quad \text{for all } v \in W_h. \quad (36)$$

Generally, a non-hybridised s -stage IMEX method can be written as

$$\mathcal{M}(Q^{(i)}, v) = \mathcal{M}(q^{(n)}, v) + \Delta t \sum_{j=1}^{i-1} a_{ij} \mathcal{N}(Q^{(j)}, v) + \Delta t \sum_{j=1}^i \tilde{a}_{ij} \mathcal{L}(Q^{(j)}, v), \quad \text{for } i = 1, \dots, s \quad (37)$$

$$\mathcal{M}(q^{(n+1)}, v) = \mathcal{M}(q^{(n)}, v) + \Delta t \sum_{i=1}^s b_i \mathcal{N}(Q^{(i)}, v) + \Delta t \sum_{i=1}^s \tilde{b}_i \mathcal{L}(Q^{(i)}, v), \quad \text{for all } v \in W_h,$$

which should be compared to Eqs. (21) and (22).

In the following section we provide further details on the configuration of the linear solvers which are used in our semi-implicit methods.

		method	order	# evaluations		# solves	
				\mathcal{L} or $\widehat{\mathcal{L}}$	\mathcal{N} or \mathcal{N}_0	\mathcal{M}	$\mathcal{M} - \alpha\widehat{\mathcal{L}}$
IMEX	{	Theta(θ)	1*	1	1	—	1
		ARS2(2,3,2)	2	2	3	2	2
		SSP2(3,2,2)	2	3	2	1	3
		ARS3(4,4,3)	3	4	4	2	4
explicit	{	Explicit Euler	1	1	1	1	—
		Heun	2	2	2	2	—
		SSPRK3	3	3	3	3	—

Table 1: Approximation order, number of function evaluations (listed separately for \mathcal{L} or $\widehat{\mathcal{L}}$ and \mathcal{N} or \mathcal{N}_0), linear solves with $\mathcal{M} - \alpha\widehat{\mathcal{L}}$ and mass solves with \mathcal{M} for all timestepping methods discussed in Section 5. For the non-hybridised versions of the IMEX integrators, the linear operator $\widehat{\mathcal{L}}$ is replaced by \mathcal{L} in $\mathcal{M} - \alpha\widehat{\mathcal{L}}$. Although for general values of θ the Theta(θ) method is first order, the linear part of the equations is integrated to second order accuracy for $\theta = 0.5$; this is indicated by the asterisk in the table above.

5.2. Linear solver

To compare the performance of different time integrators and investigate speedups achieved by using the hybridised DG discretisation we work with two different setups:

NL-LF: the non-linear SWEs with a Lax-Friedrichs flux and spatially varying bathymetry

L-Up: the linear SWEs with an upwind flux and constant bathymetry $\phi_B = 1$

5.2.1. Hybridised DG

For the semi-implicit time integrators, the hybridised system in Eq. (23) is solved with the (exact) Schur-complement reduction to flux-space described in Section 4.2. We solve the Schur-complement system defined in Eq. (29) with a Krylov-subspace method which is preconditioned with the non-nested multigrid method in Algorithm 1. For a flat bathymetry ($\phi_B = 1$), which is used in the linear setup (L-Up), the problem in Eq. (29) is symmetric and a Conjugate Gradient (CG) solver [38] can be used. For the non-linear setup (NL-LF) with spatially varying ϕ_B we use a GMRES [39] iteration since the problem is no longer symmetric. In both cases the (preconditioned) residual is reduced by a relative factor of at least ϵ , with ϵ to be specified. The bathymetry is always set to $\phi_B = 1$ in the multigrid preconditioner. As argued in Section 4.3, for the upwind flux the Schur-complement system (29) only differs from the one in [20] by a zero-order term.

It remains to specify the smoother $\text{Smooth}(\widehat{R}, \widehat{q}; n)$ in Algorithm 1 (lines 1 and 6) and the method for solving the coarse space problem $\mathcal{S}_c(q_c, v_c) = \mathcal{R}_c(v_c)$ (line 4). In our implementation we used a parallel Chebyshev iteration with SOR preconditioning and applied two pre- and post-smoothing steps, $n_{\text{pre}} = n_{\text{post}} = 2$. The coarse level problem $\mathcal{S}_c(q_c, v_c) = \mathcal{R}_c(v_c)$ is solved with a geometric- or algebraic-multigrid V-cycle, using PETSc's GAMG implementation in the latter case. Two SOR-preconditioned Chebyshev iterations are used as the smoother and coarse level solver in both AMG and the geometric multigrid solver.

5.2.2. Native DG

The linear system in the non-hybridised variants of the time-integrators in Eqs. (37) and (36) can be written in the form

$$\mathcal{A}(q, v) = \mathcal{M}(q, v) - \alpha\Delta t\mathcal{L}(q, v) = \mathcal{R}(v) \quad \text{for all } v \in W_h. \quad (38)$$

Eq. (38) is solved with a preconditioned GMRES solver, using the naive, approximate Schur-complement preconditioner which has been applied for conforming mixed finite elements in [8, 30]. As for the hybridised case, the iteration is aborted once the (preconditioned) residual has been reduced by a relative factor of at least ϵ . To construct the approximate Schur-complement preconditioner, consider the non-hybridised bilinear form \mathcal{A} which is written down for the Lax-Friedrichs flux in Eq. (A.1) and for the

upwind flux in Eq. (A.2) in Appendix A. Observe that in both cases \mathcal{A} can be decomposed into four terms as

$$\mathcal{A}(q, v) = \mathcal{A}(\phi, \mathbf{u}) = \mathcal{A}_{\phi\phi}(\phi, \psi) + \mathcal{A}_{\phi u}(\mathbf{u}, \psi) + \mathcal{A}_{u\phi}(\phi, \mathbf{w}) + \mathcal{A}_{uu}(\mathbf{u}, \mathbf{w}). \quad (39)$$

Both $\mathcal{A}_{\phi\phi}$ and \mathcal{A}_{uu} contain artificial diffusion terms which couple unknowns across adjacent cells. This is different from the hybridised case, where the term $(\mathfrak{A}q, v)_{\Omega_h}$ in Eq. (26) only contains couplings within one cell of the mesh. To construct a preconditioner, consider the matrix-representation

$$\begin{pmatrix} A_{\phi\phi} & A_{\phi u} \\ A_{u\phi} & A_{uu} \end{pmatrix} \begin{pmatrix} \Phi \\ \mathbf{U} \end{pmatrix} = \begin{pmatrix} \mathbf{R}_\phi \\ \mathbf{R}_u \end{pmatrix} \quad (40)$$

of Eq. (38) where $(\Phi, \mathbf{U})^T$ is the dof-vector of $(\phi, \mathbf{u}) = q$. $A_{\phi\phi}$, $A_{\phi u}$, $A_{u\phi}$ and A_{uu} are the (sparse) matrix representations of the four bilinear forms in Eq. (39). Given $(\mathbf{R}_\phi, \mathbf{R}_u)^T$, Eq. (40) can be solved approximately as follows:

Step 1. Approximately solve $A_{uu}\mathbf{V} = \mathbf{R}_u$ for \mathbf{V} .

Step 2. Approximately solve $S\Phi = \mathbf{R}_\phi - A_{\phi u}\mathbf{V}$ for Φ where

$$S = A_{\phi\phi} - A_{\phi u}(\text{diag}(A_{uu}))^{-1}A_{u\phi}. \quad (41)$$

Step 3. Approximately solve $A_{uu}\mathbf{U} = \mathbf{R}_u - A_{u\phi}\Phi$

Note that the approximate Schur complement S in step 2 is a sparse matrix since the diagonal matrix $\text{diag}(A_{uu})$ can be inverted exactly and is also diagonal. An ILU algorithm is applied to approximately invert the matrix A_{uu} in steps 1 and 3. One GAMG V-cycle is used to solve the Schur-complement system in step 2; two Chebyshev iterations (preconditioned with an approximate block-ILU-Jacobi method) are employed as the smoother and coarse-level solver. In our code, the PETSc “fieldsplit” preconditioner [40] was used to implement steps 1-3 of the approximate Schur-complement solver above.

The construction of the approximate Schur-complement S in Eq. (41) uses the diagonal $\text{diag}(A_{uu})$ of A_{uu} . Intuitively, a better approach appears to be to use the *block-diagonal* (i.e. the matrix $\text{blockdiag}(A_{uu})$ obtained from A_{uu} by dropping all couplings of unknowns between different grid cells) instead, since this does not change the sparsity pattern of S but uses more information from A_{uu} . However, most results in the paper were obtained by constructing the approximate Schur-complement based on the diagonal $\text{diag}(A_{uu})$, since this is the default option in the PETSc “fieldsplit” preconditioner. Potential improvements due to replacing $\text{diag}(A_{uu}) \mapsto \text{blockdiag}(A_{uu})$ in the approximate Schur-complement in Eq. (41) are quantified in Section 7.7.

5.3. Explicit methods

For reference we also consider fully explicit time integrators. In contrast to semi-implicit methods they do not require the solution of an ill-conditioned system but suffer from tighter restrictions on the time step size. Results are reported for the following methods:

Explicit Euler: The fully explicit Euler method, i.e. Eq. (19) with $\theta = 0$.

Heun’s method: An improved second order variant of the explicit Euler method.

SSPRK3: The third order strong stability preserving Runge Kutta method defined in [41].

In the context of DG discretisations, Heun’s method and SSPRK3 also also known as RKDG2 and RKDG3 which are discussed in the context of a wider framework in [42, 43]. For reference, further details on the explicit time integrators introduced in this section can be found in Appendix C.

6. Implementation

All code was implemented in the open source Firedrake library [22]. This allows the expression of weak forms in Unified Form Language (UFL) [44] in Python, from which highly efficient C-code is automatically generated. Firedrake uses PETSc [45, 40] as a linear algebra backend, which provides a wide range of iterative solvers and preconditioners. PETSc is highly configurable, and supports the composition of user-defined operators to design bespoke preconditioners [46]. Combined with the multigrid support in Firedrake, this allows the implementation of the sophisticated non-nested multigrid preconditioners described in this paper.

6.1. Local matrix algebra with Slate

The recently developed Slate package [21] for cell-wise linear algebra is crucial for the construction of the Schur-complement solver described in Section 4.2. Slate converts the UFL expression of a weak form into a so-called **Tensor** object, which stores the *locally* assembled matrix. **Tensors** can be manipulated cell-wise. To illustrate this, consider the Schur-complement system in Eq. (29) with \mathfrak{F} as in Eq. (28). Each of the operators \mathfrak{M} , \mathfrak{C} , \mathfrak{A} and \mathfrak{B} corresponds to a weak bilinear form, which is defined in Eqs. (25), (24) and (26).

For example, consider the weak form $(\mathfrak{A}q, v)_{\Omega_h}$ defined in Eq. (26). For the Lax-Friedrich flux $(\mathfrak{A}q, v)_{\Omega_h}$ is given by

$$\begin{aligned} (\mathfrak{A}q, v)_{\Omega_h} &= (\phi\psi + \mathbf{u} \cdot \mathbf{w} - c_g\alpha\Delta t [\mathbf{u} \cdot \nabla\psi + \phi\nabla \cdot \mathbf{w}])_{\Omega_h} \\ &+ c_g\alpha\Delta t \sum_{e \in \mathcal{E}_h} \left(\sqrt{\phi_B} (\mathbf{u}^+ \cdot \mathbf{w}^+ + \mathbf{u}^- \cdot \mathbf{w}^-) + \phi_B (\phi^+ \mathbf{n}^+ \cdot \mathbf{w}^+ + \phi^- \mathbf{n}^- \cdot \mathbf{w}^-) \right)_e. \end{aligned} \quad (42)$$

where $q = (\phi, \mathbf{u}), v = (\psi, \mathbf{w}) \in W_h$ are discontinuous fields. The bilinear form $(\mathfrak{A}q, v)_{\Omega_h}$ can be expressed in UFL and converted to a Slate tensor as follows:

```
a_form = (phi*psi+inner(u,w)-cg*alpha*dt*(inner(u,grad(psi))+phiB*phi*div(w))*dx
          + cg*alpha*dt*(sqrt(phiB)*(inner(u('+'),w('+'))
                                +inner(u('-'),w('-'))))
          +phiB*(phi('+')*inner(n('+'),w('+'))
                +phi('-')*inner(n('-'),w('-'))))*dS
A = Tensor(a_form)
```

Crucially, since the fields q, v in Eq. (42) are discontinuous, in this case the matrix representation of \mathfrak{A} is block-diagonal, since it only couples unknowns in the same grid cell. Hence, the inverse \mathfrak{F} can be obtained by inverting dense, cell-local matrices. The corresponding Slate tensor representation of \mathfrak{F} is given by `A.inv`. By constructing the tensors `M`, `C` and `B` which represent \mathfrak{M} , \mathfrak{C} and \mathfrak{B} , it is possible to construct the Slate tensor which corresponds to $\mathfrak{S} = \mathfrak{M} + \mathfrak{C}\mathfrak{F}\mathfrak{B}$ in Eq. (29) as

```
S = M + C*A.inv*B
```

This tensor can now be passed to a Firedrake linear variational problem solver as usual. More specifically, to solve a problem with a given right hand side `r_rhs` and return the solution in the variable `lambda`, write

```
trace_lvp = LinearVariationalProblem(a=S_expr,L=r_rhs,u=lambda)
trace_lvs = LinearVariationalSolver(trace_lvp,solver_parameters=param)
trace_lvs.solve()
```

Note that the exact solver configuration is specified by passing the dictionary `param` as the keyword argument `solver_parameters`. This is used to build the sophisticated multigrid solvers described in the following section.

We stress that the generation of all Slate expressions which are required in the HDG Schur-complement construction is automated; this is achieved by using the static condensation preconditioner `firedrake.SPC` described in the following section.

6.2. Composing the multigrid solver

A preconditioned iterative solver in Firedrake is characterised by the pair (KSP, PC) where `KSP` is the Krylov-subspace solver to be used and `PC` the preconditioner. For example, to solve the problem in the previous section with a GMRES solver, preconditioned with a Jacobi-iteration, set

```
solver_parameters={'ksp_type':'gmres','pc_type':'jacobi'}
```

when constructing the `LinearVariationalSolver` object. PETSc provides a wide range of iterative solvers and preconditioners. Crucially, the solvers can be composed [46] to construct new, problem-specific preconditioners. This is achieved by setting `'pc_type'='python'`, deriving a new python preconditioner object `MyPC` from a suitable Firedrake base class and setting `'pc_python_type'='MyPC'`.

