Upscaling of the porous shallow water equations through the use of periodic homogenization

Lisa Patrascu

June 2025 - under the supervision of Nicholas Kevlahan

Contents

List of Figures

List of Tables

1 Abstract

Due to the large and costly nature of ocean models, they are often limited in computational resolution - meaning accurate solvers must take into account boundary geometry at the subgrid-scale without explicitly modelling it. One method of characterizing subgrid-scale features is Brinkman penalization, where the solid/fluid interface is modelled as a porous medium, which yields many stability, accuracy, and efficiency benefits [10, 11]. In this work, we aim to extend the Brinkman method by testing a penalization that accounts for a position dependent tensorial permeability which will allow the model to experience friction in a directionally dependent fashion - implicitly preserving roughness that may be lost in a porosity-only approach. We simulate flow through solid/fluid and semi-permeable permeability-defined substructure configurations using the porous shallow water equations at both the subgrid-scale as well as the coarsened scale. Our findings indicate that coarsened simulations well approximate averaged fine-scale simulations in both velocity distribution and total kinetic energy. We coarsen subgrid-scale permeability using the periodic homogenization approach, which we find to be rigorous, fast, and accurate.

2 Introduction

There are many different types of coastline topographies, each with distinct characteristics that play an essential role in determining ocean–coast interactions. For example, just west of the Mississippi river, the swampy Atchafalaya Basin of Louisiana is dominated by clusters of herbaceous plants. Understanding water flow through this region is imperative for flood control and wetland conservation. The ocean dynamics in the Gulf of Mexico depend on the gulf's shape as well as its bathymetry, which is a mix of continental shelves, slopes and submarine canyons. It is important to accurately simulate since it plays a key role in circulating warm water into the global ocean system [19]. The Strait of Gibraltar is a narrow channel which connects the Atlantic Ocean to the Mediterranean Sea, which is not only significant for similar reasons to the Gulf stream, but is also widely used as a benchmark case for testing ocean models.

The diverse nature of these complex coastal geometries means that their accurate simulation not only requires a technique that properly represents their unique characteristics, but also one that accounts for their consequent numerical challenges. Simple rocky bathymetry or coastlines are reasonably well-approximated by solid, no-slip boundaries, but the introduction of any more realistic and sharply varying features make them prone to numerical instability. Areas that are not quite solid nor pure fluid, like patches of plant growth in a marsh or estuary, are more accurately modelled as a porous medium, but their small scale makes them difficult to characterize given a limited resolution. Processes that model flow through small and distinct passages, like a narrowing channel, even sometimes require ad-hoc hand tuning like the manual opening of a gap.

On top of these numerical challenges, the development of techniques that model water flow along and through complex coastlines must also strike a fine balance between efficiency and accuracy. As such, realistic ocean models are large and costly, leading to a limit in computational resolution. To mitigate any resulting instability, they might smooth sharper features and reintroduce an artificial friction term to compensate for the loss of roughness. Though both regional and global ocean models either implicitly or explicitly require some characterization of the unresolved subgrid-scale (SGS) features to ensure accuracy at a limited resolution, fulfilling this through a scalar friction term is not only arbitrary but also neglects to capture directional dependence.

An alternative approach is Brinkman penalization, described and successfully implemented on peaked bathymetry by Debreu et al. in [10]. This approach modifies the Shallow Water Equations (SWE) by the addition of a scalar permeability–porosity term that is used to penalize a region between the true bathymetry and a smoothed computational boundary that lies fully within the solid area (Figure 1). Brinkman penalization corrects for the failings of topography-conforming coordinate approaches such as spurious flows caused by staircase geometries in $z-$coordinates and the trade-off between smoothing and pressure gradient errors in $\sigma-$coordinates while increasing robustness and improving computational performance [11]. The

modelling of the interface between a solid/fluid boundary as a porous medium is what allows this method to include the impact that SGS features have on the flow, which provably maintains accuracy at a lower resolution without compromising stability (see [11] for the successful results on the Gulf Stream Separation problem).

As successful as Brinkman penalization is with porosity, we believe that the inclusion of a permeability factor would only provide further improvement. While porosity quantifies the fluid volume fraction, permeability gives a measure of the interconnectedness of pores. One computational cell may have a small amount of solid material, but arranged in such a way that completely blocks the fluid flow. For example a cell filled with clay, which is highly porous but has a low permeability. Conversely, gravel has a low porosity but a high permeability. A penalization that only relies on porosity can severely misclassify substructure that truly needs an additional tensorial permeability characterization to properly describe flow through.

Therefore, to ultimately build on Brinkman penalization, we propose to test an improvement on the scalar implementation of permeability–porosity by replacing it with a Darcy-type tensor, which would flexibly permit the independent penalization of velocity along both its principal axes [6]. In a scalar approach, the representation of SGS features at the grid-scale is done by a spatial average, which does not need to change for the upscaling of porosity (due to its inherent scalar nature as a volume fraction). However, the coarsening of an anisotropic SGS permeability matrix in the context of a differential equation is provably *not* an average [1, 17, 22], so we choose to coarsen SGS permeability using periodic homogenization, a rigorous coarsening method that is well-established in its application to conductivity and elasticity problems [18]. In fact, a tensorial permeability applied to the diffusive Darcy equation is even often referred to as a "hydraulic conductivity" in the engineering community [15].

In the limit of vanishing permeability and porosity, the porous SWE approximate no-slip solid boundary conditions. Coupled with a rigorous homogenization procedure, appropriate choices of spatially varying porosity and tensor permeability allow one to model coastlines of arbitrary complexity. Therefore, with the ultimate goal of later being able to apply this method to realistic ocean models, we test periodic homogenization's performance on diagonal and full matrices against another well known coarsening method, as well as a testing a porous Shallow Water Model with substructure configurations that mimic marsh and tunnel domains in both fine and upscaled resolutions. We test a variety of permeability matrices, both semi-permeable and solid-representing, with the hope that this method can provide improvements in modelling not only solid/fluid boundaries, but also expansive low-lying wetlands like coastal Louisiana.

