




Goal-oriented error estimation and mesh adaptation for shallow water modelling



Joseph G. Wallwork¹  · Nicolas Barral^{2,3}  · Stephan C. Kramer¹  · David A. Ham⁴  · Matthew D. Piggott¹ 

Received: 21 December 2019 / Accepted: 11 April 2020 / Published online: 10 May 2020
© The Author(s) 2020 

Abstract

This study presents a novel goal-oriented error estimate for the nonlinear shallow water equations solved using a mixed discontinuous/continuous Galerkin approach. This error estimator takes account of the discontinuities in the discrete solution and is used to drive two metric-based mesh adaptation algorithms: one which yields isotropic meshes and another which yields anisotropic meshes. An implementation of these goal-oriented mesh adaptation algorithms is described, including a method for approximating the adjoint error term which arises in the error estimate. Results are presented for simulations of two model tidal farm configurations computed using the *Thetis* coastal ocean model (Kärnä et al. in *Geosci Model Dev* 11(11):4359–4382, 2018). Convergence analysis indicates that meshes resulting from the goal-oriented adaptation strategies permit accurate QoI estimation using fewer computational resources than uniform refinement.

Keywords Adjoint methods · Mesh adaptation · Discontinuous Galerkin · Tidal turbine modelling · Firedrake

Mathematics Subject Classification 65N50

1 Introduction

Coastal modelling problems are typically multi-scale, often with a strongly direction-dependent flow. As such, anisotropic mesh adaptation is an attractive prospect for providing both accuracy and reduced computational cost; by adapting the mesh such that its anisotropy is aligned with the flow, the number of degrees of freedom (DOFs) required to yield an accurate solution is reduced. The *metric based* approach to anisotropic mesh adaptation was first introduced in [20] and uses Riemannian metric fields to control not only the size of mesh elements, but also their shape and orientation. This approach was shown to be particularly suited to multi-scale ocean modelling in [31]. For a review of recent progress in the field of anisotropic mesh adaptation, see [1].

Further, coastal modelling problems often involve a diagnostic *quantity of interest (QoI)* of greater importance than the solution itself. *Goal-oriented error estimation* represents the error accrued in computing the QoI in terms of PDE residuals and solutions of associated adjoint equations. Used within a mesh adaptation algorithm, such estimators can yield meshes permitting accurate QoI approximation. The majority of goal-oriented error estimators are based upon the pioneering work of [9, 10]. More recently, researchers have developed goal-oriented estimators which take account of discontinuous Galerkin (DG) discretisations (see [15, 22]). Integration into the metric based mesh adaptation framework has also been developed (see [12, 27, 32]).

The depth-averaged *shallow water equations* are often used in coastal ocean models, providing approximations

✉ Joseph G. Wallwork, j.wallwork16@imperial.ac.uk | ¹Earth Science and Engineering, Imperial College London, London, UK. ²University of Bordeaux, CNRS, Bordeaux INP, IMB, UMR 5251, 33400 Talence, France. ³INRIA, IMB, UMR 5251, 33400 Talence, France. ⁴Mathematics, Imperial College London, London, UK.



to the fluid velocity and elevation of the ocean surface. This work builds upon the anisotropic goal-oriented mesh adaptation research referenced above, focusing on the shallow water equations. To the best of the authors' knowledge, this work presents the first formulation of a goal-oriented error estimate for the shallow water equations discretised in a mixed discontinuous/continuous space.

The shallow water equations are solved using the *Thetis* coastal ocean modelling framework [24], which is based upon the finite element library *Firedrake* [33]. Underlying linear and nonlinear systems are solved using *PETSc* [5, 6]. As well as a 2D shallow water model [35], *Thetis* offers 3D Navier-Stokes solvers with Boussinesq and hydrostatic [24] and non-hydrostatic [30] assumptions and a tracer transport model.

Goal-oriented error estimation is introduced in Sect. 2, along with two approaches to goal-oriented mesh adaptation. Application of the error estimation and adaptation techniques to shallow water problems is considered in Sect. 3. Section 4 contains numerical experiments validating the adaptation strategy for a simple tidal turbine model. The results from Sect. 4 are discussed in Sect. 5 and conclusions are drawn in Sect. 6.

2 Goal-oriented mesh adaptation

2.1 Metric-based mesh adaptation

In this paper, *Riemannian metric fields* are used to drive the mesh adaptation process. These tensor fields, often referred to simply as *metrics*, are symmetric positive-definite (SPD) linear forms defined pointwise, which give rise to all of the geometrical quantities necessary to perform mesh adaptation. In this work, metrics are derived using the error estimates described in Sect. 2.2. For details on the metric-based approach, see [1, 7, 20, 29, 34].

The spatial domain, denoted by $\Omega \subset \mathbb{R}^n$, is assumed to have piecewise smooth boundary $\Gamma := \partial\Omega$. For a mesh \mathcal{H} of the domain, we denote mesh elements by $K \in \mathcal{H}$ and the edge set of element K by ∂K . In this work we restrict attention to triangular mesh elements, for simplicity. Denote the set of all edges which are not on the domain boundary (internal edges) by Γ_{int} .

2.2 Error estimation

The shallow water equations may be written in the 'residual form'

$$\Psi(\mathbf{q}) = \mathbf{0}, \tag{1}$$

with solution \mathbf{q} living in a space of functions denoted V . Throughout this paper, we shall refer to (1) as the *forward equation* and to \mathbf{q} as the *forward solution*. The diagnostic QoI, J , is a functional which maps members of V onto the real number line. Associated with (1) and the QoI is an adjoint equation,

$$\frac{\partial \Psi^T}{\partial \mathbf{q}} \mathbf{q}^* = \frac{\partial J^T}{\partial \mathbf{q}}. \tag{2}$$

The adjoint solution \mathbf{q}^* also lives in V and conveys the propagation of sensitivities of the QoI to perturbations in the forward solution. For a finite dimensional subspace $V_h \subset V$, (1)–(2) have Galerkin approximations given by

$$\rho(\mathbf{q}_h, \xi) = -\langle \Psi(\mathbf{q}_h), \xi \rangle = 0, \quad \forall \xi \in V_h, \tag{3}$$

$$\rho^*(\mathbf{q}_h^*, \xi) = \left\langle \frac{\partial J^T}{\partial \mathbf{q}} - \frac{\partial \Psi^T}{\partial \mathbf{q}} \mathbf{q}_h^*, \xi \right\rangle = 0, \quad \forall \xi \in V_h, \tag{4}$$

where $\langle \cdot, \cdot \rangle$ is the \mathcal{L}_2 inner product and \mathbf{q}_h and \mathbf{q}_h^* are finite element approximations to the forward and adjoint solutions. Typically, integration by parts is applied in constructing the weak residuals $\rho(\mathbf{q}_h, \cdot)$ and $\rho^*(\mathbf{q}_h^*, \cdot)$ from the inner products.

