Advances in Water Resources 164 (2022) 104189

Contents lists available at ScienceDirect

Advances in Water Resources

journal homepage: www.elsevier.com/locate/advwatres

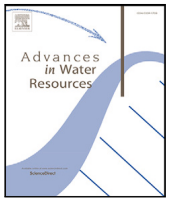

Modelling saline intrusion using dynamic mesh optimization with parallel processing

A. Hamzehloo [a,b,1], M.L. Bahlali [a,b], P. Salinas [b,c], C. Jacquemyn [b], C.C. Pain [c], A.P. Butler [a,b], M.D. Jackson [b,*]

[a] Department of Civil and Environmental Engineering, Imperial College London, UK
[b] Novel Reservoir Modelling and Simulation Group, Department of Earth Science and Engineering, Imperial College London, UK
[c] Applied Modelling and Computation Group, Department of Earth Science and Engineering, Imperial College London, UK

ARTICLE INFO

Keywords:
Coastal aquifers
Saline intrusion
Numerical modelling
Unstructured mesh
Dynamic mesh optimization
Computational cost

ABSTRACT

Saline intrusion (SI) in coastal aquifers is a global problem with the potential to contaminate groundwater used by over a billion people. Numerical modelling of SI in coastal aquifers is a key tool for risk assessment, aquifer management and resource regulation, but is extremely challenging because the mixing zone across the saline front is often very narrow, extending over metres or 10's metres, yet the saline front itself may extend laterally over a large (i.e. many km) three-dimensional (3D) domain. Moreover, the aquifer may be highly heterogeneous, further complicating the movement and geometry of the front. We test here the use of dynamic mesh optimization (DMO) in a parallel computational framework to simulate SI with higher accuracy and lower computational cost compared to fixed-mesh approaches. The framework uses a double control-volume-finite-element (DCVFE) method and is implemented in the open-source Imperial College Finite Element Reservoir SimulaTor (IC-FERST), but could be implemented in other FE-based simulators. We confirm accuracy and convergence using test cases based on the classic 'Henry' SI problem, demonstrating that solutions obtained using DMO require significantly fewer elements and therefore have much lower computational cost compared to equivalent fixed mesh solutions. We apply the framework to a realistic 3D case study simulating saline intrusion in a heterogeneous chalk aquifer, demonstrating simulation speed-up in excess of 120×. We suggest that parallelized DMO offers significant advantages over existing methods to simulate SI.

1. Introduction

Groundwater is the primary source of water for human consumption in most countries (Arfib and Charlier, 2016; Bakalowicz, 2015). Demand is particularly high in coastal regions (Custodio, 2010); over 1.2 billion people live within 100 km of a coastline and population densities in these areas are around three times the global average (Small and Nicholls, 2003). Increased groundwater abstraction in response to urbanization (Uddameri et al., 2014; Boretti and Rosa, 2019), combined with natural phenomena such as rising sea levels due to climate change (Masciopinto and Liso, 2016; Ketabchi et al., 2016) or storm surges (Cardenas et al., 2015; Klassen and Allen, 2017), can affect the balance between freshwater discharging to the sea and the inland movement of saline water (Ferguson and Gleeson, 2012). This imbalance typically results in saline intrusion (SI): a dynamic, inland movement of saline water that can lead to contamination of freshwater resources with major economic, health and social consequences (Arfib and Charlier, 2016; Morgan and Werner, 2015). SI can be particularly pronounced in fractured or highly permeable or heterogeneous aquifers (De Filippis et al., 2016), especially when subject to marked increases in groundwater abstraction, or sudden and substantial changes in sea level such as storm surges (Klassen and Allen, 2017). Such conditions can lead to highly dynamic changes in the sharp saline–freshwater interface (Graham et al., 2018), with important consequences for the management of coastal groundwater resources.

Monitoring SI using, for example, measurements of groundwater level and electrical conductivity (as a proxy for salinity) requires multiple boreholes, while geophysical monitoring requires surface and/or borehole equipment that can be logistically challenging to install and/or prohibitively expensive (MacAllister et al., 2018). Monitoring alone shows only the current extent of the intruding seawater and cannot predict future SI or its response to changes in abstraction or sea-level. Consequently, risk assessment, aquifer management and resource

* Corresponding author.
E-mail address: m.d.jackson@imperial.ac.uk (M.D. Jackson).
[1] Current address: Department of Aeronautics, Imperial College London.

https://doi.org/10.1016/j.advwatres.2022.104189
Received 13 October 2021; Received in revised form 6 March 2022; Accepted 30 March 2022
Available online 20 April 2022
0309-1708/© 2022 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

regulation, are typically based on the predictions of computational groundwater models that solve the coupled density-dependent flow and solute transport equations. Numerical simulation of SI, however, is extremely challenging, particularly where the mixing zone is very narrow (Huyakorn et al., 1996; Younes et al., 2009; Isa-Abadi et al., 2020) and dynamic (Graham et al., 2018). In such cases, the transition from saline to freshwater may occur over metres or 10's of metres, whereas the saline front and associated aquifer may extend inland and along the coast over several km. High grid or mesh resolution is essential to resolve accurately the location and movement of the saline front, but the combination of high resolution and large model domain is computationally challenging (Voss and Souza, 1987; Souza and Voss, 1987; Huyakorn et al., 1996; Roy and Datta, 2018).

One of the key constraints on accurately modelling dynamic saline intrusion is the use static or fixed meshes, whether structured (e.g. SEAWAT Langevin et al., 2007) or unstructured (e.g. FEMWATER (Lin et al., 1997); SUTRA (Voss and Provost, 1984); HydroGeoSphere (Therrien et al., 2010); PFLOTRAN (Lichtner et al., 2015); MIN3P-HPC (Su et al., 2021); COMSOL (Koohbor et al., 2019)). Although some of these codes make use of parallel processing that can speed up the computation process (e.g. HydroGeoSphere; PLOTRAN; MIN3p-HPC; COMSOL) many do not (Hwang et al., 2014). As a result, fixed mesh constraints are such that many SI simulations are simplified to two- rather than three-dimensions (3D). This can significantly reduce their practical applicability and predictive capabilities for tackling real world problems (Kerrou and Renard, 2010; Rajabi and Ataie-Ashtiani, 2014; Graham et al., 2018; Koohbora et al., 2019).

Here we test the use of dynamic mesh optimization (DMO) in conjunction with parallel computing to increase the speed and accuracy of groundwater flow and transport modelling with specific application to SI. DMO allows the geometry and resolution of an unstructured mesh to change during a simulation, typically to place higher, and often anisotropic, mesh resolution in parts of the domain where there are steep gradients in the solution, and lower resolution elsewhere. DMO allows solutions with high accuracy to be obtained at lower computational cost compared to using a fixed mesh and has been widely applied in other areas of CFD (Alauzet and Loseille, 2016; Xie et al., 2014; Hu et al., 2018). There are, however, very few applications of DMO in porous media flow modelling, because extension to heterogeneous porous media is non-trivial given the requirement to include and preserve the spatially complex material properties associated with geologic heterogeneity (Jackson et al., 2014, 2015; Jacquemyn et al., 2019; Osman et al., 2021).

The DMO algorithm reported here allows elements within an unstructured mesh to be split or amalgamated (*h*-adaptivity), or element vertices to be moved (*r*-adaptivity). It differs significantly from the related concept of adaptive grid refinement (AGR), which has previously been applied to model reservoir flows, although not to the specific problem of SI. AGR utilizes only *h*-adaptivity and Cartesian grids: grid cells are split by adding cell boundaries (Schmidt and Jacobs, 1988; Dahle et al., 1990, 1992; Hornung and Trangenstein, 1997; Trangenstein, 2002; Wolff et al., 2013) (Fig. 1). Early applications of *h*-adaptivity in reservoir modelling neglected geologic heterogeneity (Schmidt and Jacobs, 1988; Dahle et al., 1990, 1992). Later studies included heterogeneity but started with a fine-scale geologic model as input. Grid refinement in such models is straightforward as daughter cells simply inherit the hydrological properties of the parent cell. However, the minimum grid resolution remains fixed by the initial fine-scale model, which means the computational cost can only increase as new cells are added. Coarsening the initial grid to reduce cost requires the hydrological properties such as porosity and permeability to be upscaled onto the coarser grid (Hornung and Trangenstein, 1997; Trangenstein, 2002; Wolff et al., 2013). This upscaling of permeability can impose a significant computational cost (Trangenstein, 2002; Gerritsen and Lambers, 2008; Wolff et al., 2013).

In an attempt to circumvent this problem, some studies upscaled the fine-scale hydrological properties to a range of pre-selected coarse grid cell sizes prior to simulating flow (Trangenstein, 2002; Wolff et al., 2013). The pre-calculated values were then assigned to the adaptively sized coarse cells during simulation. However, the flexibility of *h*-adaptivity is compromised by the requirement to pre-select a range of cell sizes. Moreover, extension of such approaches to the more complex mesh geometries required to represent many geologic heterogeneities is non-trivial: no general methods for upscaling permeability on arbitrary meshes have yet been developed. Finally, most published studies deploying AGR for reservoir flows consider 2D models only (Trangenstein, 2002; Gerritsen and Lambers, 2008; Wolff et al., 2013). The approach adopted here allows DMO to be applied in 3D numerical simulations of flow and transport in complex geologic porous media, and is computationally efficient because it does not require the hydrological properties to be up-, cross- or downscaled each time the mesh is optimized, or to be calculated *a-priori*.

Efficient application of DMO for aquifer and reservoir modelling is facilitated by use of surface-based geologic modelling (SBGM) (Jackson et al., 2014; Jacquemyn et al., 2019; Salinas et al., 2021). SBGM is an approach whereby all geological heterogeneity, whether it is structural, stratigraphic, sedimentological or diagenetic in origin, is represented by surfaces that enclose discrete rock volumes termed geological domains (Jackson et al., 2014). Hydrogeological properties are homogeneous within these geological domains; models can therefore be created without reference to an underlying grid or mesh (Jackson et al., 2014). Previous work has shown that the correlated heterogeneity captured by these domains, and by similar concepts such as lithofacies modelling, has the most significant impact on flow; the element-to-element or cell-to-cell scale variability in hydrogeological properties often produced by point or cell-based reservoir modelling approaches can be neglected (Willis and White, 2000; Feyen and Caers, 2006; Bianchi and Zheng, 2016; Osman et al., 2021).

An initial mesh is created only when required for numerical calculations; this mesh is as coarse as possible whilst preserving the geologic surface architecture within a target error. The initial mesh defines 'master nodes' that represent the surfaces and cannot be modified by DMO. SBGM is therefore compatible with DMO because surfaces are preserved during mesh adaptivity through the use of the master nodes; moreover, because petrophysical properties are defined with domains, there is no need to up-, cross- or downscale hydrogeological properties such as permeability onto the new mesh after it has been modified by DMO (Jackson et al., 2015). New elements simply inherit the hydrogeological properties of the parent domain.

To quantify the benefits of DMO in conjunction with parallel computing for SI modelling, we extend an existing framework for porous media flow with DMO to include, for the first time, salt transport and associated density-driven flows, and demonstrate DMO for heterogeneous porous media flow in a parallel computational framework. The framework is implemented in the open-source Imperial College Finite Element Reservoir SimulaTor (IC-FERST; see http://multifluids.github.io/) but the approach we report could be implemented in other codes using the control-volume finite-element method (CVFEM). Previous work using our framework, without the parallel implementation reported here, includes simulation of two-phase flow in petroleum reservoirs (Jackson et al., 2015) and during viscous fingering (Kampitsis et al., 2020), the related problem of heat transport and associated density-driven flow in geothermal reservoirs (Salinas et al., 2021) and solid–fluid coupling for flow in fractures (Obeysekara et al., 2018). There are no previous applications of DMO to the coupled density-dependent flow and solute transport problem of interest here.