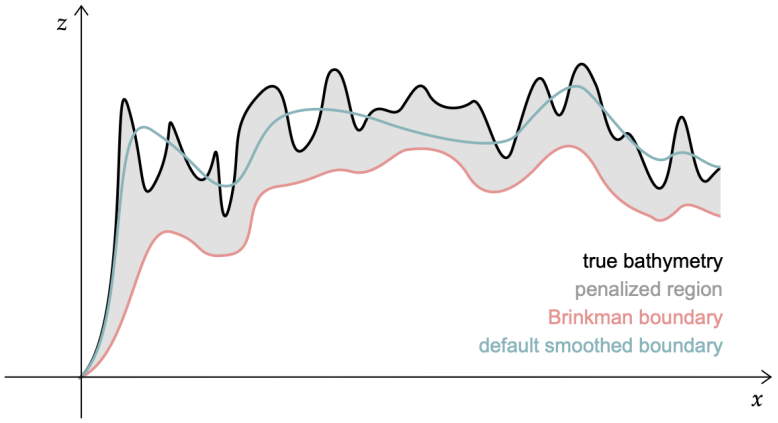

Figure 1: The bathymetry of an ocean model is often smoothed to prevent numerical instability (default bathymetry in the Figure), where an artificial bottom friction is added back in to compensate for the loss of roughness. A proven alternative approach is Brinkman penalization, which takes a smoothed boundary internal to the true bathymetry and represents the area between this boundary and the pure fluid as a porous medium.

2.1 The Darcy, Brinkman, and Shallow Water Equations

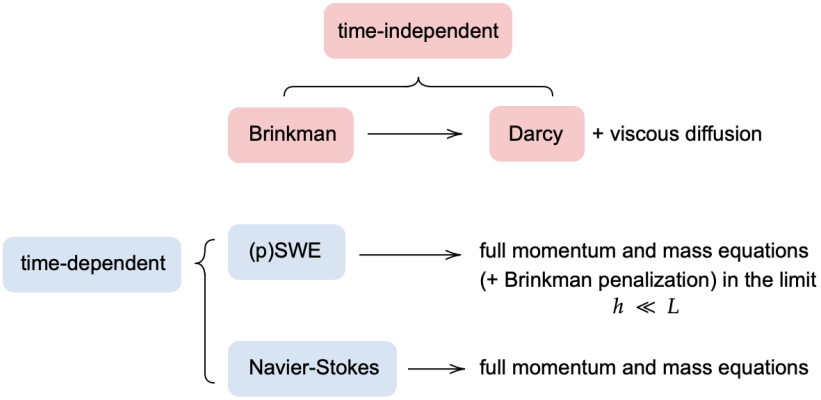

Figure 2: Fluid modelling equations from simplest (Darcy) to most involved (Navier-Stokes). Brinkman penalization was developed as a modification to the Navier-Stokes equations via the analysis of the Brinkman equations [4, 5].

The simplest model for flow through a porous medium is the Darcy Equation. In Darcy flows, velocity u at some point in space is proportional to the pressure gradient $\boldsymbol{\nabla} P$:

$$\phi u = -K \boldsymbol{\nabla} P,$$

where ϕ is the scalar porosity ($\phi = 1$ for a pure fluid) and in n dimensions, the tensor K is the $n \times n$ permeability matrix. Taking the divergence of this equation, and enforcing incompressibility, $\boldsymbol{\nabla} \cdot u = 0$, gives the Darcy equation,

$$-\boldsymbol{\nabla} \cdot (K \boldsymbol{\nabla} P) = 0. \tag{1}$$

When supplemented with Neumann boundary conditions for the pressure, (1) is an elliptic Cauchy problem (i.e. Laplace equation) for the stationary distribution of velocity and pressure in a porous medium given a permeability distribution $K(x)$. The Brinkman Equation modifies the Darcy Equation by the addition of a viscous diffusion term and allows for the modelling of time-dependent flow that is in a transitional state between the Darcy and Stokes regimes. Through the analysis of the Brinkman Equation was the idea of Brinkman penalization formed, which is the addition of a friction of the form $K^{-1}u$ to the Navier-Stokes equations [4, 5]. An assumption is made that the permeability will be symmetric positive definite to ensure invertibility.

The SWE, which are the limit of the Navier-Stokes equations for long, shallow regimes, can also be penalized by a Brinkman-type term. These are the equations we will look at in §6 when simulating flow over complex boundary geometries, but this first requires the homogenization of SGS permeability, for which now turn our focus back to the Darcy equation.

2.2 Homogenization

The discrete homogenization problem for permeability is characterized by accurately representing permeability on a coarse grid from data on a fine grid (Figure 3). We assume there is a clear separation of scales between the two grids. An accurate or physically useful homogenization can be proven or obtained in a variety of ways, for example through a mathematical limit [2, 14], in a statistical sense [12], or numerically [15].

Naively, one might think that spatially averaging the permeability over a grid-scale cell is an acceptable homogenization. However, this method may fail to preserve important quantitative properties of the porous medium flow, such as blocking effects and "turning the flow" (see Figure 4).

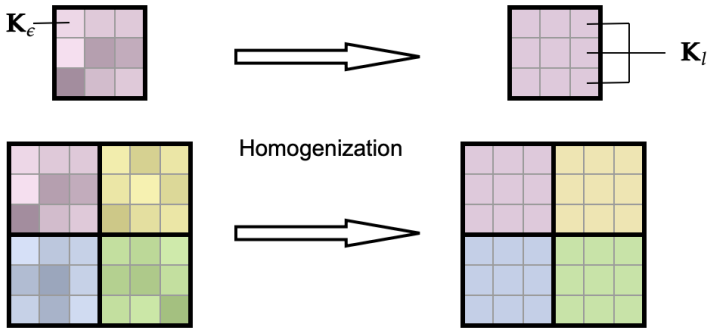

Figure 3: The process of homogenization assigns every cell in a domain, of which each had previously a unique permeability matrix K_ϵ varying on the subgrid, one representative permeability matrix K_l.

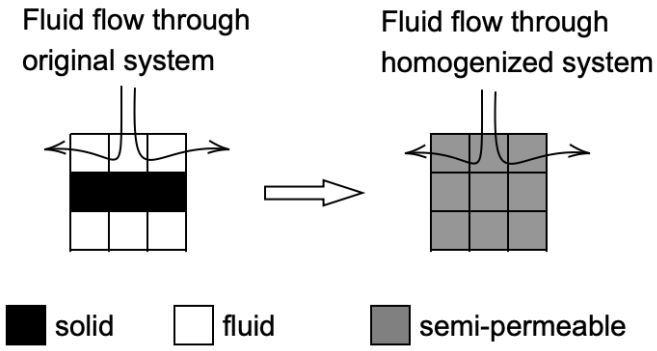

Figure 4: A homogenized system should still preserve key characteristics of the original flow, such as flow being turned by a barrier.