We consider the classical a posteriori goal-oriented error estimate known as the *dual weighted residual*, due to [9, 10]. Therein, the error result

$$J(\mathbf{q}) - J(\mathbf{q}_h) = \rho(\mathbf{q}_h, \mathbf{q}^* - \mathbf{q}_h^*) + R^{(2)}, \tag{5}$$

is presented, where the remainder term $R^{(2)}$ is quadratic in the forward and adjoint errors $\mathbf{e} := \mathbf{q} - \mathbf{q}_h$ and $\mathbf{e}^* := \mathbf{q}^* - \mathbf{q}_h^*$. Element-wise dual weighted residual error indicators and a global error estimator may be derived as

$$\mathcal{E}_K = |\rho(\mathbf{q}_h, \mathbf{e}^*)|_K, \quad \mathcal{E} = \sum_{K \in \mathcal{H}} \mathcal{E}_K. \tag{6}$$

2.3 Goal-oriented metrics

There are many potential ways to construct metric tensor fields using the scalar error indicator (6), one being to appropriately scale an identity matrix; this is referred to as an *isotropic* approach, since resulting meshes are relatively isotropic. For problems with strong directional dependence, *anisotropic* metrics can be beneficial, allowing control of the shape and orientation of mesh elements, as well as size.

In [37], an isotropic metric was compared with two approaches to constructing anisotropic metrics from goal-oriented error estimates (based on the work of [27, 32]) for

advection-diffusion problems discretised using continuous finite elements. In this subsection, we utilise a different approach, developed in [12], which straightforwardly permits discontinuous discretisations.

2.3.1 Isotropic metric

First, consider the isotropic case. Let $|\hat{K}|$ denote the reference element volume and $|K|$ denote the volume of an arbitrary element $K \in \mathcal{H}$. Constructing an element-based isotropic metric in n -dimensions amounts to the scaling [12]

$$\underline{\underline{\mathbf{M}}}_K = \frac{|\hat{K}|}{|\tilde{K}|} \underline{\underline{\mathbf{I}}}_n, \quad |\tilde{K}| = |K| \left(\frac{\sum_{K \in \mathcal{H}} \mathcal{E}_K^{\frac{1}{\alpha+1}}}{N} \right) \mathcal{E}_K^{-\frac{1}{\alpha+1}}, \quad (7)$$

where $\underline{\underline{\mathbf{I}}}_n$ is the element-wise identity metric. The desired element volume $|\tilde{K}|$ is chosen to minimise interpolation error, such that the *metric complexity* $N > 0$ is smaller than some desired complexity, N_{\max} . Metric complexity can be viewed as the continuous analogue of the mesh vertex count [1, 7]. The proof that $|\tilde{K}|$ solves this optimisation problem is given in [11]. The parameter $\alpha \geq 1$ which arises in its solution is not known a priori. However, [12] states that its influence on (7) is negligible, provided that we are sufficiently close to the optimal volume. We have found $\alpha = 1$ to be an effective choice in practice.

It is worth remarking that $|\tilde{K}|$ is mesh-dependent. This means that meshes generated using metric $\underline{\underline{\mathbf{M}}}_K$ are heavily influenced by the mesh upon which the metric was constructed. However, as detailed in Sect. 3.4, the mesh adaptation algorithm used in this work iteratively solves the PDE, evaluates (7) and adapts the mesh until convergence criteria are met. As such, the dependence on the initial mesh diminishes as the algorithm progresses. This is in agreement with what we have found in numerical experiments. For further discussion of the coupled mesh adaptation-PDE solution process, see [1].

2.3.2 Anisotropic metric

As with the approaches considered in [37], the anisotropic metric construction of [12] uses a recovered Hessian of the prognostic variables. In this instance, we compute the element-averaged Hessian $\underline{\underline{\mathbf{H}}}_K$ on an element K by solving an auxiliary finite element problem (for details, see [29]).

As a symmetric matrix, the Hessian has an orthogonal eigen-decomposition $\underline{\underline{\mathbf{H}}}_K = \underline{\underline{\mathbf{V}}}_K \underline{\underline{\Lambda}}_K \underline{\underline{\mathbf{V}}}_K^T$, with eigenvalue matrix $\underline{\underline{\Lambda}}_K = \text{diag}(\lambda_{K,1}, \dots, \lambda_{K,n})$ ordered such that $|\lambda_{K,1}| \leq \dots \leq |\lambda_{K,n}|$. The *stretching factor* associated with the element-averaged Hessian is defined by

$$s_K := \sqrt{\frac{\max_{i=1}^n |\lambda_{i,K}|}{\min_{i=1}^n |\lambda_{i,K}|}}. \quad (8)$$

We construct an element-wise metric by modifying the eigenvalues appropriately. In the two dimensional case, [12]

$$\underline{\underline{\mathbf{M}}}_K = \underline{\underline{\mathbf{V}}}_K \underline{\underline{\Lambda}}_K^{-2} \underline{\underline{\mathbf{V}}}_K^T, \quad \underline{\underline{\Lambda}}_K = \text{diag} \left(\sqrt{\frac{|\tilde{K}|}{|\hat{K}|}} s_K, \sqrt{\frac{|\tilde{K}|}{|\hat{K}|}} \frac{1}{s_K} \right). \quad (9)$$

A vertex-based metric is obtained from (9) by projection from $\mathbb{P}0$ to $\mathbb{P}1$. In this work, the projection is applied using a Galerkin projection, which amounts to averaging the element-wise values surrounding a vertex.

3 Application to the shallow water equations

3.1 Shallow water equations

In this work we consider the nonlinear shallow water equations for velocity \mathbf{u} [m s^{-1}] and surface elevation η [m] in steady-state form,

$$\mathbf{u} \cdot \nabla \mathbf{u} + g \nabla \eta + \frac{C_d \|\mathbf{u}\| \mathbf{u}}{\eta + b} = \nabla \cdot (\underline{\underline{\mathbf{D}}} \nabla \mathbf{u}), \quad \nabla \cdot ((\eta + b) \mathbf{u}) = 0, \quad (10)$$

with boundary conditions as appropriate. The fluid is modelled as having viscosity tensor $\underline{\underline{\mathbf{D}}}$ [$\text{m}^2 \text{s}^{-1}$], (unitless) quadratic drag C_d and bathymetry b [m]. We assume $g = 9.81 \text{ m s}^{-2}$.

Suppose the exact solution $\mathbf{q} = (\mathbf{u}, \eta)$ lives in a function space V with finite dimensional subspace V_h . For all test functions $\xi = (\boldsymbol{\psi}, \phi) \in V_h$, we have a Galerkin formulation of (10) given by

$$\begin{aligned} \rho_{\text{adv}}(\mathbf{q}_h, \xi) + \rho_{\text{gra}}(\mathbf{q}_h, \xi) + \rho_{\text{vis}}(\mathbf{q}_h, \xi) \\ + \rho_{\text{drg}}(\mathbf{q}_h, \xi) + \rho_{\text{cty}}(\mathbf{q}_h, \xi) = 0, \end{aligned} \quad (11)$$

where $\mathbf{q}_h \in V_h$ is a finite element approximation to \mathbf{q} . The weak form is decomposed into advection, gravity,

viscosity, drag and continuity terms. Taking all test functions, (11) forms a nonlinear system of equations which is solved using a Newton iteration. A direct solver is used in this work to solve the linear system of equations that arises in the update step of Newton’s method based on the Jacobian of (11) [3, 4].