We begin by outlining the framework and its extension to solute transport and parallel DMO. We then demonstrate accuracy and convergence of the extended framework against benchmark solutions based on the classic 'Henry' SI problem (Henry, 1964). We finish by demonstrating practical application of the framework, including parallel scaling

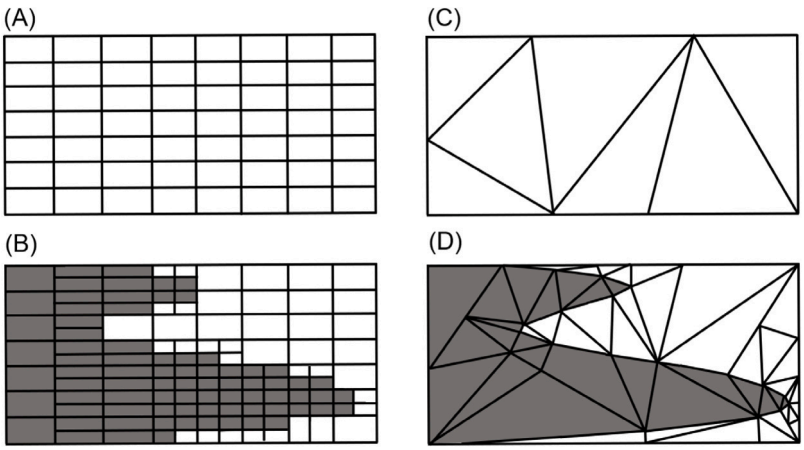

Fig. 1. Adaptive grid refinement (AGR) employing *h*-adaptivity to represent a saline front in (A) and (B), compared to the dynamic mesh optimization (DMO) approach used here, that employs *hr*-adaptivity to represent a saline front in (C) and (D). (A) and (C) show the initial mesh or grid; (B) and (D) show the grid or mesh adapted or optimized to the saline front shown as grey shading. AGR adds new cells (refining using *h*-adaptivity) to capture the front; DMO moves existing nodes and adds new nodes (*hr*-adaptivity) to capture the front.

performance, to a realistic 3D case study simulating saline intrusion in a heterogeneous chalk aquifer modelled using SBGM. The main aim of the work is to test the hypothesis that DMO in a parallel framework can yield significant computational speedup when simulating SI in realistic 3D models. The novelty lies partly in the extension of an existing computational framework for porous media flow and DMO to include salt transport in a parallel framework, but primarily in the first application of DMO to the problem of modelling SI in heterogeneous 3D aquifers.

2. Methodology

2.1. Governing equations

For a density-dependent flow in porous media, the water flux is given by Darcy's equation, which for single phase flow, where depth is positive, is written as

$$\mathbf{q} = \frac{-\mathbf{k}}{\mu}(\nabla P - \rho \mathbf{g}) \tag{1}$$

Combining this with the conservation of mass (i.e. continuity) equation gives

$$\frac{\partial(\phi\rho)}{\partial t} + \nabla \cdot (\phi\rho\mathbf{v}) = 0 \tag{2}$$

where $\mathbf{q}$ is the Darcy velocity vector [m/s], defined as $\mathbf{q} = \phi\mathbf{v}$ in which ϕ is the effective porosity [-] and $\mathbf{v}$ is the mean pore water velocity vector [m/s]. $\mathbf{k}$ denotes the intrinsic permeability tensor of the porous medium [m^2]. μ and ρ denote dynamic viscosity [Pa s] and fluid density [kg/m^3], respectively, P is fluid pressure [Pa] and $\mathbf{g}$ is the gravitational field tensor [m/s^2].

The effect of fluid and porous medium compressibility can be combined into a single parameter, the specific storage coefficient S_0 [m^{-1}], giving

$$\frac{S_0}{|\mathbf{g}|}\frac{\partial P}{\partial t} + \nabla \cdot (\phi\rho\mathbf{v}) = 0 \tag{3}$$

For saline intrusion modelling, a transport equation for the salt mass is also required

$$\frac{\partial(\phi C)}{\partial t} + \nabla \cdot (\phi\mathbf{v}C) = \nabla \cdot [\phi(D_m\mathbf{I} + \mathbf{D}) \cdot \nabla C] \tag{4}$$

where C is the solute (contaminant) concentration [kg/m^3], $\mathbf{I}$ is the identity tensor [-], D_m is the molecular diffusion coefficient [m^2/s] and $\mathbf{D}$ is the mechanical dispersion tensor [m^2/s].

The typical solute concentration of seawater (C_s) is around 35 kg/m^3 (Croucher and O'sullivan, 1995), which gives a fluid density of 1025 kg/m^3. A linear relationship is assumed between density and concentration (Croucher and O'sullivan, 1995; Fahs et al., 2016),

$$\rho = \rho_0(1 + \beta C) \tag{5}$$

where ρ_0 [1000 kg/m^3] is the corresponding reference density at $C = 0$ and $\beta = 0.7143 \times 10^{-3}$ [m^3/kg]. This means Eqs. (3) and (4) are coupled and must be solved simultaneously. We express results in terms of solute mass fraction relative to seawater (c), defined as:

$$c = \frac{C}{C_s} \tag{6}$$

An isotropic representation is assumed to approximate the mechanical dispersion component of the salt transport. Thus, $\mathbf{D}$ is a symmetric tensor defined as

$$\mathbf{D} = \begin{bmatrix} D_{xx} & D_{xy} & D_{xz} \\ D_{yx} & D_{yy} & D_{yz} \\ D_{zx} & D_{zy} & D_{zz} \end{bmatrix} \tag{7}$$

where $D_{xy} = D_{yx}$, $D_{xz} = D_{zx}$ and $D_{yz} = D_{zy}$. The diagonal components in Eq. (7) are defined as

$$D_{xx} = \frac{1}{|v|^2}(d_l v_x^2 + d_t v_y^2 + d_t v_z^2) \tag{8}$$

$$D_{yy} = \frac{1}{|v|^2}(d_t v_x^2 + d_l v_y^2 + d_t v_z^2) \tag{9}$$

$$D_{zz} = \frac{1}{|v|^2}(d_t v_x^2 + d_t v_y^2 + d_l v_z^2) \tag{10}$$

where $|v|$ denotes the magnitude of the velocity vector and v_x, v_y and v_z represent the magnitude of the velocity components in the x, y and z directions, respectively. d_l and d_t are longitudinal and transverse dispersion coefficients [m^2/s], respectively. The off-diagonal elements in (7) are

$$D_{ij} = \frac{d_l - d_t}{|v|^2}(v_i v_j) \tag{11}$$

where $i \neq j$ and $i, j = x, y, z$. The dispersion coefficients are dependent on the absolute magnitude of the local velocity and are defined as

$$d_l = \alpha_l |v| \tag{12}$$

$$d_t = \alpha_t |v| \tag{13}$$

where α_l and α_t are the longitudinal and transverse dispersivities [m] of the porous medium, respectively.

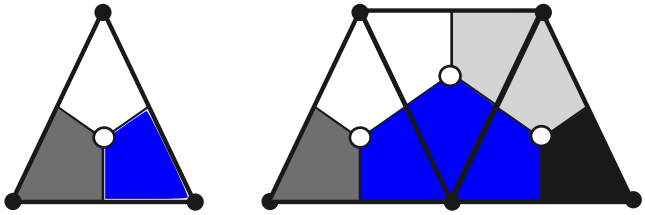

Fig. 2. $P_0(DG)P_1(CV)$ element pair in 2D. Finite-elements are represented by triangles; control-volumes are formed around nodes (solid circles) located at the element corners by connecting the barycentres of the elements. Velocity is evaluated at the centre of each finite-element (open circles); concentration and, unusually, pressure are represented in the control-volumes formed around the nodes.

2.2. Discretization

The double control-volume finite-element method (DCVFEM) used here is discussed in detail elsewhere (Salinas et al., 2017b) so we provide only a brief summary. The method is applied on fully unstructured triangular and tetrahedral (2D and 3D) meshes. Velocity is discretized conventionally using discontinuous Galerkin (DG) finite elements (FE) in the element mesh but, unlike conventional CVFEM, pressure is discretized using control volumes (CVs) in the dual mesh created by connecting the barycentres of the elements (Fig. 2). The DCVFEM was specifically developed to improve convergence rate and success in models with poor quality meshes (typically containing elements with large internal angles) that are easily created in models of heterogeneous aquifers containing high aspect ratio geological features. Extending the approach to solute transport, concentration is discretized using CVs. Rock porosity and permeability are piecewise constant within FEs, while fluid density and viscosity are piecewise constant within CVs.

Here, we use the element pair $P_0(DG)P_1(CV)$ in which velocity is discontinuous and is represented element-wise with a discretization order of 0 (constant within elements) while pressure and concentration are continuous and are represented CV-wise with a discretization order of 1 (linear variation within elements; Fig. 2). Space and time are discretized using higher order methods to increase solution accuracy. Stabilization methods are implemented to ensure monotonic solutions when dealing with high Courant–Friedrichs–Lewy or Peclet numbers.

The flux at the interface between two control volumes (CV_i and CV_j) sharing the same edge is computed as

$$\tilde{\mathbf{v}} = \tilde{\underline{\underline{\sigma}}}^{-1}\left[\frac{1}{2}\left(\underline{\underline{\sigma}}_i \mathbf{v}_i + \underline{\underline{\sigma}}_j \mathbf{v}_j\right)\right], \tag{14}$$

where $\underline{\underline{\sigma}} = \mu\mathbf{k}^{-1}$ and the tilde symbol ($\tilde{}$) denotes the value at the interface of i and j. To compute a high-order approximation of the flux at the interface, the limited value of the concentration at the interface $\tilde{C}$ is first computed using central differences and limited using a normalized variable diagram (NVD) scheme following the upwind direction. Then $\tilde{C}$ is used to compute $\tilde{\underline{\underline{\sigma}}}$ using a second-order Taylor expansion series

$$\tilde{\underline{\underline{\sigma}}} = \underline{\underline{\sigma}} + \left(\tilde{C} - C_k\right)\left(\frac{\partial \underline{\underline{\sigma}}}{\partial C}\right)_k + \frac{1}{2}\left(\tilde{C} - C_k\right)^2\left[\frac{\left(\frac{\partial \underline{\underline{\sigma}}}{\partial C}\right)_k - \left(\frac{\partial \underline{\underline{\sigma}}}{\partial C}\right)_l}{\left(C_k - C_l\right)}\right], \tag{15}$$

where k and l swap the control-volumes they represent (either i or j depending on the velocity direction). Thus, if $\mathbf{n}\cdot\mathbf{v} > 0$ then $k = i$, $l = j$; conversely, if $\mathbf{n}\cdot\mathbf{v} < 0$ then $k = j$, $l = i$. Finally, $\tilde{\mathbf{v}}\cdot\mathbf{n}$ is bounded by $\mathbf{v}_i\cdot\mathbf{n}$ and $\mathbf{v}_j\cdot\mathbf{n}$. The same flux is used for both control-volumes sharing the interface, ensuring local conservation of mass.