Homogenization, also known as upscaling or coarsening, is essentially a type of subgrid-scale (SGS) model, which is such a generalized tool that it has been implemented in many different subdomains within mathematics, physics, and engineering. Periodic homogenization (PH) is a well-established and mathematically rigorous method widely used in engineering disciplines to model qualities such as small-scale conduction or elasticity. It was introduced in the early 1990s by Allaire [3], and was soon after implemented numerically by Moulinec and Suquet [18]. This original theoretical and computational approach forms the basis of many homogenization algorithms for elliptic equations used today.

Although we will primarily use PH for two-dimensional domains, it can also be applied to three-dimensional systems with very little modification [13]. In our case, the two spatial dimensions can either be thought of as a horizontal x—y 2D-SWE model (e.g., swampy coastlines) or a vertical x—z slice (e.g., ocean models, bathymetry), as illustrated in Figure 5.

In this work, we present a new application of PH to upscale the Darcy Equation in order to provide accurate grid scale coastal or bathymetry conditions from SGS permeability data. PH generates a coarse-scale topography model that can be used in conjunction with a flow solver to create SGS-informed simulations at a limited resolution.

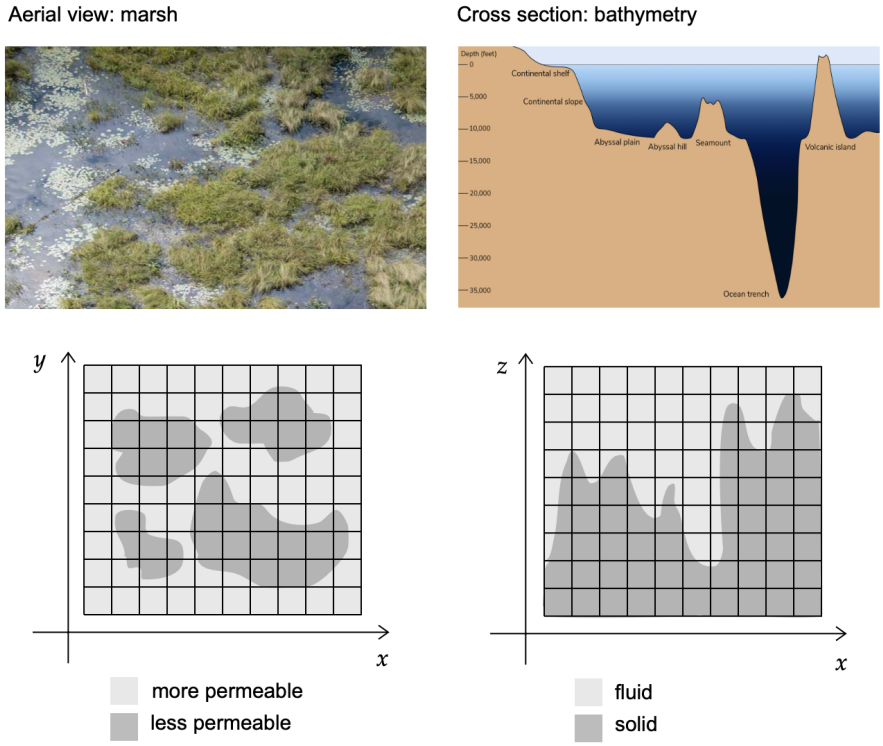

Figure 5: Horizontal and vertical homogenization problems. The two dimensions can either represent an aerial view, like a swamp with pockets of lower permeability [9], or a vertical slice with complex bathymetry. The small scale permeability can take on any value depending on the local properties of the small scale topography data.

3 Implementation of periodic homogenization (PH)

3.1 Set-up

If a global domain (say of an ocean model) Ω_L is discretized into grid-scale cells with length-scale $l \ll L$, homogenization should be performed on one of these grid-scale cells based on its contained SGS features. There should be a clear scale separation between a grid-scale cell Ω_l and subgrid cell Ω_ϵ such that $\epsilon \ll l$. Any given grid-scale cell Ω_l in the domain contains a permeability matrix K_ϵ, which varies on the subgrid scale and can be homogenized into one representative permeability matrix K_l. Our goal is to perform this homogenization using PH, which along with the assumption of scale separation also requires SGS permeabilities K_ϵ be symmetric and positive definite (i.e. flow through a porous medium must dissipate energy). Having satisfied these conditions, the following steps are performed on each grid-scale cell Ω_l to find its corresponding K_l:

1. Pose Darcy's equation as a corrector problem, with SGS permeability matrices K_ϵ and grid-scale average pressure gradient E_l as inputs.
2. Reshape corrector problem into a discrete Lippmann–Schwinger (LS) form.
3. Solve discretized corrector/LS for SGS perturbed pressure gradient $\boldsymbol{\nabla} w_\epsilon$.
4. Use Lippmann–Schwinger formulation and $\boldsymbol{\nabla} w_\epsilon$ to solve for homogenized grid-scale permeability matrix K_l.

We now describe each step of the above algorithm.

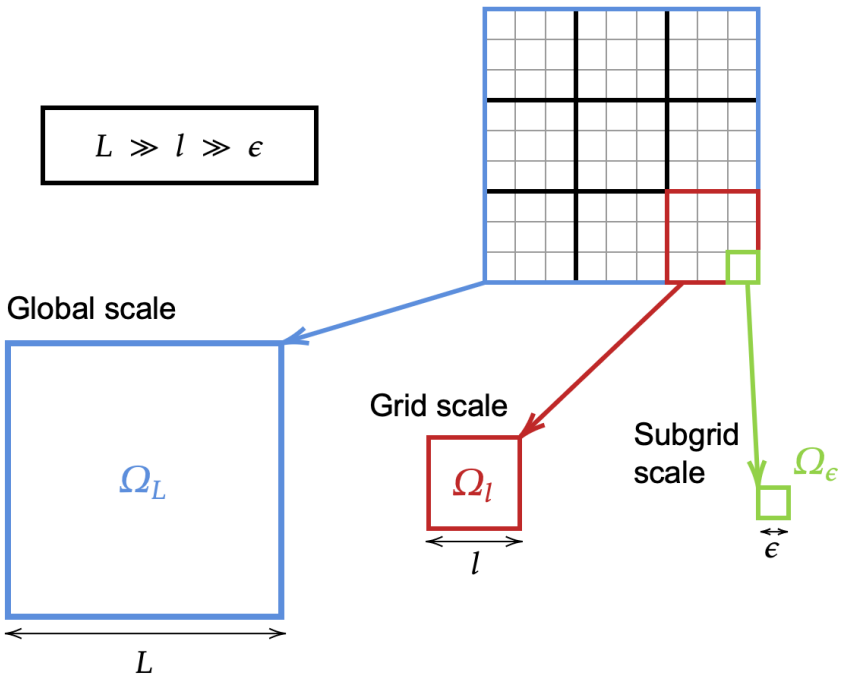

Figure 6: Graphic illustrating the separation of scales. The global domain Ω_L is split into grid-scale cells Ω_l and further into subgrid-scale cells Ω_ϵ where $L \gg l \gg \epsilon$.