For V_h , we select the mixed finite element space $\mathbb{P}1_{DG} - \mathbb{P}2$. That is, velocity is piecewise linear and discontinuous across elemental boundaries, whilst elevation is piecewise quadratic and continuous. This element pair was shown to be suitable for shallow water modelling in [13].

Clearly, functions from the velocity space are discontinuous across internal edges. Thus, internal edge integrals whose integrand involves such functions are not well defined. This necessitates the introduction of the following *restriction operators*, which assign a single value on each edge. For an edge $\gamma \in \Gamma_{int}$, arbitrarily label each side with + and −, giving rise to normal vectors $\hat{\mathbf{n}}^\pm$. Define the restriction operators given by the *average* $\{\{v\}\}_\gamma := \frac{1}{2}(v^+ + v^-)$ and *jump* $[[v]]_\gamma := v^+ - v^-$, where v can be scalar, vector or tensor valued. Note that $[[\mathbf{v} \cdot \hat{\mathbf{n}}^+]]_\gamma \equiv \mathbf{v}^+ \cdot \hat{\mathbf{n}}^+ + \mathbf{v}^- \cdot \hat{\mathbf{n}}^-$ for vector functions. For edges on the domain boundary, set both the average and jump to the function value on the boundary. Denote the outer product $\mathbf{v}_1 \mathbf{v}_2^T$ of two vector functions by $\mathbf{v}_1 \otimes \mathbf{v}_2$.

Ignoring boundary condition implementation (for brevity), Thetis uses the discretisation,

$$\begin{aligned}
 \rho_{adv}(\mathbf{q}_h, \xi) &= - \int_{\Omega} (\nabla \cdot (\mathbf{u}_h \otimes \boldsymbol{\psi})) \cdot \mathbf{u}_h \, dx \\
 &+ \int_{\Gamma_{int}} \tau [[\boldsymbol{\psi}]] \cdot [[\mathbf{u}_h]] \, dS \\
 &+ \int_{\Gamma_{int}} [[\boldsymbol{\psi}(\mathbf{u}_h \cdot \hat{\mathbf{n}}^+)]] \cdot \{\{\mathbf{u}_h\}\} \, dS, \\
 \rho_{gra}(\mathbf{q}_h, \xi) &= \int_{\Omega} g \boldsymbol{\psi} \cdot \nabla \eta_h \, dx, \\
 \rho_{vis}(\mathbf{q}_h, \xi) &= \int_{\Omega} \nabla \boldsymbol{\psi} : \underline{\underline{\mathbf{D}}} \nabla \mathbf{u}_h \, dx \\
 &+ \int_{\Gamma_{int}} \sigma [[\boldsymbol{\psi} \otimes \hat{\mathbf{n}}^+]] : \{\{\underline{\underline{\mathbf{D}}}\}\} [[\mathbf{u}_h \otimes \hat{\mathbf{n}}^+]] \, dS \\
 &- \int_{\Gamma_{int}} [[\boldsymbol{\psi} \otimes \hat{\mathbf{n}}^+]] : \{\{\underline{\underline{\mathbf{D}}} \nabla \mathbf{u}_h\}\} \, dS \\
 &- \int_{\Gamma_{int}} \{\{\nabla \boldsymbol{\psi}\}\} : \{\{\underline{\underline{\mathbf{D}}}\}\} [[\mathbf{u}_h \otimes \hat{\mathbf{n}}^+]] \, dS, \\
 \rho_{drg}(\mathbf{q}_h, \xi) &= \int_{\Omega} \boldsymbol{\psi} \cdot \frac{C_d \|\mathbf{u}_h\| \mathbf{u}_h}{\eta_h + b} \, dx, \\
 \rho_{cty}(\mathbf{q}_h, \xi) &= - \int_{\Omega} \nabla \phi \cdot ((\eta_h + b) \mathbf{u}_h) \, dx,
 \end{aligned} \tag{12}$$

where $\underline{\underline{\mathbf{S}}} : \underline{\underline{\mathbf{T}}} := \sum_{i=1}^n \sum_{j=1}^n S_{ij} T_{ij}$ for n -dimensional tensor functions $\underline{\underline{\mathbf{S}}}$ and $\underline{\underline{\mathbf{T}}}$. If the elevation was chosen from a discontinuous space (which is also a valid choice in Thetis) then the gravity term would be integrated by parts and we would need Riemann solutions for \mathbf{u} and η .

Lax–Friedrichs stabilisation [26] is applied in the advection term with parameter $\tau = \frac{1}{2} |\{\{\mathbf{u}\}\} \cdot \hat{\mathbf{n}}^+|$. The interior penalty parameter σ used in the viscous term is chosen in line with [17], depending upon the polynomial degree, variation of $\underline{\underline{\mathbf{D}}}$ and the minimal angle in each mesh element. For simplicity, we set σ as the largest element-wise value. It is important to update the penalty parameter whenever the mesh is adapted. Boundary conditions are imposed weakly, therefore contributing additional terms to (12).

3.2 Goal-oriented error estimate

As discussed in Sect. 2.2, the construction of goal-oriented error estimators from (12) typically involves integrating by parts on each element. Doing so enables us to derive from (5)–(6) error indicators of the form

$$\mathcal{E}_K = \left| \langle \boldsymbol{\Psi}(\mathbf{q}_h), \mathbf{e}^* \rangle_K + \langle \boldsymbol{\Psi}^\partial(\mathbf{q}_h), \mathbf{e}^* \rangle_{\partial K} + \mathcal{E}_K^{DG}(\mathbf{q}_h, \mathbf{e}^*) \right|, \tag{13}$$

where $\boldsymbol{\Psi}(\mathbf{q}_h)$ is the strong PDE residual (3), $\boldsymbol{\Psi}^\partial(\mathbf{q}_h)$ concatenates the residuals associated with the boundary conditions and $\mathcal{E}_K^{DG}(\mathbf{q}_h, \mathbf{e}^*)$ contains flux terms arising from the DG discretisation.

As in [10], (13) is written in this compact form so that the structure of the error indicator is transparent. This formulation has the advantage of separating out different components of the error estimator as regards the part of the problem they relate to. The first term assesses how well the PDE is solved on element interiors, the second term assesses the extent to which the boundary conditions have been weakly imposed and the final term describes the magnitude of the flux terms contributed by the DG discretisation of the velocity space. Since boundary conditions are neglected in (12), we are primarily interested in the nature of the flux term, $\mathcal{E}_K^{DG}(\mathbf{q}_h, \mathbf{e}^*)$, which conveys the smoothness of the discrete solution. Analysis of these contributions is useful in informing model development and assessing whether mesh features arise due to approximation error, boundary conditions or discretisation choice.

In order to derive an error indicator of the form (13), we need to appropriately integrate (12) by parts and substitute test functions for the adjoint error. No integration

by parts is required for the gravity or drag terms, since it is not used in their derivation from corresponding terms in (12).