Time is discretized using an implicit Θ-method, where Θ varies between 0.5 (Crank-Nicholson) and 1 (implicit Euler). Θ is controlled by a total variation diminishing (TVD) scheme to ensure that the time integration is stable regardless of the time-step size (Gomes et al., 2017).

2.3. Numerical solution

A Picard iterative method (Anderson, 1965) is used to solve the coupled system of discretized equations as described by Salinas et al. 2017a. The Picard solver iterates the coupled solutions for pressure and concentration until the concentration field has converged within a user-defined tolerance between two consecutive non-linear iterations:

1. Pressure is calculated by combining Eqs. (1) and (3) considering the solute concentration to be fixed.
2. An updated velocity field is obtained from Darcy's Law, Eq. (1).
3. From the updated pressure and velocity fields, an updated solute concentration is computed using Eq. (4).
4. Density is recomputed with the updated solute concentration using Eq. (5).

When this is done, time advances by one time-level until the final time is reached. In the examples shown here, the convergence criteria is the infinite norm of the concentration field is < 0.03, with a maximum of 15 non-linear iterations per time-level for the Henry problem test cases, and 50 non-linear iterations for the SI example case.

The system of linear equations to be solved at each step (pressure or concentration) is solved using the PETSc framework (Balay et al., 1997). A Generalized Minimal Residual (GMRes) method, with 30 restart iterations and a multigrid-based preconditioner (Hypre; Falgout, 2005), is used to reduce the initial residual by ten orders of magnitude. Use of this solver, in combination with the DCVFE formulation, ensures that the system can be solved regardless of the time-step size and mesh quality.

To reduce computational time, a Proportional–Integral–Derivative (PID) adaptive method is used to control the time-step size depending on the number of non-linear iterations and the error in the non-linear solver (Akakpo and Gildin, 2017)

$$dt_{n+1} = \left(\frac{1}{E_n}\right)^{K_i}\left(\frac{E_{n-1}}{E_n}\right)^{K_p}\left(\frac{E_{n-1}^2}{E_n E_{n-2}}\right)^{K_d}, \tag{16}$$

where n denotes the time-level, K_i, K_p and K_d are constants controlling the behaviour of the PID controller (with values 1.34, 0.001 and 0.01, respectively, taken from Akakpo & Gilden, 2017). The PID controller minimizes the overall error function E, defined as:

$$E = \frac{E_e + E_i}{2} + max(E_e, E_i), \tag{17}$$

where E_i is the ratio between the target number of non-linear iterations and that required by the non-linear solver to converge for the time level n, and E_e is the ratio of the target error and the error of the non-linear solver (E_r), which is defined as the infinite norm of the difference between two non-linear iterations:

$$E_r = \|C^m - C^{m-1}\|_{inf}, \tag{18}$$

in which m is the non-linear iteration number and C is the concentration field. The PID controller adjusts the time-step size to minimize E, thus minimizing the discrepancy between the target error and the target number of non-linear iterations. Once a new time-step is obtained, its increase or decrease is limited to less than x1.5 or x2.0, respectively, the previous time-step.

2.4. Dynamic mesh optimization

An advanced anisotropic DMO approach is used which allows mesh elements to be split or amalgamated, element edges swapped, or element vertices moved (Pain et al., 2001). The DMO procedure starts by calculating the L_2 norm of the local interpolation error introduced by the FE representation, using the weighted norm of the Hessian matrix ($\mathbf{H}$) of the specified field(s). Any CV-based solution fields are interpolated into the FE mesh using a conservative Galerkin projection (Farrell

et al., 2009; Adam et al., 2016). The quality of the mesh (I) is defined as

$$I = \sum_{i \in edges} [e_i^T \mathbf{E} e_i - 1]^2, \tag{19}$$

$$E_{ij} = \det(\mathbf{H})^{(\frac{1}{2\gamma\delta})} \frac{|H_{ij}|}{\epsilon}, \tag{20}$$

in which $\mathbf{E}$ is a matrix composed of entries E_{ij}, e_i denotes the edges connecting the vertices of the mesh, ϵ is a user-defined adaptivity tolerance to the chosen solution field(s), γ is the polynomial degree of the norm used and δ is the problem dimension (Salinas et al., 2018).

DMO applies the mesh operations sequentially to create a new mesh with an approximately uniform quality while honouring a set of user-defined criteria, such as maximum number of elements, maximum and minimum edge length, mesh gradation and mesh anisotropy. The master nodes defining the architecture of the SBGM remain fixed. Once the new mesh is generated, the solution fields from the old mesh are interpolated to the new mesh using the conservative Galerkin projection (Adam et al., 2016). Here, DMO adapts the mesh to the salt concentration field for the Henry problem test cases, and to the salt concentration and pressure fields for the SI example cases. DMO adapts the mesh to minimize error metrics defined for both fields in this latter case, maintaining solution quality for both but at the cost of using a larger number of elements if the fields are only loosely coupled such that high resolution is required in some parts of the domain for one field but not the other and *vice-versa*.

2.5. Parallel implementation

Parallelization during DMO is based on the well known distributed memory paradigm using the message passing interface (MPI) library. Halos containing duplicate nodes at the interface between two processor domains are generated to ensure a consistent communication. Parallel implementation of DMO requires additional steps compared to a fixed mesh. To balance the computational cost across different processors, dynamic domain decomposition is required in which balanced domains are created each time the mesh is adapted (Gorman et al., 2012). The approach involves the following steps:

1. The given mesh is partitioned into sub-domains using a parallel graph partitioner. In this case, ParMetis is used (Schloegel et al., 2002). The partitioning is performed with the objective of minimizing memory exchange between processors.
2. The mesh is adapted within each sub-domain while locking the nodes in the halos at the boundaries with the neighbouring sub-domains (Fig. 3A). Locking the boundary nodes yields a globally conforming mesh. The resulting mesh, however, is sub-optimal because the locked nodes have not been optimized by DMO.
3. The mesh is repartitioned to create new sub-domains (Fig. 3B) with different locked boundary nodes, so previously locked nodes can be optimized.
4. The process is repeated a number of times, typically three, to ensure that all nodes are optimized.

Applying DMO in parallel computations and heterogeneous porous media introduces additional constraints on the DMO algorithm by requiring the master nodes and the halo nodes to remain fixed. This can sometimes mean that the DMO algorithm fails to create a new mesh with a higher quality than the old mesh. If this happens, the old mesh is re-used for the next timestep and DMO is re-applied to the new solution.

3. Method evaluation

We begin by verifying that the numerical framework produces accurate results for salt transport and SI modelling by comparing results against theoretical and numerical studies. The Henry problem (Henry, 1964) considers the steady-state location of a dispersive saline wedge

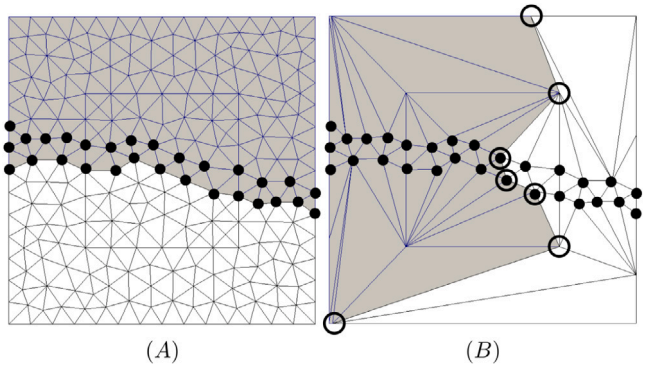

Fig. 3. Dynamic domain decomposition and DMO. Shaded and white areas denote two parallel sub-domains. (A) Partitioned mesh before DMO. Locked nodes in the halos between the two parallel sub-domains are shown as filled black circles. (B) Mesh after DMO and repartitioning. Locked nodes in halos between the new sub-domains are shown as open black circles. Previously locked nodes can now be optimized by DMO. Repeated application of this process allows all nodes to be optimized and yields parallel sub-domains on the optimized mesh.

in a confined aquifer. Henry reduced the saline intrusion problem to a 2D rectangular geometry with sealed (no flow or transport) upper and lower boundaries (Fig. 4A; for model properties, see Table 1). Hydrostatic pressure along with a Dirichlet condition for transport ($c = 1$) is imposed at the right (seawater) boundary, while freshwater discharge with salt mass fraction of zero ($c = 0$) is imposed on the left (inland) boundary. Initially, the domain is filled with freshwater. This configuration leads to a 'saline intrusion wedge' penetrating towards the freshwater boundary along the lower boundary of the aquifer (e.g. Fig. 4).

Henry provided a semi-analytical solution involving a double set of Fourier coefficients, which needed to be truncated to obtain a solution (Zidane et al., 2012). Due to numerical issues with this method, an additional numerical solution reported by Held et al. (2005) and the more recent semi-analytical solution of Fahs et al. (2016) for velocity-dependent dispersion are used here for comparison. We examine and verify the performance of the method using the original Henry problem (Case 1, with no mechanical dispersion) and five variants with different values of diffusion, longitudinal and transverse dispersion, permeability and permeability anisotropy (Table 1).

A minimum element edge length of 8.0 mm is used for all test cases, compared to the domain dimensions of 1 × 2 m (Fig. 4). Where applied, DMO imposes an absolute error tolerance of $\epsilon = 0.001$ for the salt mass fraction. The dynamic time stepping has an initial time step of $\Delta t = 0.01$ s and a maximum time step of $\Delta t = 300$ s. The time step quickly reached its maximum value with relatively limited fluctuation around this value for all test cases.

3.1. Comparison against benchmark solutions

Isolines of salt mass fraction ($c = 0.1 - 0.9$) for all test cases (Table 1) simulated with and without DMO show a good match with the published benchmark results (Fig. 4B - G). For the original Henry problem (Case 1), a relatively diffuse interface with a wide mixing zone (L_s as defined in Fig. 4B) is formed between the freshwater and saltwater, consistent with earlier observations (Croucher and O'sullivan, 1995; Held et al., 2005). Increasing the aquifer permeability in Case 2 results in a deeper intrusion of saline water (Fig. 4C), because the saltwater can move further towards the freshwater side driven by the fixed hydrostatic seawater boundary pressure (Held et al., 2005). Consequently, Case 2 has a longer wedge toe length (L_{toe} as defined in Fig. 4B) compared to the original Henry case 1 (Fig. 5A).