3.1.1 The corrector problem

The "corrector problem" is an alternate form of Darcy's equation that splits the strain field (pressure gradient) into an average gradient and a perturbation:

$$\boldsymbol{\nabla} \cdot [K_\epsilon(E_l + \boldsymbol{\nabla} w_\epsilon)] = 0$$

where E_l is the grid-scale pressure gradient and $\boldsymbol{\nabla} w_\epsilon$ is the periodic SGS pressure gradient. Scale consistency requires that the average of the total gradient be equal to the grid-scale pressure gradient ($\overline{E_l + \boldsymbol{\nabla} w_\epsilon} = E_l$). Although the corrector problem itself can be solved computationally for the SGS pressure gradient $\boldsymbol{\nabla} w_\epsilon$ with a given grid-scale pressure gradient E_l (which also provides boundary conditions for Ω_l), the need to compute derivatives makes a direct approach computationally inefficient. However, a clever transformation and discretization allows one to find a $\boldsymbol{\nabla} w_\epsilon$ that leverages the $\mathcal{O}(\epsilon)$ periodicity assumption in order to take advantage of highly efficient Fourier methods.

3.1.2 Discrete linear system for the SGS pressure gradient

This process follows directly from the work of de Geus et al. [13] as well as that from Brisard and Legoll [7], with variables scaled down a dimension (ie. tensor to vector, vector to scalar) and renamed to increase legibility in the context of this work. Defining the SGS velocity field $u_\epsilon = K_\epsilon(E_l + \boldsymbol{\nabla} w_\epsilon)$, the corrector problem can be rewritten in its weak form

$$\int_{\Omega_l} \tilde{F}(\boldsymbol{\nabla} \cdot u_\epsilon)\, \mathrm{d}\Omega_l = \int_{\Omega_l} (\boldsymbol{\nabla} \tilde{F}) u_\epsilon\, \mathrm{d}\Omega_l = 0$$

using a test function $\tilde{F}$, where the divergence theorem moves the divergence from the velocity field to a gradient on the test function.

The operator Γ_0 is defined such that $\boldsymbol{\nabla} \tilde{F} = \Gamma_0 * \delta \tilde{F}$ and using the associative property of convolution and the distributive property of the tensor product over a convolution, a new weak

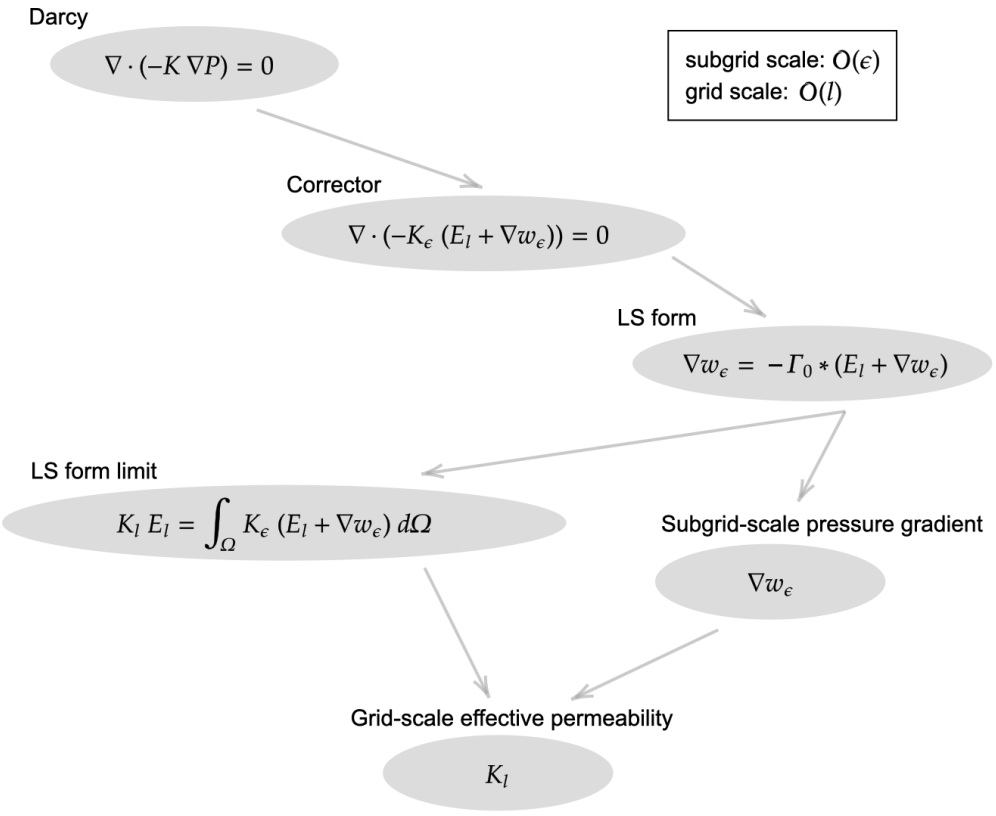

Figure 7: Flowchart depicting how the Darcy Equation is taken through the corrector problem to achieve the LS formalism, whose (numerically determined) solution is microscale pressure gradient $\boldsymbol{\nabla} w_\epsilon$, which can be used in the $\epsilon \to 0$ limit of LS to solve directly for the effective macroscale permeability matrix K_l.

integral form appears:

$$\int_{\Omega_l} (\Gamma_0 * \delta\tilde{F}) u_\epsilon \,\mathrm{d}\Omega_l = \int_{\Omega_l} \delta\tilde{F} (\Gamma_0 * u_\epsilon) \,\mathrm{d}\Omega_l = 0.$$

where $\delta\tilde{F}$ is the new test function that happens to have no position dependence, and therefore can be taken out of the integral.