The application of integration by parts on a particular element K results in a term involving an integral over its edge set, ∂K . In the DG method itself, these edge integrals must be restricted, so as to yield a unique value on each edge. However, since we seek element-based error indicators, there is no need to restrict quantities which are discontinuous across elemental boundaries.

Integrating by parts in the advection, viscosity and continuity terms yields

$$\begin{aligned}
 \rho_{\text{adv}}(\mathbf{q}_h, \boldsymbol{\xi})|_K &= \int_K \boldsymbol{\psi} \cdot (\mathbf{u}_h \cdot \nabla \mathbf{u}_h) \, dx \\
 &\quad - \int_{\partial K} ((\mathbf{u}_h \otimes \boldsymbol{\psi}) \cdot \mathbf{u}_h) \cdot \hat{\mathbf{n}} \, dS, \\
 &\quad + \int_{\partial K \setminus \Gamma} \tau \llbracket \boldsymbol{\psi} \rrbracket \cdot \llbracket \mathbf{u} \rrbracket \, dS + \int_{\partial K \setminus \Gamma} \llbracket \mathbf{u}_h \cdot \hat{\mathbf{n}}^+ \rrbracket \boldsymbol{\psi} \cdot \{\{ \mathbf{u}_h \}\} \, dS, \\
 \rho_{\text{vis}}(\mathbf{q}_h, \boldsymbol{\xi})|_K &= \int_K \boldsymbol{\psi} \cdot (\nabla \cdot (\underline{\underline{\mathbf{D}}}\nabla \mathbf{u}_h)) \, dx \\
 &\quad - \int_{\partial K} \boldsymbol{\psi} \otimes \hat{\mathbf{n}} : \underline{\underline{\mathbf{D}}}\nabla \mathbf{u}_h \, dS \\
 &\quad - \int_{\partial K \setminus \Gamma} \sigma \llbracket \boldsymbol{\psi} \otimes \hat{\mathbf{n}}^+ \rrbracket : \{\{ \underline{\underline{\mathbf{D}}}\}\} \llbracket \mathbf{u}_h \otimes \hat{\mathbf{n}}^+ \rrbracket \, dS \\
 &\quad + \int_{\partial K \setminus \Gamma} \llbracket \boldsymbol{\psi} \otimes \hat{\mathbf{n}}^+ \rrbracket : \{\{ \underline{\underline{\mathbf{D}}}\nabla \mathbf{u}_h \}\} \, dS \\
 &\quad + \int_{\partial K \setminus \Gamma} \{\{ \nabla \boldsymbol{\psi} \}\} : \{\{ \underline{\underline{\mathbf{D}}}\}\} \llbracket \mathbf{u}_h \otimes \hat{\mathbf{n}}^+ \rrbracket \, dS \\
 \rho_{\text{cty}}(\mathbf{q}_h, \boldsymbol{\xi})|_K &= \int_K \phi \nabla \cdot ((\eta_h + b)\mathbf{u}_h) \, dx \\
 &\quad - \int_{\partial K} \phi(\eta_h + b)\mathbf{u}_h \cdot \hat{\mathbf{n}} \, dS,
 \end{aligned} \tag{14}$$

where $\hat{\mathbf{n}} \in \{\hat{\mathbf{n}}^+, \hat{\mathbf{n}}^-\}$ is the outward-pointing normal to ∂K . The notation used in (14) can be simplified by application of three restriction identities, as follows. Consider a scalar function f , vector function \mathbf{g} and an edge γ shared by two adjacent elements $K^+, K^- \in \mathcal{H}$ whose outward-pointing normals are $\hat{\mathbf{n}}^+$ and $\hat{\mathbf{n}}^-$, respectively. Then

$$\begin{aligned}
 \{\{f\}\}|_\gamma &= \frac{1}{2}f|_{\partial K^+} + \frac{1}{2}f|_{\partial K^-}, \\
 \llbracket \mathbf{g} \cdot \hat{\mathbf{n}}^+ \rrbracket|_\gamma &= \mathbf{g} \cdot \hat{\mathbf{n}}^+|_{\partial K^+} + \mathbf{g} \cdot \hat{\mathbf{n}}^-|_{\partial K^-}, \\
 \llbracket \mathbf{g} \otimes \hat{\mathbf{n}}^+ \rrbracket|_\gamma &= \mathbf{g} \otimes \hat{\mathbf{n}}^+|_{\partial K^+} + \mathbf{g} \otimes \hat{\mathbf{n}}^-|_{\partial K^-}.
 \end{aligned} \tag{15}$$

Replacing test functions with the adjoint error in (14) and applying (15), we obtain error indicators of the form (13), where the flux term is given by

$$\begin{aligned}
 \mathcal{E}_K^{\text{DG}}(\mathbf{q}_h, \mathbf{e}^*) &= - \int_{\partial K} ((\mathbf{u}_h \otimes \mathbf{e}_u^*) \cdot \mathbf{u}_h) \cdot \hat{\mathbf{n}} \, dS, \\
 &\quad + \int_{\partial K \setminus \Gamma} \tau \llbracket \mathbf{e}_u^* \rrbracket \cdot \llbracket \mathbf{u} \rrbracket \, dS \\
 &\quad + \int_{\partial K \setminus \Gamma} \llbracket \mathbf{u}_h \cdot \hat{\mathbf{n}}^+ \rrbracket \mathbf{e}_u^* \cdot \{\{ \mathbf{u}_h \}\} \, dS - \int_{\partial K} \mathbf{e}_u^* \otimes \hat{\mathbf{n}} : \underline{\underline{\mathbf{D}}}\nabla \mathbf{u}_h \, dS \\
 &\quad - \int_{\partial K} e_\eta^*(\eta_h + b)\mathbf{u}_h \cdot \hat{\mathbf{n}} \, dS \\
 &\quad - \int_{\partial K \setminus \Gamma} \sigma \mathbf{e}_u^* \otimes \hat{\mathbf{n}} : \{\{ \underline{\underline{\mathbf{D}}}\}\} \llbracket \mathbf{u}_h \otimes \hat{\mathbf{n}}^+ \rrbracket \, dS \\
 &\quad + \int_{\partial K \setminus \Gamma} \mathbf{e}_u^* \otimes \hat{\mathbf{n}}^+ : \{\{ \underline{\underline{\mathbf{D}}}\nabla \mathbf{u}_h \}\} \, dS \\
 &\quad + \frac{1}{2} \int_{\partial K \setminus \Gamma} \nabla \mathbf{e}_u^* : \{\{ \underline{\underline{\mathbf{D}}}\}\} \llbracket \mathbf{u}_h \otimes \hat{\mathbf{n}}^+ \rrbracket \, dS.
 \end{aligned} \tag{16}$$

Here $\mathbf{q}^* = (\mathbf{u}^*, \eta^*)$ is the exact adjoint solution, $\mathbf{q}_h^* = (\mathbf{u}_h^*, \eta_h^*)$ is a finite element approximation thereof and $\mathbf{e}^* = (\mathbf{e}_u^*, e_\eta^*)$ is the corresponding error.

