On the use of mesh movement methods to help overcome the multi-scale challenges associated with hydro-morphodynamic modelling

Mariana C. A. Clare[a,*], Joseph G. Wallwork[a], Stephan C. Kramer[a], Hilary Weller[b], Colin J. Cotter[c], Matthew D. Piggott[a]

[a]*Department of Earth Science and Engineering, Imperial College London, U.K.*
[b]*Department of Meteorology, University of Reading, U.K.*
[c]*Department of Mathematics, Imperial College London, U.K.*

Abstract

Hydro-morphodynamic models are an important tool that can be used in the protection of coastal zones. They can be required to resolve spatial scales ranging from sub-metre to hundreds of kilometres and are computationally expensive. In this work, we apply mesh movement methods to a depth-averaged hydro-morphodynamic model for the first time, in order to tackle both these issues. Mesh movement methods are particularly suited to coastal problems as they allow the mesh to move in response to evolving flow and morphology structures. This new capability is demonstrated using test cases that exhibit complex evolving bathymetries. Some of these test cases require the use of a wetting-and-drying scheme which is known to leak sediment. To ensure sediment is conserved with this scheme, a new conservative sediment transport model is implemented. For all test cases, we demonstrate how mesh movement methods can be used to reduce discretisation error and computational cost. We also show that the optimal parameter choices in the mesh movement monitor functions are fairly predictable based upon the physical characteristics of the test case, facilitating the use of mesh movement methods on further problems.

Keywords: mesh adaptation; mesh movement; sediment transport; morphology

1. Introduction

Over the last few decades, hydro-morphodynamic models have been increasingly used to simulate erosion in both fluvial and coastal zones (see [1, 2]). These models are unfortunately generally computationally expensive, especially as they are often required to simulate very long-term morphological effects with relatively small timesteps so that hydrodynamic features such as waves and tides are appropriately resolved. The models must also often be run multiple times for calibration purposes and to quantify uncertainty as for example in [3], [4] and [5]. This means that it is infeasible to model many fluvial and coastal scenarios using standard fixed meshes of appropriate resolution. Furthermore, when applied to coastal regions, hydro-morphodynamic models must resolve problems with complex and fundamentally multi-scale domains. Unstructured mesh models, such as those based on finite element methods [6, 7], are ideally suited to providing the required mesh flexibility

*Corresponding author
Email address:* `m.clare17@imperial.ac.uk` (Mariana C. A. Clare)

and multiple tools exist for generating multi-scale fixed meshes which are suitable for a wide range of geophysical applications (such as [8]). However, in cases with significant morphology changes, the areas which require finer mesh resolution vary throughout the simulation.

Mesh adaptation provides a solution to the time-dependent multi-scale issue as well as offering the potential to improve result accuracy and/or reduce computational cost. Mesh adaptation has been used previously in a range of hydro-morphodynamic models with finite element methods, for example, [9] and [10]. These hydro-morphodynamic models are not appropriate for solving the fluvial and coastal erosion problems discussed in this work because the first simulates cohesive sediment rather than non-cohesive sediment and the latter does not simulate bedlevel changes which are required to model erosion and deposition. In [11] the authors applied mesh adaptation to an appropriate hydro-morphodynamic model solved using finite volume methods. In our work we use a discontinuous Galerkin finite element discretisation which [12] demonstrate is better at dealing with steep gradients such as those that might form due to bedlevel changes. Furthermore, all these works use so-called h-refinement, where mesh cells are subdivided, or where the mesh topology is more substantially altered with cells locally created or destroyed and mesh connectivity altered [13, 14]. h-refinement methods can be computationally expensive for hydro-morphodynamic problems such as those considered in our work, because regions warranting the most mesh resolution move with the flow. This requires a large number of expensive operations including mesh-to-mesh interpolation to track this movement [15].

Conversely, this flow dependence makes these problems well suited to mesh movement methods (also known as r-adaptation) (see [16]). For these methods, the mesh topology is fixed and the (fixed number of) mesh nodes are dynamically moved. This allows different regions of the domain to vary between low and high resolution as flow structures pass through them. Given the fixed topology, mesh movement also avoids the problem of hanging nodes, which must be considered with h-adaptation methods. In addition, from a linear algebra perspective, because the number of mesh nodes is fixed, the matrix sparsity structure remains constant throughout the simulation. This is in contrast to h-adaptation methods where complex hierarchical structures are often required in order to accommodate nodes.

Mesh movement methods have been used successfully for boundary mesh movement in 2D-vertical models in [17] and [18], where the external boundary at the bed deforms with morphological changes. In this work, we consider a depth-averaged modelling approach (otherwise known as 2D-horizontal or 2DH) and the mesh movement is thus used for the 2D-horizontal mesh. Depth-averaged models are computationally cheaper than a full 3D model, and yet unlike 2D-vertical models are still able to capture variations in the cross-stream directions, as well as changes to the bed. Specifically, we use the 2D depth-averaged coupled hydro-morphodynamic model of [19] which models both suspended sediment transport and bedload transport for non-cohesive sediment with bed updates (see Figure 1). This model uses a discontinuous Galerkin finite element discretisation on an unstructured mesh and is developed within the finite element coastal ocean modelling system *Thetis* [20], built using the *Firedrake* Python-based code generation framework [21]). This framework is well suited for the inclusion of r-adaptation methods since these often require the solution of additional complex non-linear PDEs, which Firedrake is designed to make readily solvable (see [22]). Firedrake automatically generates low level kernels that assemble the discretised equations and uses the *PETSc* solver library [23, 24, 25] to solve the resulting linear and nonlinear systems. In addition, we take advantage of the adjoint framework within Firedrake to calibrate values of uncertain parameters in our model (see [26]).
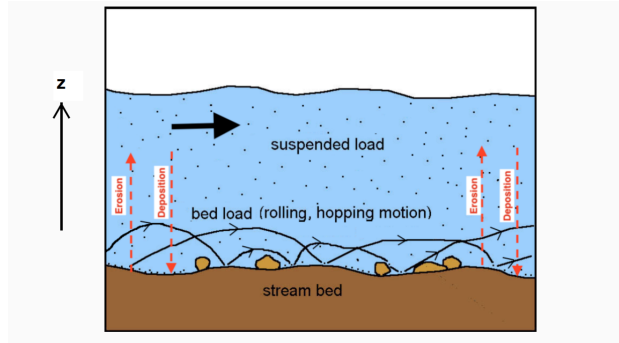
Figure 1: Diagram of sediment transport adapted from [27].

Mesh adaptation (both h and r) has already been implemented in the wider Firedrake framework, for example in [22, 28, 29, 30]. Integration of the mesh data structures with the underlying PETSc representation [31] laid the groundwork for the use of anisotropic metric-based methods in Firedrake [29]. This framework was used to apply mesh adaptation methods to the hydrodynamic component of Thetis' 2D model for the first time in [32]. The works of [22] and [28] specifically address mesh movement methods. The present work represents the first use of mesh adaptation with a coupled hydro-morphodynamic model in Thetis and the first use of mesh movement methods in Thetis.

In this work, we consider test cases with complex evolving bathymetries which are difficult to capture accurately on coarse and/or fixed meshes, which may completely miss small structures in the bathymetries. Given we are modelling sediment transport processes, the bed evolves with time, and thus a multi-scale fixed mesh is not sufficient because the areas of fine resolution must evolve with time as well. We consider test cases where the domain is fully wet as well as test cases which have a wet-dry interface – both scenarios occurring frequently in the modelling of coastal zones.

The remainder of this paper is structured as follows: in Section 2 we outline the mesh movement methods used; in Section 3 we give a brief outline of the hydro-morphodynamic model including the implementation of a conservative discretisation approach, newly developed as part of this work; in Section 4 we apply our mesh movement methods to a series of test cases and finally in Section 5 we conclude this work.

2. Mesh Movement

The mesh adaptation literature contains a number of different approaches to mesh movement, including those which impose a prescribed mesh velocity [33]; re-interpret the mesh as a structure of stiff beams [34]; and enforce mesh transformations using monitor functions [35, 36]. In this work, we consider the monitor function based approach. That is, the 'physical' mesh – upon which the prognostic equations are solved – is moved during the time period of simulation, with its 'density' prescribed by a user-provided monitor function. Moreover, the physical mesh is moved by defining it to be the image of some mapping applied to a *fixed* reference mesh. The mapping on the reference mesh determines the way in which the mesh is moved.

There are infinitely many maps which could be applied to the reference mesh and various approaches by which the map may be constructed. We focus on methods wherein the transformation is obtained from the solution of some PDE. The PDE has the role of enforcing that the moved mesh have certain properties. In this case, we enforce that the mesh evenly distributes a density function

3

across the mesh, meaning that regions of high density have elements with small areas and vice versa. As may be expected, the choice of density function – or *monitor function* – greatly impacts the geometry of the moved mesh and should therefore be chosen with care.

Both the PDE prescribing the mesh movement and the monitor functions are considered to be time-dependent, meaning that the physical mesh changes with time. As mentioned in Section 1, interpolating between meshes during the PDE time integration loop can be computationally expensive and introduces additional errors. For this reason, it may be desirable to formulate the PDE and mesh movement equation as a single system of equations to be solved *at each timestep* as in an *Arbitrary Lagrangian-Eulerian (ALE)* type approach [36, 33]. The mesh velocity due to the mesh movement equation determines the frame of reference within which here the hydro-morphodynamics equations are solved. Coupling the mesh movement to the physics in this way removes the need for mesh-to-mesh interpolation [14]. However, it is not obvious that these model improvements are worth the cost associated with being constrained to move the mesh at every timestep. As such, implementation and numerical experimentation for ALE methods within the context of hydro-morphodynamic models remains the topic of future work. In the experiments presented in this paper, fields are transferred between meshes using conservative Galerkin projection, achieved using *libsupermesh* [37].

2.1. Equidistribution and the Monge-Ampère equation

Let $\Omega_P : [0, T] \to \mathbb{R}^n$ denote the *physical* domain, which we allow to vary with time, and let $\mathcal{H}_P$ denote the associated physical mesh. This is the mesh upon which our hydro-morphodynamic model is to be solved. In addition, consider a *computational* reference domain $\Omega_C \subset \mathbb{R}^n$, which remains fixed, along with an associated computational mesh, $\mathcal{H}_C$. As with the domains, $\mathcal{H}_P = \mathcal{H}_P(t)$ is allowed to vary with time, whereas $\mathcal{H}_C$ remains fixed.

The approach to mesh movement described in this paper is concerned with obtaining a discrete representation of a sufficiently smooth map,

$$\mathbf{x} : \Omega_C \times [0, T] \to \Omega_P, \qquad \mathbf{x}(\mathcal{H}_C, t) = \mathcal{H}_P(t), \qquad t \in [0, T], \tag{1}$$

under some constraints. We aim to establish a mapping which transforms between the coordinate spaces associated with the reference mesh $\mathcal{H}_C$ and the moving mesh $\mathcal{H}_P$ upon which the hydro-morphodynamics equations are solved.

The first constraint imposed on the map (1) relates to a user-specified *monitor function*, $m : \Omega_P \times [0, T] \to (0, \infty)$. We assert that m is *equidistributed*, in the sense that

$$m \det \underline{\mathbf{J}} = \theta \qquad \text{where} \qquad \underline{\mathbf{J}} = \frac{\partial \mathbf{x}}{\partial \boldsymbol{\xi}} \tag{2}$$

is the Jacobian transform with respect to the computational coordinate $\boldsymbol{\xi} \in \Omega_C$. The normalisation coefficient $\theta : [0, T] \to (0, \infty)$ acts to conserve the domain volume. In this work, we do not consider boundary deformations and hence will always assume that $\Omega_P = \Omega_C =: \Omega$. In this case, θ is a constant.

The mesh movement problem may be stated as follows: *for a given m and θ, find a map, $\mathbf{x}$, which satisfies (2)*. In dimension $n > 1$ this problem is not well posed, meaning additional constraints must be imposed on the map (1). The hydro-morphodynamic model described in Section 3 is in $n = 2$ spatial dimensions and therefore this is indeed the case. Following the optimal transport approach advocated in [28, 36, 38, 39], we assume that the map takes the form

$$\mathbf{x}(\boldsymbol{\xi}) = \boldsymbol{\xi} + \nabla \phi_{\boldsymbol{\xi}}(\boldsymbol{\xi}), \tag{3}$$

4

for some scalar potential $\phi : \Omega_C \times [0, T] \to \Omega_P$. Substituting (3) into (2) yields

$$m \, \det\left(\underline{\mathbf{I}} + \underline{\mathbf{H}}(\phi)\right) = \theta, \tag{M-A}$$

where $\underline{\mathbf{I}}$ is the identity matrix in $\mathbb{R}^{n \times n}$ and $\underline{\mathbf{H}}(\phi)$ denotes the Hessian of the potential with respect to the computational coordinates. Equation (M-A) is a nonlinear PDE of *Monge-Ampère* type.

Under an appropriate choice of monitor function, the elliptic PDE (M-A) has been shown to be capable of producing equidistributed meshes which admit solutions of the original PDE which are of improved accuracy (with respect to particular error measures and/or diagnostics) without increasing the number of degrees of freedom or modifying the mesh topology (for example, see [28, 36, 38, 39]). The fact that mesh movement algorithms maintain a fixed mesh topology means that they can also be easily parallelised.