Reducing the diffusion coefficient but adding velocity-dependent dispersion in Cases 3 and 4 reduces the width of the mixing zone while

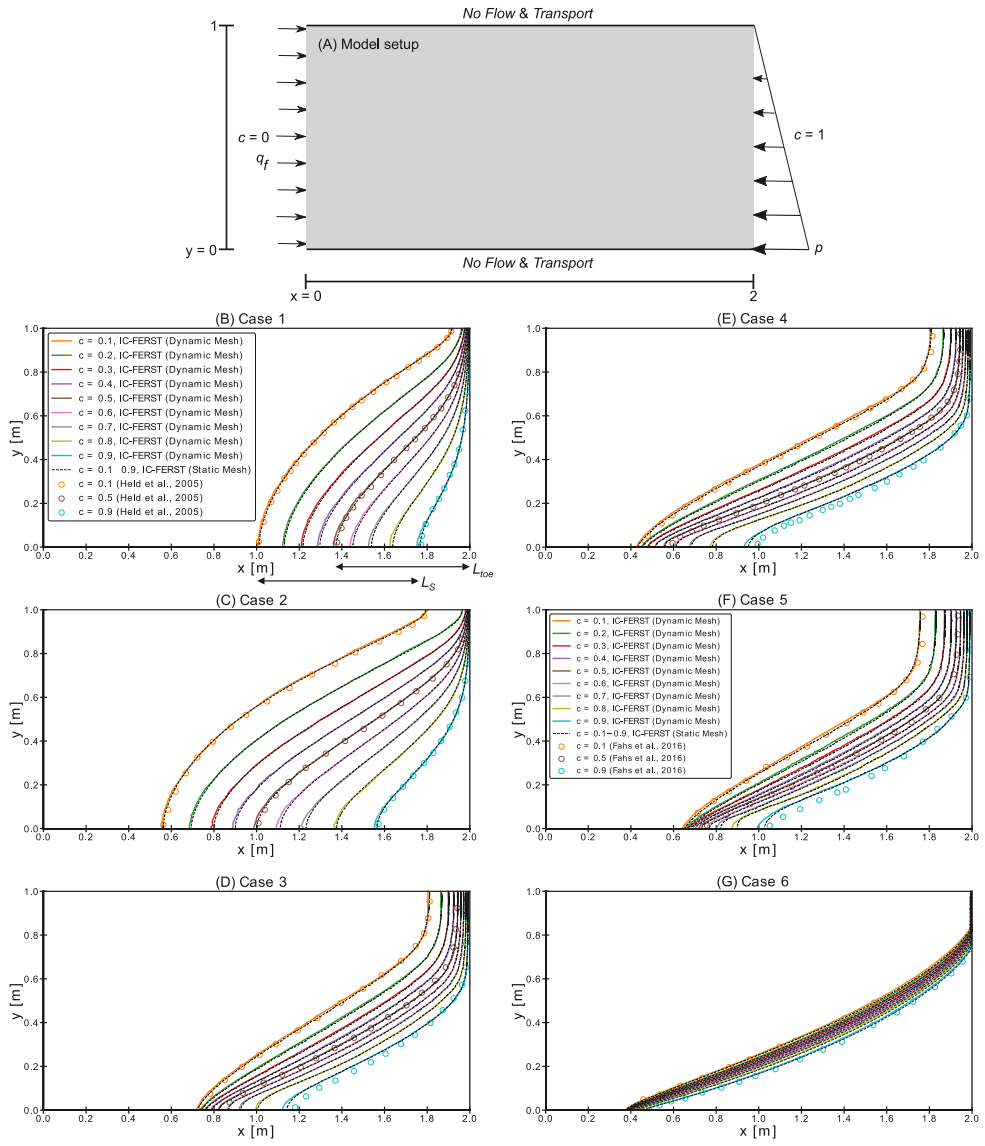

Fig. 4. Method verification using the Henry problem. (A) shows a schematic of the computational setup. (B) - (G) show solutions of all test cases plotting isolines of steady-state salt concentration compared against a computational study (Held et al., 2005) and a semi-analytical solution (Fahs et al., 2016). Case 1 (B) is the original Henry problem. Case 2 (C) has a higher permeability; Cases 3 (D) and 4 (E) have significantly lower diffusion coefficient but include velocity-dependent dispersion. Case 4 also has an anisotropic permeability. Cases 5 (F) and 6 (G) highlight the effect of dispersion (see Table 1 for model properties). The saltwater wedge toe, L_{toe} and the mixing zone, L_S, are defined as the distances along the x-axis between the sea boundary and $c = 0.5$ isoline and between the $c = 0.1$ and $c = 0.9$ isolines, respectively, as illustrated in (B).

increasing the wedge toe length (Fig. 5A). Case 4 has a longer wedge toe length compared to Case 3 because the former has anisotropic permeability. Fig. 4 shows our approach captures the salt concentration in cases with velocity-dependent dispersion and anisotropic permeability with good accuracy, with and without DMO, compared to the numerical predictions of Held et al. (2005).

Reducing further the diffusion coefficient in Case 5 and then reducing the longitudinal and transverse dispersivities by two orders of magnitude in Case 6 yields an increase in the wedge toe length (Fig. 5a) and a narrower mixing zone, which becomes extremely narrow for the lowest dispersivities tested in Case 6 (Fig. 4G). These observations are in very close agreement with those of Fahs et al. (2016). In addition, the parameter values chosen for Case 6 are comparable with and, for longitudinal dispersivity the same as, those used by Robinson et al. (2015) to simulate observations of saline intrusion obtained from a benchtop-scale experimental representation of the Henry problem.

The salt concentration (Fig. 4) and wedge toe length results (Fig. 5a) demonstrate that simulations with DMO produce comparable results to those obtained using a fixed mesh and reported in the literature. However, DMO requires significantly fewer elements so the computational cost is much lower (Fig. 5b). All test cases with DMO start from an extremely coarse initial grid with only 44 elements, because there is no geological heterogeneity that must be captured (Fig. 6). As the saline front moves into the domain, the mesh refines to capture the front, and then moves with the front, coarsening behind it where high resolution is no longer required.

Fig. 6 shows the time evolution of Case 6, which is of particular interest since its relatively narrow ('sharp') mixing zone is comparable to those typically observed in real-life saline intrusion problems, where $L_s \ll$ the scale of the aquifer domain. As a result, the elements in the mixing zone become highly anisotropic, to capture the interface geometry and associated steep gradients in salt concentration. The DCVFEM formulation used here ensures very good numerical accuracy and avoids spurious oscillations despite the presence of these highly anisotropic elements.

Table 1
Parameters for the Henry problem test cases.

	Case 1	Case 2	Case 3	Case 4	Case 5	Case 6
Molecular diffusion coefficient [m^2/s]	1.88571×10^{-5}			5×10^{-7}	9.43×10^{-8}	
Longitudinal disp. [m]	0			0.075	0.1	0.001
Transverse disp. [m]	0			0.015	0.01	0.0001
Permeability $\times 10^{-9}$ [m^2]	1.01936	1.67539	1.01958	k_{xx}= 1.29119 k_{yy}= 0.79272	1.0204	
Porosity [m^3/m^3]	0.35					
Freshwater density [kg/m^3]	1.0×10^{-3}					
Saltwater density [kg/m^3]	1.025×10^{-3}					

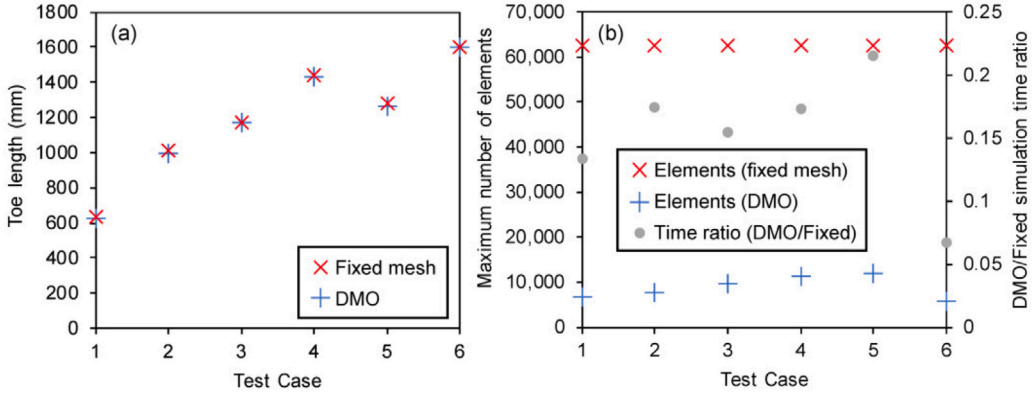

Fig. 5. (a) SI toe length for each of the test cases, for both fixed and DMO mesh cases. (b) Maximum number of elements, and ratio of DMO to fixed mesh simulation times, for each of the test cases.

3.2. Convergence and speedup

The comparisons between our simulation results and the benchmark solutions shown in the previous section were qualitative. Here we provide a quantitative convergence test for the most challenging test (Case 6) with the sharpest saline front. All simulations here were conducted for a duration of 3600 s with a fixed Courant number of 0.9 to allow direct comparison between cases with different spatial resolution. The solution error is measured using the Relative L2 Norm R_N, defined as

$$R_N = \sqrt{\frac{\sum_{n=1}^{n_p}(c_0(n) - c(n))^2}{\sum_{n=1}^{n_p} c_0(n)^2}} \tag{21}$$

where n_p denotes the number of test points and '0' denotes the reference value of the salt concentration at a given test point, obtained from a very fine, fixed mesh simulation using triangular elements with an edge length of 4 mm.

Fixed (i.e. static) mesh cases show approximately linear convergence (reduction in error with increasing resolution), similar to that found previously for the DCVFEM applied to other problems (Salinas et al. 2017) (Fig. 7A). DMO, however, shows much more rapid convergence and an equivalent simulation error using a much lower number of elements than required with a fixed mesh. The use of fewer elements with DMO yields a reduction in simulation time, with speed-up of order 5× for the finest mesh (Fig. 7B). The offset between the ratio of cell numbers and simulation times in Fig. 7B represents the cost of DMO, and is of order 10 percent of the cost of the equivalent fixed mesh case. The cost of DMO is significantly outweighed by its benefit, even though we generate a new mesh and new matrices. It is clear that DMO is more effective in reducing the simulation time when the domain contains a larger number of elements, which suggests DMO will yield even better speed-up for 3D cases.

4. Application: saline intrusion modelling

This section demonstrates the value of DMO for modelling SI in large (kilometre) scale groundwater models with realistic boundary conditions and heterogeneous geology. The test case is based on the UK Chalk aquifer, an important groundwater resource for southern and eastern England, including London. The area of interest is located on the south coast of the UK, where the Chalk provides groundwater to the city of Brighton and surrounding area. The data used in the modelling were obtained from the Saltdean Observation Borehole (OBH), which is located approximately 1.8 km from the coast. A map of the Saltdean study area and schematic of the Saltdean OBH can be found in Graham et al. (2018).

The Saltdean OBH is now used only for observation. The borehole extends to around 60 m below ground level and intersects a series of chalk, marl and hardground layers within the Cretaceous Seaford and Lewes Nodular Chalks (Fig. 8(a)). The strata dip towards the coast at an angle of approximately 5 ° (Graham et al., 2018). The hydraulic conductivity in the aquifer generally decreases with depth. Recharge estimates given by Graham et al. (2018) for the Saltdean area are about 490 mm/year.

Here, we make use of two years' monitoring data acquired by MacAllister et al. (2018). Vertical profiling of fluid electrical conductivity in the borehole over several tidal cycles during a period of SI showed upwards and downwards movement of a sharp saline interface, with saline water interpreted to be entering and leaving the borehole via a high permeability zone at its base (MacAllister et al., 2018). This high permeability feature most likely reflects the presence of a fracture horizon (Jones and Robins, 1999; MacAllister et al., 2018; Graham et al., 2018).

Measured self-potential (SP) in the borehole could only be matched by models with a very sharp saline interface and, overall, the borehole data are consistent with the presence of a sharp saline interface in the vicinity of the borehole (MacAllister et al., 2018; Graham et al., 2018). Graham et al. (2018) conducted numerical simulations using a widely applied code but the inability to simulate at high resolution in a large domain meant that they were not able to produce a saline–freshwater wedge that was sharp enough to match the data. The interface had to be artificially sharpened in a post-processing step; essentially, the numerical model results were adjusted to give the desired concentration field in the vicinity of the interface (Graham et al., 2018).