3.3 Approximation of the adjoint error

Note that (13) and (16) contain the (unknown) exact adjoint solution, \mathbf{q}^* . In practice, it suffices to approximate it in an enriched space, $V_h^+ \supset V_h$. We again choose V_h^+ as $\mathbb{P}1_{\text{DG}} - \mathbb{P}2$, but defined on a uniformly refined mesh. That is, a single, global *iso- $\mathbb{P}2$ refinement* is made, which amounts to inserting vertices wherever a quadrature node would exist in a quadratic element. The adjoint equation is then derived using a linearisation about the projected forward solution, $\Pi_h^+ \mathbf{q}_h \in V_h^+$. Whilst this approach is shown to be effective in Sect. 4.3, it implies an additional computational cost associated with solving the adjoint equation on a mesh with four times as many elements. In future work, a more efficient approach will be implemented, such as the one described on pp.590–593 of [15], which solves local PDEs to approximate \mathbf{q}^* .

3.4 Implementation details

The anisotropic metric formulation described in Sect. 2.3 relies on the provision of an approximate Hessian. Hessian recovery techniques typically seek second derivatives for scalar fields. In this case, there are a number of options for fields to recover a Hessian from—the free surface elevation, velocity components and the fluid speed being the most obvious candidates. Individual Hessians may be combined, using a strategy such as metric averaging or metric superposition (see pp. 131–138 of [7] for details and an investigation of the differences). In this work, we superpose the free surface elevation and velocity component

Hessians, so that the anisotropy of each field is accounted for. The eigenvalues of the resulting metric are then modified according to (9) in order to account of the error indicator (13), giving rise to a goal-oriented metric.

Mesh adaptation is performed using *Pragmatic* [8, 34], which is a mesh optimisation library primarily performing *h*-like adaptive operations, with some local Laplacian smoothing. The mesh adaptation workflow used is identical to that described in Algorithm 1 of [37]. The adaptation loop is terminated if the QoI value, number of mesh elements or error estimator (13) change by less than 0.1% from one iteration to the next.

Previously published studies, such as [19], use the *dolfin-adjoint* package [18], which automates the solution of the discrete adjoint of a finite element problem expressed in *Unified Form Language (UFL)* [2]. In this work, we wish to be able to discretise the adjoint problem in a different way than the forward problem. Thus, we exploit the automatic differentiation capabilities of UFL in order to avoid an error-prone manual calculation.

Given a weak form PDE $\mathbb{F} = 0$ with finite element solution q and QoI J , the adjoint solution, q_{star} , may be computed in only a few lines of code:

```
dFdq = derivative(F, q, TrialFunction(q.function_space()))
dFdq_transpose = adjoint(dFdq)
dJdq = derivative(J, q, TestFunction(q.function_space()))
solve(dFdq_transpose == dJdq, q_star, solver_parameters = {...}).
```

Writing the adjoint equation using Firedrake solve calls enables us to solve the adjoint equation in V_h^+ , as opposed to V_h . To do this, we replace q by $q \in V_h^+$ and \mathbb{F} by \mathbb{F}_+ in the above, where \mathbb{F}_+ is defined by prolonging the variables used in \mathbb{F} from V_h to V_h^+ .

The Firedrake installation used for all simulations documented in Sect. 4 is archived at [23, 36], with all simulation code archived at [38].

4 Numerical experiments

4.1 Tidal turbine modelling

Marine renewable energy is an active area of research in coastal ocean modelling (for example, see [14, 16, 19, 28, 35]). In particular, tidal power presents an opportunity to generate large amounts of low-carbon electricity in coastal countries such as the UK. One major advantage of tidal power over other renewable energy sources is that tides

are highly predictable, meaning that power is generated reliably.

The shallow water depth in which tidal turbines are deployed has a significant impact on wake recovery and hydrodynamic blockage effects, meaning that the positioning of turbines can be very important. By modifying the shallow water equations to account for tidal turbines positioned within the domain, recent research formulated tidal array positioning as a PDE-constrained optimisation problem [19]. Solving this problem gives the configuration with maximum power, with the potential to incorporate penalties in the optimisation functional to account for financial [14] and environmental impact factors [16]. A parametrisation based approach is used, meaning that the turbines are modelled using a density function $d = d(\mathbf{x})$.

Modifying (10) to account for a set of tidal turbines \mathcal{T} amounts to choosing an appropriate drag coefficient C_d . Suppose turbine T has thrust coefficient c_T , area A_T and footprint indicated by $\mathbb{1}_T$. For a background drag $C_b = 0.0025$,

$$C_d := C_b + C_t, \quad C_t := \sum_{T \in \mathcal{T}} \frac{1}{2} d c_T A_T \mathbb{1}_T. \tag{17}$$

The binary footprint function $\mathbb{1}_T : \Omega \rightarrow \{0, 1\}$ is unity in the region where turbine T is deployed and zero elsewhere.

As such, it specifies the spatial region where the turbine drag is active. The thrust coefficient used in (17) is based on an upstream velocity, whereas \mathbf{u} is the depth-averaged velocity at the turbine. Hence, we correct the thrust coefficient using the rescaling recommended in [25].

We use the following proxy for the power output of the tidal array:

$$J(\mathbf{u}, \eta) := \int_{\Omega} C_t \|\mathbf{u}\|^3 dx. \tag{18}$$

This provides a QoI for goal-oriented error estimation and has units of Watts.

4.2 Problem setup

Whilst realistic tidal turbine applications are inherently time-dependent, we consider a steady-state test case to highlight the impact of the turbine position on power

output. For a simple tidal farm with two turbines, T_1 and T_2 , we consider two configurations: one where T_1 is directly upstream of T_2 and another where the turbines are offset by one turbine diameter to south and north, respectively. Our numerical experiments involve applying the goal-oriented mesh adaptation approaches described in Sect. 2 to provide accurate approximations to the total tidal farm power output in each case.

For a channel domain $\Omega = [0, \ell_1] \times [0, \ell_2]$ with $\ell_1 = 1.2 \text{ km}$ and $\ell_2 = 500 \text{ m}$ and uniform bathymetry $b = 40 \text{ m}$, flow is driven by an inflow condition, $\mathbf{u}|_{x=0} = (5, 0) \text{ ms}^{-1}$. These depths and fluid speeds are representative of the Pentland Firth, Scotland - one of the UK's greatest tidal resources [16]. Viscosity is set to be the isotropic constant $\underline{\underline{\mathbf{D}}} = \nu \underline{\underline{\mathbf{I}}}_2$ with $\nu = 0.5 \text{ m}^2 \text{ s}^{-1}$, meaning we have a moderately advection-dominated problem. Free-slip conditions are imposed on the channel walls, along with a Dirichlet condition $\eta|_{x=\ell_1} = 0 \text{ m}$ on the outflow which acts to close the system. Turbines of diameter 18m are centred at $\{(456, 250), (744, 250)\}$ in the aligned case and $\{(456, 232), (744, 218)\}$ in the offset case. We use J_0 and J_1 to denote the total power output (18) for each of the two array configurations considered, where the subscript indicates the offset in terms of number of turbine diameters.