To solve the problem $\widehat{\mathcal{A}}(q, \widehat{q}, v, \widehat{v}) = \mathcal{R}(v)$ defined in Eq. (23), two bespoke Python preconditioners are constructed and composed

fire Drake.SCP C: Given Eq. (23), this preconditioner implements steps 1-3 of the Schur-complement solver described at the beginning of Section 4.2. It uses Slate to reduce the problem to the Schur-complement system $\widehat{\mathcal{S}}(\widehat{q}, \widehat{v}) = \widehat{\mathcal{R}}(\widehat{v})$ defined in Eq. (29) (steps 1 and 2). After solving this system with `fire Drake.GTMGPC` (step 2), the DG field q is reconstructed with Eq. (30) (step 3).

fire Drake.GTMGPC: This solves the system $\widehat{\mathcal{S}}(\widehat{q}, \widehat{v}) = \widehat{\mathcal{R}}(\widehat{v})$ in the space $\widehat{\Lambda}_h$ by using the non-nested multigrid method described in Algorithm 1 in Section 4.4. The user can specify the smoother in the DG space and the coarse level solver by passing options to the solver.

The `fire Drake.GTMGPC` object requires additional, problem specific information to construct the bilinear form in the P1 or Raviart-Thomas space defined in Eqs. (33) and (35). This information is passed via an application-context (`appctx`) keyword when constructing the `LinearVariationalSolver`.

Schematically, given the bilinear form `a_form` written down in the code listing in Section 6.1, the problem $\widehat{\mathcal{A}}(q, \widehat{q}, v, \widehat{v}) = \mathcal{R}(v)$ is solved in Firedrake as follows:

```
lvp = LinearVariationalProblem(a=a_form, L=r_rhs, u=q_qhat)
param = {'mat_type': 'matfree',
         'ksp_type': 'preonly',
         'pc_type': 'python',
         'pc_python_type': 'fire Drake.SCP C',
         'pc_sc_eliminate_fields': '0, 1',
         'condensed_field': {'ksp_type': 'cg',
                             'mat_type': 'matfree',
                             'pc_type': 'python',
                             'pc_python_type': 'fire Drake.GTMGPC',
                             'gt': {...}}
        }
lvs = LinearVariationalSolver(lvp, solver_parameters=param, appctx=appctx)
lvs.solve()
```

Note that since the reduction to the Schur-complement in Section 4.2 is exact, the outer problem is solved by applying the preconditioner once (this is achieved by setting `'ksp_type'='preonly'`). The dictionary `gt` is used to specify the smoother and coarse level solver of the `fire Drake.GTMGPC` object. For reference, the full configuration is written down in Appendix D.1.

As explained in Section 5.2.2, the approximate Schur complement preconditioner for the native DG method, is implemented with the PETSc `fieldsplit` preconditioner. The exact solver options for this are written down in Appendix D.2. All code developed for this paper is openly available as a git repository at <https://bitbucket.org/em459/hybridizeddg>.

7. Results

The numerical methods developed in the first part of the paper are now applied to different model problems. For this, we first consider an idealised, smooth stationary flow in a flat domain, which exhibits the typical separation between advective- and gravity-wave velocities observed in atmospheric flows. The problem setup is described in Section 7.1 and results on the accuracy, performance and parallel scalability are presented in sections 7.3 to 7.6. Results for an alternative approximate DG Schur-complement are shown in Section 7.7 and a comparison to explicit time integrators is presented in Section 7.8.

Following this, we test our methods for several well-established test cases in spherical geometry from the Williamson et al. suite in [23] in Section 7.9. After verifying expected spatial convergence rates for a smooth stationary test case (zonal flow with compact support), we run our code for two non-stationary examples: (i) flow over an isolated mountain and (ii) the Rossby-Haurwitz wave. In the latter two cases we compare the performance of our new multigrid algorithms to a direct solver, i.e. a setup which is similar to the one used in [11]; we find that our preconditioners are competitive with a highly-optimised direct method.

7.1. Problem setup for flat geometry

7.1.1. Stationary vortex

To assess the correctness of our code and quantify numerical accuracy we consider a smooth, time-independent rotationally invariant solution of the shallow water equations. As will be shown below, the

numerical parameters can be adjusted such that the relative size of the gravity wave speed and advective velocities is similar to what is encountered in atmospheric models. To construct this stationary solution, all fields are assumed to be radially symmetric, and the bathymetry ϕ_B is described by

$$\phi_B(r) = \begin{cases} 1 - \delta_B \exp \left[\frac{1}{r-r_+} + \frac{4}{r_+-r_-} - \frac{1}{r-r_-} \right] & \text{for } r_- \leq r \leq r_+, \\ 1 & \text{otherwise.} \end{cases} \quad (43)$$

An exact solution of the SWEs is given by the height-perturbation

$$\phi(r) = -\delta \cdot \begin{cases} 1 & \text{for } r < r_-, \\ \frac{1}{2} \left(1 + \tanh \left[\frac{\sigma}{r-r_-} + \frac{\sigma}{r-r_+} \right] \right) & \text{for } r_- \leq r \leq r_+, \\ 0 & \text{for } r_+ < r, \end{cases} \quad (44)$$

and purely tangential velocity $\mathbf{u}(r) = u(r)\mathbf{e}_\theta$ in the tangential direction. For the non-linear SWEs in Eq. (1) the function $u(r)$ is given by

$$u(r) = \frac{r}{2L_R}(\phi + \phi_B) \left(-1 + \sqrt{1 + 4L_R^2 \frac{\phi'}{r}} \right). \quad (45)$$

where $L_R = c_g/f$ is the Rossby radius of deformation, measured in units of the reference length scale \bar{L} . For the linear SWEs in Eq. (3), the tangential velocity simplifies to

$$u(r) = L_R(\phi + \phi_B)\phi'. \quad (46)$$

The solution is smooth ($\phi, u \in C^\infty$) and has compact support, i.e. ϕ and \mathbf{u} are non-zero only on a disk of radius r_+ around the origin. In particular, if we set $r_+ < \frac{1}{2}$, the computational domain can be set to the unit square $\Omega = [-\frac{1}{2}, \frac{1}{2}] \times [-\frac{1}{2}, \frac{1}{2}]$. Both fields are visualised together with the bathymetry ϕ_B in Fig. 4. An important parameter is the relative size of the advection velocity $\mathbf{U}^{(\text{adv})} = \mathbf{u}/(\phi_B + \phi)$ and the gravity wave speed $\phi_B^{1/2}$. This ratio is given by the spatially varying non-dimensional Froude number

$$F := \frac{|\mathbf{U}^{(\text{adv})}|}{\phi_B^{1/2}} = \begin{cases} \frac{r}{2L_R}\phi_B^{-1/2} \left(-1 + \sqrt{1 + 4L_R^2 \frac{\phi'}{r}} \right) & \text{(non-linear),} \\ L_R\phi_B^{-1/2}\phi' & \text{(linear).} \end{cases} \quad (47)$$

It can be shown that for sufficiently large values of $(r_+ - r_-)/\sigma$ the function ϕ' has a single maximum at $r = \frac{1}{2}(r_+ + r_-)$ and in both cases the Froude number can be bounded as

$$F \leq L_R\phi_B^{-1/2}\phi' \leq \frac{4\sigma L_R}{(r_+ - r_-)^2} (1 + \delta_B) \delta. \quad (48)$$

To obtain the results in the rest of this paper we use the following numerical values: The reference depth is set to $\bar{H}_B = 2 \cdot 10^3 \text{m}$, the reference length is chosen to be the radius of the earth $\bar{L} = R_{\text{earth}} = 6.37 \cdot 10^6 \text{m}$ and the time scale is $\bar{T} = 1 \text{day} = 8.64 \cdot 10^4 \text{s}$. The Coriolis parameter is set to $2 \cos(\vartheta)\Omega$, where $\Omega = 2\pi \text{day}^{-1}$ is the angular speed of rotation of the earth and ϑ is the latitude (measured from the pole). For simplicity we further use the constant value of f at the pole, i.e. set $\vartheta = 0$. This results in the following non-dimensional values

$$c_g = 1.89, \quad f = 4\pi \approx 12.567, \quad L_R = 0.15.$$

We further set

$$\delta = \delta_B = 0.1, \quad r_- = 0.05, \quad r_+ = 0.45, \quad \sigma = 0.25.$$

The Froude number is plotted for this choice of parameters in the range $r_- \leq r \leq r_+$ in Fig. 5. In all numerical experiments below the initial condition $q_0 = (\phi_0, \mathbf{u}_0)$ was chosen such that ϕ_0 is given by Eq. (44) and $\mathbf{u}_0 = u(r)\mathbf{e}_\theta$ with $u(r)$ given by either Eq. (45) for the non-linear SWEs or Eq. (46) for the linearised equations. With this choice the exact solution $q_{\text{exact}} = (\phi_{\text{exact}}, \mathbf{u}_{\text{exact}}) = (\phi_0, \mathbf{u}_0)$ is stationary, which allows the easy quantification of numerical errors.

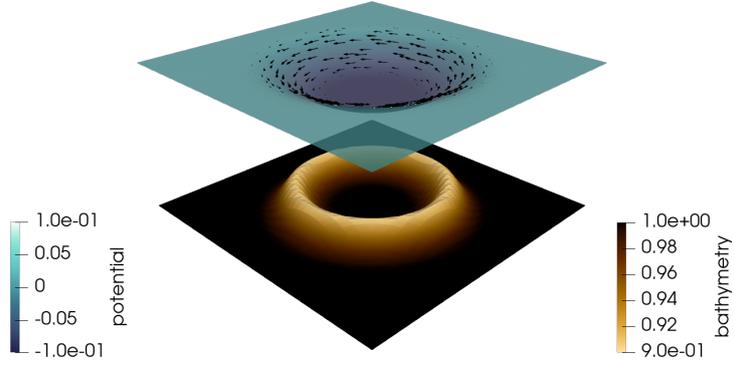


Figure 4: Stationary vortex solution ϕ , \mathbf{u} as defined in Eqs. (44) and (45). The bathymetry given in Eq. (43) is shown at the bottom of the figure.

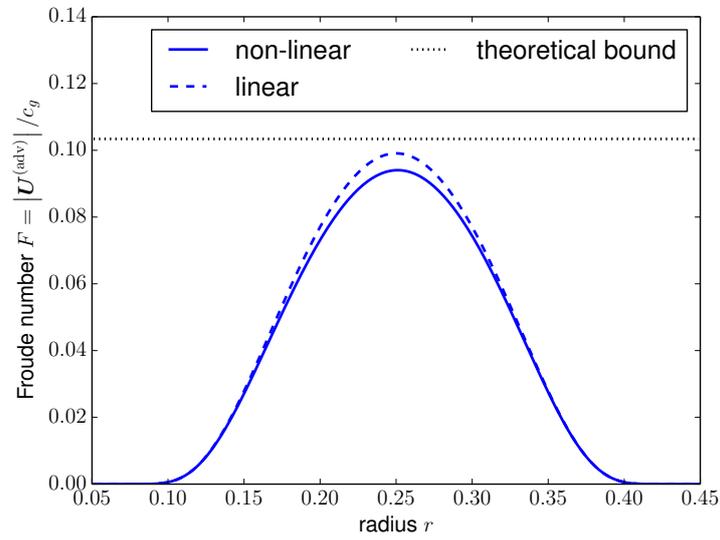


Figure 5: Froude number defined in Eq. (47) for the linear and non-linear SWEs. The theoretical upper bound in Eq. (48) is also shown as a dotted line.

r	N_{cell}	$h = 2^{-r}$	$p = 1$		$p = 3$		$p = 5$	
4	512	0.0625	4 608	/ 1 536	15 360	/ 3 072	32 256	/ 4 608
5	2 048	0.0313	18 432	/ 6 144	61 440	/ 12 288	129 024	/ 18 432
6	8 192	0.0156	73 728	/ 24 576	245 760	/ 49 152	516 096	/ 73 728
7	32 768	0.0078	294 912	/ 98 304	983 040	/ 196 608	2 064 384	/ 294 912
8	131 072	0.0039	1 179 648	/ 393 216	3 932 160	/ 786 432		

Table 2: Grid layout and numbers of unknowns for different polynomial degrees p . Each of the final three columns shows the total number of unknowns associated with cells and facets as calculated by Eq. (50) for the scalar-valued flux problem ($\nu^{(\text{flux})} = \nu^{(\text{up})} = 1$) in the form $N_{\text{dof}}^{(\text{cell})}/N_{\text{dof}}^{(\text{facet})}$.

7.1.2. Time step size

For a given grid resolution h and polynomial degree p of the DG spaces defined in Eqs. (5) and (8) the time step size is chosen as follows. First, a marginally stable explicit timestep size Δt_{expl} is found by requiring that the Courant number is expressed as

$$\frac{c_g \Delta t_{\text{expl}}}{h} = \frac{\rho}{2p+1}. \quad (49)$$

Here the dependency on p is chosen as proposed by [47]. Numerical experiments confirm that explicit time integrators are marginally stable, i.e. the timestep Δt_{expl} is close to the stability limit, if $\rho = 0.2$ in Eq. (49). We then set the timestep for the semi-implicit IMEX methods to $\Delta t = \gamma \Delta t_{\text{expl}}$ where $\gamma > 1$ is chosen such that the *advective* Courant number $|\mathbf{U}^{(\text{adv})}| \Delta t/h$ is sufficiently small, but again close to the stability limit for advective transport. This can be achieved by requiring that $F\gamma \leq 1$ (where F is the Froude number defined in Eq. (47)). In atmospheric models the gravity wave speed is typically one order of magnitude larger than advective velocities, and setting $\gamma = 10$ in the numerical experiments therefore replicates this scale separation in an idealised scenario (as can be seen from Fig. 5, the Froude number is bound by approximately 0.1, so this choice of γ guarantees that $F\gamma \lesssim 1$). The final time was always set to $T = \frac{1}{2}$.

7.1.3. Grid layout

The grid was constructed by first partitioning the domain Ω into n^2 squares of width $h = 1/n$ and then subdividing each cell into two triangles, resulting in $N = 2n^2$ triangular cells overall. The number n is chosen such that $n = 2^r$ where $r \in \mathbb{N}$ is referred to as the *refinement level* in the following. Tab. 2 shows the grid sizes used. To highlight the relative size of the original DG system and the hybridised system, we also show the total number of unknowns for varying polynomial degrees p . For a fully periodic triangular grid with N_{cell} cells the number of unknowns for storing the fields ϕ and \mathbf{u} on cells is given by $N_{\text{dof}}^{(\text{cell})}$; storing the flux unknowns on the facets requires $N_{\text{dof}}^{(\text{facet})}$ variables with

$$N_{\text{dof}}^{(\text{cell})} = \frac{3}{2}(p+1)(p+2)N_{\text{cell}}, \quad N_{\text{dof}}^{(\text{facet})} = \nu^{(\text{flux})} \frac{3}{2}(p+1)N_{\text{cell}}. \quad (50)$$

$\nu^{(\text{flux})}$ depends on the numerical flux. Since only scalar valued unknowns are stored on the facets for the upwind flux, whereas vectors are used in the Lax-Friedrichs flux, we have $\nu^{(\text{LF})} = 2$ and $\nu^{(\text{up})} = 1$. The Schur-complement flux-system in Eq. (29) is a factor $\nu^{(\text{flux})}/(p+2)$ smaller than the original DG system in the variables ϕ, \mathbf{u} .