This integral is then discretized using shape functions N that sample discrete values of the continuous fields $\delta\tilde{F}$ and u_ϵ at nodes X_m,

$$N_k(X_m) = \delta_{km}. \tag{2}$$

Additionally, since this relation must be valid for any $\delta\tilde{F}$, the integral further simplifies to

$$\int_{\Omega_l} N(\Gamma_0 * u_\epsilon) \,\mathrm{d}\Omega_l = 0.$$

Using the trapezoid rule to approximate the integral, where all nodal quantities have an equal weight of one assigned by the shape function (2), we have

$$\sum_k N\,(\Gamma_0 * u_\epsilon) = 0. \tag{3}$$

Although this integral represents a convolution in physical space, it is a tensor product in Fourier space. In Fourier space (with periodic SGS cells Ω_ϵ), the projection $\hat{\Gamma}_0$ is a diagonal matrix where

$$\hat{\Gamma}_0(q) = \begin{cases} 0, & \|q\| = 0, \\ \frac{q_i q_j}{\|q\|^{1/2}}, & \|q\| \neq 0, \end{cases}$$

where q is the wave number. The definition of the operator Γ_0 enforces a zero mean component such that

$$\int_{\Omega_l} \sum_k N(\Gamma_0 * u_\epsilon)\, \mathrm{d}\Omega_l = 0 \tag{4}$$

In the above sums, k indexes the nodes of N.

Since K_ϵ is a linear operator, solving equation 4 for SGS velocity u_ϵ is akin to solving

$$\mathcal{F}^{-1}[\hat{\Gamma}_0 \mathcal{F}(K_\epsilon(E_l + \boldsymbol{\nabla} w_\epsilon))] = 0 \tag{5}$$

for SGS pressure gradient $\boldsymbol{\nabla} w_\epsilon$.

To use a nonlinear operator K_ϵ, a linearization process is required with initialization such that $u_\epsilon^{(0)} = K_\epsilon E_l$ and $u_\epsilon^{(1)} = u_\epsilon^{(0)} + \delta u = K_\epsilon(E_l + \boldsymbol{\nabla} w_\epsilon)$. With Equation 4, this results in the same target equation to solve as the already linear system.

The code used for this section is largely based on that accompanying [13]. Written in Python, we have modified it use the notation introduced here and with the additional goal of solving for the homogenized matrix K_l instead of only the solution for $\boldsymbol{\nabla} w$ in both two and three dimensions. The process of solving for K_l is covered in the following section.

In the script associated with Table 1, equation 5 is represented in a single line of code

```
GK_fun(grad_w) = -GK_fun(E),
```

which is solved for the SGS pressure $\boldsymbol{\nabla} w_\epsilon$ using the conjugate gradient method. A summary of code objects, their types, and their descriptions is provided below.

Code Object	Type	Description
`ndim`	scalar	number of physical dimensions (ie. 1D, 2D, 3D)
`N`	scalar	number of nodes/cells in one dimension
`prodN`	scalar	number of entries in `E` and `grad_w` (`ndim` x `N` x `N`)
`K`	(`ndim`,`ndim`,`N`,`N`) vector	symmetric invertible positive definite microscale permeability matrix K_ϵ
`E`	(`ndim` x `N` x `N`) vector	macroscale pressure gradient E_l
`grad_w`	(`ndim` x `N` x `N`) vector	microscale periodic pressure gradient $\boldsymbol{\nabla} w_\epsilon$ (found using conjugate gradient method)
`K_l`	variable	2×2 symmetric positive definite homogenized macroscale permeability matrix K_l
`G_hat (v)`	function	projection operator $\hat{\Gamma}_0 v$
`K_fun(v)`	function	applies permeability matrix $K_\epsilon v$
`G_fun(v)`	function	applies $\mathcal{F}^{-1}(\hat{\Gamma}_0 \mathcal{F}(v))$
`GK_fun(v)`	function	applies $\mathcal{F}^{-1}[\hat{\Gamma}_0 \mathcal{F}(K_\epsilon v)]$

Table 1: Description of key code variables and functions used to solve for the microscale pressure gradient $\boldsymbol{\nabla} w_\epsilon$ in the function `get_K_l(ndim,N,A)`.

3.1.3 The Lippmann–Schwinger formulation for homogenization

Using the Green's function, the corrector problem can be transformed to the Lippmann–Schwinger equation [7][16]. This equation takes the form

$$\boldsymbol{\nabla} w_\epsilon = -\Gamma_0 * (E_l + \boldsymbol{\nabla} w_\epsilon),$$

where in the limit as $\epsilon \to 0$ a homogeneous form

$$K_l E_l = \int_{\Omega_l} K_\epsilon (E_l + \boldsymbol{\nabla} w_\epsilon)\, \mathrm{d}\Omega_l$$

is obtained. K_l is the homogenized permeability coefficient, which is constant over grid-scale cells Ω_l, and E_l is the prescribed grid-scale pressure gradient.

Choosing appropriate values for E_l and pre-factor T allows one to solve for each of the three distinct entries of the symmetric grid-scale permeability matrix K_l using the relationship

$$T K_l E_l = T \left(\frac{1}{\texttt{prodN}} \sum_k K_k (E_{l\,k} + \boldsymbol{\nabla} w_{\epsilon\,k}) \right).$$

As a two-dimensional example, to find the entry $K_{l\,11}$, we prescribe a macroscale gradient pressure $E_l = [1\ 0]^T$ and a pre-factor $T = [1\ 0]$. The left hand side of the above equation then becomes

$$T\, K_l E_l = [1\ 0]\ K [1\ 0]^T = [1\ 0]\ [K_{l\,11}\ K_{l\,21}]^T = K_{\iota 11}.$$

The other components of the homogenized permeability matrix can be found similarly, using the given E_l in Table 2.

K_l component	Choice of macro-strain	Pre-factor
$K_{l\,11}$	$E_l = [1\ 0]^T$	$T = [1\ 0]$
$K_{l\,12}$	$E_l = [0\ 1]^T$	$T = [1\ 0]$
$K_{l\,21}$	$E_l = [1\ 0]^T$	$T = [0\ 1]$
$K_{l\,22}$	$E_l = [0\ 1]^T$	$T = [0\ 1]$

Table 2: Imposed macroscale pressure gradients E_l and pre-factors T for computing components of the homogenized macroscale permeability matrix K_l in 2D. Note that the expected symmetry of the permeability matrix can be used as a sanity check by verifying $K_{l\,12} = K_{l\,21}$.

The same method can be applied to extract the components of a homogenized permeability matrix in 3D by making the appropriate dimensionality respecting choices for the macro-strain and pre-factor (Table 3).