Coarse initial meshes which take account of the tidal turbines are generated using *gmsh* [21] and shown in Fig. 1 a, b. Figure 1 c, d show the magnitude of the velocity (i.e. the speed) given by solving the test case on meshes which have been refined three times using iso-P2 refinement.

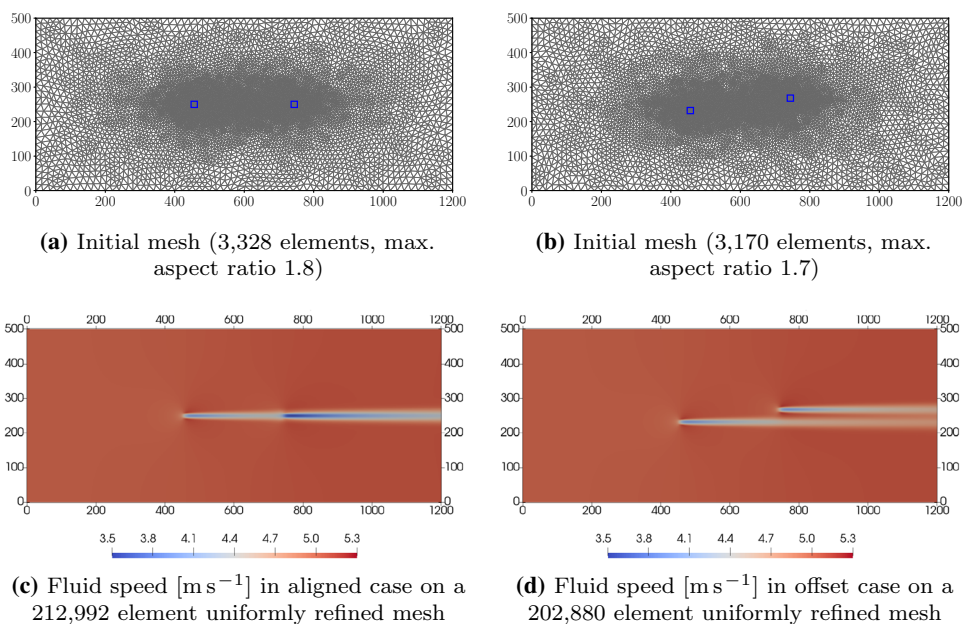
Observe that, in the aligned case, the momentum deficit at T_2 is more significant than at T_1 . This is because the wake of T_1 has not fully recovered and so the fluid speed meeting T_2 is lower than that which meets the first. That is, T_2 is fully in the wake of T_1 and thus overall experiences slower flow and consequently generates less power [cf. (18)]. In the offset case, the momentum deficit is similar at each turbine, suggesting that the presence of T_1 has less impact upon the fluid speed meeting T_2 . The result is that the total power output of the aligned array configuration is lower than that of the offset configuration, as shown in the QoI values displayed on the fourth row of Table 1.

Table 1 illustrates the convergence of QoI values under iso-P2 (uniform) refinement to four significant figures. Final values, $J_0 = 19.7174 \text{ kW}$ and $J_1 = 23.1905 \text{ kW}$, present benchmark values to approximate using adaptive meshes. In agreement with expectations, the converged power output in the aligned configuration is 15% lower than in the offset case. This illustrates how the

Table 1 Convergence of QoIs J_0 and J_1 evaluated at finite element solutions on a sequence of meshes generated by uniform refinement of the initial mesh

Aligned		Offset			
Elements	DOFs	J_0 (kW)	Elements	DOFs	J_1 (kW)
3328	26,711	20.7460	3170	25,447	24.1583
13,312	106,669	19.9454	12,680	101,613	23.4034
53,248	426,329	19.7451	50,720	406,105	23.2172
212,992	1,704,625	19.7201	202,880	1,623,729	23.1923
851,968	6,817,121	19.7174	811,520	6,493,537	23.1905

Fig. 1 a, b Initial meshes for the aligned and offset cases with turbine footprint regions indicated by blue squares. **c, d** Fluid speed as computed on meshes generated by three uniform refinements of the meshes shown in **a, b**



positioning of turbines within an array can significantly impact power output.

4.3 Convergence analysis under goal-oriented mesh adaptation

Having obtained benchmark values for the power output under uniform refinement in each turbine configuration, we apply mesh adaptation in an attempt to achieve convergence to these values using fewer DOFs.

Figure 2 shows meshes generated by the two adaptation strategies, each with low resolution downstream of the second turbine. This should be expected since we have an advection-dominated problem, meaning the power output of the array is independent of the downstream dynamics. We observe high mesh resolution surrounding and upstream of the turbines.

The moderate advection-dominance manifests in the anisotropy of the meshes shown in Fig. 2c, d. The maximum aspect ratio is reported to be higher in the aligned configuration, in which case almost all of the dynamics to which the QoI value is sensitive lie in a narrow horizontal band through the middle of the domain. The moderate mesh resolution near the domain boundaries in the tidal

farm and upstream regions is due to residual contributions from the weakly enforced boundary conditions.

Figure 3 is useful in understanding how each component of the error indicator (13) contributes towards the mesh adaptation. We observe that the PDE residual is most significant surrounding the turbines, whilst the error indicator contributions due to flux terms are generally in the upstream region. As has already been noted, the fact that we weakly impose boundary conditions means that there are significant flux term contributions near to the boundary.

Figure 4 illustrates the convergence of evaluated QoI values to the benchmark values established in Table 1 under each adaptation strategy with increasing DOF count. This validates the adaptive solution strategies. The DOF count is increased by increasing the target mesh complexity in (7).

For a given number of DOFs, isotropic goal-oriented adaptation generally offers an improvement in QoI approximation accuracy over uniform refinement. Anisotropic goal-oriented adaptation is shown to yield yet more accurate power output estimates, with the improvement over uniform refinement of the initial mesh being significant.

Fig. 2 Example meshes generated using goal-oriented mesh adaptation. In the aligned case, adaptation is applied to the mesh shown in Fig. 1a in order to generate **a, c**. In the offset case, adaptation is applied to the mesh shown in Fig. 1b in order to generate **b, d**

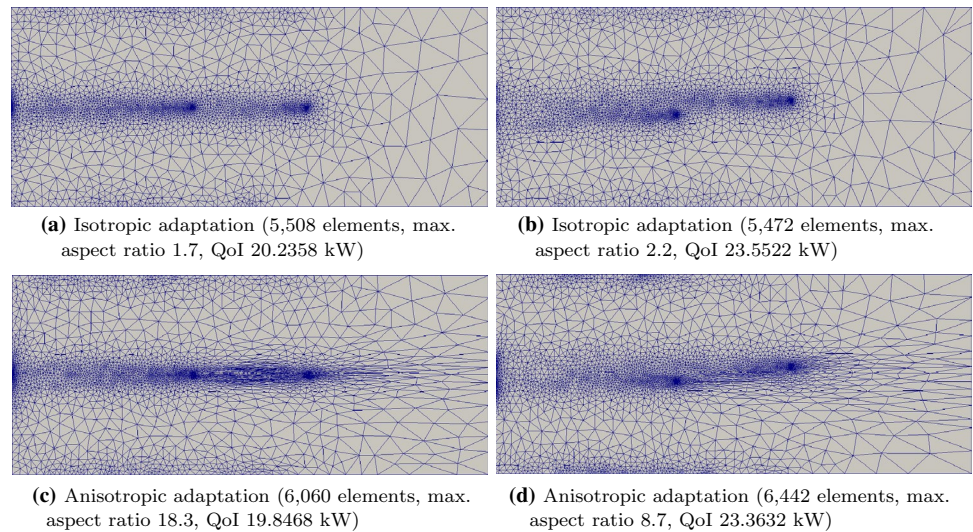



Fig. 3 Element-wise error indicator contributions resulting from **a** the strong residual evaluated on element interiors and **b** flux terms due to boundary conditions, integration by parts and the DG velocity discretisation. Both fields are evaluated on the mesh displayed in Fig. 2d