Many mesh movement algorithms come with the risk of mesh tangling, where one or more elements become inverted which will oftentimes crash the PDE solver. This is accounted for by the equidistribution relation (2): the fact that both the monitor function and normalisation coefficient are positive forces the determinant of the Jacobian transform to remain positive. This implies that no element orientation changes occur (assuming we have solved (M-A) to a reasonable accuracy).

A major criticism of mesh movement methods which are driven by solutions of Monge-Ampère type equations is that such equations require a degree of convexity of domain geometry which is not apparent in coastal ocean domains bounded by coasts. However, for many realistic coastlines, the use of a wetting-and-drying scheme in the hydrodynamics solver means that domain convexity can be ensured, because the domain boundary is no longer constrained to topographic contours and may be chosen as a convex superset of the domain. On the other hand, it is only the theory which requires domain convexity; it may be that this is not a strict requirement in practice.

2.2. Solution Strategy

Mesh adaptation based on (M-A) can be conducted as follows. First, we choose a computational mesh $\mathcal{H}_C$, an initial physical mesh $\mathcal{H}_P$ with an identical topology (usually chosen to coincide with $\mathcal{H}_C$), a normalisation coefficient θ and a monitor function m. The choice of monitor function is very important, as will become apparent in Section 4; it determines the way in which mesh resolution is to be distributed over the domain. Given the definitions above, mesh movement is achieved in an iterative fashion. Each iteration involves the numerical solution of (M-A) for the scalar potential on the current physical mesh, followed by the transformation of this mesh according to (3).

As mentioned above, (M-A) is an elliptic PDE, with two sources of nonlinearity: one from the determinant and another from the product with the monitor function, via the dependence of the physical coordinates on the computational coordinates. One solution strategy is to parabolise the Monge-Ampère equation, as in Section 2.3 of [28]:

$$-\frac{\partial}{\partial \tau} \Delta \phi = m \, \det\left(\underline{\mathbf{I}} + \underline{\mathbf{H}}(\phi)\right) - \theta, \tag{4}$$

where τ denotes 'pseudotime' and $\Delta \tau > 0$ is a 'pseudotimestep'. This expression is closely related to the earlier work of [36] on mesh movement under parabolic Monge-Ampère formulations. Suppose we apply a time integration scheme to (4). As ϕ approaches a steady state, the derivative term on the LHS converges to zero, whereby the solution of (4) tends to the solution of the elliptic form (M-A). It is proved in [40] that the sequence of solutions given by solving (4) converges to the solution of (M-A), provided the pseudotimestep is sufficiently small and the initial guess for ϕ is

sufficiently close to the solution. A pseudotimestep of $\Delta\tau = 0.1$ was found to be sufficient in this work.

Repeated solution of (4) can be computationally expensive, because each solve involves establishing a new map from Ω_C to Ω_P. For time-dependent PDEs, however, computational cost can be reduced by making an appropriate initial guesses for the potential, ϕ. The final value of ϕ computed during one iteration provides an appropriate initial guess for the next, provided that the hydro-morphodynamics have not changed too significantly. This is dependent on the choice of timestep and also mesh movement frequency.

Whilst the exact equidistribution provided by (M-A) is mathematically desirable, it is often unnecessary in practice. In order to reduce computational cost, we propose solving the parabolised form (4) to relative tolerance $\mathcal{O}(10^{-3})$ or even $\mathcal{O}(10^{-2})$. As shown in Section 4, it can also be excessive to apply mesh movement at every timestep, especially if the modifications to mesh coordinates are only minor. Reduction of the mesh movement frequency provides another method for controlling computational cost.

We follow [28] in solving the parabolised Monge-Ampère equation using a mixed finite element method. Let Ψ be a scalar function space in which the potential is approximated and let Σ be a tensor function space in which its Hessian is approximated. At pseudotimestep k, first apply the forward Euler scheme as

$$\langle \nabla\psi, \nabla\phi_{k+1}\rangle_{\mathcal{L}_2(\Omega)} = \langle \nabla\psi, \nabla\phi_k\rangle_{\mathcal{L}_2(\Omega)} + \Delta\tau\langle\psi, m\det(\underline{\mathbf{I}} + \underline{\mathbf{H}}_k) - \theta\rangle_{\mathcal{L}_2(\Omega)}, \quad \forall\psi \in \Psi. \tag{5}$$

The Hessian $\underline{\mathbf{H}}_k$ is initialised to zero for the first iteration (assuming that the initial physical mesh coincides with the computational mesh). It is recovered via an $\mathcal{L}_2$ projection in iterations thereafter [41]:

$$\langle\underline{\boldsymbol{\sigma}}, \underline{\mathbf{H}}_{k+1}\rangle_{\mathcal{L}_2(\Omega)} = -\langle\nabla\cdot\underline{\boldsymbol{\sigma}}, \nabla\phi_{k+1}\rangle_{\mathcal{L}_2(\Omega)} + \langle\underline{\boldsymbol{\sigma}}\cdot\widehat{\mathbf{n}}, \nabla\phi_{k+1}\rangle_{\mathcal{L}_2(\partial\Omega)}, \quad \forall\underline{\boldsymbol{\sigma}} \in \Sigma. \tag{6}$$

3. Hydro-Morphodynamic Model

In this work, we use the hydro-morphodynamic model derived in [19]. This model is able to simulate both suspended sediment and bedload transport taking into account gravitational and helical flow effects. To maximise stability, the time derivatives of our model equations are approximated using a fully-implicit backward Euler timestepping scheme. Full details of the model and its development are provided in [19].

So far this model has been used for test cases in which the whole domain is wet. However, in coastal problems, wetting-and-drying plays an important part. We use the inbuilt wetting-and-drying process in Thetis, which follows the approach detailed in [42], wherein the total water depth, H, is modified using the expression

$$\tilde{H} := \eta - h + f(H), \tag{7}$$

where η is the elevation, h is the bed height, and

$$f(H) := \frac{1}{2}\left(\sqrt{H^2 + \delta^2} - H\right), \tag{8}$$

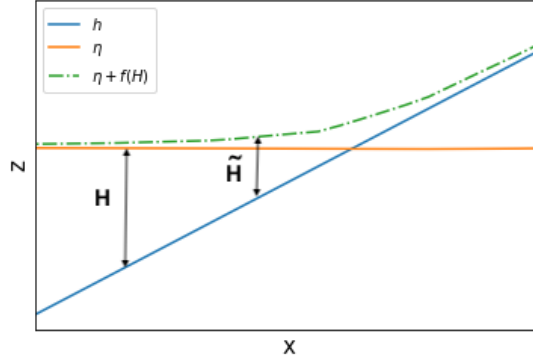as shown in Figure 2. Here δ is the wetting-and-drying parameter. In [42], it is recommended that

Figure 2: Wetting and drying scheme diagram showing the relationship between the total water depth H and its modified form $\tilde{H}$.

it is set approximately equal to $d||\nabla h||$ where d is the typical length scale of the mesh. Thus, the hydrodynamic equations in our model become

$$\frac{\partial \tilde{H}}{\partial t} + \nabla \cdot (\tilde{H}\mathbf{U}) = 0, \tag{9}$$

$$\frac{\partial \mathbf{U}}{\partial t} + \mathbf{U} \cdot \nabla \mathbf{U} + \mathbf{g}\nabla\eta = \nu\nabla^2\mathbf{U} - \frac{C_h(\tilde{H})}{\tilde{H}}||\mathbf{U}||\mathbf{U}, \tag{10}$$

where $\mathbf{U}$ is the depth-averaged velocity, $\mathbf{g}$ is gravity, ν the viscosity parameter and C_h the bed friction. Following [43], to avoid non-differentiable functions, we make the following smooth approximation to the norm operator in (10)

$$||\mathbf{U}|| \approx \sqrt{||\mathbf{U}||^2 + \beta^2}. \tag{11}$$

In [43], β is set equal to the wetting-and-drying parameter. However, the issue of non-differentiability exists independent of the wetting-and-drying formulation and so here we do not follow this choice. Rather, β is chosen such that it is large enough to have the desired effect, but not too large that the friction term spuriously affects model results. In the test case in Section 4.3, we experimented with values of β between 0.1 and 1 and found that a value of 0.4 provided good results.

To simulate suspended sediment transport in the fluid, we use an advection-diffusion equation for the sediment concentration. The non-conservative form

$$\frac{\partial C}{\partial t} + \nabla \cdot (F_{\text{corr}}\mathbf{U}C) = \epsilon_s\nabla^2 C + \frac{E_b - D_b}{H}, \tag{12}$$

is used in [19] where C is the depth-averaged sediment concentration, ϵ_s the diffusivity coefficient, E_b the erosion flux, D_b the deposition flux, H the depth and F_{corr} a correction factor which accounts for the fact that depth-averaging the product of two variables is not equivalent to multiplying two depth-averaged variables (see [19]). However, the wetting-and-drying scheme used is known to leak sediment. Thus, in order to ensure sediment is conserved when using this wetting-and-drying scheme, in this work we add the choice to solve the following conservative form into Thetis:

$$\frac{\partial}{\partial t}(\tilde{H}C) + \nabla \cdot (F_{\text{corr}}\mathbf{U}\tilde{H}C) = \epsilon_s\nabla^2(\tilde{H}C) + E_b - D_b, \tag{13}$$
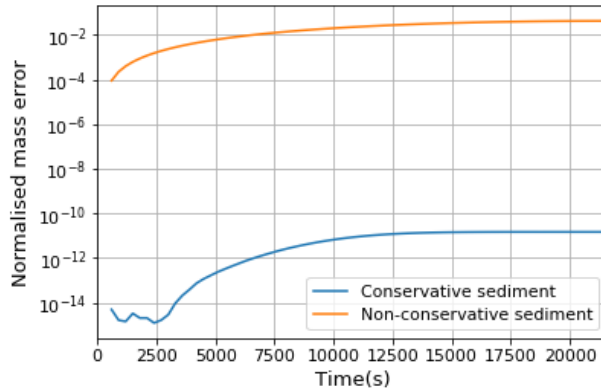
Figure 3: Sediment mass conservation error when using wetting-and-drying with non-conservative sediment equation (12) and with conservative sediment equation (13).

where $\tilde{H}$ is the modified water depth given by (7). Instead of solving for C, we solve for the depth-integrated concentration, $\tilde{H}C$, which allows us to use the same finite element formulation that is used for the non-conservative sediment equation. To verify that our conservative scheme has improved the sediment conservation, we consider the simple Thacker test case of oscillations in a paraboloid bowl with diameter $430.620\,\mathrm{m}$. The hydrodynamic version of this test case is presented in [44] and we refer the reader here for more details. The free surface is initially a paraboloid of revolution with a depth of approximately $50\,\mathrm{m}$ which oscillates inside the bowl with no forcing. No fluid leaves the bowl and therefore the problem is closed hydrodynamically, which means sediments should not leave at the boundaries. We introduce a Gaussian blob of sediment in the wet part of the domain defined by the expression

$$C(t_0) = \begin{cases} 100 \exp^{-(x^2+y^2)/100000} & H \geq 0 \\ 0 & H < 0 \end{cases} \tag{14}$$

and run the simulation for $6\,\mathrm{h}$. The principal differences between the conservative and non-conservative schemes are due to the advection term and therefore to avoid unnecessary complications we set the erosion and deposition terms to 0. At each timestep, t_n, we calculate the normalised mass error using the following formula

$$\frac{\int C(t_n)\tilde{H}(t_n)\,dx - \int C(t_0)\tilde{H}(t_0)\,dx}{\int C(t_0)\tilde{H}(t_0)\,dx}. \tag{15}$$

Figure 3 shows that the normalised mass error is $\mathcal{O}(10^{-2})$ for the non-conservative sediment equation (12) but $\mathcal{O}(10^{-12})$ for the conservative sediment equation (13), which is close to the numerical precision of the model. Thus, the conservative sediment equation conserves sediment much better.

To complete our hydro-morphodynamic model, we use the Exner equation, which governs how the suspended sediment and bedload transport affect the bed. This is unaffected by wetting-and-drying and thus has the same form as that given in [19]:

$$\frac{(1-p')}{m_f}\frac{\partial z_b}{\partial t} + \nabla \cdot \mathbf{Q}_b = D_b - E_b, \tag{16}$$

8

where p' represents the porosity, m_f a morphological scale factor, z_b the bathymetry (also known as the bed profile) and $\mathbf{Q}_b$ the bedload transport. The morphological scale factor is used to artificially increase the rate of bedlevel changes compared with the underlying hydrodynamics and thus save computational time when simulating long-term morphodynamic change (see [19] for more details).

In summary, the wetting-and-drying hydro-morphodynamic model is given by equations (9), (10), (13) and (16). This system may be interpreted as a vector PDE and three scalar PDEs, or five scalar PDEs. These equations are discretised and solved using finite element methods. Due to the complexity and nonlinearity of the combined system of equations, the triple formed by (9)–(10), (13) and (16) is solved alternately, as opposed to as a monolithic system. The mesh upon which the finite element methods are defined is moved from timestep to timestep using the methods described in Section 2. However, it is not enforced that this occurs *every* timestep, meaning that we are free to choose the frequency with which mesh movement is applied.