K_l component	Choice of macro-strain	Pre-factor
$K_{l\,11}$	$E_l = [1\ 0\ 0]^T$	$T = [1\ 0\ 0]$
$K_{l\,12}$	$E_l = [0\ 1\ 0]^T$	$T = [1\ 0\ 0]$
$K_{l\,13}$	$E_l = [0\ 0\ 1]^T$	$T = [1\ 0\ 0]$
$K_{l\,21}$	$E_l = [1\ 0\ 0]^T$	$T = [0\ 1\ 0]$
$K_{l\,22}$	$E_l = [0\ 1\ 0]^T$	$T = [0\ 1\ 0]$
$K_{l\,23}$	$E_l = [0\ 0\ 1]^T$	$T = [0\ 1\ 0]$
$K_{l\,31}$	$E_l = [1\ 0\ 0]^T$	$T = [0\ 0\ 1]$
$K_{l\,32}$	$E_l = [0\ 1\ 0]^T$	$T = [0\ 0\ 1]$
$K_{l\,33}$	$E_l = [0\ 0\ 1]^T$	$T = [0\ 0\ 1]$

Table 3: Imposed macroscale pressure gradients E_l and pre-factors T for computing components of the homogenized macroscale permeability matrix K_l in 3D.

In order to determine whether PH is a viable means of coarsening data, it needs to pass a series of validation tests. Most importantly, a valid method should be able to reproduce comparable results for cases that can be solved analytically and cases that will have known and observable physical effects. In the following section we present validation results in two and three dimensions.

4 A quick validation of PH

4.1 Laminate tests

Domains that have laminated structures have known analytical homogenized solutions where coarsened matrix components are expected to follow the arithmetic mean of the corresponding SGS components in the direction parallel to the laminate and the harmonic mean in the direction perpendicular to the laminate [6]. We introduce a simple case where the domain Ω_l is filled with permeability matrix K_1 except for a "blocking" band with lower permeability K_2.

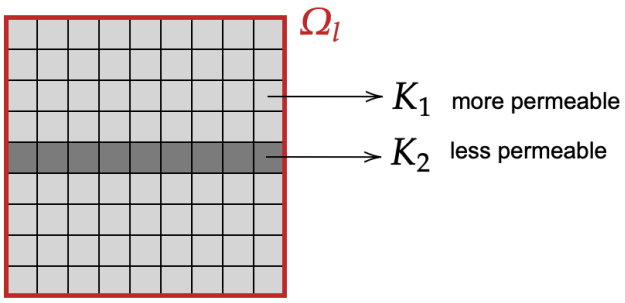

Figure 8: Blockage set-up where the computational homogenization domain with permeability K_1 is split by a blockage of less permeable (K_2) cells. Vertically-directed pressure gradients should theoretically turn flow horizontally and homogenization of the entire domain should aim to preserve this characteristic.

Tag	Set-up	K_1	K_2	K_l
B1	blockage (5x5)	$\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$	$\begin{bmatrix} 0.0101 & 0.01 \\ 0.01 & 0.0101 \end{bmatrix}$	$\begin{bmatrix} 0.8584 & 0.0094 \\ 0.0094 & 0.0660 \end{bmatrix}$
B2	blockage (5x5)	$\begin{bmatrix} 0.8333 & 0 \\ 0 & 0.4394 \end{bmatrix}$	$\begin{bmatrix} 1.0389 & 0 \\ 0 & 0.2678 \end{bmatrix}$	$\begin{bmatrix} 0.8627 & 0 \\ 0 & 0.4025 \end{bmatrix}$
B3	blockage (5x5)	$\begin{bmatrix} 0.5 & 0.2 \\ 0.2 & 0.3 \end{bmatrix}$	$\begin{bmatrix} 0.93 & 0.4 \\ 0.4 & 0.6 \end{bmatrix}$	$\begin{bmatrix} 0.5526 & 0.2154 \\ 0.2154 & 0.3231 \end{bmatrix}$

Table 4: Test set-ups and their corresponding computed homogeneous permeability values. K_1 and K_2 are permeabilities of the blockage domain (Figure 8). K_l is the calculated homogenized permeability matrix. Diagonal cases reproduce analytically expected results to machine precision. All values are rounded to 4 decimal places.

Table 4 lists a sample of specific SGS permeability choices and their corresponding homogenized results using the blockage configuration. For diagonal subgrid permeabilities, analytical results are easily recovered to machine precision. The results of the non-diagonal cases, however, require more analysis.

Theoretically, in the Darcy framework, non-zero off-diagonal components of subgrid permeabilities are responsible for converting some part of a unidirectional pressure gradient into an orthogonally directed velocity. This "turning" of the flow (which in this case is not meant in a time-dependent sense) is an important observed and confirmed characteristic of an anisotropic porous medium (e.g. soil or layered media [1, 22]) and therefore homogenization should preserve such a behaviour. We do observe that in non-diagonal cases (B1 and B2, Table 4), homogenized permeability matrices successfully preserve the presence of off-diagonal components. We also test a fully laminate configuration (Figure 9) and recover results of a similar nature to the blockage substructure.

Furthermore, say we attempt to represent a pure solid/fluid system as a porous medium; a perfect fluid will be represented by the identity matrix while a solid must be represented by a permeability matrix that properly penalizes pressure-incited Darcy flow in a way that

reflects the physical nature of a completely impermeable blockage. Take for example case B1 in Table 4, where we choose to define a solid blockage by a full matrix of small coefficients. We would expect a vertical pressure gradient to allow little to no flow in the vertical direction (as it should be blocked) and we would expect almost unpenalized flow in the horizontal direction provided a horizontal pressure gradient. We should also expect that the blockage may turn a vertical pressure gradient into a horizontal velocity - which would be indicated by an off-diagonal component. Looking at the calculated effective permeability components for case B1 does in fact reflect these effects.

Though the requirement of symmetry does reduce flexibility in capturing turning capabilities of some media, the maintenance of a non-diagonal structure through homogenization is a positive result regardless. For future testing, there may also be ways to homogenize based on a permeability that has been transformed into a diagonal matrix through the finding of its principal axes [6], which could potentially allow for a relaxation of the symmetry requirement.

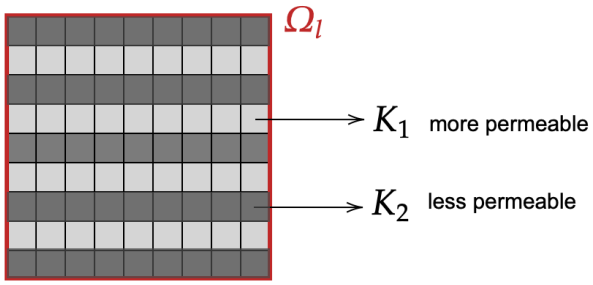

Figure 9: Laminate set-up where the domain is filled with bands of alternate permeability. Homogenized permeability matrix components should follow the arithmetic mean in the band-parallel direction and the harmonic mean in the perpendicular direction.