7.1.4. Linear solver configurations

The following solver configurations are used in our numerical experiments:

HDG+GTMG+GMG : As described in Section 4, the implicit system for the hybridised DG formulation in Eq. (23) is solved with a Schur-complement approach, which requires the solution of a linear system in the flux-space Λ_h . The resulting problem (which is given in Eq. (29)) is solved with GMRES or CG and preconditioned with the non-nested multigrid method described in Section 4.4; this multigrid algorithm is written down as a two-level method in Alg. 1. The coarse level system $\mathcal{S}_c(q_c, v_c) = \mathcal{R}_c(v_c)$ is solved with geometric multigrid (see Sections 5.2.2 and 6.2 for details on the implementation and Appendix D.1 for the PETSc solver options).

HDG+GTMG+AMG : This configuration is very similar to HDG+GTMG+GMG, but the coarse level problem in line 4 of Alg. 1 is solved with AMG instead of a geometric multigrid method.

HDG+AMG : For reference, the flux system in Eq. (29) is also solved directly with AMG instead of the non-nested two-level method in Alg. 1 (again, see Appendix D.1 for details on the PETSc solver options).

DG+ApproxSchur : The non-hybridised (ϕ, \mathbf{u}) system in Eq. (38) is solved with GMRES. The preconditioner is the approximate Schur complement approach with AMG for the ϕ -system described in Section 5.2.2 (see Appendix D.2 for PETSc solver options).

7.2. Computational resources and code version

All runs were carried out on the “Balena” supercomputer at the University of Bath. Each node of the machine contains two Intel Xeon E5-2650v2 (IvyBridge) processors with eight cores each, resulting in 16 cores per node. The results in the paper can be reproduced with the code released as [48]. All runtimes reported in Sections 7.4 and 7.5 were obtained by running on a single fully occupied node (16 cores).

7.3. Accuracy

We begin by verifying the accuracy of the DG discretisation, which serves as an important test of the correctness of the implementation. Theoretically [11] we expect the spatial discretisation error (measured in a suitable norm defined in [11]) to be $\mathcal{O}(h^{p+\frac{1}{2}})$, where p is the polynomial degree of the DG discretisation. To quantify the error, both the fully non-linear SWEs with bathymetry (the NL-Up configuration) and the linearised equations with flat bathymetry (L-Up) are solved with an ARS(2,3,2) IMEX timestepper. We verified that the DG+ApproxSchur and the HDG+GTMG+AMG solver give the same solution in both cases. For a given grid spacing h and polynomial degree p the L_2 error is defined as

$$\Delta_2 := \|q - q_{\text{exact}}\|_2 = \left(\int_{\Omega} ((\phi - \phi_{\text{exact}})^2 + (\mathbf{u} - \mathbf{u}_{\text{exact}})^2) dx \right)^{1/2}. \quad (51)$$

Here $q_{\text{exact}} = (\phi_{\text{exact}}, \mathbf{u}_{\text{exact}})$ is the known exact solution of the stationary vortex problem described in Section 7.1.1, and this is used as the initial condition. $q = (\phi, \mathbf{u})$ denotes the numerical solution at the final time $T = \frac{1}{2}$. The implicit time step size is chosen as in Section 7.1.2. Fig. 6 shows the error Δ_2 for varying grid refinement r and different polynomial degrees p . The error decreases at a rate of at least $\mathcal{O}(h^{p+\frac{1}{2}})$ (and empirically no more than $\mathcal{O}(h^{p+1})$).

7.4. Comparison of linear solvers

Next, we verify that the non-nested HDG solver described in Sections 4.2 and 4.4 is indeed superior to the naive DG solver in Section 5.2.2 and demonstrate that the non-nested multigrid method developed in this paper (see Alg. 1) is competitive with a black-box AMG solver for the hybridised system in Eq. (29). Fig. 7 compares the performance of the four solver setups listed in Section 7.1.4 for different polynomial degrees p and increasing spatial resolution. The linear SWEs with flat bathymetry (L-Up) are solved with a Theta($\theta = 0.5$) timestepper and again the timestep size is chosen as in Section 7.1.2. The plot demonstrates that for sufficiently large problems there is very little difference between the performance of the three hybridised solvers, namely HDG+GTMG+AMG, HDG+GTMG+GMG and HDG+AMG. To understand this, the total runtime for the simulations with refinement $r = 7$ in Fig. 7 is broken down into the time spent in the solver for the Schur-complement system in Eq. (29) and the combined time spent both in the fine-level smoother and residual calculation in Tab. 3. The results show that for high polynomial degrees the performance of the three solvers (GTMG+GMG, GTMG+AMG and AMG) differs by around 10%. This small variation in performance is readily explained by the fact that the solve time is dominated by the fine space smoother and residual calculation, which is the same for all methods.

All solvers based on the hybridised DG discretisation perform significantly better the native DG method with an approximate Schur complement preconditioner (DG+ApproxSchur). At the highest considered polynomial degree ($p = 5$) the speedup is approximately a factor of 3. Fig. 8 shows the runtimes achieved with the DG+ApproxSchur and HDG+GTMG+AMG method both for the linear

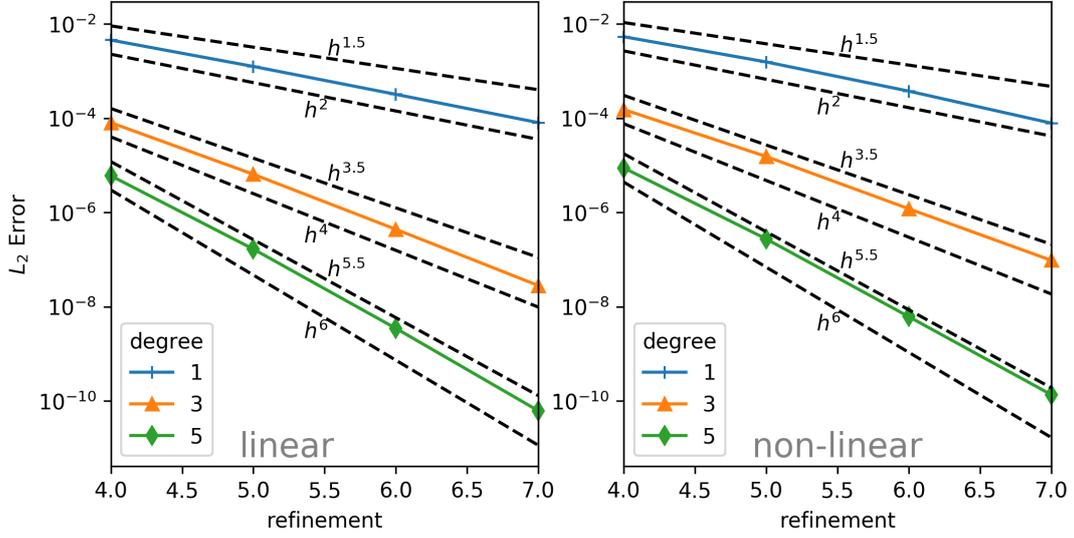


Figure 6: L_2 error Δ_2 as defined in EQ. (51) for the linearised equations with flat bathymetry (left) and the fully non-linear SWEs with spatially varying bathymetry (right) as a function of the grid refinement r for different polynomial degrees p . All results are obtained with the ARS2(2,3,3) timestepper. The Lax-Friedrichs flux was used in the fully non-linear case and the upwind flux for the linear problem. The grid spacing is $h = 2^{-r}$, where r is the refinement shown on the horizontal axis.

	solver	t_{total}	t_{solve}	t_{fine}
$p = 1$	GTMG+GMG	21.0	6.9	1.6
	GTMG+AMG	17.7	3.5	1.5
	AMG	16.9	2.7	1.5
$p = 3$	GTMG+GMG	207.3	38.3	23.0
	GTMG+AMG	198.9	31.1	23.4
	AMG	200.7	32.3	26.4
$p = 5$	GTMG+GMG	1691.7	121.1	90.3
	GTMG+AMG	1680.5	112.2	91.2
	AMG	1675.8	108.2	93.0

Table 3: Breakdown of runtime for the HDG-based solvers for refinement $r = 7$ and different polynomial degrees p . The tolerance in the iterative solver is set to $\epsilon = 10^{-8}$. In addition to the total runtime t_{total} , the time t_{solve} spent in the solver for the Schur-complement system (29) and the combined time t_{fine} spent in the fine level smoother and residual calculation are shown. All results are given in seconds.

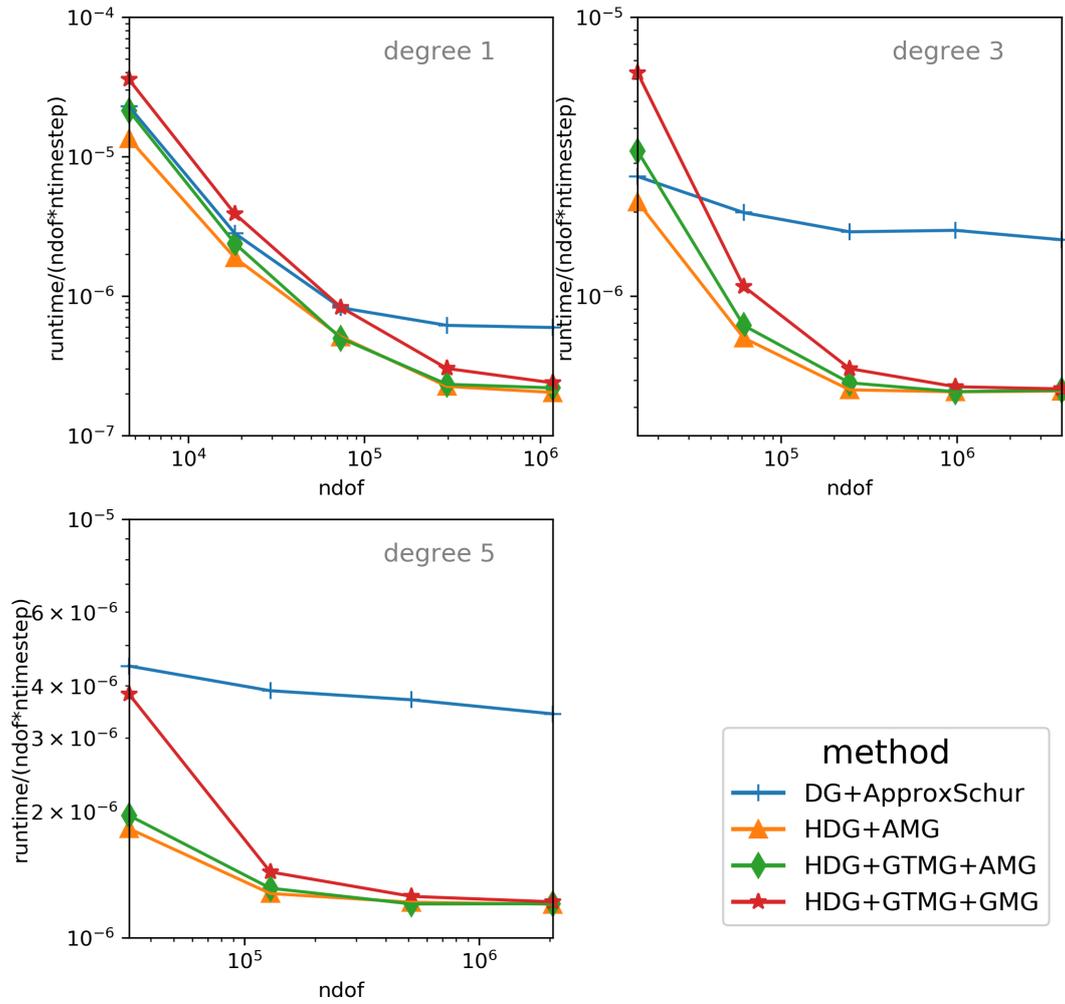


Figure 7: Comparison of the performance of different linear solvers for the linearised SWEs without bathymetry (L-Up testcase). The runtime per timestep and per unknown is shown as a function of the number of degrees of freedom for different polynomial degrees p . The ARS2(2,3,2) timestepper is used in all cases and results are given in seconds.

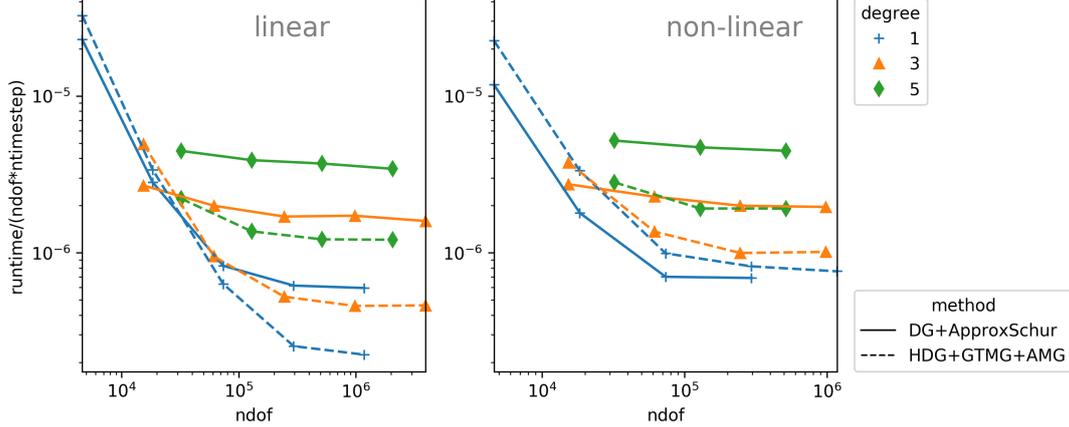


Figure 8: Runtime per unknown per timestep as a function of the total number of DG unknowns $N_{\text{dof}}^{(\text{cell})}$ for different solvers. Results for the non-linear setup with spatially varying bathymetry (NL-LF) are shown on the right and times for the linear SWEs with constant bathymetry (L-Up) on the left. All results are given in seconds.

r	$h = 2^{-r}$	$p = 1$		$p = 3$		$p = 5$	
		DG	HDG	DG	HDG	DG	HDG
4	0.0625	14.1	8.1	15.5	8.0	17.3	8.0
5	0.0313	13.8	8.1	14.9	7.0	16.5	8.0
6	0.0078	13.0	8.0	14.0	7.0	15.9	8.0
7	0.0078	12.9	8.0	14.0	7.0	15.0	8.0
8	0.0039	12.0	8.0	13.0	7.0	—	—

Table 4: Average number of GMRES/CG iterations for solving the linear problem arising in the linearised SWEs with flat bathymetry (L-Up) to a tolerance of $\epsilon = 10^{-8}$. For the DG+SchurApprox solver the number of GMRES iterations over the outer (ϕ, \mathbf{u}) system in Eq. (38) are given in the “DG” column. The number of CG iterations required for solving the flux problem in Eq. (29) is reported for the HDG+GTMG+AMG method in the “HDG” column.

system (L-Up, left) and the non-linear problem with spatially varying bathymetry (NL-LF, right). The relative speedup of the HDG+GTMG+AMG method is quantified in both cases in Fig. 9. As can be seen from Fig. 8 the drop in speedup for higher refinements can be explained by the fact that the runtime of the DG+ApproxSchur decreases slightly more than expected as r is increased from 7 to 8 for $p = 3$ and from 6 to 7 for $p = 5$.