4. Test Cases

4.1. Migrating Trench

As a first test case, we consider the case of a migrating trench, for which there exists experimental data in [45]. We choose this test case because it has already been used in [19] with a fixed uniform mesh for validation of our hydro-morphodynamic model configuration without wetting-and-drying. This non-wetting-and-drying configuration is equivalent to (9)–(10) with $\delta = 0$ in (8), coupled with the non-conservative form of the sediment equation (12), as well as (16) which remains unchanged.

The set-up for the migrating trench considered here is the same as that in [19], modelling both suspended and bedload transport with the magnitude and angle corrections from the slope effect, using the Nikuradse friction formula for C_h in (10), and following guidance in that work using a morphological scale factor, m_f, of 100 for computational efficiency. In [19], it was shown that this test case is very sensitive to the value specified for diffusivity. Using mesh adaptation on this problem will change mesh resolution and this will change the effective numerical diffusivity of the model. This makes assessing the performance of different mesh movement strategies against experimental data challenging, due in part to the issue of "getting the right answer for the wrong reasons". We therefore also choose to compare against a numerical result obtained using a uniform fixed mesh with $\Delta x = 0.05\,\mathrm{m}$ (320 mesh elements in the x-direction). In order for this benchmark result to best reflect the experimental data we choose here to calibrate a value for diffusivity on this mesh that will then be held constant for all subsequent experiments. We use a gradient-based optimisation routine to invert, or perform optimal parameter estimation, for diffusivity. This involves the minimization of the reduced functional

$$J(\epsilon_s) = \frac{1}{2} \int_0^T \int_\Omega (h(\epsilon_s) - h^{\mathrm{obs}})^2 \; \mathrm{d}x \; \mathrm{d}t, \tag{17}$$

subject to the forward version of the hydro-morphodynamic model being satisfied. The gradients are computed using the adjoint method. Here ϵ_s (diffusivity coefficient) is the parameter being optimised for, h^{obs} is the experimental data from [45] and h is the bathymetry obtained from the model, which in (17) is assumed to be a function of ϵ_s alone, but is understood to result from a run of the full hydro-morphodynamic model. More detail on the use of the adjoint framework with the hydro-morphodynamic model can be found in [46]. Figure 4a shows the reduction in the functional value for each iteration and Figure 4b shows how this corresponds to the convergence of

ϵ_s to an optimum parameter value. We find that the optimal value of the diffusivity coefficient is $\epsilon_s = 0.18011$ to five significant figures. Throughout this section, this value is used as the diffusivity parameter.
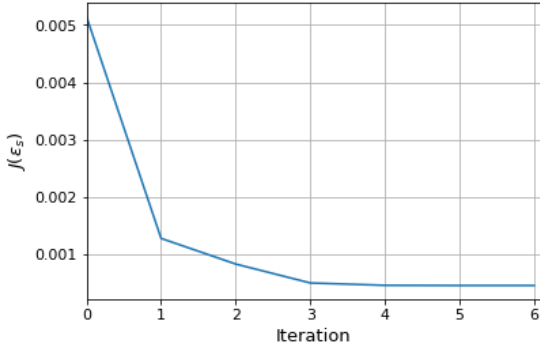

Figure 4a: Minimization of the reduced functional for determining the diffusivity parameter in the migrating trench test case.
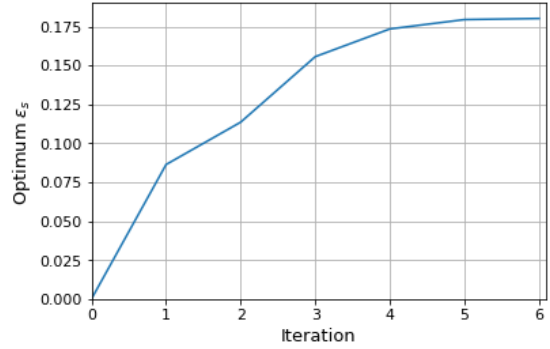
Figure 4b: Convergence of the diffusivity parameter in the inversion performed for the migrating trench test case.

To begin with, consider a series of fixed meshes varying from 320 mesh elements in the x-direction down to only 8 elements in the x-direction. Throughout this work, the error in the bed profile is calculated at the end of the simulation. The existence of experimental data for this test case means that we are able to calculate the total error. Total error can be thought of as being model error (the error due to using a simplified model to approximate a real world problem), plus discretisation error (the error arising from using a finite mesh to solve the model equations). In this section, the error is calculated using a pointwise ℓ_2 error norm at the location of the data points in the experiment.

Figure 5a shows the total error between the final bed profile from the experimental data and from the model output for a series of fixed uniform meshes. In this figure, the points in orange are the results where the number of mesh elements in the x-direction $= [10, 20, 40, 80]$ and in green are those where the number of mesh elements in the x-direction $= [8]$. This is because the initial profile of the trench is incorrectly defined on meshes with these numbers of mesh elements, due to the complex structure of the profile. Furthermore, the orange points are incorrectly defined in a different way to the green ones. For example, instead of the second slope starting at the correct location of $x = 9.5\,\mathrm{m}$, for the orange points it starts at $x = 9.6\,\mathrm{m}$, whereas for the green points it starts at $10\,\mathrm{m}$. These differences affect the results on the fixed uniform mesh. For the remaining blue points (number of mesh elements in the x-direction $= [32, 64, 160, 320]$), the initial profile is correctly defined. This difference in the initial trench profile definition results in fluctuations in the error trend. For the orange points, the error increases as the number of mesh elements increases, whereas for the green point the error is higher (relative to the final converged outcome). It is surprising that with 10 elements, the total error is lower than on finer meshes. This is likely due to this test case's sensitivity to diffusivity and the fact that the effective diffusivity is different at the coarser mesh resolutions. Despite these fluctuations, the general trend is that the value of the error norm begins to plateau around 20 mesh elements and by 32 mesh elements is virtually unchanged as the number of mesh elements further increases. This is because, for larger numbers of mesh elements, the initial trench profile is well-defined on the mesh and as was noted in [19] for this test

10

case the model is then quite insensitive to changes in Δx and Δt; the total error in the solution is thus dominated by the error from the model rather than the discretisation error from the mesh. Figure 5a shows that the total error in the bed profile using a series of fixed uniform meshes clearly converges to a value of approximately $0.0244\,\mathrm{m}$ for all values and thus we can proceed with the test case.

To better understand the discretisation error, we calculate the error norm between the final bed profile for different mesh resolutions compared to the final bed profile obtained with the finest fixed uniform mesh with 320 mesh elements in the x-direction ($\Delta x = 0.05\,\mathrm{m}$). Figure 5b shows that as the mesh becomes finer the error decreases with an order of approximately $(\Delta x)^4$. We again separate out visually number of mesh elements in x-direction $= [10, 20, 40, 80]$ marked in orange, number of mesh elements in x-direction $= [8]$ marked in green and the remaining marked in blue. All sets clearly converge to the same value at slightly different rates, and thus we can proceed.
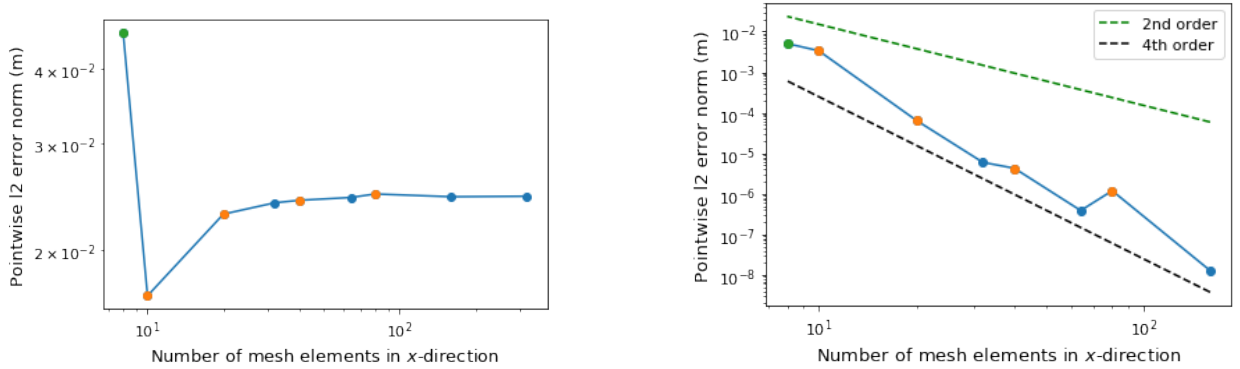
(a) Total error between final bed profile from experiment data and from model output.

(b) Discretisation error between final bed profiles where reference solution has 320 mesh elements ($\Delta x = 0.05\,\mathrm{m}$).

Figure 5: Errors for a series of fixed uniform meshes for the migrating trench test case. The colours distinguish differences in the definition of the initial profile on the mesh e.g. for orange, green and blue the initial second trench slope starts at $9.6\,\mathrm{m}$, $10\,\mathrm{m}$ and $9.5\,\mathrm{m}$ respectively.

4.1.1. Monitor function

The main source of discretisation error in this test case is due to the fact that the initial trench profile is not well defined on the mesh. Even if a mesh is chosen so the profile is initially well defined, as soon as the simulation starts, the bed begins to move so the profile may quickly become ill-defined (hence the differences in total error even for the blue points in Figure 5a). Furthermore, there are large sections of the trench profile which are flat and are relatively unchanged during the simulation, especially at the inflow. Therefore, a good choice of monitor function for this test case is one that results in refinement of the mesh in regions where the bed gradient changes and coarsening of the mesh in regions where there is little bed gradient change. We thus consider here the monitor function

$$m(x, y) = 1 + \alpha \, \frac{\|\boldsymbol{H}(z_b)\|_F}{\max\limits_{x,y} \|\boldsymbol{H}(z_b)\|_F}, \tag{18}$$

where $\boldsymbol{H}(z_b)$ represents the Hessian of the bathymetry, $\|\cdot\|_F$ the Frobenius norm and α is a user-defined parameter, which we determine later in this section. Note that the Frobenius norm is taken on an element-by-element basis, giving rise to a spatially varying field. This field is projected into the $\mathbb{P}1$ space for use in the monitor function.

11

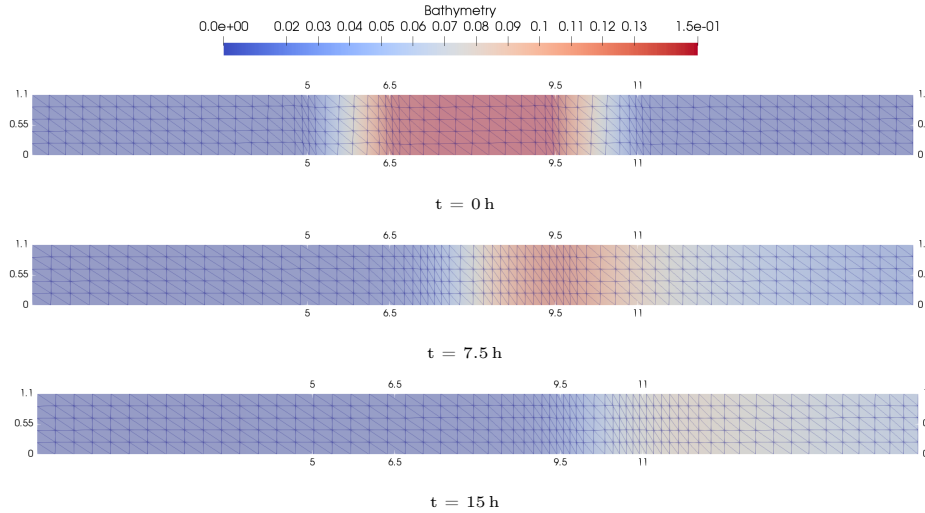Figure 6 shows an example of how the mesh is moved using (18) for this particular test case.

Figure 6: Mesh movement using (18) for the migrating trench test case with 64 mesh elements in the x-direction and $\alpha = 2$. Results are shown at three time levels demonstrating mesh movement to capture bed evolution to the right.

4.1.2. Results

Below 64 equally sized mesh elements, Figures 5a and 5b show the total error and discretisation error respectively are near convergence. Thus we apply our adaptive mesh technique only for meshes with less than or equal to 64 mesh elements in the x-direction. Figures 5a shows that there is a distinction at 20 mesh elements where the total error decreases significantly. Thus we have chosen to consider the number of mesh elements < 20 as one group which we call the smaller number of elements hereafter, and $20 \leq$ number of mesh elements ≤ 64 another group which we call the larger number of elements hereafter.
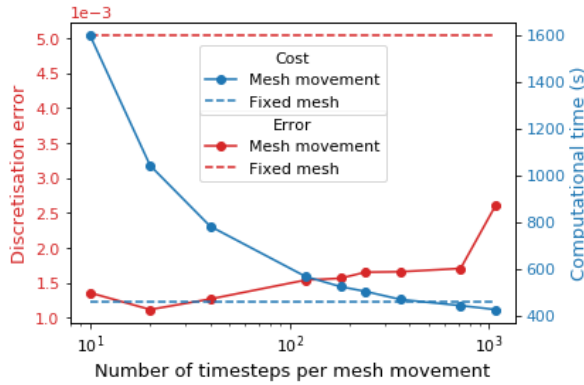
(a) For $\alpha = 4$ and with 8 mesh elements. Dashed lines represent baseline accuracy and computational cost achieved with fixed uniform mesh of 8 elements.
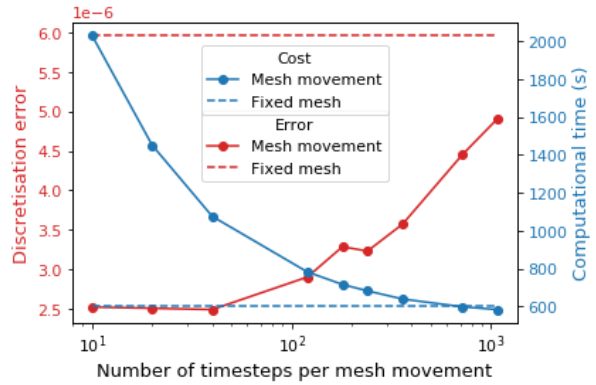
(b) For $\alpha = 1.5$ and with 32 mesh elements. Dashed lines represent baseline accuracy and computational cost achieved with fixed uniform mesh of 32 elements.