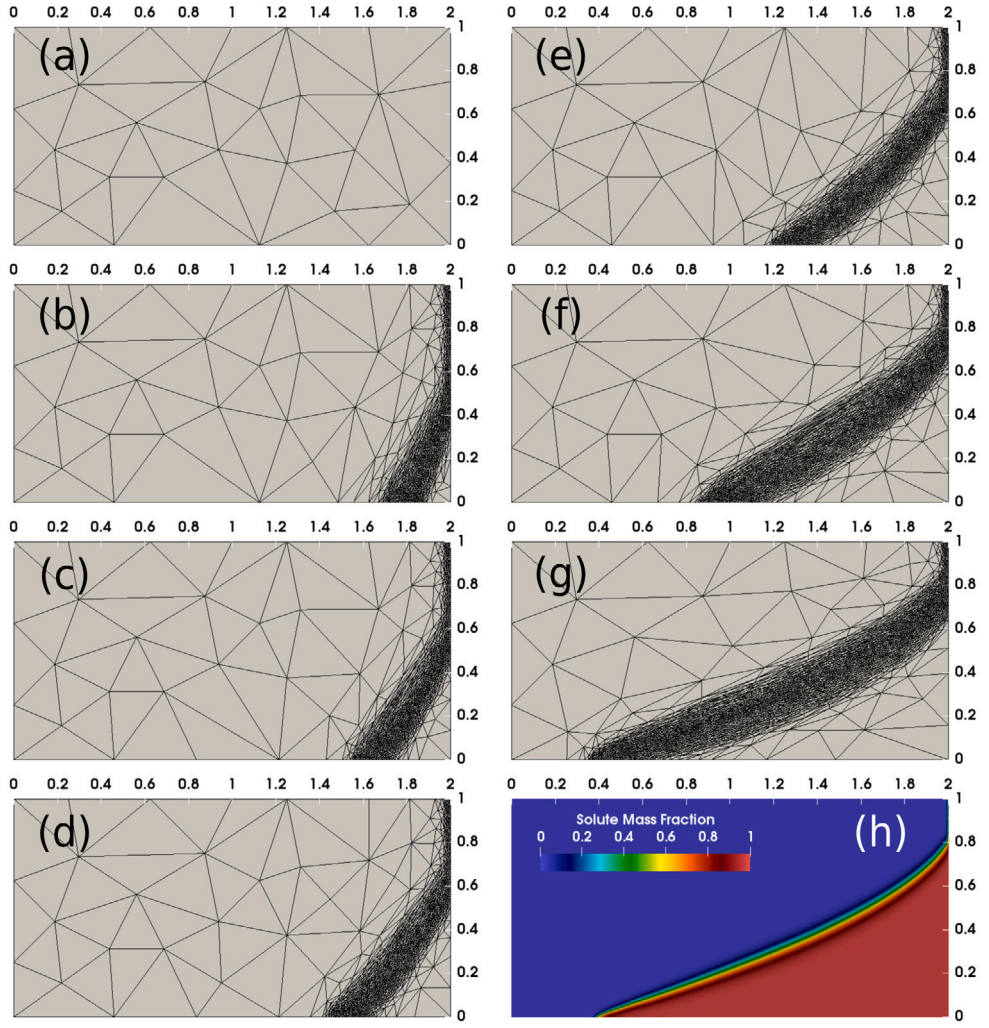

Fig. 6. Snapshots of the time evolution of the DMO mesh of Test Case 6 from (a) t=0 s to (g) t=20,000 s. Plot (h) shows the corresponding salt mass fraction at t=20,000 s.

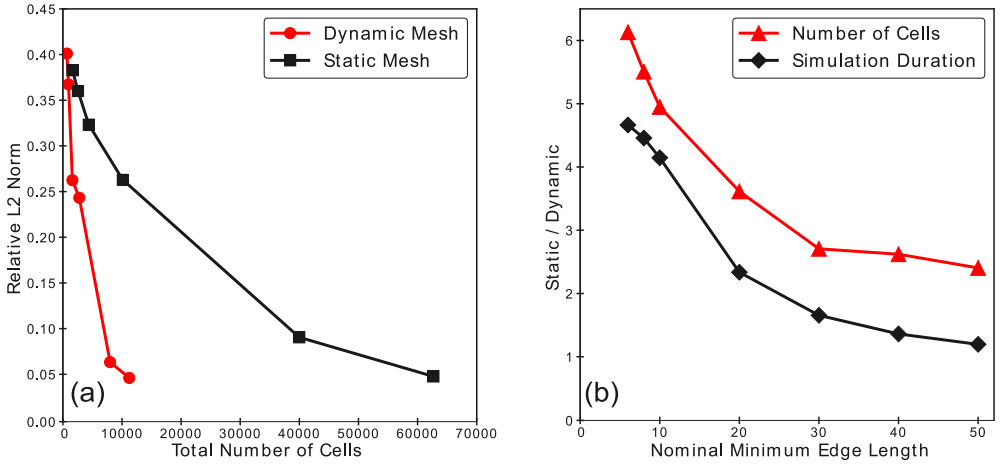

Fig. 7. Metrics for mesh adaptivity in Test Case 6. (a) Relative L2 Norm as a function of number of elements for static (i.e. fixed) mesh and DMO simulations with varying mesh resolution as defined by the minimum edge length. (b) Corresponding ratio of element numbers and simulation times between fixed and DMO cases with the same minimum edge length. Timings are for serial runs on 1 CPU and a simulation (model time) duration of 3600 s with a fixed Courant number of 0.9.

Table 2
Geological and hydrodynamic model parameters for the Saltdean SI model. Modified from Graham et al. (2018).

Parameter	Seaford Chalk	Lewes Chalk, marls & hardgrounds
Porosity	0.39	0.39
Permeability (m^2)	$k_{xx} = 3 \times 10^{-10}$ $k_{yy} = 3 \times 10^{-10}$ $k_{zz} = 1 \times 10^{-11}$	$k_{xx} = 3 \times 10^{-11}$ $k_{yy} = 3 \times 10^{-11}$ $k_{zz} = 1 \times 10^{-12}$
Longitudinal disp. (m)	0.05	0.05
Transverse disp. (m)	0.005	0.005
Fluid viscosity (Pa s)	0.001	
Molecular diffusion coefficient (m^2/s)	5×10^{-10}	
Freshwater recharge (m/s)	1.553×10^{-8}	
Freshwater density (kg/m^3)	1000	
Saltwater density (kg/m^3)	1025	

Such an approach has little value for prediction or management of the aquifer response to SI.

We consider two models of the Saltdean OBH and associated saline interface, based on the model presented by Graham et al. (2018). The first is a pseudo 2D model in the x (horizontal) and z (vertical) directions, as there is only one element (of 20 m) in the y (orthogonal horizontal) direction and no variation of any parameter or variable in that direction. This model is small enough that DMO can be benchmarked against fixed mesh cases. The second is a fully 3D model and is computationally tractable only through use of DMO and/or parallel computation.

A simple implementation of SBGM was used to create the models. The chalk, marl and hardground layers penetrated by the Saltdean OBH are interpreted to be laterally extensive and the model captures their interpreted seaward dip (Jones and Robins, 1999; MacAllister et al., 2018; Graham et al., 2018) (Fig. 8(a)). There are known to be other heterogeneous layers present in the stratigraphy above and below the interval penetrated by the Saltdean OBH, but these were not included by Graham et al. (2018) because the model was designed to focus on the movement of the saline front in the vicinity of the Saltdean OBH.

Surfaces were created using the SBGM approach of Jacquemyn et al. (2019) to represent the boundaries between the modelled layers. Petrophysical properties, based on core and well data, were assigned to each layer (Table 2). The permeability tensor in each geological layer is anisotropic with vertical permeability 30× lower than horizontal permeability (MacDonald and Allen, 2001; Butler et al., 2009). The upper boundary condition comprises freshwater recharge with a salt mass fraction of zero ($c = 0$) imposed along the top right of the domain (dotted line in Fig. 8a) while the top left boundary (continuous line) is a hydrostatic pressure distribution of the overlying seawater with a Dirichlet condition for transport ($c = 1$). The left (seaward), right (inland) and bottom boundaries are zero flux for both flow and transport. Initially, the model is filled with seawater.

4.1. Pseudo 2D model of saline intrusion

The effect of fixed and DMO mesh resolution was tested in the pseudo 2D model using three different minimum edge lengths (1.5 m, 1 m and 0.5 m). The initial mesh used in the DMO cases contains 1 378 elements and defines the master nodes that represent the geologic surfaces. DMO was applied every two time steps with an absolute tolerance of $\epsilon = 0.001$ for salt concentration and a relative tolerance of $\epsilon = 50\,000$ Pa for pressure. Dynamic time stepping was used with no minimum or maximum time step enforced; typical values were approximately 45 min and 200 days, respectively. The model was run for a simulation (model) time of 1 000 years; pseudo equilibrium was obtained after approximately 350 years.

Fig. 8b shows salt mass fraction isolines for both fixed and DMO meshes, simulated using a minimum edge length of $z_{min} = 0.5$ m. The mixing zone (here defined as the distance between $c = 0.1$ and $c = 0.9$ normal to the slope of $c = 0.5$) has a width of just 4 m, consistent in both the fixed and DMO simulations, confirming the challenge of resolving such a sharp saline front in a large (km-scale) domain. A slight change in the width of the mixing zone occurs at the boundary between the Seaford Chalk and the Lewes Nodular Chalk, consistent with the change in permeability across these units (Table 2). The results with and without DMO are in close agreement. The accuracy of the match is demonstrated in Fig. 8c, which shows the variation of the mean proportional absolute error E between fixed and DMO simulation results versus the minimum edge length. E is defined as

$$E = \frac{1}{n_p} \sum_{n=1}^{n_p} \frac{|c_{0n} - c_n|}{c_{0n}} \tag{22}$$

where n_p denotes the number of test locations and the '0' subscript denotes the reference value of the salt concentration. The reference values are taken from the finest static grid simulation with $z_{min} = 0.5$ m. The saline wedge is captured by DMO with high accuracy when compared to the high resolution fixed grid, as the maximum error (obtained for $z_{min} = 1.5$ m) remains below 3%. This accuracy is achieved using significantly fewer elements compared to the equivalent fixed mesh, yielding a simulation speed-up in excess of 11× for the most accurate (finest mesh) case (Fig. 8d).

4.2. 3D model of saline intrusion

Here we test the use of DMO in 3D modelling of saline intrusion. The pseudo-2D model is extruded into the 3rd dimension so the model width becomes 50 m. We also explicitly model the Saltdean OBH and associated adit (see schematic of the Saltdean OBH in Graham et al. (2018)). The borehole domain is modelled using a porosity of 1 and a permeability of 7.92×10^{-6} m^2, following Graham et al. (2018).

Only DMO simulations were conducted for this 3D case because, even with DMO, the number of elements required was approximately 260 000. A fixed mesh with the resolution required to deliver comparable solution accuracy contains over 9 000 000 elements; numerical simulation of this case would require several 100s of cores and a simulation time of order weeks. In common with many groups modelling saline intrusion, we did not have access in this research to the high performance computing facilities required to run such a large, long time varying 3D case on a fixed mesh. The value of DMO in facilitating numerical simulation of such a case is clear. All model settings were the same as in the pseudo-2D case, except DMO was triggered every 10 time steps and the minimum time step was $t = 86\,400$ s (1 day). The minimum edge length was $z_{min} = 1$ m in the vertical direction, in the middle of the range tested in the pseudo-2D model. The initial mesh contained approximately 51 000 elements. The simulation was run in parallel on 30 cores, thus demonstrating the use of DMO in conjunction with parallel processing. We discuss the parallel scaling performance in the next section.