This speedup can be partly attributed to the larger number of GMRES iterations in the “outer” solve of the (ϕ, \mathbf{u}) system for the DG+ApproxSchur solver and the relatively smaller number of Krylov-solver iterations required to solve the flux-system with the non-nested multigrid preconditioner for the HDG+GTMG+AMG method; in both cases the residual was reduced by a factor of $\epsilon = 10^{-8}$. This number of iterations is shown for the linear L-Up problem (which uses a CG-solver for the symmetric positive definite flux system) in Tab. 4 and for the non-linear NL-LF problem in Tab. 5. Compared to the native DG method, the hybridised DG solver requires significantly fewer iterations and is robust with respect to the polynomial degree p . Another advantage of the HDG+GTMG+AMG solver, in particular for higher polynomial degree p , is that the flux system, whose solution dominates the runtime, has significantly fewer unknowns than the (ϕ, \mathbf{u}) -system which is solved by the DG+ApproxSchur solver. As discussed at the end of section 7.1.4, the (ϕ, \mathbf{u}) system is $p + 2$ times larger than the flux-system for the linearised SWEs (L-Up). In the fully non-linear case (NL-LF) the relative size is only $\frac{1}{2}(p + 2)$ since the Lax-Friedrich flux will result in two flux unknowns per facet. This partly explains the smaller speedup in this case.

7.5. Comparison of IMEX timesteppers

To assess the overall performance of our methods when solving the time-dependent shallow water equations, we compare different IMEX timesteppers introduced in Section 5. Tab. 6 gives the runtime and

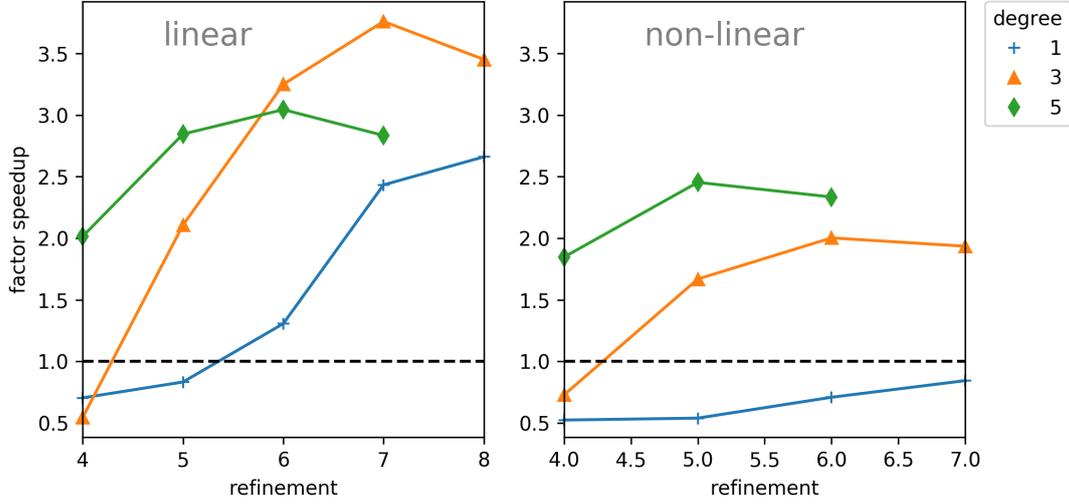


Figure 9: Speedup of the hybridised HDG+GTMG+AMG solver relative to the DG+ApproxSchur method as a function of the total number of DG unknowns $N_{\text{dof}}^{(\text{cell})}$ given in Eq. (50). The Theta($\theta = 0.5$) timestepper is used in both cases. Results for the non-linear setup with spatially varying bathymetry (NL-LF) are shown on the left and speedups for the linear SWEs with constant bathymetry (L-Up) on the right.

r	$h = 2^{-r}$	$p = 1$		$p = 3$		$p = 5$	
		DG	HDG	DG	HDG	DG	HDG
4	0.0625	15.9	10.1	18.0	8.9	21.0	8.8
5	0.0313	15.0	10.2	17.3	9.0	20.4	8.8
6	0.0078	14.1	10.1	16.3	9.0	19.4	9.0
7	0.0078	14.0	10.1	16.0	8.9	—	—
8	0.0039	—	10.0	—	—	—	—

Table 5: Average number of GMRES iterations for the linear problem arising in the fully non-linear SWEs with bathymetry (NL-LF) to a tolerance of $\epsilon = 10^{-8}$. For the DG+SchurApprox solver the number of iterations over the outer (ϕ, \mathbf{u}) system in Eq. (38) are given in the “DG” column. The number of iterations required for solving the flux problem in Eq. (29) is reported for the HDG+GTMG+AMG method in the “HDG” column.

timestepper	tolerance	runtime [s]		n_{iter}		error Δ_2
	ϵ	DG	HDG	DG	HDG	
Theta($\theta = 0.5$)	10^{-8}	107.4	56.1	16.3	9.0	$5.29 \cdot 10^{-6}$
	10^{-6}	78.5	49.7	12.0	7.0	$5.14 \cdot 10^{-6}$
	10^{-4}	53.0	49.3	7.7	5.0	$4.77 \cdot 10^{-5}$
Theta($\theta = 0.55$)	10^{-8}	117.0	56.9	17.9	8.9	$5.38 \cdot 10^{-8}$
	10^{-6}	84.5	49.7	13.0	7.0	$5.40 \cdot 10^{-8}$
	10^{-4}	59.3	47.1	8.8	5.0	$2.79 \cdot 10^{-7}$
ARS2(2,3,2)	10^{-8}	184.4	112.0	14.0	8.6	$5.39 \cdot 10^{-8}$
	10^{-6}	137.7	102.7	10.0	7.0	$5.39 \cdot 10^{-8}$
	10^{-4}	95.2	97.0	6.4	5.0	$2.43 \cdot 10^{-7}$
SSP2(3,2,2)	10^{-8}	327.0	158.7	17.0	8.0	$7.36 \cdot 10^{-6}$
	10^{-6}	235.8	146.1	12.0	6.0	$7.36 \cdot 10^{-6}$
	10^{-4}	169.6	140.6	8.0	5.0	$7.36 \cdot 10^{-6}$
ARS3(4,4,3)	10^{-8}	456.0	212.7	17.0	8.0	$5.38 \cdot 10^{-8}$
	10^{-6}	357.4	205.9	13.0	7.0	$5.38 \cdot 10^{-8}$
	10^{-4}	230.6	195.0	8.0	5.0	$1.10 \cdot 10^{-7}$

Table 6: Comparison of runtime, number of solver iterations and L_2 error for different IMEX timesteppers. Results are reported for solver tolerances of $\epsilon = 10^{-8}$, 10^{-6} and 10^{-4} . The polynomial degree is $p = 3$ and the refinement $r = 6$.

L_2 error Δ_2 (defined in Eq. (51)) for the fully non-linear setup with bathymetry (NL-LF). For all IMEX methods the average number of iterations is reported both for DG and hybridised DG discretisations. The polynomial degree is $p = 3$ and the refinement $r = 6$. Results are given for linear solver tolerances of $\epsilon = 10^{-8}$, 10^{-6} and 10^{-4} . For the NL-LF setup the HDG+GTMG+AMG method is used to solve the Schur-complement problem for the flux variables. For the timesteppers considered here the HDG solver is typically twice as fast as the corresponding DG method for $\epsilon = 10^{-8}$, which confirms the findings in Section 7.4. For looser tolerances the advantage of the HDG method is less pronounced. This is readily explained by the following observation: for the DG method the cost is roughly proportional to the number of outer solver iterations, which becomes smaller as ϵ increases. For the HDG method, however, a significant amount of time is spent in the construction of the right hand side of the Schur-complement flux system and the reconstruction of ϕ and \mathbf{u} (it is well known that the current Slate implementation of those parts of the algorithm is sub-optimal and will be improved in a future version of Firedrake). Since the solution of the flux system is relatively cheap, any reduction in the number of solver iterations with looser tolerances has a smaller impact on the overall runtime. Off-centering the Theta(θ)-method slightly by setting $\theta = 0.55$ instead of $\theta = 0.5$ reduces the error by two orders of magnitude at the cost of a small increase in the number of linear solver iterations and overall runtime (this increase is not observed for the hybridised solver). Looking back at Tab. 1, the significantly higher cost of the ARS2(2,3,2), SSP2(3,2,2) and ARS4(4,4,3) integrators (which are not really more accurate than the Theta($\theta = 0.55$) method for this particular testcase) is largely explained by the fact that they require two, three and four linear solves per timestep respectively. In virtually all cases the HDG based solver with our non-nested multigrid algorithm performs significantly better than the naive DG based method.

7.6. Parallel scalability

Finally, the parallel scalability of the code is quantified for different polynomial degrees p . In a strong scaling experiment the global problem size is kept fixed at $2.95 \cdot 10^5$ ($p = 1$), $9.83 \cdot 10^5$ ($p = 3$) and $2.06 \cdot 10^6$ ($p = 5$) unknowns while increasing the number of processor cores. The chosen global problem sizes correspond to a refinement of $r = 7$ for all polynomial degrees. As will be discussed below, a higher refinement is required for the lowest polynomial degree to achieve good parallel scalability. To illustrate this, results with $r = 10$ are also reported for $p = 1$ and the linearised SWEs (L-Up setup). The local problem size for all considered processor counts are shown in Tab. 7. Fig. 10 shows the strong scaling of the time per timestep. As can be seen from this plot, scalability is reasonably good for $p = 3, 5$ until the local problem size reaches a few thousand unknowns on 256 cores. For the lowest polynomial degree ($p = 1$), however, the code antiscals if more than two nodes are used. This issue can only be overcome

$p \setminus \#procs$	16	32	64	128	256
1	$1.84 \cdot 10^4$	$9.22 \cdot 10^3$	$4.61 \cdot 10^3$	$2.30 \cdot 10^3$	$1.15 \cdot 10^3$
3	$6.14 \cdot 10^4$	$3.07 \cdot 10^4$	$1.54 \cdot 10^4$	$7.68 \cdot 10^3$	$3.84 \cdot 10^3$
5	$1.29 \cdot 10^5$	$6.45 \cdot 10^4$	$3.23 \cdot 10^4$	$1.61 \cdot 10^4$	$8.06 \cdot 10^3$
1*	$1.18 \cdot 10^6$	$5.90 \cdot 10^5$	$2.95 \cdot 10^5$	$1.47 \cdot 10^5$	$7.37 \cdot 10^4$

Table 7: Local problem sizes for strong scaling experiments for different polynomial degrees p and numbers of processor cores. The refinement level is $r = 7$ for the first three rows, the final row (marked by an asterisk $*$) shows the local problem sizes for $p = 1$ with refinement $r = 10$.

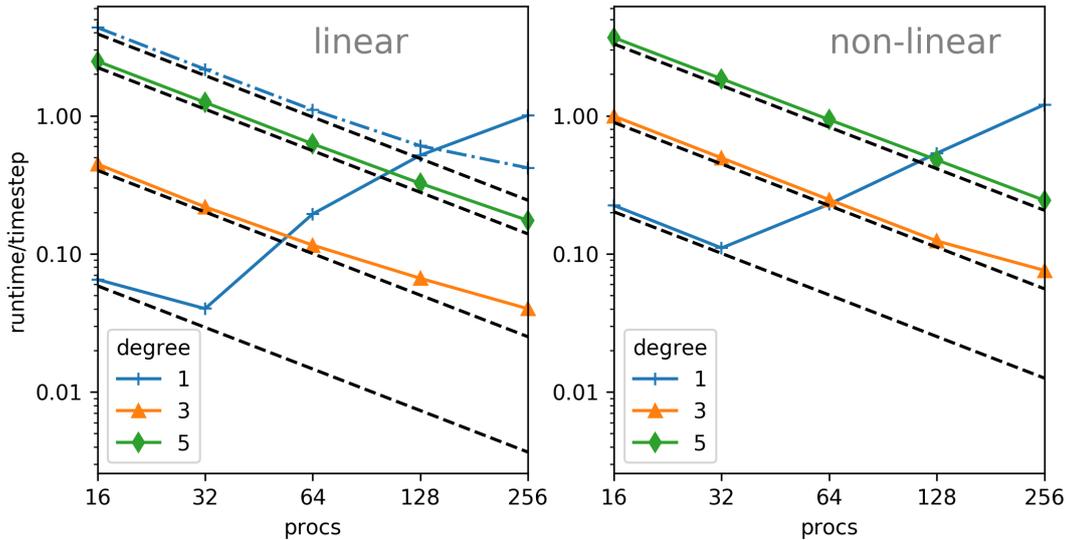


Figure 10: Strong scaling of the time per timestep for the HDG+GTMG+AMG solver on 16 to 256 processors. Results are shown for both for the linearised problem with constant bathymetry (L-Up, left) and the fully non-linear SWEs with spatially varying bathymetry (NL-LF, right). The dot-dashed curve in the left plot shows the results for degree $p = 1$ with refinement $r = 10$, the solid curves use a refinement of $r = 7$. The black dashed lines correspond to optimal speedup for which the runtime per timestep would be inversely proportional to the number of processors.

by significantly increasing the global problem size by using a refinement level of $r = 10$ instead.

In summary, the results clearly demonstrate that increasing the polynomial degree improves scalability. This can be explained by the fact that, compared to the cost of computations, the time spent communicating data decreases as the polynomial degree increases.

For the weak scaling experiment the total problem size is increased in proportion with the number of processors while keeping the local problem size fixed. On one processor the number of unknowns is $4.61 \cdot 10^3$ ($p = 1$), $1.54 \cdot 10^4$ ($p = 3$) and $3.23 \cdot 10^4$ ($p = 5$); this corresponds to a refinement of $r = 4$ for all polynomial degrees. Tab. 8 lists the global problem sizes. As this table shows, problems with more than a million unknowns are solved on 256 cores. Analogously to the strong scaling experiment, for degree $p = 1$ we also report results for a higher refinement ($r = 7$ in this case). The time per timestep for increasing problem sizes is shown in Fig. 11. Again, we find that scalability is very good for high polynomial degrees ($p = 3, 5$). Scaling breaks down for $p = 1$ once inter-node communication is required. As for the strong scaling experiment, this is caused by the poorer calculation-to-communication ratio in this case. Increasing the global problem size for $p = 1$ by setting $r = 7$ improves this, but the code still does not scale as well as for higher polynomial degrees. As usual, interpreting results within a single node (i.e. for 16 or less cores) is very challenging since the variation of the runtime can be influenced by saturation of the memory bandwidth and the complexities of the cache hierarchy.