Figure 7: Effect of changing the number of timesteps of the hydro-morphodynamic model performed between each mesh movement (ie. the frequency of mesh movement) for the migrating trench test case.

With the mesh movement technique used in this work, we can choose how frequently the mesh is moved. If the modification to the mesh coordinates proposed by (M-A) is only minor then solving this equation at every timestep of our hydro-morphodynamic model may prove an unnecessary expense. Figures 7a and 7b show the effect of changing the frequency of mesh movement for two particular choices for the number of mesh elements and the α parameter in (18). For consistency, we use the same Δt value in our hydro-morphodynamic model throughout. In order to have confidence in our conclusions, we have chosen a number of elements from the smaller number of elements group (Figure 7a) where the initial trench profile is not well-defined on a uniform mesh and the larger number of elements group (Figure 7b) where the initial trench profile is well-defined on a uniform mesh.

Figures 7a and 7b show that the larger the number of timesteps per mesh movement the greater the discretisation error but the lower the computational cost, which is to be expected as mesh movement methods have a non-negligible computational cost. Note that the total number of timesteps in this simulation is 2160 and therefore for 1080, the largest number of timesteps per mesh movement considered here, the mesh is only moved twice during the simulation. Significantly, both figures show that if a large enough number of timesteps per mesh movement is chosen, we can reduce the computational cost below that of using a fixed mesh and still be more accurate. This supports our argument, at least for this simple test case, that mesh movement can not only reduce error but also reduce the computational cost of the simulation. In the remainder of this section, we choose to move the mesh after every 40 timesteps of the hydro-morphodynamic model (equivalent to moving the mesh 54 times during the simulation), as this provides a good balance between computational cost and accuracy.

We must also define the value of α in 18, which controls the mesh movement due to the magnitude of the Hessian of the bathymetry. In order to determine an optimum value of α, we conduct a small sensitivity study, but in future work will seek a gradient-based approach (see Section 5). Figures 9a and 9b show the discretisation error for a range of α values for smaller and larger number of elements in the x-direction respectively. Comparing these two figures shows that we are right to make the distinction as particularly for the discretisation error, the relationship to α is very different. For the smaller number of elements, the discretisation error is significantly smaller than the discretisation error for the fixed mesh ($\alpha = 0$) for all values of α. The minimum discretisation error is achieved when $\alpha = 4$. For larger number of elements, the discretisation error is minimized with $\alpha \approx 1.5$, which is a smaller value than that for the small number of elements. As the number of elements increases the effect of mesh movement on the error decreases. This is likely because for larger number of elements, the discretisation error is already very small ($\mathcal{O}(10^{-5})$) and at this scale discretisation error due to interpolation between meshes begins to become important, rather than just discretisation error linked to inaccurate definition of the bathymetry.

Figures 8a and 8b show the total error for a range of α values for smaller and larger number of elements respectively. The error is always smaller when using mesh movement methods compared to a fixed uniform mesh ($\alpha = 0$) and in all cases the total error is minimized when α is at its largest. Unfortunately, we are unable to test for larger α than the values presented because, as α becomes larger, the ratio between the maximum and minimum mesh spacing increases. Thus the hydro-morphodynamic model diverges for the timestep size used in this experiment due to CFL number constraints (note we have confirmed with experiments that if a smaller Δt is used the model no longer diverges). Figure 8a shows that, for a small number of mesh elements in the x-direction where for a fixed uniform mesh the initial profile would be ill-defined, the total error seems to be

13

converging to a value of around 0.015 m. However, for larger number of elements in Figure 8b, the total error seems to converge to a different value for each number of elements. Figure 10a shows that if we choose values of α where the discretisation error is minimized then the total error is approximately the same for all large number of elements and approximately the same as the value which the uniform mesh converges to. This suggests that the differences in total error at large values of α are due to using a simplified model to approximate a real world problem (i.e. model error). We have already shown that this problem is very sensitive to changes in diffusivity, and hence this increase in model error may be due to changes in the effective diffusivity in the simulation.

(a) Small number of mesh elements.

(b) Large number of mesh elements.
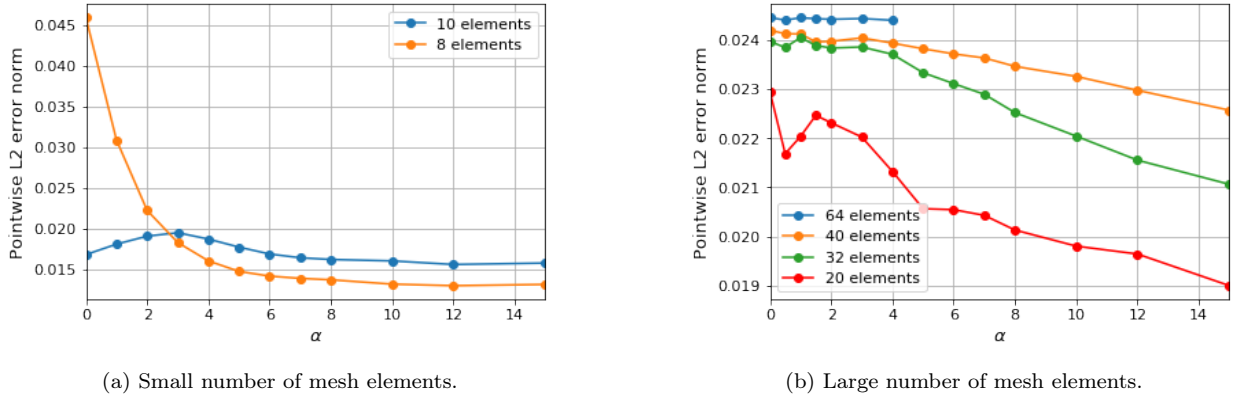
Figure 8: Total error between model and experiment outputs for varying α values for the migrating trench test case.

(a) Small number of mesh elements.

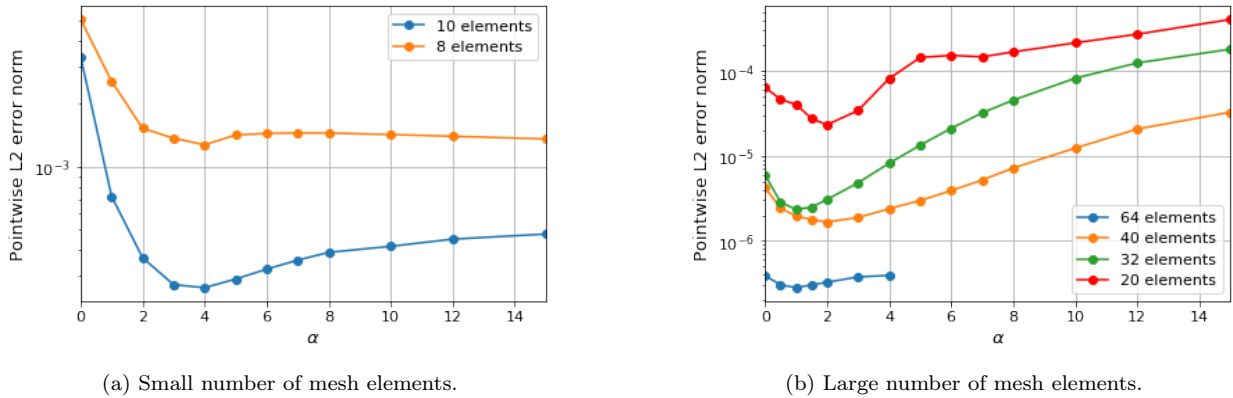(b) Large number of mesh elements.

Figure 9: Discretisation error for varying α values for the migrating trench test case.

Figure 10b presents a summary of our discretisation error results and shows both the minimum error achieved (optimum parameter set) and the error achieved by a generalised parameter set ($\alpha = 4$ for small number of elements and $\alpha = 2$ for large number of elements). Moving mesh methods consistently result in a lower error than fixed mesh methods and furthermore our general parameter choice produces very similar error results to the optimum values. This leads us to suggest that for this monitor function (18), if the bathymetry is ill-defined on the uniform mesh $\alpha = 4$ represents a good choice and if the bathymetry is well-defined then $\alpha = 2$ represents a good choice.

14

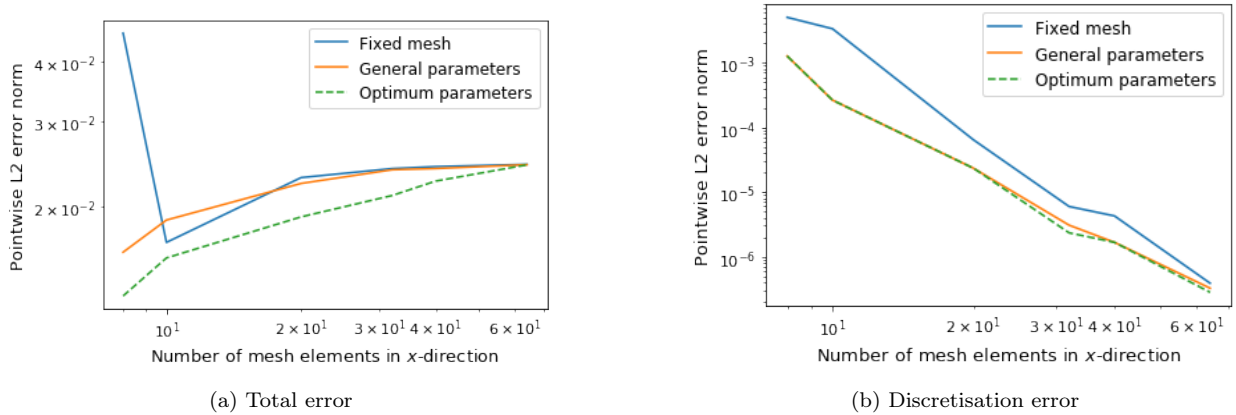(a) Total error   (b) Discretisation error

Figure 10: Comparing errors for a fixed uniform mesh and for mesh movement methods with general and optimal parameters, for the migrating trench test case.

Using a simple test case, the experiments in this section show that mesh movement methods can provide a means of controlling both the discretisation error and the computational cost of a simulation. We have demonstrated that the choice of monitor function and the mesh movement frequency are two factors which can allow us to achieve this control. There exist other such parameters (for example, the relative tolerances used when solving (M-A)). However, these are not discussed here, for brevity.

4.2. Migrating Trench 2D

The test case considered in Section 4.1 is effectively 1D with very little variation in the bed or the flow in the y-direction. However, most hydro-morphodynamic problems are 2D in nature and thus we modify the test case in Section 4.1 to make it 2D by introducing a slope of 0.1 in the y-direction. All other set-up parameters are kept the same and the initial profile is shown in Figure 11.
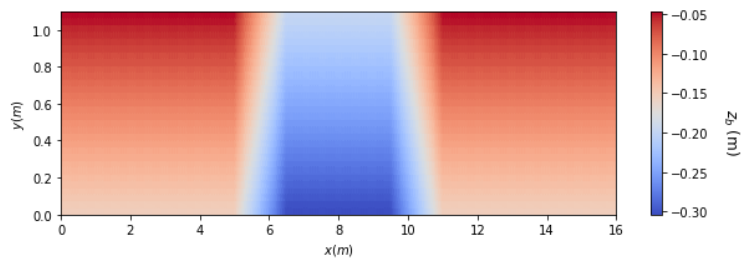

Figure 11: Set-up of the 2D migrating trench test case.

As before, we start by considering a series of fixed uniform meshes, choosing 8, 16, 32, 40, 64, 80 and 128 mesh elements in the x-direction. As the length in the y-direction is one sixteenth that in the x-direction, we choose the number of mesh elements in the y-direction to be one sixteenth of the number in the x-direction (rounding up to the nearest integer where necessary). Note that for the eight elements in the x-direction we use one element in the y-direction and for sixteen elements in the x-direction we use two elements in the y-direction, to allow for some mesh movement in the y-direction.