4.2 Extension of testing to three dimensions

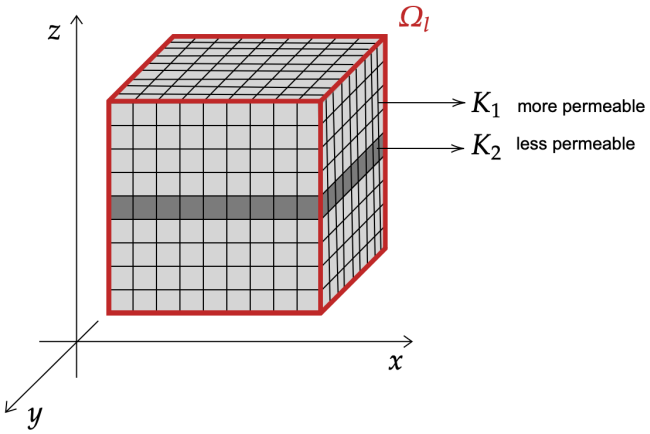

Figure 10: Three-dimensional laminated blockage test case.

In order to confirm that the modification of the code to account for three-dimensional structure calculates permeability correctly, we expand the laminated blockage into a third dimension. If the previous two-dimensional case represented the xz plane, the computational domain extends in the y dimensions with layers of the same laminated structure (Figure 10). Given that the domain remains cubic, we expect flow behaviour through the newly created yz faces to mimic

the flow through the xz faces. In other words, there should be permeability symmetry in the two coordinate directions that display physical symmetry. This behaviour is confirmed when all blockage and laminate configurations described in the previous section are expanded into a third dimension with the above symmetry.

4.3 A note about using conservation of total dissipated power as a validation criterion

At the beginning of this study, we thought that we might be able to validate homogenization results by using a measure of dissipated power over the computational domain to compare fine-scale and coarsened representations of the Darcy equation. One might think that along with preserving the flow characteristics of a porous medium represented with a fine grid, a homogenized block should also conserve the total dissipated power (TDP). TDP over a domain Ω takes the form

$$TDP = \int_\Omega u^T \nabla P \, \mathrm{d}\Omega = \int_\Omega u^T K^{-1} u \, \mathrm{d}\Omega$$

where u is the Darcy velocity. However, any type of averaging is dissipative in nature and having calculated coarse and fine TDP (with proper scaling and equivalent velocities) generated by systems with both diagonal and non-diagonal permeability matrices yielded very poor comparisons. Figure 11 shows probability density function estimates of relative error in coarse and fine TDP for randomly generated blockage, laminate, and Gaussian type permeability matrices (with both diagonal and non-diagonal set-ups) using 500 samples each. Both cases are found to have very large TDP consistency errors: between 90 and 100%. The Gaussian case represents a domain in which every fine-scale permeability matrix was sampled from a normal distribution with a given mean and standard deviation.

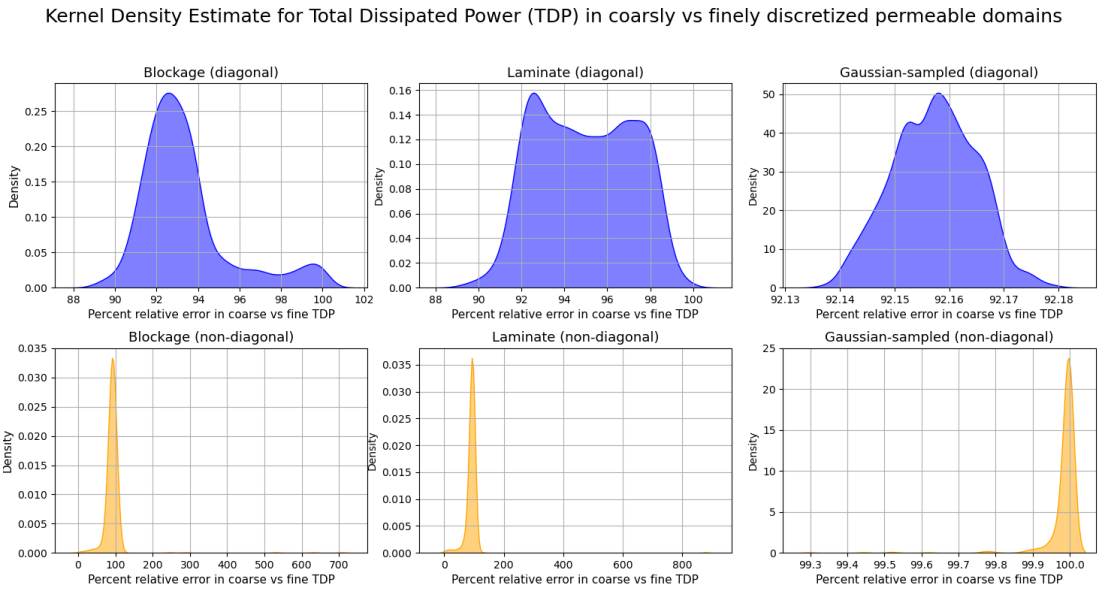

Figure 11: Kernel density estimates for coarse total dissipated power as compared to fine. Run with equivalent velocities.

Though this measure of TDP is not a reputable method of validating homogenized permeabilities, we do go on to use it as a measure of permeability-dissipated energy in the analysis of our shallow water models later in this study, where it proves to be an important contributor to balancing energy levels in systems that reach a steady state.

5 Comparison of PH to a well known intuitive approach

PH is not the only technique known to coarsen simulations of porous media. Given the nature of upscaling, it is not surprising that studies that require the use of a lower resolution often rely on some form of averaging, whether the average is applied to fluxes, gradients, or another variable. Such methods tend to be empirically derived and lack a rigorous mathematical foundation. However, they have practical advantages since they require fewer constraints and assumptions (e.g., periodicity, symmetry, positive-definiteness) and can be easier to implement than a method like PH. To judge how PH performs compared to an averaging approach, we implement the volume averaging method (VA) described in [20] and will run a series of tests. VA also uses the Darcy equation to find a homogenized permeability tensor K_{NSF}. As indicated by the subscript, this method is derived by requiring that the homogenized coarse-scale net surface flux (NSF) matches the fine-scale (SGS) NSF.

5.1 Testing methods

The test cases that we have chosen to compare the PH and VA approaches upon can either be solved analytically or have well established empirical results. The computational domain is an $n \times n$ square grid on which four fine-scale structure configurations are tested: blockage, laminate, Gaussian, and homogeneous.