Fig. 9 shows the time evolution of salt concentration along with the adapting mesh during the early stages of wedge growth ($t = 0$ s

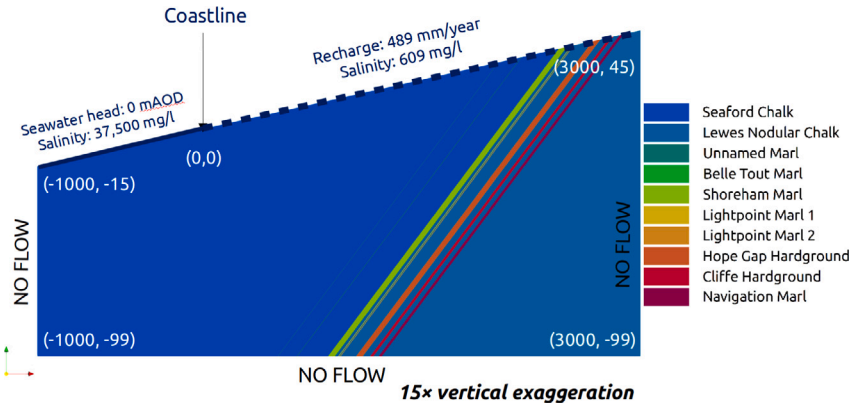

(a)

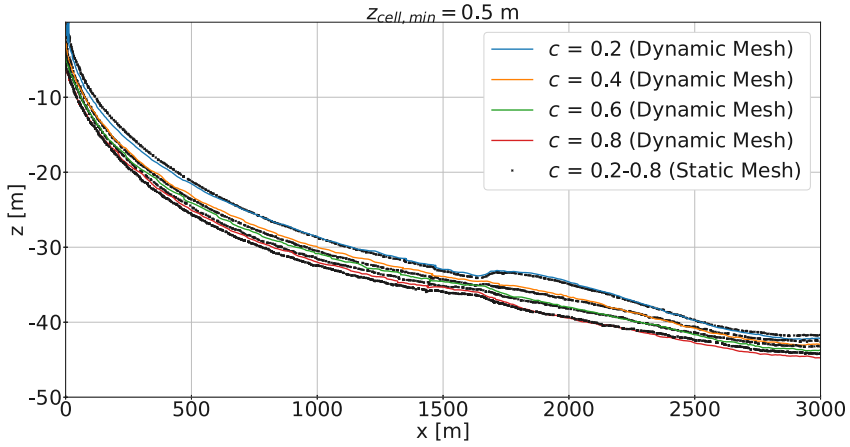

(b)

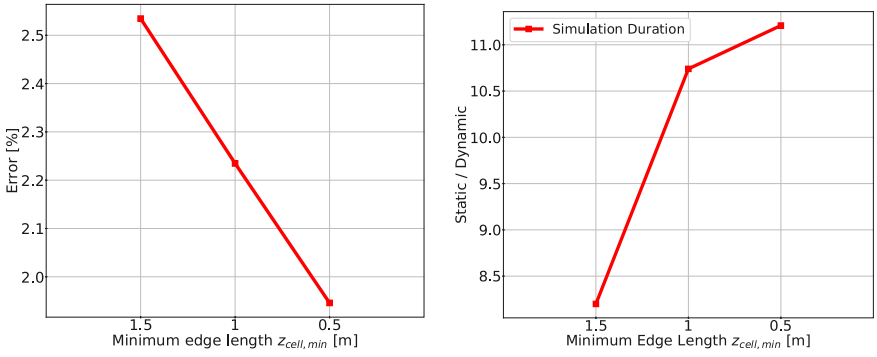

(c) (d)

Fig. 8. (a) Schematic of the pseudo 2D Saltdean model, including boundary conditions. The inclined layers represent different marl and hardground layers. Numbers in brackets denote coordinates in metres: the seaward-to-inland extent of the model is 4 km, and the vertical extent is 84 m at the seaward end and 144 m at the landward end. A 15× vertical exaggeration is used on the vertical scale. (b) Simulation results showing concentration isolines at t = 350 years, for both fixed and DMO cases with a minimum edge length of z_{min} = 0.5 m. (c) Error (as defined in Eq. (22)) between DMO simulations and the finest fixed mesh simulation as a function of the minimum edge length used in the DMO simulations. (d) Corresponding ratio of simulation times between fixed and DMO mesh cases with the same minimum edge length. Timings are for serial runs on 1 CPU.

to t = 40 years). DMO captures the steep concentration gradient at the moving saline–freshwater interface, but maintains a coarse mesh elsewhere. The narrow and highly dynamic saline front has a width of ca. 5 m consistent with the value obtained in the pseudo-2D model for the same minimum edge length. The front is, therefore, well captured. We discuss the computational cost in the next section.

4.3. Parallel performance

The parallel computing capability of the computational framework was tested using the same 3D SI model presented in the previous section (Fig. 10). Purely to compare the scaling on a consistent number of elements, the model was discretized using 2.5M elements for both fixed

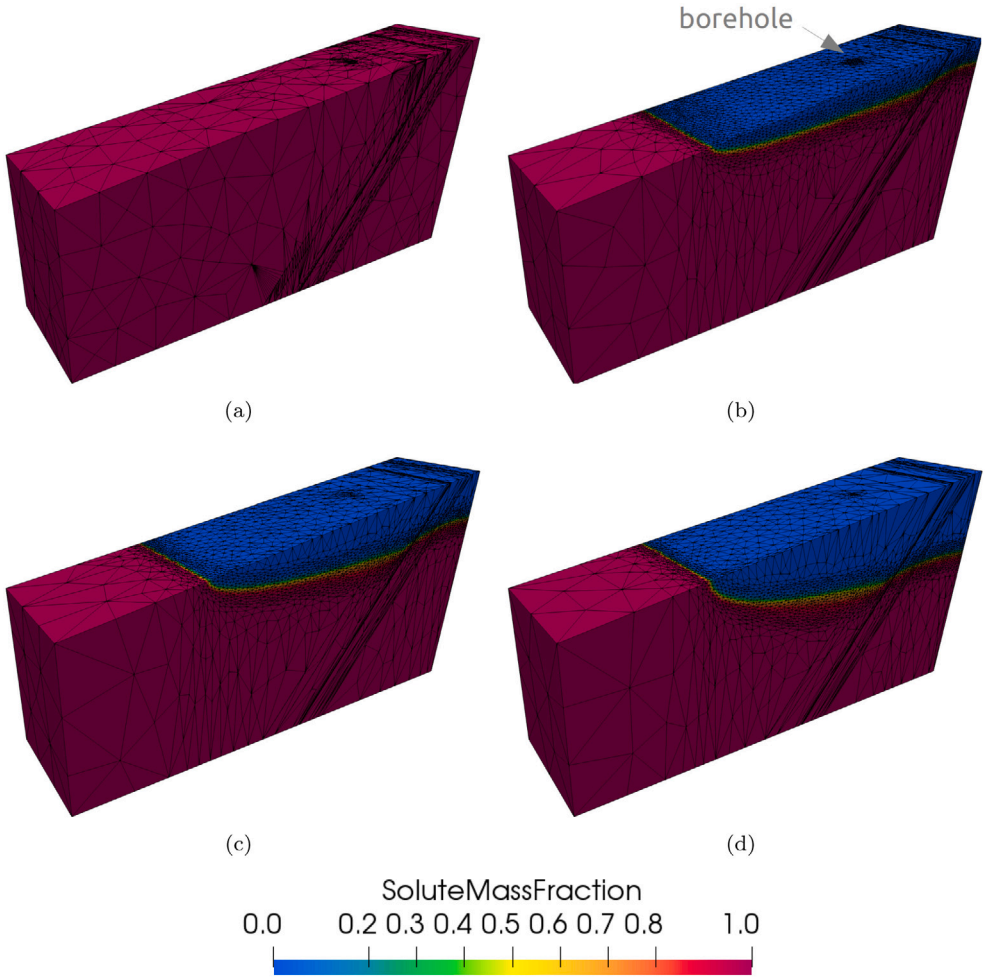

Fig. 9. (a)–(d) Time evolution of the dynamic grid of 3D SI simulation from $t = 0$ s to $t = 40$ years. Small and highly anisotropic elements are constructed by DMO at the vicinity of the freshwater–saltwater interface as well as near the borehole. The vertical exaggeration is 15×.

and DMO cases. The scaling was tested against a reference case with 2 cores (CPU) to account for the added cost of distributed memory. The number of elements per core is always larger than the minimum required for good parallel performance, which is of order 15k. The parallel performance was tested by executing 10 time-steps with a fixed time step of 1 month, with DMO applied every two timesteps. All other settings were the same as in the previous section. The strong-scaling was tested, with the number of cores increasing from two to 30. The scaling performance was evaluated using:

$$\eta = \frac{w_{ref} \times N_{ref}}{w_n \times N_n} \tag{23}$$

where w_{ref} and N_{ref} are the walltime of the reference simulation and the reference number of cores used (here $N_{ref} = 2$); likewise, w_n denotes the walltime using N_n cores.

The scaling in both fixed and DMO cases is very similar and shows good performance, falling to 0.7 at 30 cores (i.e. the speed-up is x0.7 of the ideal (perfect) speed-up). The results confirm that application of DMO does not compromise the efficiency of parallel scaling. In these scaling results, DMO is used to change the mesh, but the number of elements remains fixed to ensure a fair comparison with the fixed mesh results. In "real-world" simulations, the DMO performance varies more because the number of elements varies, so the parallel partitioning may become more challenging as the mesh changes owing to changes in the size of the halos, and the element size varies which impacts on the dynamic time-stepping. We tested the "real-world" scaling of the

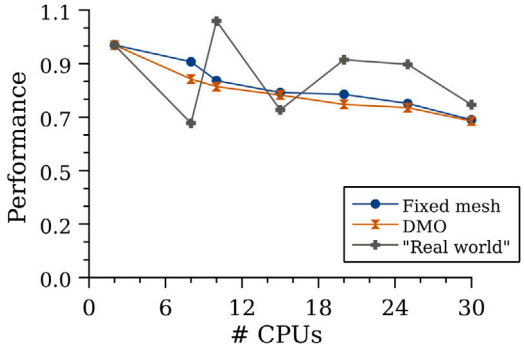

Fig. 10. Strong scaling results for the 3D Saltdean model for a fixed mesh and DMO, both with a fixed number (2.5M) of elements, and the "real world" scaling using DMO in which the number of elements varies. The ideal strong scaling performance is a fixed value of 1. Non-ideal behaviour is always observed.

Saltdean model using DMO and allowing the number of elements to vary over a simulation period of ten years. As expected, the scaling is less stable (oscillates more) than the cases with a fixed number of elements, but the overall scaling performance is comparable, showing that DMO in parallel yields significant speed-up compared to using a fixed mesh in serial.

5. Discussion

We demonstrate here practical use of dynamic mesh optimization, in a parallel computational framework, to reduce simulation cost while maintaining solution accuracy in modelling saline intrusion. DMO allows key solution features such as the saline front to be resolved at high resolution, in a large 3D model domain, at much lower computational cost than fixed mesh simulations with comparable accuracy. We measure the computational advantage of DMO by comparing against equivalent fixed mesh solutions. Here we could do this only in 2D models. In the 3D example shown, we could not simulate the equivalent fixed mesh case at sufficient resolution because of the high computational cost. We can, however, estimate the computational advantage of simulating this 3D model in parallel and with DMO, versus a single core and a fixed mesh with the same resolution.