15

Since this is a theoretical test case, we no longer have experimental data to compare against but can still analyse the discretisation error by considering results obtained on a fine fixed uniform mesh of 320 mesh elements in the x-direction ($\Delta x = 0.05$ m) and 20 mesh elements in the y-direction ($\Delta y = 0.055$ m) to be truth. Furthermore, as this is a 2D case, we can no longer just consider the pointwise error at specific location, but instead calculate the $\mathcal{L}_2$ error norm over the whole domain. Figure 12 shows that as the mesh becomes finer the error decreases approximately linearly.
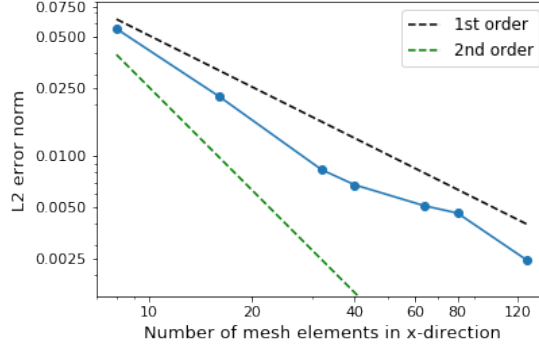
Figure 12: Discretisation error for the bed profile in the 2D migrating trench test case using a series of fixed uniform meshes. The x-axis lists the number of elements in the x direction in the experiment, the y-direction is refined commensurately. The reference solution compared against is that obtained using a fine fixed unifom mesh with 320 mesh elements in x-direction ($\Delta x = 0.05$ m) and 20 mesh elements in y-direction ($\Delta y = 0.055$ m).

4.2.1. Monitor function

As with the previous test case, the main source of discretisation error is due to the initial trench profile being ill-defined on the mesh. Given that the profile is now more complex with a linear slope, we construct a more complex monitor function than (18) which involves not only the second order derivatives but also the first order ones:

$$m(x,y) = 1 + \max_{x,y}\left(\alpha\frac{\|\boldsymbol{H}(z_b)\|_F}{\max\limits_{x,y}\|\boldsymbol{H}(z_b)\|_F},\ \beta\frac{\|\nabla z_b\|_2}{\max\limits_{x,y}\|\nabla z_b\|_2}\right). \tag{19}$$

This introduces two user-defined parameters which control the effect of the underlying bathymetry on mesh movement; α controls the effect of the second order derivative (curvature), whilst β controls the effect of the first-order derivatives (slope). As with the Frobenius norm, the 2-norm of the gradient term in (19) is taken on an element-by-element basis and the resulting piecewise constant field is projected into $\mathbb{P}1$ space.

Figure 13 shows an example of how the mesh moves based upon (19). To emphasise the differences between the monitor functions (18) and (19), we have picked example parameters so that the monitor function is only dependent on the first order derivatives (i.e. $\alpha = 0$). Figure 13 shows that the mesh moves in both the x and y-direction with the element height in the y-direction being generally smaller at the top of the domain than at the bottom. This is especially visible between $x = 6.5$ m and $x = 9.5$ m.
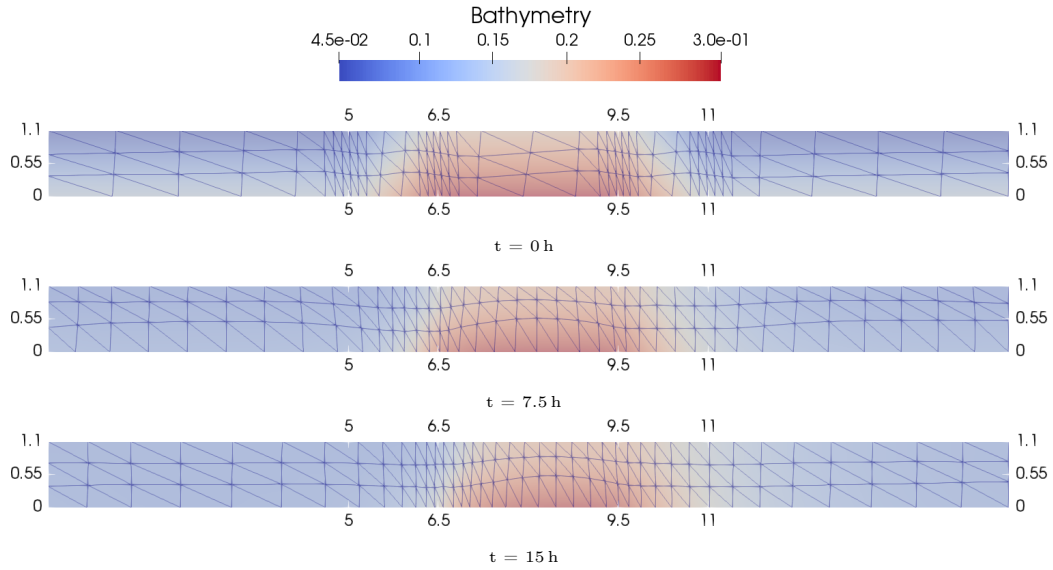
16

Figure 13: Mesh movement using (19) for the 2D migrating trench test case with 40 elements in x-direction, 3 elements in y-direction, $\alpha = 0$ and $\beta = 10$. Results are shown at three time levels demonstrating mesh movement to capture bed evolution to the right.

4.2.2. Results

The other user-defined parameter is the frequency of the mesh movement. Given our results for the previous test case, we could make an educated guess that a value of $\mathcal{O}(5 \times 10^2)$ would be a good frequency for this type of problem, providing a good balance between computational cost and error. Indeed Figures 14a and 14b confirm this guess and show that as the number of model timesteps per mesh movement increases, the computational cost decreases and the discretisation error increases (note that the results in Figure 14a are produced using a timestep of $0.25\,\mathrm{s}$, whilst for Figure 14b a timestep of $0.125\,\mathrm{s}$ is used for reasons of Courant number stability discussed later). This suggests that it is possible to predict a good mesh movement frequency without performing extensive sensitivity analysis. Importantly, the discretisation error using mesh movement methods always remains lower than the error with a fixed uniform mesh so whatever frequency is chosen, mesh movement always represents an improvement. Using this analysis, as before, we choose to move the mesh every 40 timesteps of the hydro-morphodynamic model when this timestep is $0.25\,\mathrm{s}$, as this represents a good balance between computational cost and accuracy. For consistency, when this timestep is $0.125\,\mathrm{s}$, we move the mesh every 80 timesteps.

17

(a) For $\alpha = 7$, $\beta = 7$ and 8 mesh elements in x-direction.   (b) For $\alpha = 5$, $\beta = 2.5$ and 128 mesh elements in x-direction.
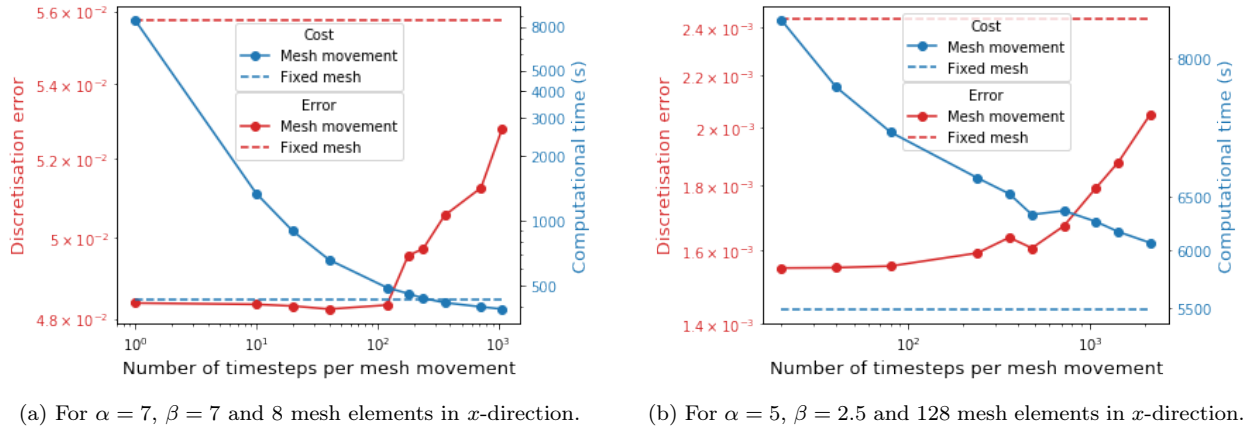
Figure 14: Effect of changing the number of timesteps of the hydro-morphodynamic model performed between each mesh movement (ie. the frequency of mesh movement) for the 2D migrating trench test case.

Because the monitor function (19) has user-defined parameters, α and β, we conduct a sensitivity study to investigate the impact on discretisation error of these parameters and determine optimum values. Consider Figure 15 and note the different scales on the y-axes. This figure shows that when the number of mesh elements in the x-direction is less than or equal to 64 the error is minimized by choosing $\beta = \alpha$, corresponding to a monitor function where the first- and second-order derivatives of the bathymetry have equal weighting. In contrast, for larger numbers of mesh elements shown in the top left quadrant of Figure 15, the importance of the first derivative decreases and the discretisation error is lowest for $\beta = 0$ i.e. when there is no contribution from the first derivative. This is likely because, as the mesh becomes finer, any areas where the bathymetry has high first-order derivatives are adequately captured, but the more difficult areas with high second-order derivatives are not. We note that, in order to run our model at these fine resolutions, it is necessary to reduce the timestep from $\Delta t = 0.25\,\mathrm{s}$ to $\Delta t = 0.125\,\mathrm{s}$. Nevertheless, for 128 mesh elements, $\beta = 0$ and $\alpha > 4$, the hydro-morphodynamic model still becomes unstable (at the timestep used in this experiment) and diverges. This instability is likely why the error increases from $\alpha = 2$ to $\alpha = 4$ when $\beta = 0$. This prevents us from drawing complete conclusions from our results for this number of mesh elements, but the trend is nevertheless clear.

In all cases, the relation between α and β is more significant than the value of either parameter. This is a useful result as it is much easier to predict an appropriate relationship between β and α than their value. However, for all numbers of elements, $\alpha = 7$ appears to be a good choice. This choice of α is confirmed in Figure 16, which shows the error achieved by an optimum parameter choice and the general parameter choice of $\alpha = 7$ and $\beta = 7$ for small numbers of mesh elements and $\alpha = 7$ and $\beta = 0$ if large numbers. This general optimum choice for α shows the influence of the slope in the y-direction because, for the 1D case in Section 4.1, the optimum α is approximately half this value. Thus, we have shown that mesh movement methods can be used to improve accuracy and reduce computational cost for 2D cases with non-trivial bathymetries. We have also been able to use this test case to draw some conclusions about general good parameter choices for (19). In the next section, we test these general optimum parameter choices on a completely different 2D problem to see if they are again optimal or close to optimal choices.
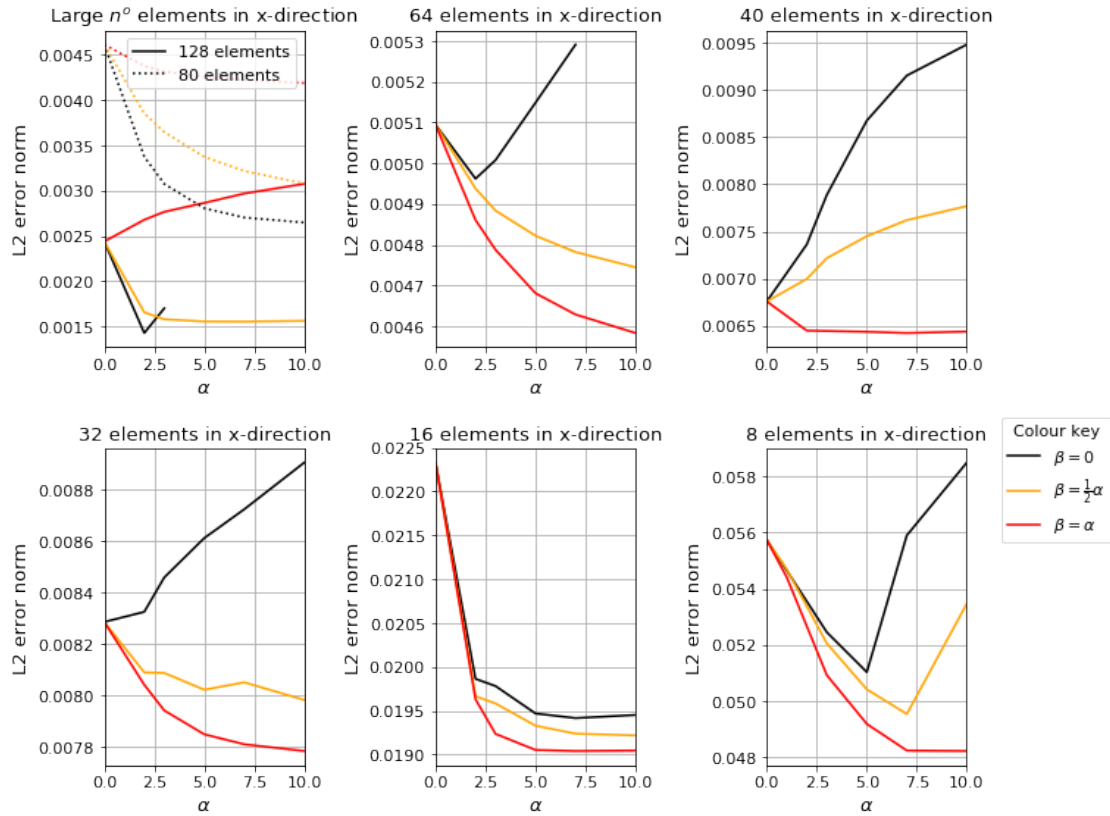
18

Figure 15: Discretisation error with mesh movement methods for varying α, β and number of mesh elements for the 2D migrating trench test case.
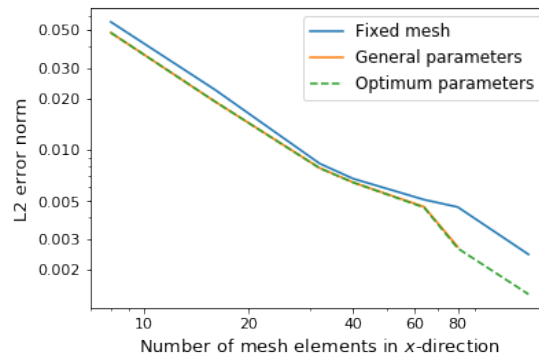

Figure 16: Comparing discretisation error for a fixed uniform mesh and for mesh movement methods with general and optimal parameters, for the 2D migrating trench test case.