$p \setminus \#procs$	1	4	16	64	256
1	$4.61 \cdot 10^3$	$1.84 \cdot 10^4$	$7.37 \cdot 10^4$	$2.95 \cdot 10^5$	$1.18 \cdot 10^6$
3	$1.54 \cdot 10^4$	$6.14 \cdot 10^4$	$2.46 \cdot 10^5$	$9.83 \cdot 10^5$	$3.93 \cdot 10^6$
5	$3.23 \cdot 10^4$	$1.29 \cdot 10^5$	$5.16 \cdot 10^5$	$2.06 \cdot 10^6$	$8.26 \cdot 10^6$
1*	$2.95 \cdot 10^5$	$1.18 \cdot 10^6$	$4.72 \cdot 10^6$	$1.89 \cdot 10^7$	$7.55 \cdot 10^7$

Table 8: Global problem sizes for weak scaling experiments for different polynomial degrees p and numbers of processor cores. The refinement level is $r = 4$ for the first three rows, the final row (marked by an asterisk $*$) shows the local problem sizes for $p = 1$ with refinement $r = 7$.

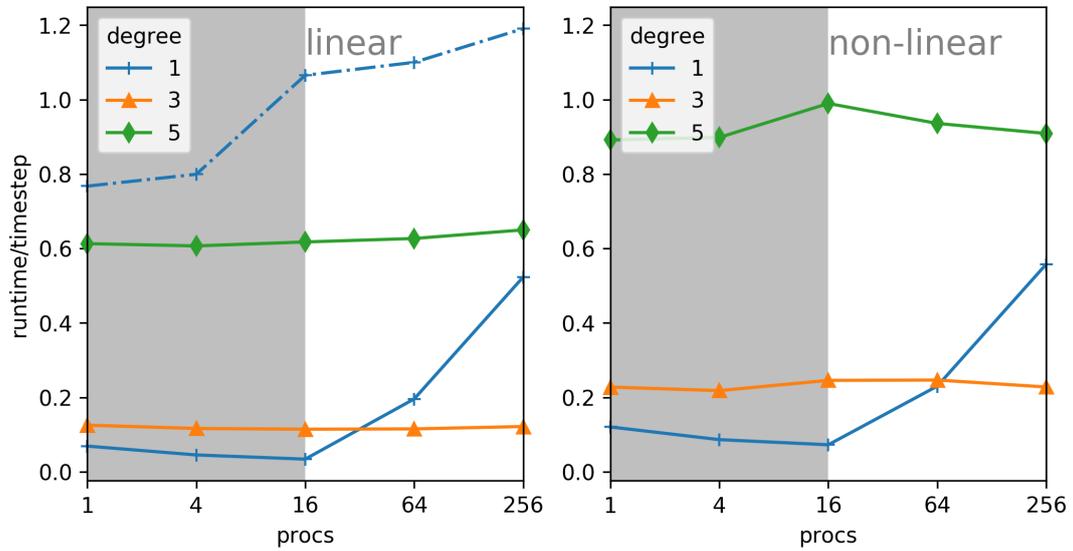


Figure 11: Weak scaling of the time per timestep for the HDG+GTMG+AMG solver on 16 to 256 processors. Results are shown for both the fully non-linear SWEs with spatially varying bathymetry (NL-LF, left) and for the linearised problem with constant bathymetry (L-Up, right). The gray region shows scaling within a node.

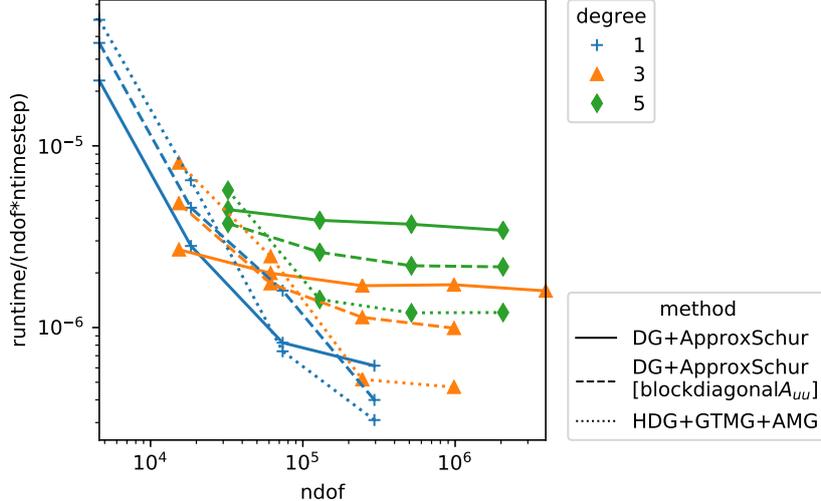


Figure 12: Runtime per unknown per timestep as a function of the total number of DG unknowns $N_{\text{dof}}^{(\text{cell})}$ for the HDG+GTMG+AMG solver and different approximate Schur complements, using either the diagonal (DG+ApproxSchur) or the block-diagonal (DG+ApproxSchur [blockdiagonal A_{uu}]) of the matrix A_{uu} . All results are given in seconds and were obtained for the linear SWEs with constant bathymetry (L-Up). The tolerance in the linear solver is set to $\epsilon = 10^{-8}$.

timestepper	runtime [s]	error Δ_2
Explicit Euler	32.0	$6.25 \cdot 10^{-2}$
Heun	35.3	$5.41 \cdot 10^{-8}$
SSPRK3	59.1	$5.41 \cdot 10^{-8}$

Table 9: Comparison of runtime and L_2 error for different explicit timesteppers. The polynomial degree is $p = 3$ and the refinement $r = 6$.

7.7. Block-diagonal approximation in Schur complement

As discussed at the end of Section 5.2.2, the approximate Schur-complement preconditioner for the native DG method (the DG+ApproxSchur solver) can be potentially improved by using the block-diagonal instead of the diagonal of the matrix A_{uu} . To quantify any improvements due to this modification, the runs for the linearised SWE testcase (L-Up) from Section 7.4 were repeated for a variant of the DG+ApproxSchur preconditioner, in which $\text{diag}(A_{uu})$ has been replaced by $\text{blockdiag}(A_{uu})$ in Eq. (41). Fig. 12 shows those results, together with the ones for DG+ApproxSchur and HDG+GTMG+AMG which are already plotted in Fig. 8 (left). As those results demonstrate, the approximate Schur-complement preconditioner can be improved by this modification, but is still not competitive with solvers based on the hybridised DG method.

7.8. Comparison to explicit time integrators

For completeness, we also present a tentative comparison of our IMEX-HDG solvers with fully explicit time integrators. We stress, however, that this is not the main focus of this paper, and that the comparison of absolute runtimes should be interpreted with caution: the simplified considered setup here does not replicate the conditions encountered for realistic flows and a more reliable comparison should be carried out for example for the spherical testcase in [23] to assess the relative merits of explicit and semi-implicit methods.

We consider the three explicit time integrators introduced in Section 5.3: the lowest order explicit (forward) Euler method, Heun’s method and the third order accurate SSPRK3 integrator. The absolute runtime and errors for the same setup as used in Tab. 6 are shown in Tab. 9. Since the Froude number is around $F = 10$, those results were obtained with an explicit timestep which is a factor ten shorter than the implicit timestep size used for the IMEX integrators in Section 7.5 to guarantee stability.

While the forward Euler integrator is extremely simple and leads to the shortest runtime, not very surprisingly the error is very large in this case; potentially the method is also affected by a time-stepping instability, which is seeded by the imperfect balance of the initial state. The Heun method, which requires an additional function evaluation and two instead of one mass solves, reduces the error significantly, while only increasing the runtime by around 10%. Comparing the most efficient IMEX method (Theta($\theta = 0.55$)) to the fastest explicit method with the same L_2 error (Heun), we conclude that for this particular setup using hybridisation allows the design of IMEX timesteppers which (for a linear solver tolerance of 10^{-6}) are only $1.4\times$ slower than the significantly simpler explicit integrators. The difference in performance can be understood by looking at the number of function evaluations and linear solves given in Tab. 1. The Theta(θ) method requires one expensive linear solve and one function evaluation per timestep while two function evaluations and two mass-solves are needed for the Heun-method.

Although here we restrict ourselves to the case where the implicit time step size Δt is not more than ten times larger than the explicit step size Δt_{expl} , the authors of [11] also consider scenarios where this ratio is larger. They find that for $\Delta t/\Delta t_{\text{expl}} = 20$ their IMEX HDG-ARS2 code is nearly as fast as an RKDG2 (SSPRK3) integrator, whereas for $\Delta t/\Delta t_{\text{expl}} = 200$ their HDG-based IMEX solver is ten times faster than the explicit integrator, albeit with a larger error.

7.9. Spherical geometry

We now apply our methods to well-established test cases in spherical geometry [23]. Care has to be taken when extending the DG discretisation to non-Euclidean geometries since in general the cell-normals of two neighbouring cells (see Fig. 2) do not necessarily satisfy $\mathbf{n}^- = -\mathbf{n}^+$. To address this issue, we followed the approach described in [49]. As can be seen for example in Eq. (A.1), this introduces additional ‘‘curvature’’ terms in the linear- and non-linear standard DG-fluxes. It is not necessary to modify the HDG flux since in this case unknowns in neighbouring cells only couple indirectly through the hybridised variables on the facets.

The main aim of the following numerical experiments is to compare the performance of our non-nested multigrid methods to direct solvers which were also used in [11].

7.9.1. Problem setup

We consider three configurations from [23] which are labelled W3, W5 and W6 for future reference:

W3 : Steady state zonal geostrophic flow with compact support described in [23, Sec. 3.3]. Since this test case has an analytic solution, we use it to verify the correctness of our code and study convergence under h - and p -refinement. Note also that the setup is conceptually very similar to the steady state vortex in flat geometry described in Section 7.1.1.

W5 : Zonal flow over an isolated mountain, as described in [23, Sec. 3.5]. Although the initial condition is the same as for W3, in this case the bathymetry is not constant and describes an isolated conical mountain in the Northern hemisphere.

W6 : Rossby-Haurwitz waves described in [23, Sec 3.6] are analytical solutions of the nonlinear barotropic vorticity equation (but now of the SWEs) on the sphere.

The parameters used in our numerical experiments are the same as in [23]; for reference they are collected in Tab. D.13 in Appendix D.3; in the following all quantities are non-dimensionalised as described in Section 2. For this choice of parameters the maximal value of the Froude number $F := \left| \mathbf{U}^{(\text{adv})} \right| \phi_B^{-1/2}$ at the initial time is approximately 0.23 for W3, 0.17 for W5 and 0.34 for W6, and hence the scale separation is slightly less pronounced than for the steady-state vortex in Section 7.1.1. Since the test cases considered are non-linear, the Lax-Friedrichs flux formulation of the discretised SWEs is used in all cases. The computational grid is obtained by recursively sub-dividing the cells of an icosahedral coarse mesh consisting of 20 triangles, taking into account the curvature of the sphere. If r denotes the number of refinement steps, the average side length of a triangular grid cell is (in units of the radius of the Earth)

$$h = 2^{-r} h_0 \quad \text{with } h_0 = \sqrt{\frac{4\pi}{5\sqrt{3}}} \approx 1.20459. \quad (52)$$

r	N_{cell}	h	$p = 3$		$p = 5$	
2	320	0.3011	9 600 /	1 920	20 160 /	2 880
3	1 280	0.1506	38 400 /	7 680	80 640 /	11 520
4	5 120	0.0753	153 600 /	30 720	322 560 /	46 080
5	20 480	0.0376	614 400 /	122 880	1 290 240 /	184 320
6	81 920	0.0188	2 457 600 /	491 520	5 160 960 /	737 280
7	327 680	0.0094	9 830 400 /	1 966 080	20 643 840 /	2 949 120

Table 10: Grid layout and numbers of unknowns for different polynomial degrees p for the numerical experiments in Section 7.9. Each of the final two columns shows the total number of unknowns associated with cells and facets as calculated by Eq. (50) for the scalar-valued flux problem ($\nu^{(\text{flux})} = \nu^{(\text{up})} = 1$) in the form $N_{\text{dof}}^{(\text{cell})}/N_{\text{dof}}^{(\text{facet})}$.

Tab. 10 shows the grid layout and number of unknowns for a range of refinement levels and polynomial degrees. Integrals over cells and facets in the weak form require evaluation of geometric quantities to capture the local curvature of the sphere, such as the Jacobian of the pullback to a flat reference cell. In the following the coordinate field is approximated by a polynomial of degree 3 in each grid cell. For all numerical experiments we use the ARS2 time integrator and choose the timestep size Δt such that

$$\frac{\max\{|\mathbf{U}_0^{(\text{adv})}|\}\Delta t}{h} = \frac{\rho}{2p+1} \quad (53)$$

where $\mathbf{U}_0^{(\text{adv})}$ is the initial advection velocity and ρ is a parameter that is adjusted such that the semi-implicit timestepper (which treats advective processes explicitly) is stable. We find that the choice $\rho = 0.1/\sqrt{2} \approx 0.0707$ works well for the majority of considered test cases, larger values lead to instabilities. The only exception are the numerical experiments for polynomial degree $p = 5$ in the (non-linear) W5- and W6- setups; in this case we had to reduce the timestep size by an additional factor of two to ensure stability over the entire simulation window of 15 days.

7.9.2. Linear solvers

We use three different methods to solve the linear problem that arises in the semi-implicit time integrator:

HDG+GTMG+AMG : The non-nested multigrid approach with AMG solver for the system in the Raviart-Thomas subspace, as described in more detail in Section 7.1.4.

HDG+LU : In this configuration the hybridised linear system in Eq. (23) is solved with a direct method from the MUMPS suite [24, 25]; this is similar to the setup in [11].

DG+ApproxSchur : As described in Section 7.1.4, the non-hybridised (ϕ, \mathbf{u}) system in Eq. (38) is solved with an iterative solver preconditioned with an approximate Schur complement.

Recall that the nature of the hybridised system in Eq. (23) is very different if the upwind flux is used: in contrast to the Lax-Friedrichs flux, the hybridised unknowns on each facet are scalar-valued and the coarse-level problem is formulated in the P1-subspace in this case. To also explore the performance of our solvers in this setup, we constructed linearised variants W3^(lin), W5^(lin) and W6^(lin) of the test cases in [23]. Those are obtained by dropping all terms in the initial conditions which are quadratic in the velocity and using the linear shallow water equations with upwind flux. Note that with those modifications the solution of W3^(lin) is still stationary and known analytically.