We obtained a 11.5× speed-up in the 3D model through use of multiple cores and DMO (Fig. 10). In the equivalent 2D model, we obtained 10.75× speed-up using DMO compared to an equivalent fixed mesh (Fig. 9d). It is reasonable to assume that use of DMO yields a similar speed-up in the 3D model compared to the equivalent fixed-mesh; indeed, the advantage of DMO is typically greater in 3D, as the volume of the model domain (and hence the total number of fixed-mesh elements required) scales as L^3 (where L is the model size), but the area of the saline front where DMO is required scales as L^2. Combining these two speed-ups suggests that DMO and parallelization yielded a reduction in computational cost in excess of 120× for the 3D saline intrusion problem tested here, compared to a fixed mesh case with equivalent resolution simulated on a single core, as is typical in current SI intrusion models. This clearly demonstrates the significant computational opportunities that the parallelized adaptive meshing framework presented here can bring to the field of groundwater modelling.

We tested the parallel scaling of DMO only to 30 cores. We are aware that high performance computing (HPC) allows access to many more cores, and the parallel scaling of numerical codes that use fixed or adaptive meshes has been tested to many more cores (Hammond et al., 2014; Wei et al., 2015), although we are not aware of any studies that have tested the parallel scaling of numerical methods and codes used for SI. We restricted the scaling to 30 cores because, as we show here, DMO for SI modelling provides similar accuracy to fixed mesh codes using of order 10× fewer elements. A DMO model with sufficient elements to merit parallel simulation on several thousand cores would require several tens of millions of elements, and yield similar accuracy to a fixed mesh simulation with several hundreds of millions of elements. We do not anticipate a demand for such large models of SI. Rather, we argue that DMO enables practitioners to simulate heterogeneous, 3D SI problems on a desktop multi-core machine, which would otherwise require hundreds of millions of fixed elements on a HPC facility.

We have demonstrated DMO in parallel using a 3D model of a heterogeneous chalk aquifer containing laterally extensive marls and hardgrounds that are just a few metres in thickness. The layers have very high aspect ratio and contrasting permeability; hence, this is a challenging example case. We have also demonstrated benefits of DMO for flow problems (although not SI) in other aquifer types, using SBGM to model geological features such as inclined, intersecting faults, growth faults, folded strata and complex stratigraphy, and complex intersecting fractures (Jackson et al. 2015; Jacquemyn et al. 2019; Salinas et al. 2021). In the models presented here, smaller scale variability within the layers is neglected, but this is not essential to SBGM and DMO: it is possible to include, for example, finer-scale stratigraphy in SBGMs (Jacquemyn et al. 2019). DMO works efficiently in SBGMs that neglect the cell-to-cell scale variability often observed in pixel- or grid-based models, but this has been shown to have a negligible impact on flow (Osman et al. 2021). It is capturing and preserving the spatially correlated heterogeneity and connectivity of geological flow zones and flow barriers that is key to predict accurately flow and transport. Note that this does not mean that SBGMs must be deterministic; surfaces representing key correlated heterogeneity can be created stochastically (Jacquemyn et al. 2019) and conditioned to borehole data (Titus et al., 2021).

We see significant value in DMO for modelling SI in reducing simulation cost such that computational resources can be used to run many models or cases. Simulations of many models or cases are often required to explore and quantify the impact of uncertain aquifer architecture and hydrological properties on SI, to optimize well placement, completion depth and operation in aquifers at risk of SI, and for model calibration against measured data (Kerrou et al., 2013; Rajabi and Ataie-Ashtiani, 2014; Rajabi et al., 2015; Koohbora et al., 2019; Yang et al., 2021). DMO can also ensure sufficient resolution is maintained in models irrespective of the dynamics of the saline front: when the aquifer architecture or properties change in a sensitivity analysis or optimization process, the saline front geometry or thickness may vary such that a fixed mesh resolution previously identified to have sufficient resolution may no longer be appropriate.

6. Conclusions

Groundwater is an increasingly important water resource but its long-term viability is under threat due to over extraction, contamination or both. In many cases, these threats are being exacerbated by climate change. Understanding of subsurface geological structures, along with their parameterization and associated uncertainties, is steadily improving due to the plethora of geophysical tools now available and there is a need to maximize the benefits these additional data provide. To achieve this, a new generation of numerical models, which can readily accommodate geological complexity in a highly efficient computational framework, is required. Here we present a framework for use of dynamic mesh optimization (DMO), which respects geological heterogeneity, to simulate complex flow and transport problems in a parallel framework, exemplified here through saline intrusion modelling, in a highly efficient manner. We suggest that parallelized DMO provides opportunities for ultra-high resolution modelling of complex contaminant transport problems with sufficient rapidity that multiple 3D realizations to assess and quantify model and data uncertainty are not only desirable but also achievable.

CRediT authorship contribution statement

A. Hamzehloo: Methodology, Software, Visualization, Validation, Writing – original draft. **M.L. Bahlali:** Methodology, Software, Visualization, Validation, Writing – original draft. **P. Salinas:** Conceptualization, Methodology, Software, Writing – review & editing. **C. Jacquemyn:** Methodology, Software. **C.C. Pain:** Methodology, Software. **A.P. Butler:** Writing – original draft, Writing – review & editing, Supervision, Funding. **M.D. Jackson:** Conceptualization, Writing – original draft, Writing – review & editing, Supervision, Funding.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

The authors would like to thank the UK Engineering and Physical Sciences Research Council (EPSRC) for supporting this work by grant EP/R01938X/1 under project SALINA (Saline Intrusion in coastal Aquifers: Hydrodynamic Assessment and Prediction of Dynamic Response). The authors would also like to thank members of the Novel Reservoir Modelling and Simulation (NORMS) Group at Imperial College London for many valuable discussions. The authors have no conflicts of interest to declare. All data reported here were created using the Imperial College Finite Element Reservoir Simulator (IC-FERST) which is available under the Open Source Affero General Purpose Licence v.3.0 (AGPL v.3.0) at the IC-FERST Github repository (http://multifluids.github.io/).

References

Adam, A., Pavlidis, D., Percival, J.R., Salinas, P., Xie, Z., Fang, F., Pain, C.C., Muggeridge, A.H., Jackson, M.D., 2016. Higher-order conservative interpolation between control-volume meshes: Application to advection and multiphase flow problems with dynamic mesh adaptivity. J. Comput. Phys. 321, 512–531.

Akakpo, D., Gildin, E., 2017. A control perspective to adaptive time-stepping in reservoir simulation. In: SPE Reservoir Simulation Conference. Day 3 Wed, February 22, 2017, http://dx.doi.org/10.2118/182601-MS.

Alauzet, F., Loseille, A., 2016. A decade of progress on anisotropic mesh adaptation for computational fluid dynamics. Comput. Aided Des. 72, 13–39.

Anderson, D.G., 1965. Iterative procedures for nonlinear integral equations. J. Assoc. Comput. Mach. 12, 547–560.

Arfib, B., Charlier, J.-B., 2016. Insights into saline intrusion and freshwater resources in coastal karstic aquifers using a lumped rainfall–discharge–salinity model (the Port-Miou brackish spring, SE France). J. Hydrol. 540, 148–161.

Bakalowicz, M., 2015. Karst and karst groundwater resources in the mediterranean. Environ. Earth Sci. 74 (1), 5–14.

Balay, S., Gropp, W., McInnes, L., Smith, B., 1997. Efficient Management of Parallelism in Object Oriented Numerical Software Libraries. Birkha User Press, pp. 163–202.

Bianchi, M., Zheng, C., 2016. A lithofacies approach for modeling non-fickian solute transport in a heterogeneous alluvial aquifer. Water Resour. Res. 52 (1), 552–565.

Boretti, A., Rosa, L., 2019. Reassessing the projections of the world water development report. Npj Clean Water 2 (15).

Butler, A., Mathias, S., Gallagher, A., Peach, D., Williams, A., 2009. Analysis of flow processes in fractured chalk under pumped and ambient conditions (UK). Hydrogeol. J. 17, 1849—1858.

Cardenas, M.B., Bennett, P.C., Zamora, P.B., Befus, K.M., Rodolfo, R.S., Cabria, H.B., Lapus, M.R., 2015. Devastation of aquifers from tsunami-like storm surge by supertyphoon haiyan. Geophys. Res. Lett. 42 (8), 2844–2851.

Croucher, A., O'sullivan, M., 1995. The henry problem for saltwater intrusion. Water Resour. Res. 31 (7), 1809–1814.

Custodio, E., 2010. Coastal aquifers of Europe: an overview. Hydrogeol. J. 18 (1), 269–280.

Dahle, H., Espedal, M., Ewing, R., Sævereid, O., 1990. Characteristic adaptive subdomain methods for reservoir flow problems. Numer. Methods Partial Differential Equations 6, 279–309.

Dahle, H., Espedal, M., Ewing, R., Sævereid, O., 1992. Characteristic, local grid refinement techniques for reservoir flow problems. Int. J. Numer. Eng. 34, 1051.

De Filippis, G., Foglia, L., Giudici, M., Mehl, S., Margiotta, S., Negri, S.L., 2016. Seawater intrusion in karstic, coastal aquifers: Current challenges and future scenarios in the taranto area (southern Italy). Sci. Total Environ. 573, 1340–1351.

Fahs, M., Ataie-Ashtiani, B., Younes, A., Simmons, C.T., Ackerer, P., 2016. The Henry problem: New semi-analytical solution for velocity-dependent dispersion. Water Resour. Res. 52 (9), 7382–7407.

Farrell, P.E., Piggott, M.D., Pain, C.C., Gorman, G.J., Wilson, C.R., 2009. Conservative interpolation between unstructured meshes via supermesh construction. Comput. Methods Appl. Mech. Engrg. 198, 2632–2642.

Ferguson, G., Gleeson, T., 2012. Vulnerability of coastal aquifers to groundwater use and climate change. Nature Clim. Change 2, 342–345. http://dx.doi.org/10.1038/nclimate1413DO.

Feyen, L., Caers, J., 2006. Quantifying geological uncertainty for flow and transport modeling in multi-modal heterogeneous formations. Adv. Water Resour. 29 (6), 912–929.

Gerritsen, M., Lambers, J., 2008. Integration of local–global upscaling and grid adaptivity for simulation of subsurface flow in heterogeneous formations. Comput. Geosci. 12, 193–208.

Gomes, J., Pavlidis, D., Salinas, P., Xie, Z., Percival, J.R., Melnikova, Y., Pain, C.C., Jackson, M.D., 2017. A force-balanced control volume finite element method for multi-phase porous media flow modelling. Internat. J. Numer. Methods Fluids 83 (5), 431–445.

Gorman, G.J., Southern, J., Farrell, P.E., Piggott, M., Rokos, G., Kelly, P.H., 2012. Hybrid openMP/MPI anisotropic mesh smoothing. Procedia Comput. Sci. 9, 1513–1522.

Graham, M., MacAllister, D., Vinogradov, J., Jackson, M.D., Butler, A., 2018. Self-potential as a predictor of seawater intrusion in coastal groundwater boreholes. Water Resour. Res. 54 (9), 6055–6071.

Hammond, G., Lichtner, P., Mills, R., 2014. Evaluating the performance of parallel subsurface simulators: An illustrative example with PFLOTRAN. Water Resour. Res. 50, 208–228.

Held, R., Attinger, S., Kinzelbach, W., 2005. Homogenization and effective parameters for the Henry problem in heterogeneous formations. Water Resour. Res. 41 (11).