4.3. Beach Profile

We have shown that moving mesh methods can improve accuracy and efficiency for test cases with steep gradients in a fully wet domain. Coastal problems often have a wetting-and-drying interface, for example as a wave or tide moves up a beach. These problems have been historically difficult to solve because of their computational expense, but mesh refinement methods provide a

19

way to retain accuracy, whilst improving efficiency. Mesh movement methods have been used previously to successfully simulate a wetting-and-drying interface, for example in [47]). However to the best of our knowledge they have not previously been used to solve coupled hydro-morphodynamic cases with wetting and drying interfaces.

In this section, we consider the test case of a wave running up a beach slope, whose setup is shown in Figure 17. The incoming wave is governed by

$$\mathbf{u} = \cos\left(-\frac{1}{2}t\right) \quad \text{and} \quad h = \frac{1}{4}\cos\left(-\frac{1}{2}t\right) \quad \text{at } x = 0\,\text{m}. \tag{20}$$

The simulation is run for $3600\,\text{h}$ (approximately equivalent to 5 months) with a morphological scale factor of 10,000. This means the full wave is generated over one hundred times during the simulation.

For this test case, we use a similar hydrodynamic set-up to that used for the Balzano test case in [42]. Following that work, instead of using the Nikuradse friction formula as in previous test cases, the Manning friction formula is used to determine the bed friction

$$C_h = g\frac{n^2}{H^{1/3}}, \tag{21}$$

where n is the manning drag coefficient, set to $n = 0.02\,\text{s}\,\text{m}^{1/3}$, as in [42], a standard value for sandy coasts. However, whereas [42] model only the hydrodynamics, we model the morphodynamic effects as well, simulating both suspended sediment and bedload transport using the wetting-and-drying conservative form of the hydro-morphodynamic model discussed in Section 3. The magnitude corrections to the slope to account for gravitational effects are also applied (see [19] for more details). The parameters used in the simulation are summarised in Table 1.

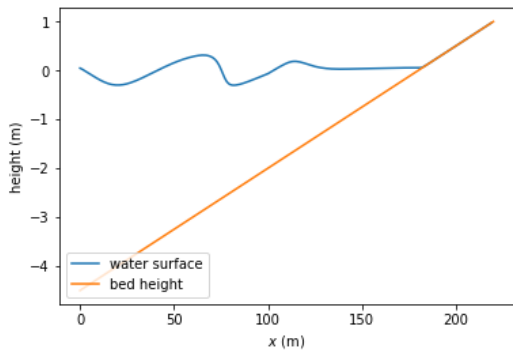Table 1: Parameter values for beach profile test case.

Figure 17: Set-up of beach profile test case.

Variable Name	Variable Value
Length in x-direction	$220\,\text{m}$
Length in y-direction	$10\,\text{m}$
Morphological simulation time	$3600\,\text{h}$
Morphological scale factor	10000
Median particle size (d_{50})	$2\times 10^{-4}\,\text{m}$
Sediment density (ρ_s)	$2650\,\text{kg}\,\text{m}^{-3}$
Water density (ρ_f)	$1000\,\text{kg}\,\text{m}^{-3}$
Kinematic viscosity (ν)	$0.5\,\text{m}^2\,\text{s}^{-1}$
Bed sediment porosity (p')	0.4
Diffusivity (ϵ_s)	$100\,\text{m}^2\,\text{s}^{-1}$
Manning friction coefficient (n)	$0.02\,\text{s}\,\text{m}^{1/3}$
Wetting-and-drying parameter (α)	$0.025\,\text{m}$
Norm smoother parameter (β)	$0.4\,\text{m}\,\text{s}^{-1}$

Note our model does not currently have the ability to simulate shoaling waves and thus we have used a viscosity "sponge" in the dry area to cause wave dissipation defined by

$$\text{viscosity sponge} = \begin{cases} \nu & x < 100, \\ \nu(-399 + 4x) & x \geq 100. \end{cases} \tag{22}$$

20

Given the test case set-up, we can predict that the incoming wave will produce two different types
of sediment transport behaviour in the domain. At the input of the domain, we expect large scale
erosion as the wave entering the domain acts as a breaking wave would in a normal coastal set-up.
Further along the domain, we would expect smaller ripple effects to occur as the wave passes over
the bed. We could make an educated assumption for where the location of the breaking wave effects
ends and the ripple effects starts. However, Figure 18 shows that, even for the coarsest fixed uniform
mesh considered in this section (44 mesh elements in the x-direction), we can deduce the limit is
approximately $x = 70$ m. To confirm that this limit does not change when the mesh becomes finer,
Figure 18 also shows the evolution profile for the finest mesh considered in this section (440 mesh
elements in the x-direction).

The erosion at the input of the domain clearly dominates the sediment transport behaviour in
the domain and any mesh movement will likely focus on this part. Thus, in order to be able to
improve the accuracy of the predictions of the ripple like behaviour, throughout this section we also
consider the error over the subdomain separately to the error over the whole domain. The part the
user is more interested in modelling will determine the optimum monitor function.

To begin with, we consider a series of fixed meshes varying from a uniform spacing of 440 mesh
elements in the x-direction ($\Delta x = 0.5$ m) to 44 mesh elements ($\Delta x = 5$ m). In order to keep the
elements in the fixed uniform mesh roughly isosceles, where the number of mesh elements in the
x-direction is less than 100, we use five elements in the y-direction ($\Delta y = 2$ m), between 100 and
400 elements in the x-direction we use 10 elements in the y-direction ($\Delta y = 1$ m) and for 440 mesh
elements we use 20 mesh elements in the y-direction ($\Delta y = 0.5$ m). Note that we set $\Delta t = 0.5$ s for
all the simulations in this section, in order to ensure that the wave is correctly defined. Because this
is a theoretical test case for which there exists no real data, we use the model solution at 440 mesh
elements in the x-direction ($\Delta x = 0.5$ m) and 20 mesh elements in the y-direction ($\Delta y = 0.5$ m)
as the real solution and are thus showing the estimates of the discretisation error. Figures 19 and
21 show that, as the mesh becomes finer, the error decreases approximately linearly over both
the whole domain and the subdomain. Already we see a difference in the error between the two
domains, with the error converging more uniformly in the subdomain than in the whole domain.
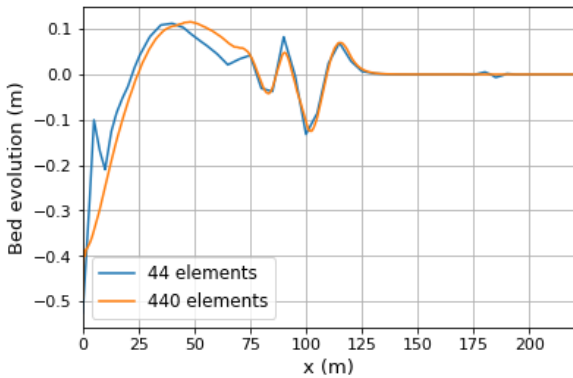

Figure 18: Total bed evolution of beach profile over
whole domain after 3600 h for fixed uniform mesh with
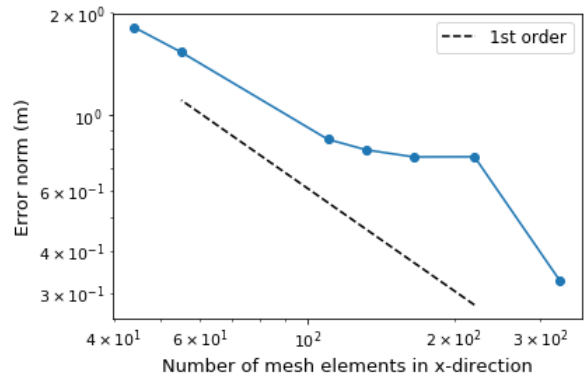440 mesh elements ($\Delta x = 0.5$ m).

Figure 19: Discretisation error for whole domain, using
a fixed uniform mesh where reference solution has 440
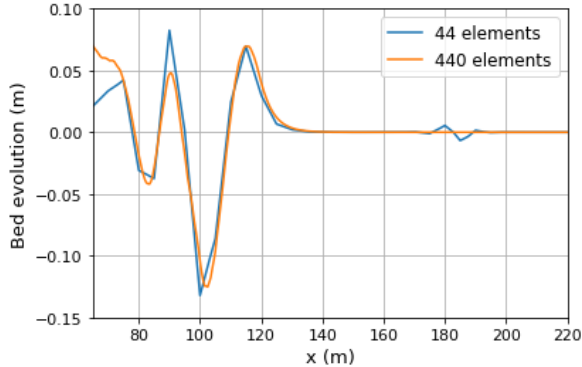mesh elements ($\Delta x = 0.5$ m) for beach test case.

21

Figure 20: Total bed evolution of beach profile for $x >$ 70 m after 3600 h for fixed uniform mesh with 440 mesh elements ($\Delta x = 0.5$ m).

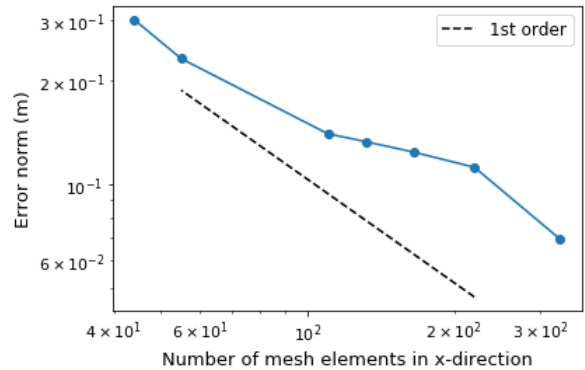
Figure 21: Discretisation error for $x > 70$ m, using a fixed uniform mesh where reference solution has 440 mesh elements ($\Delta x = 0.5$ m) for beach test case.

4.3.1. Monitor function

To apply mesh movement methods to this test case, we use the same monitor function as in Section 4.2 (19). We are principally interested in bed changes which occur in the direction perpendicular to the shoreline (x-axis) and thus smooth the monitor function in the cross-shore direction (y-direction) by applying Laplacian smoothing,

$$\frac{m_{\text{smooth}} - m}{\Delta t} = K \frac{\partial^2}{\partial y^2} \left(m_{\text{smooth}} \right), \quad \text{where} \quad K = N \frac{\Delta x^2}{4 \Delta t}. \tag{23}$$

Here N is equivalent to the number of applications of a $(1, -2, 1)$ filter, which we choose to be 40.

4.3.2. Results

As in the previous sections, we can choose the mesh movement frequency. In previous simulations, we used a frequency of $\mathcal{O}(5 \times 10^2)$ and could reasonably expect again that this would be appropriate. (Note that the frequency must exactly divide the total number of model timesteps, so we cannot use exactly the same value every time.) However, upon running a sensitivity study, Figure 22a shows that when we consider the whole domain, moving the mesh four times during the simulation (*i.e.* at least every 648 timesteps) results in a significant reduction in the discretisation error compared to the fixed mesh. Furthermore, using any frequency greater than this also results in a similar error reduction. This is predictable because we know the error for the whole domain case is focused at the domain input, where the bed evolution pattern remains steady throughout the simulation, *i.e.* the areas at the input with the greatest erosion rates when the bed begins to evolve are the same areas with the greatest erosion rates at the end of the simulation. Thus, once the mesh movement algorithm has identified the region of interest, it only needs to move the mesh a couple of times in order to optimise the error reduction. A frequency of $\mathcal{O}(5 \times 10^2)$ would still result in an optimum error reduction, just not an optimum cost reduction.

If only the subdomain is considered, Figure 22b shows that the number of mesh movements during the simulation does affect the error over this subdomain. Here a frequency of $\mathcal{O}(5 \times 10^2)$ timesteps per mesh movement would be appropriate, and when we run the sensitivity study we see that the discretisation error is marginally more minimised by moving the mesh every 72 timesteps, but that every value of this order results in an error value close to the minimum. Surprisingly,

22

moving the mesh more frequently than this increases the error. This may be because the mesh becomes too concentrated around areas of complex bathymetry, meaning the incoming wave is ill-defined. Another cause may be because more frequent mesh movement induces greater interpolation error over the whole simulation. Moving the mesh every 16 timesteps (the smallest number of timesteps considered here) corresponds to a total of 324 mesh movements. This is almost five times more times than if it is moved every 72 timesteps and thus we would expect the mesh interpolation error to be much greater. Alternatively, because of the lack of experimental data, we do not know the true solution for this problem and thus, it may be that with more frequent mesh movement, the model captures an aspect of the problem which cannot be captured using a fixed uniform mesh no matter the resolution. Thus, in cases where both the hydrodynamics and morphodynamics change frequently throughout the simulation, it may be more appropriate to use a slightly larger frequency than that used in the tests cases in Sections 4.1 and 4.2. However, our beach profile test case suggests that any frequency of $\mathcal{O}(5 \times 10^2)$ would still result in a close to optimum accuracy improvement compared to fixed uniform meshes, and a close to optimum cost improvement.