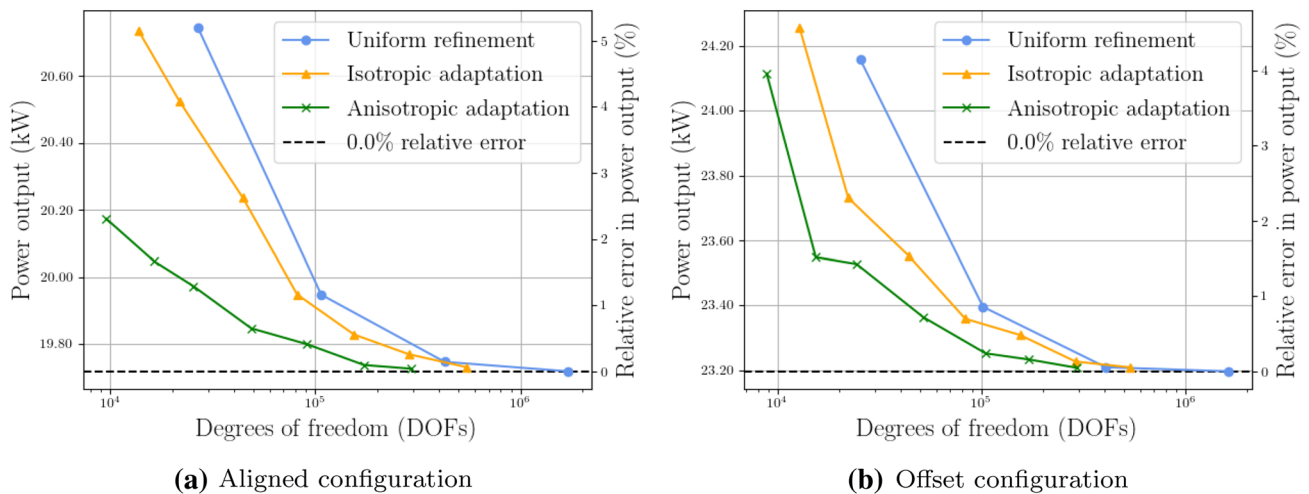


Fig. 4 QoI convergence analysis under uniform refinement, isotropic adaptation and anisotropic adaptation for the **a** aligned and **b** offset array configurations

5 Discussion

The numerical experiments conducted in Sect. 4 validate the two goal-oriented mesh adaptation strategies considered. The experimental results show that the anisotropic approach in particular is able to achieve notably higher QoI estimation accuracy than is obtained using the same number of DOFs under uniform refinement. The isotropic approach also offers some improvement over uniform refinement for these test cases.

Consider again the goal-oriented meshes shown in Fig. 2. Observe that, in each case, the tidal farm and other well-resolved regions correspond to a small proportion of the domain, with low resolution used away from these regions. One can imagine that extending the problem size to something akin to a realistic tidal turbine modelling application would imply high resolution deployed in an even smaller proportion of the domain. Goal-oriented mesh adaptation is most effective when the QoI is insensitive to the dynamics in the majority of the domain and this becomes increasingly true as the domain size increases relative to the tidal farm. As remarked above, goal-oriented adaptation strategies deploy high resolution mostly in upstream regions for advection-dominated problems. Any combination of widening the channel, extending the downstream region and increasing the advection-dominance of the problem (as per the advection/viscosity relationship determined by the Reynolds number) would imply a leftwards shift of the goal-oriented convergence curves shown in Fig. 4, relative to the uniform refinement curve. It is likely that all three of these would be present in a realistic application, meaning that goal-oriented mesh

adaptation would require relatively few DOFs to provide an accurate power output estimate.

Mixed finite element methods are becoming increasingly popular discretisation choices for shallow water problems [13]. Through the derivation shown in Sect. 3.2, we illustrate how to perform goal-oriented error estimation for mixed discretisations with discontinuous components. The derivation of this estimate follows the procedure used in [22], which considers DG discretisations for advection-diffusion-reaction problems.

The anisotropic mesh adaptation algorithm used in this work is based on that described in [12], which also focuses on advection-diffusion-reaction problems. This element-based approach is advantageous because it straightforwardly permits the incorporation of discontinuous goal-oriented error indicators, such as arise from DG discretisations. As in the numerical experiments considered in [12], we observe convergence of the QoI to reference values computed on a high resolution mesh. In agreement with the numerical experiments considered in [22], we found that, for a given number of DOFs, the relative error in evaluating the QoI was smaller under anisotropic adaptation than under isotropic adaptation.

6 Conclusion

The main achievement of this paper is the formulation of a goal-oriented error estimate for the nonlinear shallow water equations solved using a mixed discontinuous/continuous finite element method, along with the implementation of isotropic and anisotropic mesh adaptation algorithms using this estimate.

Convergence analysis is performed in the context of the power output of a steady-state tidal turbine problem. This analysis illustrates that fewer DOFs are required to achieve a certain QoI error threshold using goal-oriented approaches than uniform refinement.

In future work, we intend to use the goal-oriented adaptation framework discussed in this paper for modelling proposed tidal farms, with the aim of accurately approximating the associated power output.

Acknowledgements Thank you to the Applied Modelling and Computation Group (AMCG) at Imperial College London, the Firedrake and Thetis development teams and staff and students of the Mathematics of Planet Earth Centre for Doctoral Training (MPE CDT).

Funding This work was funded by the Engineering and Physical Sciences Research Council (EPSRC) under Grants EP/L016613/1 and EP/R029423/1.