7.9.3. Convergence

To verify the accuracy and expected convergence rates of the error as the grid is refined we solved the stationary W3 problem and compared the numerical result to the exact solution. Fig. 13 shows the potential ϕ , absolute momentum $|\mathbf{u}|$ and the relative error $(\phi - \phi^{(\text{exact})})/(\phi_B + \phi^{(\text{exact})})$ of the total potential $\phi_B + \phi$ after 12 days for polynomial degree $p = 3$ and refinement level $r = 3$ for the HDG+GTMG+AMG solver. As the figure shows, the absolute value of the relative error of the total potential is of the order 10^{-4} . Fig. 14 shows the L_2 norm of the total error as defined in Eq. (51)

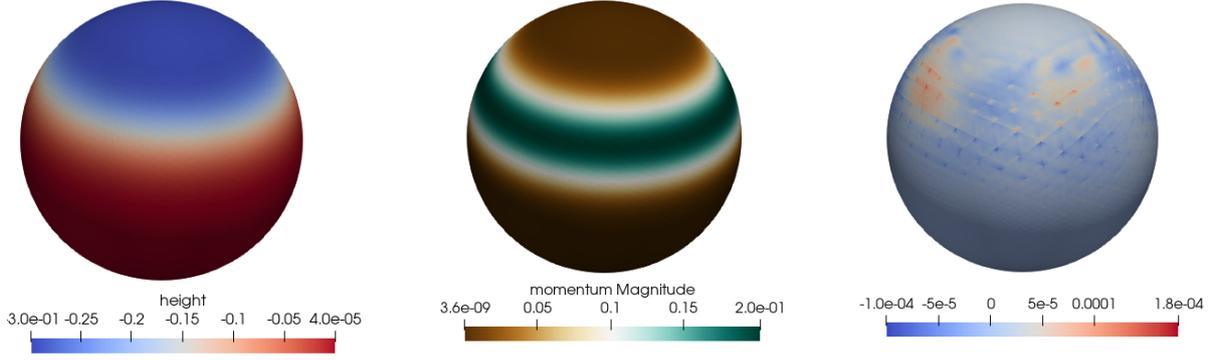


Figure 13: Potential ϕ (left), absolute value of the momentum $|\mathbf{u}|$ (center) and relative error in the potential $\phi + \phi_B$ (right) for the non-linear W3 (steady-state zonal flow) testcase at $t = 12\text{days}$. The polynomial degree is $p = 3$ and a refinement level of $r = 3$ was used. The timestep size in physical units is 650.6s. The linear flux system is solved with the HDG+GTMG+AMG solver.

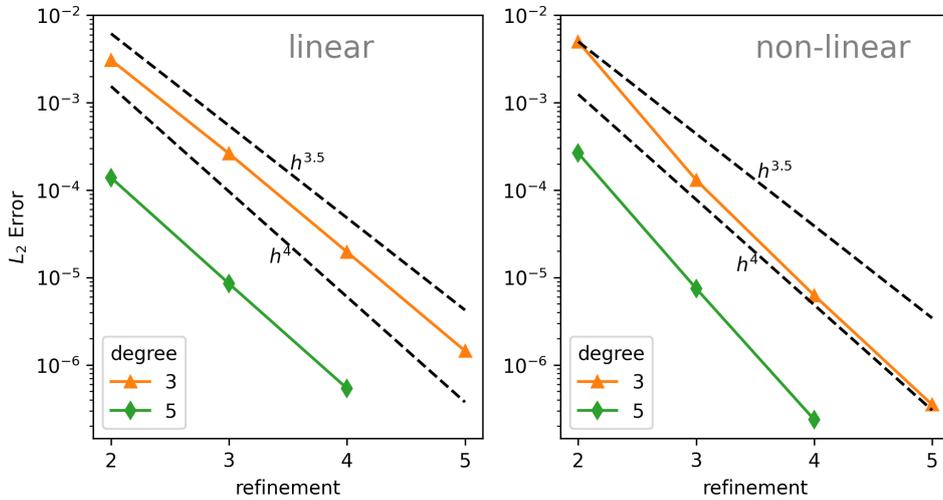


Figure 14: Total L_2 error defined in Eq. (51) as a function of the refinement level for $p = 3$ and $p = 5$. The code was run to the final time $T = 5$, using both W3 (right) and W3^(lin) (left).

for increasing refinement levels and two different polynomial degrees p . For each value of h and p the timestep size Δt was adjusted according to Eq. (53). For $p = 3$ the error decreases with a rate between $h^{3.5}$ and h^4 , which is consistent with the results in Section 7.3. Although the error is over an order of magnitude smaller for $p = 5$, asymptotically it does not appear to decrease with the naively expected rate between $h^{p+1/2}$ and h^{p+1} . This, however, is readily explained by the fact that in our implementation the coordinate field is approximated by a cubic polynomial in each grid cell, which introduces an additional error that limits the asymptotic rate of convergence. Nevertheless, even in this context, the benefit of using $p = 5$ instead of $p = 3$ is evident.

7.9.4. Computational performance

The performance of the different solver configurations discussed in Section 7.9.2 is compared by running the isolated mountain (W5) and Rossby-Haurwitz (W6) test. Fig. 15 shows the potential ϕ and momentum \mathbf{u} for the W5 test after 15 days for polynomial degree $p = 3$ and refinement $r = 3$; the corresponding results for the W6 test at the same resolution and after 14 days are shown in Fig. 16. The figures should be compared for example to Figs. 9 and 11 in [11]. For each solver configuration the time $t_{\text{dof,iter}}$ per timestep and per (DG) unknown is measured over the first 15 days of the run for

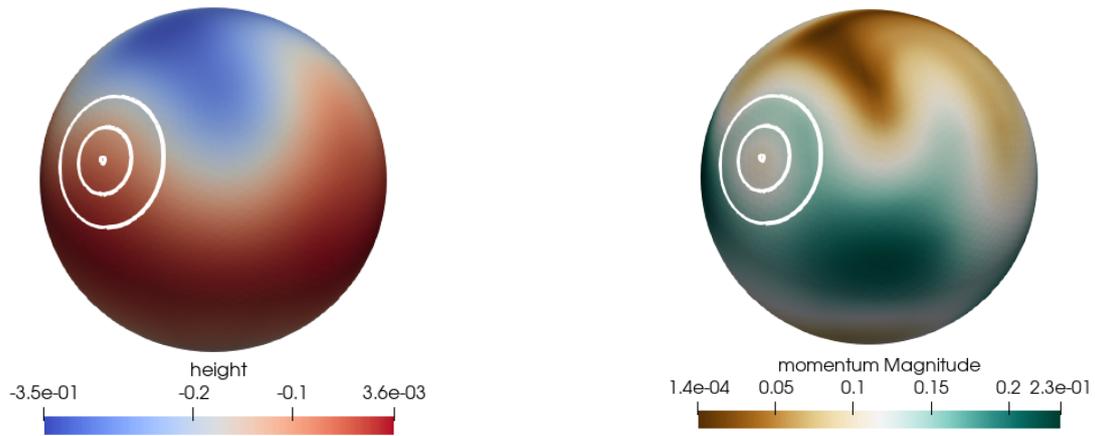


Figure 15: Potential ϕ (left) and absolute value of the momentum $|\mathbf{u}|$ (right) for the non-linear W5 (isolated mountain) testcase at $t = 15$ days, with the location of the mountain marked by concentric circles. The polynomial degree is $p = 3$ and a refinement level of $r = 3$ was used. The timestep size in physical units is 783.2s. The linear flux system is solved with the HDG+GTMG+AMG solver.

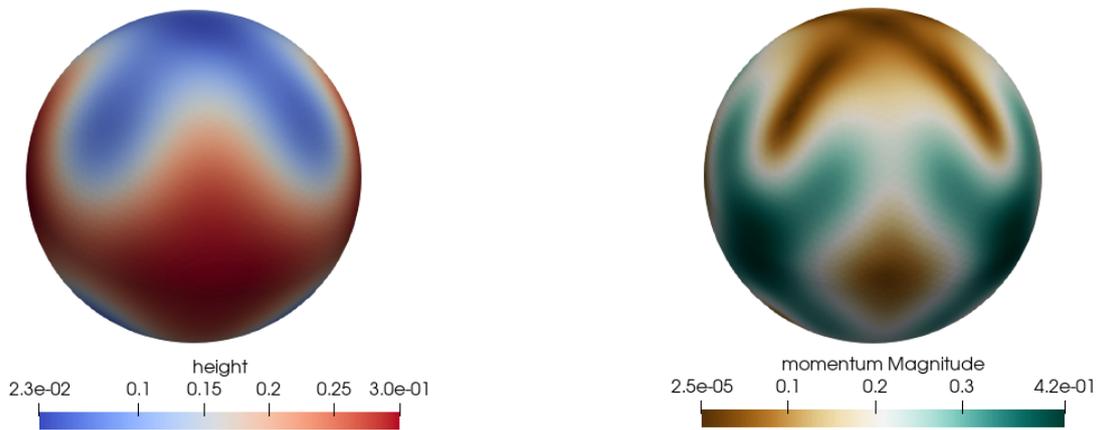


Figure 16: Potential ϕ (left) and absolute value of the momentum $|\mathbf{u}|$ (right) for the non-linear W6 (Rossby-Haurwitz) testcase at $t = 14$ days. The polynomial degree is $p = 3$ and a refinement level of $r = 3$ was used. The timestep size in physical units is 306.8s. The linear flux system is solved with the HDG+GTMG+AMG solver.

	W5 (mountain)				W6 (Rossby-Haurwitz)			
	$p = 3$ ($\Delta t = 391.5\text{s}$)		$p = 5$ ($\Delta t = 124.6\text{s}$)		$p = 3$ ($\Delta t = 153.4\text{s}$)		$p = 5$ ($\Delta t = 48.8\text{s}$)	
solver	n_{iter}	$t_{\text{dof,iter}}$	n_{iter}	$t_{\text{dof,iter}}$	n_{iter}	$t_{\text{dof,iter}}$	n_{iter}	$t_{\text{dof,iter}}$
DG+ApproxSchur	11.0	3.99	9.8	5.61	9.0	3.48	8.9	5.23
HDG+LU	—	5.45	—	10.86	—	5.24	—	10.73
HDG+GTMG+AMG	8.0	6.29	8.0	11.78	8.0	6.12	8.0	11.52

Table 11: Average number n_{iter} of solver iterations and time $t_{\text{dof,iter}}$ (measured in μs) per time step and per unknown for the non-linear W5 (isolated mountain) and W6 (Rossby-Haurwitz) testcases. In all cases the refinement level is $r = 4$ and the SWEs are integrated to a final time of $T = 15\text{days}$ with the timestep size Δt given in the relevant column. The iterative solver tolerance is 10^{-8} .

	W5 ^(lin) (mountain)				W6 ^(lin) (Rossby-Haurwitz)			
	$p = 3$ ($\Delta t = 391.5\text{s}$)		$p = 5$ ($\Delta t = 249.2\text{s}$)		$p = 3$ ($\Delta t = 153.4\text{s}$)		$p = 5$ ($\Delta t = 97.6\text{s}$)	
solver	n_{iter}	$t_{\text{dof,iter}}$	n_{iter}	$t_{\text{dof,iter}}$	n_{iter}	$t_{\text{dof,iter}}$	n_{iter}	$t_{\text{dof,iter}}$
DG+ApproxSchur	10.0	3.49	11.0	6.17	9.0	3.27	9.0	5.21
HDG+LU	-	*4.42	-	9.13	-	4.23	-	8.95
HDG+GTMG+AMG	*8.0	*4.42	8.0	9.13	8.0	4.20	7.3	8.95

Table 12: Average number n_{iter} of solver iterations and time $t_{\text{dof,iter}}$ (measured in μs) per time step and per unknown for the linear W5^(lin) (isolated mountain) and W6^(lin) (Rossby-Haurwitz) testcases. In all cases the refinement level is $r = 5$ and the SWEs are integrated to a final time of $T = 15\text{days}$ with the timestep size Δt given in the relevant column. The iterative solver tolerance is 10^{-8} . Results marked with an asterisk (*) were compiled with slightly different optimisation flags to avoid a known Intel compiler bug.

both W5 and W6. More specifically, $t_{\text{dof,iter}}$ is defined by

$$t_{\text{dof,iter}} = \frac{\text{total runtime}}{(\# \text{ unknowns}) \times (\# \text{ timesteps})}. \quad (54)$$

For the DG+ApproxSchur and HDG+GTMG+AMG solvers the average number n_{iter} of solver iterations required to reduce the relative residual by a factor of at least 10^{-8} is also recorded. Again, the timestep size is adjusted following Eq. (53) in all cases. Tab. 11 shows n_{iter} and $t_{\text{dof,iter}}$ for the non-linear test-cases W5 and W6, the corresponding results for the linear setups W5^(lin) and W6^(lin) are shown in Tab. 12. As the results show, the performance of the different solvers (as measured by $t_{\text{dof,iter}}$) differs by no more than a factor of around two. In contrast to what we found in Section 7.4 the non-hybridised DG+ApproxSchur solver gives the best overall performance. This might be due to inefficiencies in the Slate implementation in the spherical case and further investigation is required to understand this. Note, however, that the hybridised HDG+GTMG+AMG solver consistently requires less iterations to reduce the residual by the same factor. While for the parameters chosen here the iterative GTMG+AMG solver is not faster than the direct solver, several points should be taken into account when comparing those two methods: firstly, MUMPS [24, 25] is a highly optimised sparse direct solver, provided as a monolithic compiled library, whereas there are likely to be inefficiencies due to Python overheads in our Firedrake implementation of GTMG+AMG. Secondly, reducing the relative residual by at least eight orders of magnitude is probably overly cautious. In fact, we find that for the chosen tolerance the linear problem is solved to almost machine precision. Since the performance of the GTMG+AMG solver is directly proportional on the number of iterations, we expect it to perform better for tolerances of, say, 10^{-4} which are more relevant to atmospheric simulations. Finally, the runs presented in this section were carried out on a single compute node. While there is no reason why iterative solvers and multigrid preconditioners would not scale to very large problem sizes on massively parallel machines (see the review in [33]), this is less clear for direct solvers. Overall, we believe that our results show that our new methods provide a competitive alternative to existing methods.

8. Conclusion

Since IMEX methods are able to remove the tight timestep size restrictions due to fast modes, they are of particular interest when simulating large-scale geophysical flow. Combining IMEX time integrators with a hybridised DG discretisation in space leads to an elliptic system of equations for the flux-unknowns on the facets of the spatial grid. In this paper we described an efficient non-nested multigrid preconditioner for this elliptic problem by extending the algorithm in [20]. In contrast to the direct solver described in [11], we expect our approach to allow the efficient simulation of much larger problems. The first coarsening step projects the problem onto a conforming P1- or lowest-order Raviart-Thomas- space, where it can be solved with established geometric- or algebraic- multigrid methods. Compared to the standard DG discretisation with a naive Schur-complement preconditioner, we show numerically that our method is algorithmically robust under both p - and h -refinement. Since the problem in pressure space has fewer unknowns and converges in a smaller number of iterations, it leads to better performance overall for an idealised stationary test case in a flat domain. Compared to the standard DG approach, the hybridised method with our non-nested multigrid algorithm approximately halves the overall runtime for a linear solver tolerance of $\epsilon = 10^{-8}$ in this idealised setup.

For more complex flow problems in spherical geometry [23] the performance of our bespoke multigrid solver is comparable to that of the highly-optimised state-of-the-art direct MUMPS solver [24, 25] on a single compute node. However, we expect that further improvements to the implementation will eventually lead to superior performance of the non-nested multigrid when simulating at high resolution for large core counts, especially if the solver tolerance is reduced to levels that are more common in atmospheric simulations. In contrast to the simplified Euclidean setup, we find that for the considered test cases from the Williamson et al. suite the non-hybridised DG solver with an approximate Schur-complement preconditioner is faster than the HDG based method.