In both the whole domain and subdomain case, the number of mesh movements may be chosen so that the computational cost is less than the fixed mesh algorithm, whilst still being more accurate. Following our observations, when considering the error over the whole domain, the mesh is moved every 648 timesteps of the hydro-morphodynamic model, because this gives a low computational cost and a low discretisation error. When considering the error over the subdomain, the mesh is moved every 72 timesteps of the hydro-morphodynamic model as this minimizes the discretisation error and results in a low computational cost.
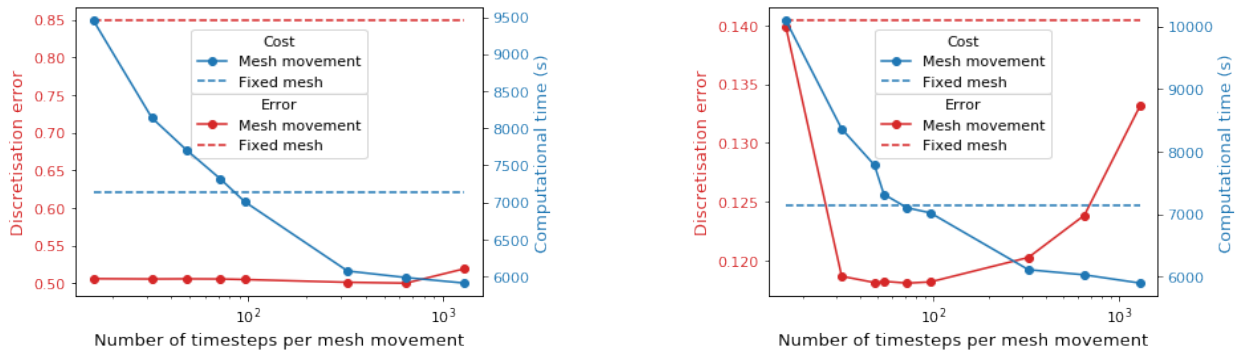
(a) On whole domain for $\alpha = 5, \beta = 0$, 110 mesh elements in x-direction.

(b) On the subdomain for $\alpha = 0, \beta = 7$, 110 mesh elements in x-direction.

Figure 22: Effect of changing the number of timesteps of the hydro-morphodynamic model performed between each mesh movement (ie. the frequency of mesh movement) for beach test case.

We can also choose an α and β in (19) to minimize the discretisation error. In the previous section, we found that a value of α or $\beta = 7$ provided a good general optimisation of the discretisation error and thus we conduct our sensitivity study using similar values. We consider first the discretisation error over the whole domain. Figure 23 shows that $\alpha = 7$ does indeed correspond to an either optimal or very close to optimal minimization in the discretisation error over the whole domain. Like in the test case in Section 4.2, the relation between α and β is more important, though, and Figure 23 shows that the error over the whole domain is minimized when $\beta = \mu$ and $\alpha = 0$. This corresponds to a monitor function which is dependent on the first derivative of the bed. This is what we would expect from the evolution profile in Figure 18 because most of the error is

concentrated at the input where the second derivative is small but the first derivative is large. For these values of α and β, the discretisation error is significantly lower than the discretisation error for the fixed uniform mesh (equivalent to $\mu = 0$), for all the mesh resolutions considered.
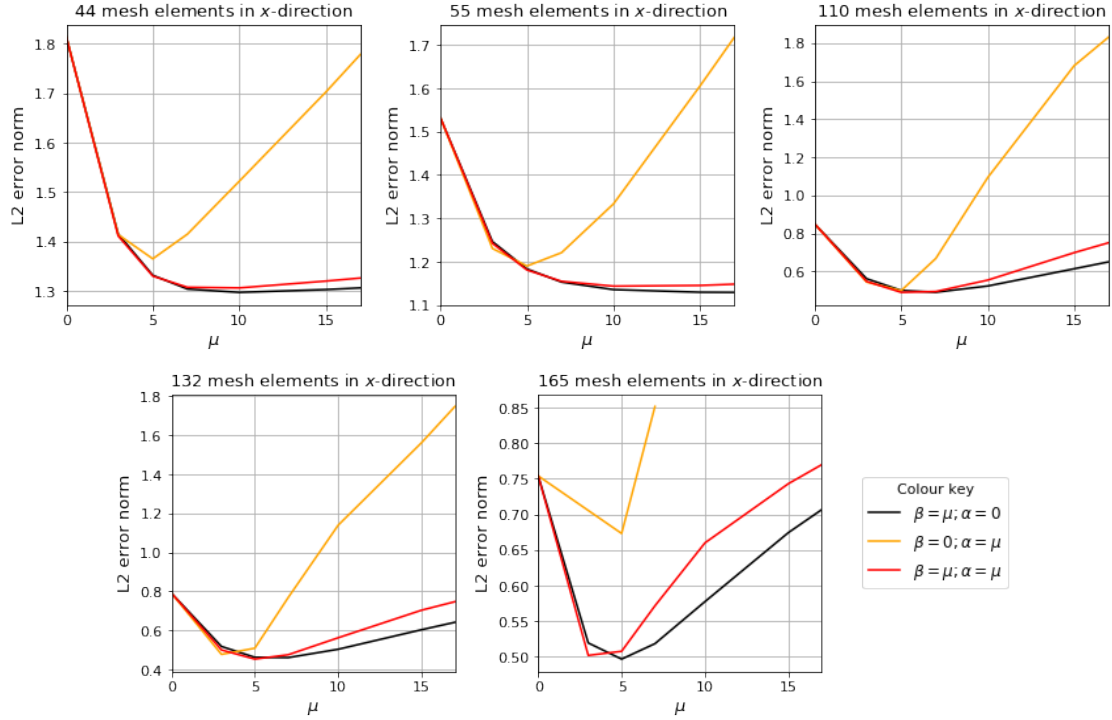

Figure 23: Discretisation error with mesh movement methods for varying α, β and number of mesh elements for whole domain in beach test case.

We next consider the error over the subdomain, $x > 70\,\mathrm{m}$. Figure 24 shows that the error is in general minimized by $\alpha = \mu, \beta = 0$, with $\mu \approx 7$ again seeming to correspond to the greatest minimization. The exception is when there are 55 mesh elements, where the error is minimized by $\alpha = 0$ and $\beta = \mu$. This may be because for this small number of elements, there is a particular structure that is better represented by the first derivative. Nevertheless, for $\alpha = 7$ and $\beta = 0$ and all mesh resolutions considered, the discretisation error using the mesh movement method is significantly lower than the discretisation error using the fixed mesh method. Setting $\alpha = \mu$ and $\beta = 0$ corresponds to a monitor function which is dependent only on the second-order derivatives of the bed. This is again what we would expect from the subdomain evolution profile in Figure 20, where there are large changes in the second-order derivatives in regions of interest. The difference between the optimum relationship between α and β for the whole domain case and the subdomain case shows it was correct to consider them separately.
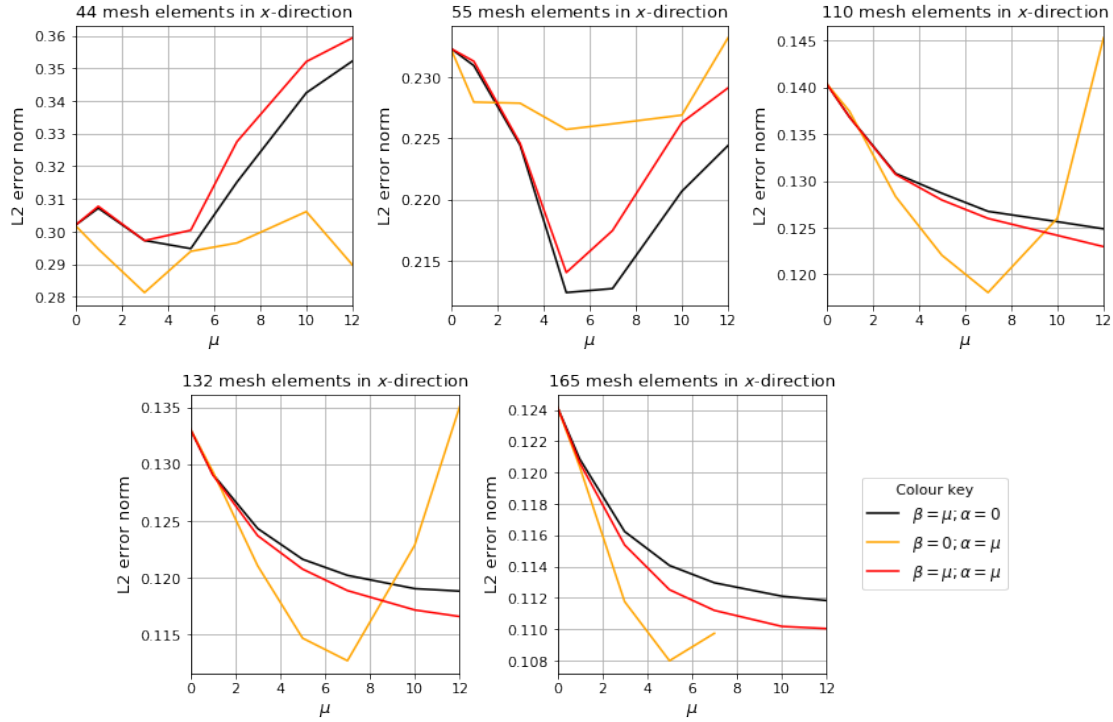
24

Figure 24: Discretisation error with mesh movement methods for varying α, β and number of mesh elements for subdomain $x > 70\,\mathrm{m}$ in beach test case.
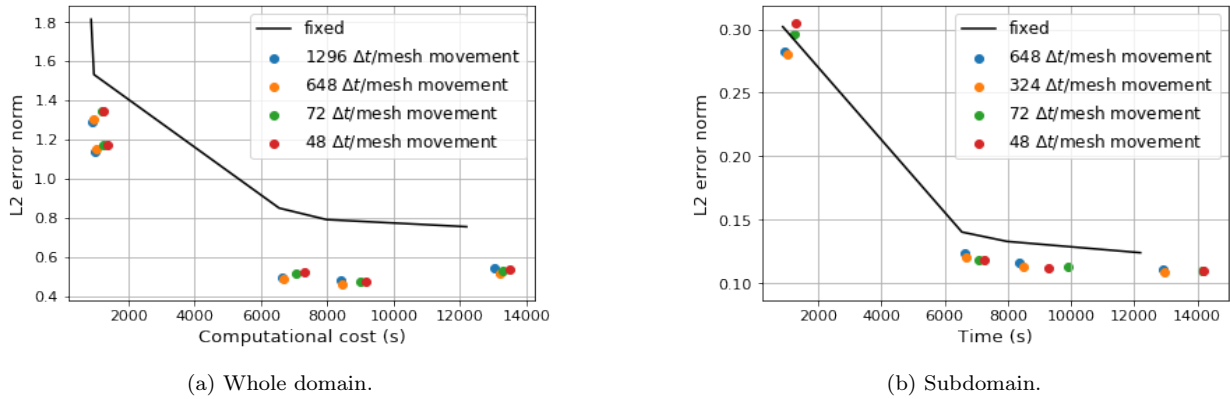

(a) Whole domain.

(b) Subdomain.

Figure 25: Computational cost vs error for fixed mesh and mesh movement methods with general parameters for the beach test case.

We have thus shown that mesh movement methods can accurately represent both large changes to the bed where the waves enter the domain, as well as smaller ripple-like structures. We have confirmed that - just as in the test case in Section 4.2 - a value of $\mu = 7$ is a good general parameter for optimising the discretisation error in both the subdomain and the whole domain. The relationship between α and β has again been shown to be physically intuitive, with $\alpha = 0, \beta = \mu$ in the whole domain and $\alpha = \mu, \beta = 0$ in the subdomain being shown to be the optimal relationships. Using these general parameters, in Figure 25 we consider the crucial relationship of simulation

25

accuracy versus computational cost. This shows that using moving mesh methods results in greater accuracy for the same computational cost as that of a fixed uniform mesh for all mesh movement frequencies considered (with a slight exception when the number of mesh elements is low in the subdomain where the accuracy is roughly the same). Figure 25 also shows that as the total cost of the simulation increases, the cost of each mesh movement increases, too. This is to be expected because the total cost here is a proxy for the number of mesh elements and increasing the number of mesh elements increases the cost of solving (4). Importantly, the proportion of the actual cost of the simulation stays roughly the same and thus using mesh movement methods with large numbers of mesh elements is still computationally efficient when compared to using fixed uniform meshes.