Henry, H.R., 1964. Effects of dispersion on salt encroachment in coastal aquifers, in"Seawater in Coastal Aquifers". US Geol. Surv. Water Supply Pap. 1613, C70–C80.

Hornung, R., Trangenstein, J., 1997. Adaptive mesh refinement and multilevel iteration for flow in porous media. J. Comput. Phys. 136, 522–545.

Hu, R., Fang, F., Salinas, P., Pain, C., 2018. Unstructured mesh adaptivity for urban flooding modelling. J. Hydrol. 560, 354–363.

Huyakorn, P., Wu, Y., Park, N., 1996. Multiphase approach to the numerical solution of a sharp interface saltwater intrusion problem. Water Resour. Res. 32, 93–102.

Hwang, H.-T., Park, Y.-J., Sudicky, E., Forsyth, P.A., 2014. A parallel computational framework to solve flow and transport in integrated surface–subsurface hydrologic systems. Environ. Model. Softw. 61, 39–58.

Isa-Abadi, A.R., Fontaine, V., Ghafouri, H.-R., Younes, A., Fahs, M., 2020. A fully interior penalty discontinuous Galerkin method for variable density groundwater flow problems. Comput. & Fluids 213, 104744.

Jackson, M.D., Hampson, G.J., Saunders, J.H., El-Sheikh, A., Graham, G.H., Massart, B.Y.G., 2014. Surface-based reservoir modelling for flow simulation. Geol. Soc. Lond. Special Publ. 387 (1), 271–292.

Jackson, M.D., Percival, J., Mostaghimi, P., Tollit, B., Pavlidis, D., Pain, C., Gomes, J., Elsheikh, A.H., Salinas, P., Muggeridge, A., Blunt, M.J., 2015. Reservoir modeling for flow simulation by use of surfaces, adaptive unstructured meshes, and an overlapping-control-volume finite-element method. SPE Reserv. Eval. Eng. 18 (02), 115–132.

Jacquemyn, C., Jackson, M.D., Hampson, G.J., 2019. Surface-based geological reservoir modelling using grid-free NURBS curves and surfaces. Math. Geosci. 51 (1), 1–28.

Jones, H., Robins, N., 1999. The Chalk Aquifer of the South Downs. British Geological Survey.

Kampitsis, A., Adam, A., Salinas, P., Pain, C., Muggeridge, A., Jackson, M.D., 2020. Dynamic adaptive mesh optimisation for immiscible viscous fingering. Comput. Geosci. 24 (3), 1221–1237.

Kerrou, J., Renard, P., 2010. A numerical analysis of dimensionality and heterogeneity effects on advective dispersive seawater intrusion processes. Hydrogeol. J. 18, 55–72.

Kerrou, J., Renard, P., Cornaton, F., Perrochet, P., 2013. Stochastic forecasts of seawater intrusion towards sustainable groundwater management: application to the Korba aquifer (Tunisia). Hydrogeol. J. 21, 425–440.

Ketabchi, H., Mahmoodzadeh, D., Ataie-Ashtiani, B., Simmons, C.T., 2016. Sea-level rise impacts on seawater intrusion in coastal aquifers: Review and integration. J. Hydrol. 535, 235–255.

Klassen, J., Allen, D., 2017. Assessing the risk of saltwater intrusion in coastal aquifers. J. Hydrol. 551, 730–745.

Koohbor, B., Fahs, M., Ataie-Ashtiani, B., Belfort, B., Simmons, C.T., Younes, A., 2019. Uncertainty analysis for seawater intrusion in fractured coastal aquifers: Effects of fracture location, aperture, density and hydrodynamic parameters. J. Hydrol. 571, 159–177.

Koohbora, B., Fahsa, M., Ataie-Ashtiani, B., Belfort, B., Simmons, C.T., Younes, A., 2019. Uncertainty analysis for seawater intrusion in fractured coastal aquifers: Effects of fracture location, aperture, density and hydrodynamic parameters. J. Hydrol. 571, 159–177.

Langevin, C., Thorne, D., Dausman, A., Sukop, M., Guo, W., 2007. SEAWAT version 4: a computer program for simulation of multi-species solute and heat transport. In: Tech Methods Book 6 (Chapter A 22:39). US Geological Survey.

Lichtner, P.C., Hammond, G.E., Lu, C., Karra, S., Bisht, G., Andre, B., Mills, R., Kumar, J., 2015. PFLOTRAN user manual.

Lin, H.-C., Rechards, D., Talbot, C., Yeh, G.-T., Cheng, J.-R., Cheng, H.-P., Jones, N., 1997. FEMWATER: A Three-Dimensional Finite Element Computer Model for Simulating Density-Dependent Flow and Transport in Variable Saturated Media. Technical Rep, No CHL-97-12 Vicksburg, MS, US Army Engineer Waterways Experiment Station Coastal and Hydraulics Laboratory.

MacAllister, D., Jackson, M.D., Butler, A., Vinogradov, J., 2018. Remote detection of saline intrusion in a coastal aquifer using borehole measurements of self-potential. Water Resour. Res. 54 (3), 1669–1687.

MacDonald, A.M., Allen, D.J., 2001. Aquifer properties of the chalk of England. Q. J. Eng. Geol. Hydrogeol. 34 (4), 371–384.

Masciopinto, C., Liso, I.S., 2016. Assessment of the impact of sea-level rise due to climate change on coastal groundwater discharge. Sci. Total Environ. 569, 672–680.

Morgan, L.K., Werner, A.D., 2015. A national inventory of seawater intrusion vulnerability for Australia. J. Hydrol. Reg. Stud. 4, 686–698.

Obeysekara, A., Lei, Q., Salinas, P., Pavlidis, D., Xiang, J., Latham, J.-P., Pain, C., 2018. Modelling stress-dependent single and multi-phase flows in fractured porous media based on an immersed-body method with mesh adaptivity. Comput. Geotech. 103, 229–241.

Osman, H., Graham, G.H., Moncorge, A., Jacquemyn, C., Jackson, M.D., 2021. Is cell-to-cell scale variability necessary in reservoir models? Math. Geosci. 53 (4), 571–596.

Pain, C., Umpleby, A., De Oliveira, C., Goddard, A., 2001. Tetrahedral mesh optimisation and adaptivity for steady-state and transient finite element calculations. Comput. Methods Appl. Mech. Engrg. 190 (29–30), 3771–3796.

Rajabi, M.M., Ataie-Ashtiani, B., 2014. Sampling efficiency in Monte Carlo based uncertainty propagation strategies: Application in seawater intrusion simulations. Adv. Water Resour. 67, 46–64.

Rajabi, M.M., Ataie-Ashtiani, B., Simmons, C.T., 2015. Polynomial chaos expansions for uncertainty propagation and moment independent sensitivity analysis of seawater intrusion simulations. J. Hydrol. 520, 101–122.

Roy, D.K., Datta, B., 2018. A review of surrogate models and their ensembles to develop saltwater intrusion management strategies in coastal aquifers. Earth Syst. Environ. 2, 193–211.

Salinas, P., Pavlidis, D., Xie, Z., Adam, A., Pain, C.C., Jackson, M.D., 2017a. Improving the convergence behaviour of a fixed-point-iteration solver for multiphase flow in porous media. Internat. J. Numer. Methods Fluids 84 (8), 466–476.

Salinas, P., Pavlidis, D., Xie, Z., Jacquemyn, C., Melnikova, Y., Jackson, M.D., Pain, C.C., 2017b. Improving the robustness of the control volume finite element method with application to multiphase porous media flow. Internat. J. Numer. Methods Fluids 85 (4), 235–246.

Salinas, P., Pavlidis, D., Xie, Z., Osman, H., Pain, C., Jackson, M.D., 2018. A robust mesh optimisation method for multiphase porous media flows. Comput. Geosci. 22 (5), 1389–1401.

Salinas, P., Regnier, G., Jacquemyn, C., Pain, C.C., Jackson, M.D., 2021. Dynamic mesh optimisation for geothermal reservoir modelling. Geothermics 94, 102089.

Schloegel, K., Karypis, G., Kumar, V., 2002. Parallel static and dynamic multi-constraint graph partitioning. Concurr. Comput.: Pract. Exper. 14, 219–240.

Schmidt, G., Jacobs, F., 1988. Adaptive local grid refinement and multi-grid in numerical reservoir simulation. J. Comput. Phys. 77, 140–165.

Small, C., Nicholls, R.J., 2003. A global analysis of human settlement in coastal zones. J. Coast. Res. 19 (3), 584–599.

Souza, W., Voss, C., 1987. Analysis of an anisotropic coastal aquifer system using variable-density flow and solute transport simulation. J. Hydrol. 92, 17–41.

Su, D., Mayer, K.U., MacQuarrie, K.T., 2021. MIN3P-HPC: a high-performance unstructured grid code for subsurface flow and reactive transport simulation. Math. Geosci. 53 (4), 517–550.

Therrien, R., McLaren, R., Sudicky, E., Panday, S., 2010. HydroGeoSphere: A Three-Dimensional Numerical Model Describing Fully-Integrated Subsurface and Surface Flow and Solute Transport. Vol. 35. Groundwater Simulations Group, University of Waterloo, Waterloo, on.

Titus, Z., Heaney, C., Jacquemyn, C., Salinas, P., Jackson, M.D., Pain, C., 2021. Conditioning surface-based geological models to well data using artificial neural networks. Comput. Geosci. http://dx.doi.org/10.1007/s10596-021-10088-5.

Trangenstein, J., 2002. Multi-scale iterative techniques and adaptive mesh refinement for flow in porous media. Adv. Water Resour. 25, 1175–1213.

Uddameri, V., Singaraju, S., Hernandez, E.A., 2014. Impacts of sea-level rise and urbanization on groundwater availability and sustainability of coastal communities in semi-arid South Texas. Environ. Earth Sci. (ISSN: 1866-6299) 71, 2503–2515.

Voss, C.I., Provost, A.M., 1984. SUTRA. In: US Geological Survey Water Resources Investigation Reports. pp. 84–4369.

Voss, C., Souza, W., 1987. Variable density flow and solute transport simulation of regional aquifers containing a narrow freshwater-saltwater transition zone. Water Resour. Res. 23, 1851–1866.

Wei, X., Li, W., Tian, H., Li, H., Xu, H., Xu, T., 2015. THC-MP: High performance numerical simulation of reactive transport and multiphase flow in porous media. Comput. Geosci. 80, 26–37.

Willis, B., White, C., 2000. Quantitative outcrop data for flow simulation. J. Sediment. Res. 70, 788–802.

Wolff, M., Flemisch, B., Helmig, R., 2013. An adaptive multiscale approach for modeling two-phase flow in porous media including capillary pressure. Water Resour. Res. 49, 8139–8159.

Xie, Z., Pavlidis, D., Percival, J., Gomes, J., Pain, C., Matar, O., 2014. Adaptive unstructured mesh modelling of multiphase flows. Int. J. Multiphase Flow 67, 104–110.

Yang, Y., Song, J., Simmons, C.T., Ataie-Ashtiani, B., Wu, J., Wang, J., Wu, J., 2021. A conjunctive management framework for the optimal design of pumping and injection strategies to mitigate seawater intrusion. J. Environ. Manag. 282, 111964.

Younes, A., Fahs, M., Ahmed, S., 2009. Solving density driven flow problems with efficient spatial discretizations and higher-order time integration methods. Adv. Water Resour. 32, 340–352.

Zidane, A., Younes, A., Huggenberger, P., Zechner, E., 2012. The Henry semi-analytical solution for saltwater intrusion with reduced dispersion. Water Resour. Res. 48 (6).