All code was implemented in the Firedrake library. In particular we used the recently developed Slate package [21] to realise the Schur-complement reduction to an elliptic system in flux-space. The PETSc library allowed the easy composition of our complex solver hierarchy. For an idealised flow problem good parallel scalability was demonstrated on up to 256 cores, in particular for higher-order DG discretisations. Strong scalability was good until the size of the local problem was reduced to a few thousand unknowns per core. The largest problems solved in the weak scaling runs had several million degrees of freedom.

While the numerical results imply that solving the elliptic flux system with a black-box AMG method gives approximately the same performance as our sophisticated non-nested multigrid preconditioner, the latter can have several advantages:

1. It does not suffer from the high setup costs, in particular for the coarse level operator construction via the Galerkin triple matrix-product.
2. If matrix-free smoothers for the flux-space are combined with a geometric multigrid solver for the coarse space problem, the memory requirements are significantly reduced, allowing the solution of much larger problems.

There are several ways of extending the work presented here. First of all, the performance of the non-nested multigrid preconditioner can be improved further by tailoring the individual components for the problem at hand. For example, recall that the bilinear form \hat{S} in Eq. (29) couples the unknowns on each facet and unknowns on facets which lie on the boundary of the same cell. Hence, the corresponding matrix has a block-structure. To exploit this structure, which is particularly important at higher polynomial degree p , one could use a block-Jacobi smoother in which the diagonal blocks are inverted exactly via LU-factorisation. In the future this could be replaced by the approximate matrix-free iterative inversion described in [50] which can lead to further speedups for high polynomial orders. The coarse level grad-div problem for the Lax-Friedrichs flux (the corresponding bilinear form is given in Eq. (35)) is currently solved with AMG. In the future the bespoke patch-based smoothers [51] could be used for this instead, which would result in a fully geometric method that can again potentially be optimised with a matrix-free implementation. While the numerical results presented in this work imply that our method is efficient and robust under both h - and p -refinement, this should be confirmed theoretically by adapting the proof in [20], as suggested in Section 4.3. Comparing our approach to time-explicit integrators, we find that IMEX methods are only slightly slower if the methods are chosen such that they have the same accuracy for our simplified setup. Furthermore, it is well known that the construction of the HDG

Schur-complement system and the reconstruction of the momentum- and height-perturbation field in each time step is currently not implemented very efficiently in Slate; improvements in future versions of Firedrake will likely remove this bottleneck and make the IMEX-HDG methods more competitive. Finally, the non-nested multigrid approach should also be applied to more complicated physical problems, such as the compressible Euler- and Navier-Stokes equations, for which a hybridised DG discretisation has been developed in [52]. Ultimately it is of course desirable to use the computationally most efficient discretisation. For this, it is necessary to compare the performance of the IMEX-HDG methods with the multigrid preconditioners described in this paper to other approaches such as spectral- and finite-volume discretisations. While such a comparison is of course extremely important, it is a major undertaking which is well beyond the scope of this paper and which should be pursued in a future publication.

Acknowledgments

This research made use of the Balena High Performance Computing (HPC) Service at the University of Bath. The PhD of Jack Betteridge was funded by EPSRC as part of the SAMBa CDT in Bath, grant number EP/L015684/1. We would like to thank all Firedrake developers for their continuous support, in particular Lawrence Mitchell for help with the multigrid implementation. We are grateful to Patrick Farrell (Oxford) for helpful comments. Part of this work was carried out during a short research visit (SRV) to Imperial College in May 2018, funded by the UK-Fluids network (EPSRC grant EP/N032861/1).

Appendix A. Bilinear forms for standard DG

The standard DG discretisation of the bilinear form $\widehat{\mathcal{A}}(q, v) = \mathcal{M}(q, v) - c_g \alpha \Delta t \mathcal{L}(q, v)$ is given for the Lax-Friedrichs flux in Eq. (A.1) and for the upwind flux in Eq. (A.2). Those two expressions, which only differ by one term ($\langle\langle \mathbf{u} \rangle\rangle \cdot \langle\langle \mathbf{w} \rangle\rangle$ vs. $[[\mathbf{u}]][\mathbf{w}]$), should be compared to the corresponding HDG versions in Eqs. (25) and (24).

Lax-Friedrichs flux:

$$\begin{aligned} \mathcal{A}^{(\text{LF})}(q, v) = \mathcal{A}^{(\text{LF})}(\phi, \mathbf{u}, \psi, \mathbf{w}) &= (\phi\psi + \mathbf{u} \cdot \mathbf{w})_{\Omega_h} - c_g \alpha \Delta t \left[(\mathbf{u} \cdot \nabla \psi + \phi \nabla \cdot \mathbf{w})_{\Omega_h} + \mathcal{K}(\mathbf{u}, \mathbf{w}) \right. \\ &\quad \left. - \left(\frac{1}{2} \langle\langle \mathbf{u} \cdot \mathbf{n} \rangle\rangle \langle\langle \psi \rangle\rangle + \frac{1}{2} \sqrt{\phi_B} (\langle\langle \phi \rangle\rangle \langle\langle \psi \rangle\rangle + \langle\langle \mathbf{u} \rangle\rangle \cdot \langle\langle \mathbf{w} \rangle\rangle) + \phi_B \{\{\phi\}\} [[\mathbf{w}]] \right) \right]_{\varepsilon_h} \end{aligned} \quad (\text{A.1})$$

with the ‘‘curvature’’ term

$$\mathcal{K}(\mathbf{u}, \mathbf{w}) = - \left(\frac{1}{2} \sqrt{\phi_B} ((\mathbf{u}_+ \cdot \mathbf{n}_+) (\mathbf{n}_+ + \mathbf{n}_-) \cdot \mathbf{w}_- + (\mathbf{u}_- \cdot \mathbf{n}_-) (\mathbf{n}_+ + \mathbf{n}_-) \cdot \mathbf{w}_+) \right)_{\varepsilon_h}$$

which can be derived following [49]. Note that $\mathcal{K}(\mathbf{u}, \mathbf{w}) = 0$ for flat geometries where $\mathbf{n}_- = -\mathbf{n}_+$.

Upwind flux:

$$\begin{aligned} \mathcal{A}^{(\text{up})}(q, v) = \mathcal{A}^{(\text{up})}(\phi, \mathbf{u}, \psi, \mathbf{w}) &= (\phi\psi + \mathbf{u} \cdot \mathbf{w})_{\Omega_h} - c_g \alpha \Delta t \left[(\mathbf{u} \cdot \nabla \psi + \phi \nabla \cdot \mathbf{w})_{\Omega_h} \right. \\ &\quad \left. - \left(\frac{1}{2} \langle\langle \mathbf{u} \cdot \mathbf{n} \rangle\rangle \langle\langle \psi \rangle\rangle + \frac{1}{2} \sqrt{\phi_B} (\langle\langle \phi \rangle\rangle \langle\langle \psi \rangle\rangle + [[\mathbf{u}]][\mathbf{w}]) + \phi_B \{\{\phi\}\} [[\mathbf{w}]] \right) \right]_{\varepsilon_h} \end{aligned} \quad (\text{A.2})$$

Appendix B. Linear operators

Let

$$\mathcal{A}_0(\phi, \mathbf{u}; \psi, \mathbf{w}) = (\phi\psi + \mathbf{u} \cdot \mathbf{w})_{\Omega_h} - c_g \alpha \Delta t [(\mathbf{u} \cdot \nabla \psi + \phi \nabla \cdot \mathbf{w})_{\Omega_h}].$$

The linear operators defined in Eq. (26) are given the following expressions for the two fluxes considered in this paper:

Lax-Friedrichs flux

$$\begin{aligned}
(\mathfrak{A}^{(\text{LF})} q, v)_{\Omega_h} &= \mathcal{A}_0(\phi, \mathbf{u}; \psi, \mathbf{w}) + c_g \alpha \Delta t \left(2\sqrt{\phi_B} \{\!\{ \mathbf{u} \cdot \mathbf{w} \}\!\} + \phi_B \llbracket \phi \mathbf{w} \rrbracket \right)_{\Omega_h} \\
(\widehat{q}, (\mathfrak{B}^{(\text{LF})})^\top v)_{\mathcal{E}_h} &= c_g \alpha \Delta t \left(\widehat{\mathbf{u}} \cdot \left(\llbracket \psi \rrbracket - 2\sqrt{\phi_B} \{\!\{ \mathbf{w} \}\!\} \right) \right)_{\mathcal{E}_h} \\
(\mathfrak{C}^{(\text{LF})} q, \widehat{v})_{\mathcal{E}_h} &= \left(\widehat{\mathbf{w}} \cdot \left(\phi_B \llbracket \phi \rrbracket + 2\sqrt{\phi_B} \{\!\{ \mathbf{u} \}\!\} \right) \right)_{\mathcal{E}_h} \\
(\mathfrak{M}^{(\text{LF})} \widehat{q}, \widehat{v})_{\mathcal{E}_h} &= (\widehat{\mathbf{w}} \cdot \widehat{\mathbf{u}})_{\mathcal{E}_h}
\end{aligned}$$

upwind flux

$$\begin{aligned}
(\mathfrak{A}^{(\text{up})} q, v)_{\Omega_h} &= \mathcal{A}_0(\phi, \mathbf{u}; \psi, \mathbf{w}) + c_g \alpha \Delta t \left(2\sqrt{\phi_B} \{\!\{ \phi \psi \}\!\} + \llbracket \mathbf{u} \psi \rrbracket \right)_{\Omega_h} \\
(\widehat{q}, (\mathfrak{B}^{(\text{up})})^\top v)_{\mathcal{E}_h} &= c_g \alpha \Delta t \left(\widehat{\phi} \left(\phi_B \llbracket \mathbf{w} \rrbracket - 2\sqrt{\phi_B} \{\!\{ \psi \}\!\} \right) \right)_{\mathcal{E}_h} \\
(\mathfrak{C}^{(\text{up})} q, \widehat{v})_{\mathcal{E}_h} &= \left(\widehat{\psi} \cdot \left(\llbracket \mathbf{u} \rrbracket + 2\sqrt{\phi_B} \{\!\{ \phi \}\!\} \right) \right)_{\mathcal{E}_h} \\
(\mathfrak{M}^{(\text{up})} \widehat{q}, \widehat{v})_{\mathcal{E}_h} &= \left(\widehat{\psi} \widehat{\phi} \right)_{\mathcal{E}_h}
\end{aligned}$$

Appendix C. Runge Kutta methods

Generally, an explicit s -stage Runge Kutta (RK) method can be written as

$$\begin{aligned}
\mathcal{M}(Q^{(i)}, v) &= \mathcal{M}(q^{(n)}, v) + \Delta t \sum_{j=1}^{i-1} \bar{a}_{ij} \left(\mathcal{N}(Q^{(j)}, v) + \mathcal{L}(Q^{(j)}, v) \right), \quad \text{for } i = 1, \dots, s \text{ and all } v \in W_h, \\
\mathcal{M}(q^{(n+1)}, v) &= \mathcal{M}(q^{(n)}, v) + \Delta t \sum_{i=1}^s \bar{b}_i \left(\mathcal{N}(Q^{(i)}, v) + \mathcal{L}(Q^{(i)}, v) \right), \quad \text{for all } v \in W_h.
\end{aligned} \tag{C.1}$$

Since the $\widehat{\mathcal{L}}$ in (19) does not make a contribution in the case of explicit integrators, it is not necessary to use hybridisation, i.e. only fields in the standard DG space W_h are required. In general, s evaluations of $\mathcal{N} + \mathcal{L}$ and $s + 1$ mass-solves on each element are required per timestep. For reference, the matrices \bar{a} and \bar{b} for the explicit RKDG methods used in this paper are given in the following equation:

$$\begin{aligned}
\text{Heun's method:} \quad \bar{a} &= \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix}, \quad \bar{b} = \begin{pmatrix} 1/2 \\ 1/2 \end{pmatrix} \\
\text{SSPRK3:} \quad \bar{a} &= \begin{pmatrix} 0 & 0 & 0 \\ 1 & 0 & 0 \\ 1/4 & 1/4 & 0 \end{pmatrix}, \quad \bar{b} = \begin{pmatrix} 1/6 \\ 1/6 \\ 2/3 \end{pmatrix}
\end{aligned}$$

Appendix D. Solver configuration

Appendix D.1. Hybridised DG

To solve the problem in Eq. (23), with the HDG+GTMG+AMG solver configuration described in Section 5.2.1, the following parameter dictionary is passed to the Firedrake `LinearVariationalSolver` object:

```

param = {'mat_type': 'matfree',
         'ksp_type': 'preonly',
         'pc_type': 'python',
         'pc_python_type': 'firedrake.SPCP',
         'pc_sc_eliminate_fields': '0, 1',
         'condensed_field': {'ksp_type': inner_ksp_type,
                             'ksp_rtol': rtol,
                             'mat_type': 'aij',
                             'pc_type': 'python',
                             'pc_python_type': 'firedrake.GTMGPC',

```

```

        'gt': {'mat_type': 'aij',
              'mg_levels': {'ksp_type': 'chebyshev',
                           'pc_type': 'bjacobi',
                           'sub_pc_type': 'sor',
                           'ksp_max_it': 2},
              'mg_coarse': {'ksp_type': 'preonly',
                           'pc_type': "gamg",
                           'ksp_rtol': rtol,
                           'pc_mg_cycles': 'v',
                           'mg_levels':
                               {'ksp_type': 'chebyshev',
                                'ksp_max_it': 2,
                                'pc_type': 'sor'},
                           'mg_coarse':
                               {'ksp_type': 'chebyshev',
                                'ksp_max_it': 2,
                                'pc_type': 'sor'}}
        }
    }
}

```

Here `rtol` is the relative tolerance (on the preconditioned residual) to which the linear system is solved. The solver for the Schur-complement system (29) is either GMRES for the non-linear shallow-water equations or CG in the linearised case. This is controlled by setting the variable `inner_ksp_type` to `gmres` or `cg`. To use the HDG+AMG solver instead, the `'condensed_field'` dictionary in the code above is replaced by:

```

'condensed_field': {'ksp_type': inner_ksp_type,
                   'ksp_rtol': rtol,
                   'mat_type': 'aij',
                   'pc_type': 'gamg',
                   'mg_levels': {'ksp_type': 'chebyshev',
                                'ksp_max_it': 2,
                                'pc_type': 'bjacobi',
                                'sub_pc_type': 'sor'}}
}

```

To use the HDG+GMG solver (for the upwind flux) this becomes

```

'condensed_field': {'ksp_type': inner_ksp_type,
                   'ksp_rtol': rtol,
                   'mat_type': 'aij',
                   'pc_type': 'python',
                   'pc_python_type': 'firedrake.GTMGPC',
                   'gt': {'mat_type': 'aij',
                          'mg_levels': {'ksp_type': 'chebyshev',
                                       'ksp_max_it': 2,
                                       'pc_type': 'bjacobi',
                                       'sub_pc_type': 'sor'}}
                   }
}

```

Appendix D.2. Native DG

To solve the native DG system $\mathcal{A}(q, v) = \mathcal{R}(v)$ in Eq. (38) with the DG+ApproxSchur solver, the following parameters are used:

```

solver_param = {'ksp_type': 'gmres',
               'ksp_rtol': rtol,
               'pc_type': 'fieldsplit',
               'pc_fieldsplit_type': 'schur',
               'pc_fieldsplit_schur_fact_type': 'FULL',
               'pc_fieldsplit_schur_precondition': 'selfp',
               'fieldsplit_0': {'ksp_type': 'preonly',
                               'pc_type': 'gamg',
                               'mg_levels': {'ksp_type': 'chebyshev',
                                             'ksp_max_it': 2,
                                             'pc_type': 'bjacobi',
                                             'sub_pc_type': 'sor'}}
}

```

```

        'ksp_max_it': 2,
        'pc_type': 'bjacobi',
        'sub_pc_type': 'sor'}
    'fieldsplit_1': {'ksp_type': 'preonly',
                    'pc_type': 'bjacobi',
                    'sub_pc_type': 'ilu'},
                    }
}

```

As above, the system is solved to a relative tolerance `rtol`.