Compliance with ethical standards

Conflicts of interest The authors declare that they have no conflict of interest.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Alauzet F, Loseille A (2016) A decade of progress on anisotropic mesh adaptation for computational fluid dynamics. *Comput Aided Des* 72:13–39
- Alnæs MS, Logg A, Ølgaard KB, Rognes ME, Wells GN (2014) Unified form language: a domain-specific language for weak formulations of partial differential equations. *ACM T Math Softw* 40(2):9
- Amestoy PR, Duff IS, L'Excellent JY, Koster J (2001) A fully asynchronous multifrontal solver using distributed dynamic scheduling. *SIAM J Matrix Anal A* 23(1):15–41
- Amestoy PR, Guermouche A, L'Excellent JY, Pralet S (2006) Hybrid scheduling for the parallel solution of linear systems. *Parallel Comput* 32(2):136–156
- Balay S, Abhyankar S, Adams MF, Brown J, Brune P, Buschelman K, Dalcin L, Dener A, Eijkhout V, Gropp WD, Karpeyev D, Kaushik D, Knepley MG, May DA, McInnes LC, Mills RT, Munson T, Rupp K, Sanan P, Smith BF, Zampini S, Zhang H, Zhang H (2019) PETSc users manual. Technical Report ANL-95/11—Revision 3.12, Argonne National Laboratory. <https://www.mcs.anl.gov/petsc>
- Balay S, Gropp WD, McInnes LC, Smith BF (1997) Efficient management of parallelism in object oriented numerical software libraries. In: Arge E, Bruaset AM, Langtangen HP (eds) *Modern software tools in scientific computing*. Birkhäuser Press, Boston, pp 163–202
- Barral N (2015) Time-accurate anisotropic mesh adaptation for three-dimensional moving mesh problems. Ph.D. thesis, Université Pierre et Marie Curie
- Barral N, Knepley MG, Lange M, Piggott MD, Gorman GJ (2016) Anisotropic mesh adaptation in Firedrake with PETSc DMplex. In: 25th International meshing roundtable
- Becker R, Rannacher R (1996) A feed-back approach to error control in finite element methods: basic analysis and examples. *IWR*
- Becker R, Rannacher R (2001) An optimal control approach to a posteriori error estimation in finite element methods. *Acta Numer* 10:1–102
- Braack M (1998) An adaptive finite element method for reactive-flow problems. Ph.D. thesis, IWR
- Carpio J, Prieto JL, Bermejo R (2013) Anisotropic “goal-oriented” mesh adaptivity for elliptic problems. *SIAM J Sci Comput* 35(2):A861–A885
- Cotter CJ, Ham DA, Pain CC (2009) A mixed discontinuous/continuous finite element pair for shallow-water ocean modelling. *Ocean Model* 26(1–2):86–90
- Culley DM, Funke SW, Kramer SC, Piggott MD (2016) Integration of cost modelling within the micro-siting design optimisation of tidal turbine arrays. *Renew Energy* 85:215–227
- Dolejší V, Roskovec F (2017) Goal-oriented error estimates including algebraic errors in discontinuous Galerkin discretizations of linear boundary value problems. *Appl Math* 62(6):579–605
- du Feu RJ, Funke SW, Kramer SC, Hill J, Piggott MD (2019) The trade-off between tidal-turbine array yield and environmental impact: a habitat suitability modelling approach. *Renew Energy* 143:390–403
- Epshteyn Y, Rivière B (2007) Estimation of penalty parameters for symmetric interior penalty Galerkin methods. *J Comput Appl Math* 206(2):843–872
- Farrell PE, Ham DA, Funke SW, Rognes ME (2013) Automated derivation of the adjoint of high-level transient finite element programs. *SIAM J Sci Comput* 35(4):C369–C393
- Funke SW, Farrell PE, Piggott M (2014) Tidal turbine array optimisation using the adjoint approach. *Renew Energy* 63:658–673
- George P, Hecht F, Vallet M (1991) Creation of internal points in Voronoi's type method. *Control and adaptation. Adv Eng Softw* 13(5–6):303–312
- Geuzaine C, Remacle JF (2009) Gmsh: a 3-d finite element mesh generator with built-in pre- and post-processing facilities. *Int J Numer Methods Eng* 79(11):1309–1331
- Houston P, Georgoulis EH, Hall E (2006) Adaptivity and a posteriori error estimation for DG methods on anisotropic meshes. In: *International conference on boundary and interior layers*
- Kärnä T, Kramer S, Mitchell L, Angeloudis A, Ham DA, Wallwork JG, Barral N, Warder S, Clare M, McRae ATT (2020) thetisproject/thetis: Thetis coastal ocean model. <https://doi.org/10.5281/zenodo.3689727>
- Kärnä T, Kramer SC, Mitchell L, Ham DA, Piggott MD, Baptista AM (2018) Thetis coastal ocean model: discontinuous Galerkin discretization for the three-dimensional hydrostatic equations. *Geosci Model Dev* 11(11):4359–4382
- Kramer SC, Piggott MD (2016) A correction to the enhanced bottom drag parameterisation of tidal turbines. *Renew Energy* 92:385–396

26. Lax PD (1954) Weak solutions of nonlinear hyperbolic equations and their numerical computation. *Commun Pure Appl Math* 7(1):159–193
27. Loseille A, Dervieux A, Alauzet F (2010) Fully anisotropic goal-oriented mesh adaptation for 3D steady euler equations. *J Comput Phys* 229(8):2866–2897
28. Neill SP, Angeloudis A, Robins PE, Walkington I, Ward SL, Masters I, Lewis MJ, Piano M, Avdis A, Piggott MD et al (2018) Tidal range energy resource and optimization-past perspectives and future challenges. *Renew Energy* 127:763–778
29. Pain C, Umpheby A, De Oliveira C, Goddard A (2001) Tetrahedral mesh optimisation and adaptivity for steady-state and transient finite element calculations. *Comput Method Appl M* 190(29–30):3771–3796
30. Pan W, Kramer SC, Piggott MD (2019) Multi-layer non-hydrostatic free surface modelling using the discontinuous Galerkin method. *Ocean Model* 134:68–83
31. Piggott M, Farrell P, Wilson C, Gorman G, Pain C (2009) Anisotropic mesh adaptivity for multi-scale ocean modelling. *Philos Trans R Soc A Math Phys Eng Sci* 367(1907):4591–4611
32. Power P, Pain CC, Piggott M, Fang F, Gorman GJ, Umpheby A, Goddard AJ, Navon I (2006) Adjoint a posteriori error measures for anisotropic mesh optimisation. *Comput Math Appl* 52(8–9):1213–1242
33. Rathgeber F, Ham DA, Mitchell L, Lange M, Luporini F, McRae AT, Bercea GT, Markall GR, Kelly PH (2016) Firedrake: automating the finite element method by composing abstractions. *ACM T Math Softw* 43(3):24
34. Rokos G, Gorman G (2013) Pragmatic–parallel anisotropic adaptive mesh toolkit. In: Keller R, Kramer D, Weiss JP (eds) *Facing the multicore-challenge III*. Springer, Berlin, pp 143–144
35. Vouriot CVM, Angeloudis A, Kramer SC, Piggott MD (2019) Fate of large-scale vortices in idealized tidal lagoons. *Environ Fluid Mech* 19(2):329–348
36. Software used in ‘Goal-oriented error estimation and mesh adaptation for shallow water modelling’ (2019). <https://doi.org/10.5281/zenodo.3568997>
37. Wallwork JG, Barral N, Ham DA, Piggott MD (2020) Anisotropic goal-oriented mesh adaptation in Firedrake. In: 28th international meshing roundtable, Zenodo. <https://doi.org/10.5281/zenodo.3653101>
38. Wallwork JG, Clare M (2020) Simulation code for ‘Goal-oriented error estimation and mesh adaptation for shallow water modelling’. <https://doi.org/10.5281/zenodo.3748366>

Publisher’s Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.