For this relatively complex wetting-and-drying test case, we have thus shown mesh movement methods not only improve accuracy for the same number of mesh elements but can also be used to reduce computational cost, even when using general parameters which have not been tuned for this specific test case.

5. Conclusion

In this work we have implemented a mesh movement scheme as part of a hydro-morphodynamic model for the first time. We have shown that these mesh movement methods can be used to improve accuracy in hydro-morphodynamic test cases with complex bathymetries and with moving wetting-and-drying interfaces. They can also be used to reduce the computational cost of solving the problem to a given accuracy. As part of the work on wetting-and-drying test cases, we have implemented a new conservative sediment transport model which is useful for dealing with coastal hydro-morphodynamic problems. In future work, we will consider using mesh movement methods to resolve the evolution of the wetting-and-drying interface and thus maximise the accuracy with which the faster hydrodynamic scale is simulated, rather than the slower morphodynamic scale which was the focus here through the construction of a monitor function based on the bed geometry alone.

For the mesh movement method considered, we present a monitor function for which the scaling factor that optimises the error reduction is fairly predictable from the physical characteristics of the test case. This will facilitate using this monitor function on further problems. In this work, we have used small scale sensitivity studies to obtain these optimum mesh parameters. In future work, we will seek to develop a gradient-based approach using the adjoint framework within Firedrake to allow us to more rigorously determine the optimum values for the scaling parameters. Section 4.1 showed that our hydro-morphodynamic model can be combined with the adjoint framework to determine an optimum diffusivity coefficient, and thus this is a promising line for follow-up research.

We will also consider mesh movement methods which couple the coordinate transformation with the solution of the prognostic equations, in an arbitrary Lagrangian Eulerian (ALE) type approach [33]. A major advantage of such an approach is that there is no longer any need to interpolate between meshes as solution fields and mesh coordinates are solved together.

26

Computer code availability

References

[1] L. O. Amoudry, A. J. Souza, Deterministic coastal morphological and sediment transport modeling: A review and discussion, Reviews of Geophysics 49 (2) (2011).

[2] A. T. N. Papanicolaou, M. Elhakeem, G. Krallis, S. Prakash, J. Edinger, Sediment transport modeling review - current and future developments, Journal of Hydraulic Engineering 134 (1) (2008) 1–14.

[3] D. L. Harris, A. Rovere, E. Casella, H. Power, R. Canavesio, A. Collin, A. Pomeroy, J. M. Webster, V. Parravicini, Coral reef structural complexity provides important coastal protection from waves under rising sea levels, Science Advances 4 (2) (2018).

[4] D. P. Callaghan, R. Ranasinghe, D. Roelvink, Probabilistic estimation of storm erosion using analytical, semi-empirical, and process based storm erosion models, Coastal Engineering 82 (2013) 64–75.

[5] C. Villaret, R. Kopmann, D. Wyncoll, J. Riehme, U. Merkel, U. Naumann, First-order uncertainty analysis using Algorithmic Differentiation of morphodynamic models, Computers & geosciences 90 (2016) 144–151.

[6] M. D. Piggott, G. J. Gorman, C. C. Pain, P. A. Allison, A. S. Candy, B. T. Martin, M. R. Wells, A new computational framework for multi-scale ocean modelling based on adapting unstructured meshes, International Journal for Numerical Methods in Fluids 56 (8) (2008) 1003–1015.

[7] M. D. Piggott, C. C. Pain, G. J. Gorman, D. P. Marshall, P. D. Killworth, Unstructured Adaptive Meshes for Ocean Modeling, in: M. W. Hecht, H. Hasumi (Eds.), Ocean Modeling in an Eddying Regime, American Geophysical Union, 2008, pp. 383–408.

[8] A. Avdis, A. S. Candy, J. Hill, S. C. Kramer, M. D. Piggott, Efficient unstructured mesh generation for marine renewable energy applications, Renewable Energy 116 (2018) 842–856.

[9] D. A. Mayne, A. S. Usmani, M. Crapper, An Adaptive Finite Element Solution for Cohesive Sediment Transport, in: Proceedings in Marine Science, Vol. 5, Elsevier, 2002, pp. 627–641. `doi:10.1016/S1568-2692(02)80044-4`.

[10] P. Delandmeter, Discontinuous Galerkin finite element modelling of geophysical and environmental flows, Ph.D. thesis, Universite catholique de Louvain (2017).

[11] F. Benkhaldoun, I. El Mahi, S. Sari, M. Seaid, An unstructured finite volume method for coupled models of suspended sediment and bed-load transport in shallow water flows, International journal for numerical methods in fluids 72 (9) (2013) 967–993. `doi:10.1002/d.3771`.

[12] E. Hanert, D. Y. Le Roux, V. Legat, E. Deleersnijder, Advection schemes for unstructured grid ocean modelling, Ocean Modelling 7 (1-2) (2004) 39–58.

[13] M. D. Piggott, P. E. Farrell, C. R. Wilson, G. J. Gorman, C. C. Pain, Anisotropic mesh adaptivity for multi-scale ocean modelling, Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences 367 (1907) (2009) 4591–4611. `arXiv:https://royalsocietypublishing.org/doi/pdf/10.1098/rsta.2009.0155`, `doi:10.1098/rsta.2009.0155`.

[14] M. D. Piggott, C. C. Pain, G. J. Gorman, P. W. Power, A. J. H. Goddard, h, r, and hr adaptivity with applications in numerical ocean modelling, Ocean Modelling 10 (1-2 SPEC. ISS.) (2006) 95–113. `doi:10.1016/j.ocemod.2004.07.007`.

[15] P. E. Farrell, M. D. Piggott, C. C. Pain, G. J. Gorman, C. R. Wilson, Conservative interpolation between unstructured meshes via supermesh construction, Computer Methods in Applied Mechanics and Engineering 198 (33) (2009) 2632 – 2642. `doi:https://doi.org/10.1016/j.cma.2009.03.004`.

[16] C. J. Budd, W. Huang, R. D. Russell, Adaptivity with moving grids, Acta Numerica (2009) 1–131`doi:10.1017/S0962492906400015`.

[17] J. M. Nunez Rattia, J. R. Percival, S. J. Neethling, M. D. Piggott, Modelling local scour near structures with combined mesh movement and mesh optimisation, Journal of Computational Physics 375 (2018) 1220–1237. `doi:10.1016/j.jcp.2018.09.026`.

[18] M. Grave, J. J. Camata, A. L. G. A. Coutinho, Residual-based variational multiscale 2D simulation of sediment transport with morphological changes, Computers and Fluids 196 (2020). `doi:10.1016/j.compfluid.2019.104312`.

[19] M. C. A. Clare, J. Percival, A. Angeloudis, C. J. Cotter, M. D. Piggott, Hydro-morphodynamics 2D modelling using a discontinuous Galerkin discretisation (2020). `doi:10.31223/OSF.IO/TPQVY`.

[20] T. Kärnä, S. C. Kramer, L. Mitchell, D. A. Ham, M. D. Piggott, A. M. Baptista, Thetis coastal ocean model: discontinuous galerkin discretization for the three-dimensional hydrostatic equations, arXiv preprint arXiv:1711.08552 (2017).

[21] 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) (2016) 1–27.

[22] T. M. McManus, J. R. Percival, B. A. Yeager, N. Barral, G. J. Gorman, M. D. Piggott, Moving mesh methods in Fluidity and Firedrake, Tech. Rep. July, Imperial College London (2017). `doi:10.13140/RG.2.2.27670.24648`.

[23] S. Balay, S. Abhyankar, M. F. Adams, J. Brown, P. Brune, K. Buschelman, L. Dalcin, A. Dener, V. Eijkhout, W. D. Gropp, D. Karpeyev, D. Kaushik, M. G. Knepley, D. A. May, L. C. McInnes, R. T. Mills, T. Munson, K. Rupp, P. Sanan, B. F. Smith, S. Zampini, H. Zhang, PETSc Web page, `https://www.mcs.anl.gov/petsc` (2019).

[24] S. Balay, S. Abhyankar, M. F. Adams, J. Brown, P. Brune, K. Buschelman, L. Dalcin, A. Dener, V. Eijkhout, W. D. Gropp, D. Karpeyev, D. Kaushik, M. G. Knepley, D. A. May, L. C. McInnes, R. T. Mills, T. Munson, K. Rupp, P. Sanan, B. F. Smith, S. Zampini, H. Zhang, H. Zhang, PETSc users manual, Tech. Rep. ANL-95/11 - Revision 3.13, Argonne National Laboratory (2020).

[25] 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.

[26] P. E. Farrell, D. A. Ham, S. W. Funke, M. E. Rognes, Automated derivation of the adjoint of high-level transient finite element programs, SIAM Journal on Scientific Computing 35 (4) (2013) C369–C393.

[27] Geology Cafe, Chapter 11 - rivers, streams, and water underground, `http://geologycafe.com/class/chapter11.html`, accessed on 2020-09-11 (2015).

[28] A. T. T. McRae, C. J. Cotter, C. J. Budd, Optimal-transport-based mesh adaptivity on the plane and sphere using finite elements, SIAM Journal on Scientific Computing 40 (2) (2018) 1121–1148. `doi:10.1137/16M1109515`.

[29] N. Barral, M. G. Knepley, M. Lange, M. D. Piggott, G. J. Gorman, Anisotropic mesh adaptation in Firedrake with PETSc DMPlex, arXiv preprint arXiv:1610.09874 (2016).

[30] J. G. Wallwork, N. Barral, D. A. Ham, M. D. Piggott, Anisotropic goal-oriented mesh adaptation in Firedrake, in: 28th Intl Meshing Roundtable, Zenodo, 2020. `doi:10.5281/zenodo.3653101`.

[31] M. Lange, L. Mitchell, M. G. Knepley, G. J. Gorman, Efficient mesh management in Firedrake using PETSc-DMPlex, SIAM Journal on Scientific Computing 38 (5) (2016) S143–S155. `arXiv:1506.07749, doi:10.1137/15M1026092`.

[32] J. G. Wallwork, N. Barral, S. C. Kramer, D. A. Ham, M. D. Piggott, Goal-oriented error estimation and mesh adaptation for shallow water modelling, Springer Nature Applied Sciences (to appear) (2020).

[33] J. Donea, A. Huerta, J.-P. Ponthot, A. Rodríguez-Ferran, Arbitrary l agrangian–e ulerian methods, Encyclopedia of Computational Mechanics Second Edition (2017) 1–23.

[34] C. Farhat, C. Degand, B. Koobus, M. Lesoinne, Torsional springs for two-dimensional dynamic unstructured fluid meshes, Computer methods in applied mechanics and engineering 163 (1-4) (1998) 231–245.

[35] W. Huang, R. D. Russell, Adaptive mesh movement—the mmpde approach and its applications, Journal of Computational and Applied Mathematics 128 (1-2) (2001) 383–398.

[36] C. J. Budd, J. Williams, Moving mesh generation using the parabolic Monge–Ampère equation, SIAM Journal on Scientific Computing 31 (5) (2009) 3438–3465.

[37] J. R. Maddison, I. Panourgias, P. E. Farrell, libsupermesh version 1.0. (2017).

[38] C. J. Budd, M. J. P. Cullen, E. J. Walsh, Monge–Ampére based moving mesh methods for numerical weather prediction, with applications to the eady problem, Journal of Computational Physics 236 (2013) 247–270.

[39] H. Weller, P. Browne, C. Budd, M. Cullen, Mesh adaptation on the sphere using optimal transport and the numerical solution of a Monge–Ampère type equation, Journal of Computational Physics 308 (2016) 102–123.

[40] G. Awanou, Quadratic mixed finite element approximations of the Monge–Ampère equation in 2d, Calcolo 52 (4) (2015) 503–518.

[41] O. Lakkis, T. Pryer, A finite element method for nonlinear elliptic problems, SIAM Journal on Scientific Computing 35 (4) (2013) A2025–A2045.

28

[42] T. Kärnä, B. de Brye, O. Gourgue, J. Lambrechts, R. Comblen, V. Legat, E. Deleersnijder, A fully implicit wetting–drying method for DG-FEM shallow water models, with an application to the Scheldt Estuary, Computer Methods in Applied Mechanics and Engineering 200 (5-8) (2011) 509–524.

[43] S. Funke, P. Farrell, M. Piggott, Reconstructing wave profiles from inundation data, Computer Methods in Applied Mechanics and Engineering 322 (2017) 167–186.

[44] A. Balzano, Evaluation of methods for numerical simulation of wetting and drying in shallow water flow models, Coastal Engineering 34 (1-2) (1998) 83–107. `doi:10.1016/S0378-3839(98)00015-5`.

[45] L. C. Van Rijn, Storm surge barrier Oosterschelde-computation of siltation in dredged trenches: Semi-empirical model for the flow in dredged trenches, Deltares, Delft, The Netherlands, 1980.

[46] M. C. A. Clare, S. C. Kramer, C. J. Cotter, M. D. Piggott, Calibration, inversion and sensitivity analysis for hydro-morphodynamic models through the application of adjoint methods, in preparation.

[47] F. Zhou, G. Chen, Y. Huang, J. Z. Yang, H. Feng, An adaptive moving finite volume scheme for modeling flood inundation over dry and complex topography, Water Resources Research 49 (4) (2013) 1914–1928. `doi:10.1002/wrcr.20179`.