Appendix D.3. Numerical parameters for Williamson et al. testcases

Tab. D.13 show the parameters used for the numerical experiments in Section 7.9.

parameter	value
earth radius	$R_{\text{earth}} = 6.37122 \cdot 10^6 \text{m}$
angular rotation frequency	$\Omega = 7.292 \cdot 10^{-5} \text{s}^{-1}$
reference height (see Section 2)	$\bar{H}_B = \begin{cases} 3000\text{m} & \text{for W3} \\ 5960\text{m} & \text{for W5} \\ 8000\text{m} & \text{for W6} \end{cases}$
velocity	$u_0 = \begin{cases} 2\pi R_{\text{earth}}/(12\text{days}) & \text{for W3} \\ 20\text{ms}^{-1} & \text{for W5} \end{cases}$
zonal flow parameters for W3	$\theta_b = \pi/6, \theta_e = \pi/2, x_e = 0.3$
mountain height and radius for W5	$h_{s_0} = 2000\text{m}, R = \pi/9$
latitude and longitude of mountain for W5	$\lambda_c = 3\pi/2, \theta_c = \pi/6$
parameters for W6	$\omega = K = 7.848 \cdot 10^{-6} \text{s}^{-1}$

Table D.13: Numerical parameters used for numerical experiments in Section 7.9.

References

- [1] N. Wood, A. Staniforth, A. White, T. Allen, M. Diamantakis, M. Gross, T. Melvin, C. Smith, S. Vosper, M. Zerroukat, et al., An inherently mass-conserving semi-implicit semi-Lagrangian discretization of the deep-atmosphere global non-hydrostatic equations, *Quarterly Journal of the Royal Meteorological Society* 140 (682) (2014) 1505–1520.
- [2] U. M. Ascher, S. J. Ruuth, B. T. Wetton, Implicit-explicit methods for time-dependent partial differential equations, *SIAM Journal on Numerical Analysis* 32 (3) (1995) 797–823.
- [3] U. M. Ascher, S. J. Ruuth, R. J. Spiteri, Implicit-explicit Runge-Kutta methods for time-dependent partial differential equations, *Applied Numerical Mathematics* 25 (2-3) (1997) 151–167.
- [4] L. Pareschi, G. Russo, Implicit-explicit Runge-Kutta schemes for stiff systems of differential equations, *Recent trends in numerical analysis* 3 (2000) 269–289.
- [5] C. A. Kennedy, M. H. Carpenter, Additive Runge–Kutta schemes for convection–diffusion–reaction equations, *Applied Numerical Mathematics* 44 (1-2) (2003) 139–181.
- [6] H. Weller, S.-J. Lock, N. Wood, Runge–Kutta IMEX schemes for the horizontally explicit/vertically implicit (HEVI) solution of wave equations, *Journal of Computational Physics* 252 (2013) 365–381.
- [7] A. Staniforth, J. Thuburn, Horizontal grids for global weather and climate prediction models: a review, *Quarterly Journal of the Royal Meteorological Society* 138 (662) (2012) 1–26.
- [8] L. Mitchell, E. H. Müller, High level implementation of geometric multigrid solvers for finite element problems: Applications in atmospheric modelling, *Journal of Computational Physics* 327 (2016) 1–18.
- [9] S. Marras, J. F. Kelly, M. Moragues, A. Müller, M. A. Kopera, M. Vázquez, F. X. Giraldo, G. Houzeaux, O. Jorba, A review of element-based Galerkin methods for numerical weather prediction: Finite elements, spectral elements, and discontinuous Galerkin, *Archives of Computational Methods in Engineering* 23 (4) (2016) 673–722.
- [10] S. Müthing, M. Piatkowski, P. Bastian, High-performance implementation of matrix-free high-order discontinuous galerkin methods, *arXiv preprint arXiv:1711.10885* (2017).
- [11] S. Kang, F. X. Giraldo, T. Bui-Thanh, IMEX HDG-DG: A coupled implicit hybridized discontinuous Galerkin and explicit discontinuous Galerkin approach for shallow water systems, *Journal of Computational Physics* 401 (2020) 109010.
- [12] B. Cockburn, B. Dong, J. Guzmán, M. Restelli, R. Sacco, A hybridizable discontinuous Galerkin method for steady-state convection-diffusion-reaction problems, *SIAM Journal on Scientific Computing* 31 (5) (2009) 3827–3846.
- [13] B. Cockburn, J. Gopalakrishnan, R. Lazarov, Unified hybridization of discontinuous Galerkin, mixed, and continuous Galerkin methods for second order elliptic problems, *SIAM Journal on Numerical Analysis* 47 (2) (2009) 1319–1365.
- [14] H. Egger, J. Schöberl, A hybrid mixed discontinuous Galerkin finite-element method for convection–diffusion problems, *IMA Journal of Numerical Analysis* 30 (4) (2009) 1206–1234.
- [15] N. C. Nguyen, J. Peraire, B. Cockburn, An implicit high-order hybridizable discontinuous Galerkin method for linear convection–diffusion equations, *Journal of Computational Physics* 228 (9) (2009) 3232–3254.
- [16] T. Bui-Thanh, From Godunov to a unified hybridized discontinuous Galerkin framework for partial differential equations, *Journal of Computational Physics* 295 (2015) 114–146.
- [17] T. Bui-Thanh, Construction and analysis of HDG methods for linearized shallow water equations, *SIAM Journal on Scientific Computing* 38 (6) (2016) A3696–A3719.

- [18] I. S. Duff, A. M. Erisman, J. K. Reid, *Direct methods for sparse matrices*, Oxford University Press, 2017.
- [19] J. Gopalakrishnan, S. Tan, A convergent multigrid cycle for the hybridized mixed method, *Numerical Linear Algebra with Applications* 16 (9) (2009) 689–714.
- [20] B. Cockburn, O. Dubois, J. Gopalakrishnan, S. Tan, Multigrid for an HDG method, *IMA Journal of Numerical Analysis* 34 (4) (2014) 1386–1425.
- [21] T. H. Gibson, L. Mitchell, D. A. Ham, C. J. Cotter, [Slate: extending Firedrake’s domain-specific abstraction to hybridized solvers for geoscience and beyond](#). (2020). [arXiv:1802.00303](#), [doi: 10.5194/gmd-13-735-2020](#).
URL <https://arxiv.org/abs/1802.00303>
- [22] F. Rathgeber, D. A. Ham, L. Mitchell, M. Lange, F. Luporini, A. T. McRae, G.-T. Bercea, G. R. Markall, P. H. Kelly, Firedrake: automating the finite element method by composing abstractions, *ACM Transactions on Mathematical Software (TOMS)* 43 (3) (2017) 24.
- [23] D. L. Williamson, J. B. Drake, J. J. Hack, R. Jakob, P. N. Swarztrauber, A standard test set for numerical approximations to the shallow water equations in spherical geometry, *Journal of Computational Physics* 102 (1) (1992) 211–224.
- [24] P. Amestoy, I. S. Duff, J. Koster, J.-Y. L’Excellent, A Fully Asynchronous Multifrontal Solver Using Distributed Dynamic Scheduling, *SIAM Journal on Matrix Analysis and Applications* 23 (1) (2001) 15–41.
- [25] P. Amestoy, A. Buttari, J.-Y. L’Excellent, T. Mary, Performance and Scalability of the Block Low-Rank Multifrontal Factorization on Multicore Architectures, *ACM Transactions on Mathematical Software* 45 (2019) 2:1–2:26.
- [26] T. Wildey, S. Muralikrishnan, T. Bui-Thanh, Unified geometric multigrid algorithm for hybridized high-order finite element methods, *arXiv preprint arXiv:1811.09909* (2018).
- [27] S. Muralikrishnan, T. Bui-Thanh, J. N. Shadid, A Multilevel Approach for Trace System in HDG Discretizations, *arXiv preprint arXiv:1903.11045* (2019).
- [28] V. V. Rusanov, The calculation of the interaction of non-stationary shock waves with barriers, *Zhurnal Vychislitel’noi Matematiki i Matematicheskoi Fiziki* 1 (2) (1961) 267–279.
- [29] R. J. LeVeque, *Finite volume methods for hyperbolic problems*, Vol. 31, Cambridge university press, 2002.
- [30] T. Melvin, T. Benacchio, B. Shipway, N. Wood, J. Thuburn, C. Cotter, A mixed finite-element, finite-volume, semi-implicit discretization for atmospheric dynamics: Cartesian geometry, *Quarterly Journal of the Royal Meteorological Society* 145 (724) (2019) 2835–2853.
- [31] J. S. Hesthaven, T. Warburton, *Nodal discontinuous Galerkin methods: algorithms, analysis, and applications*, Springer Science & Business Media, 2007.
- [32] J. Betteridge, *Efficient elliptic solvers for higher order DG discretisations on modern architectures and applications in atmospheric modelling*, Ph.D. thesis, University of Bath (2019).
- [33] E. H. Müller, R. Scheichl, Massively parallel solvers for elliptic partial differential equations in numerical weather and climate prediction, *Quarterly Journal of the Royal Meteorological Society* 140 (685) (2014) 2608–2624.
- [34] S. Sandbach, J. Thuburn, D. Vassilev, M. G. Duda, A semi-implicit version of the MPAS-Atmosphere dynamical core, *Monthly Weather Review* 143 (9) (2015) 3838–3855.
- [35] P.-A. Raviart, J.-M. Thomas, A mixed finite element method for 2-nd order elliptic problems, in: *Mathematical aspects of finite element methods*, Springer, 1977, pp. 292–315.

- [36] J. Crank, P. Nicolson, A practical method for numerical evaluation of solutions of partial differential equations of the heat-conduction type, in: *Mathematical Proceedings of the Cambridge Philosophical Society*, Vol. 43, Cambridge University Press, 1947, pp. 50–67.
- [37] L. Pareschi, G. Russo, Implicit–explicit Runge–Kutta schemes and applications to hyperbolic systems with relaxation, *Journal of Scientific computing* 25 (1) (2005) 129–155.
- [38] M. R. Hestenes, E. Stiefel, *Methods of conjugate gradients for solving linear systems*, Vol. 49, NBS Washington, DC, 1952.
- [39] Y. Saad, M. H. Schultz, GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems, *SIAM Journal on scientific and statistical computing* 7 (3) (1986) 856–869.
- [40] S. Balay, S. Abhyankar, M. F. Adams, J. Brown, P. Brune, K. Buschelman, L. Dalcin, V. Eijkhout, W. D. Gropp, D. Kaushik, M. G. Knepley, L. C. McInnes, K. Rupp, B. F. Smith, S. Zampini, H. Zhang, *PETSc users manual*, Tech. Rep. ANL-95/11 - Revision 3.6, Argonne National Laboratory (2015).
- [41] C.-W. Shu, S. Osher, Efficient implementation of essentially non-oscillatory shock-capturing schemes, *Journal of computational physics* 77 (2) (1988) 439–471.
- [42] B. Cockburn, C.-W. Shu, TVB Runge-Kutta local projection discontinuous Galerkin finite element method for conservation laws. II. General framework, *Mathematics of computation* 52 (186) (1989) 411–435.
- [43] B. Cockburn, C.-W. Shu, The Runge–Kutta discontinuous Galerkin method for conservation laws V: multidimensional systems, *Journal of Computational Physics* 141 (2) (1998) 199–224.
- [44] M. S. Alnæs, A. Logg, K. B. Ølgaard, M. E. Rognes, G. N. Wells, Unified Form Language: A Domain-specific Language for Weak Formulations of Partial Differential Equations, *ACM Transactions on Mathematical Software* 40 (2) (2014) 9:1–9:37. [arXiv:1211.4047](https://arxiv.org/abs/1211.4047).
- [45] S. Balay, W. D. Gropp, L. C. McInnes, B. F. Smith, Efficient management of parallelism in object oriented numerical software libraries, in: E. Arge, A. M. Bruaset, H. P. Langtangen (Eds.), *Modern Software Tools in Scientific Computing*, Birkhäuser Press, 1997, pp. 163–202.
- [46] R. C. Kirby, L. Mitchell, Solver composition across the PDE/linear algebra barrier, *SIAM Journal on Scientific Computing* 40 (1) (2018) C76–C98.
- [47] B. Cockburn, C.-W. Shu, Runge–Kutta discontinuous Galerkin methods for convection-dominated problems, *Journal of scientific computing* 16 (3) (2001) 173–261.
- [48] J. Betteridge, I. Graham, T. H. Gibson, E. H. Mueller, [Code Release: Multigrid preconditioners for the hybridized Discontinuous Galerkin discretisation of the shallow water equations](https://arxiv.org/abs/2010.12345) (Oct. 2020). [doi:10.5281/zenodo.4081310](https://doi.org/10.5281/zenodo.4081310). URL <https://doi.org/10.5281/zenodo.4081310>
- [49] P.-E. Bernard, J.-F. Remacle, R. Comblen, V. Legat, K. Hillewaert, High-order discontinuous Galerkin schemes on general 2D manifolds applied to the shallow water equations, *Journal of Computational Physics* 228 (17) (2009) 6514–6535.
- [50] P. Bastian, E. H. Müller, S. Müthing, M. Piatkowski, Matrix-free multigrid block-preconditioners for higher order Discontinuous Galerkin discretisations, *Journal of Computational Physics* (2019).
- [51] P. E. Farrell, L. Mitchell, F. Wechsung, An Augmented Lagrangian Preconditioner for the 3D Stationary Incompressible Navier–Stokes Equations at High Reynolds Number, *SIAM Journal on Scientific Computing* 41 (5) (2019) A3073–A3096.
- [52] J. Peraire, N. Nguyen, B. Cockburn, A hybridizable discontinuous Galerkin method for the compressible Euler and Navier-Stokes equations, in: *48th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition*, 2010, p. 363.