The homogeneous configuration assigns every grid cell in the domain the same permeability matrix K_1. Therefore, when homogenized the coarse-scale permeability matrix should also be K_1. The blockage configuration is shown in Figure 8. The domain has cells with permeability matrix K_1, except for a single row of cells with a lower permeability matrix K_2. The laminate configuration is shown in Figure 9. The domain alternates between bands of permeability matrices K_1 and K_2. As mentioned earlier, the blockage and the laminated configurations have known analytical solutions: the effective permeability in the parallel and perpendicular directions are exactly the arithmetic and harmonic means respectively [6]. The Gaussian test case samples randomly generated permeability matrices for each fine-scale grid cell from a normal distribution with a given mean and standard deviation. It has been shown to homogenize into an effective permeability matrix whose entries reflect the geometric mean, which is valued slightly less than the arithmetic mean [17].

Blockage configurations were tested with 11×11, 31×31, and 101×101 cells to quantify accuracy in terms of the upscaling factor. These cases are referred to in the analysis of results as low, medium, and high resolutions. The other configurations were tested with 31×31 grids. All configurations were tested with both diagonal and non-diagonal SGS permeability matrices. A full summary of configurations and results is provided in Table 5.

5.2 Results of PH and VA test cases

The volume averaging method is not stable for non-diagonal SGS inputs: the effective permeability components calculated get extremely large. This is indicative of the conjugate gradient method failing to converge. No orientation nor resolution is able to produce reliable results. There would be an argument to modify this method by first diagonalizing the matrix before homogenizing and then transforming back to its original basis were not diagonal trials of PH to outperform VA in both accuracy and computational time by at least an order of magnitude.

All laminated and blockage cases tested with PH produce the expected analytical results within a couple orders of machine precision, and the homogeneous test cases further produce the exact result (reproducing the fine-scale permeability). Higher grid resolutions (greater upscaling ratios) increase accuracy, but also run time. However, all run times (including high resolution non-diagonal) are fast: 0.01 to 0.1 s. It should be noted that PH is by nature limited to grids with odd-numbered side lengths and the method cannot be used for grids of even dimensions.

Figure 12 compares the accuracy and computational time of all configurations for VA and PH, where accuracy is measured against the known analytical results described earlier. Every single VA case ranks below the worst performing PH case in both metrics. In fact, every non-diagonal VA case lies in the worst quadrant: slow and inaccurate (or unstable). In fact, the

highest resolution (101×101) non-diagonal VA case is not even included on the plot as it did not even reach the iteration ceiling without timing out. Configurations which produce unexpected results, or use other methods of inferring accuracy are marked with an asterisk and are detailed in the following section discussing "special cases".

The high computational efficiency of PH compared to VA is an not unexpected. VA must make a conjugate gradient calculation with a complex derivative structure, whereas PH reduces the same system to a linear one via the Lippmann–Schwinger approach and is consequently a much cheaper approach. The same computational complexity may also explain VA's lack of accuracy in diagonal cases, but its instability for non-diagonal cases rules it out as a practically useful method in its current state.

The two methods also differ significantly in their assumptions and enforced conditions, though we find PH's requirements easily met. Although PH assumes periodicity on the fine-scale, this is not necessarily a drawback since our goal is to apply this method to homogenize SGS structure in order to provide better physical models at a fixed computational resolution. In other words, taking permeability matrices that are constant for a single fine-scale cell satisfies the periodicity assumption, which is automatically fulfilled by the nature of the subgrid's $\mathcal{O}(\epsilon)$ structure. The fact that PH can only be applied to grids of odd dimension is no more of a serious limitation on the method. The only requirement that must be treated with more awareness is maintaining a separation of scales.

Because PH lacks an explicit conservation of net surface flux between coarse and fine flows, we look to test its more physical effects on stationary and non-stationary flows modelled by the porous shallow water equations (pSWE). The following sections will ultimately apply periodic homogenization to upscale permeability for the pSWE, where permeability enters via a Brinkman-type penalization term and porosity is set equal to one.

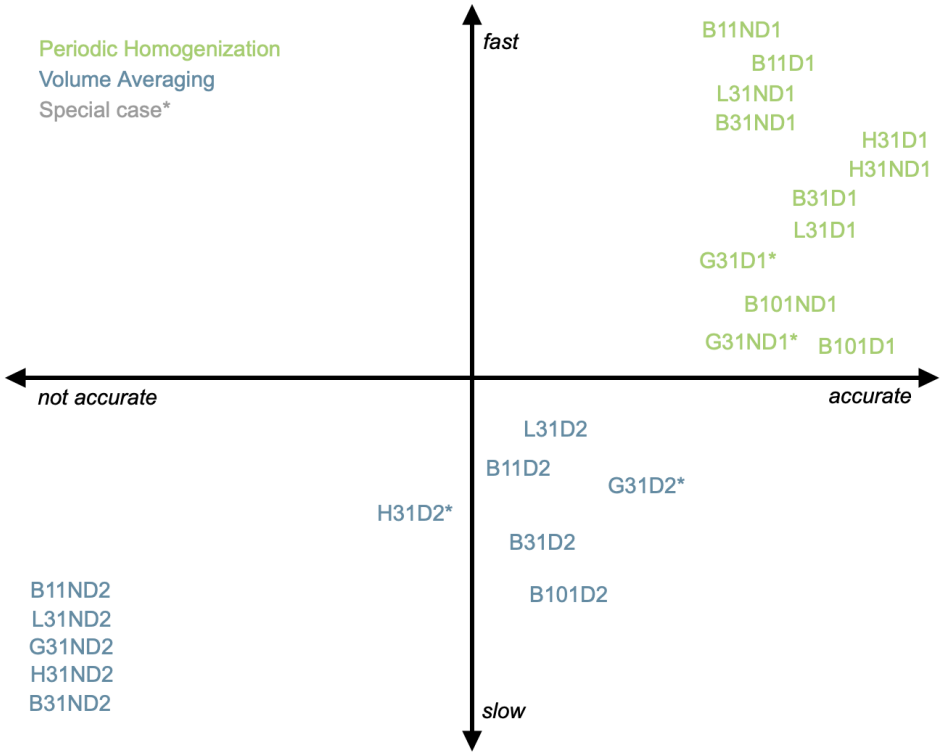

Figure 12: A comparison in speed and accuracy of each case in Table 5 where cases suffixed with 1 refer to PH and cases suffixed with 2 refer to VA. All cases that use PH perform faster than those that use VA and all but one perform more accurately. In most cases, accuracy is measured against corresponding analytical or known solutions, however, cases with asterisks have particular unphysical properties, such as unexpected negative entries or non-zero entries (see §5.2.1).

Tag	Diag	Method	Run Time	K_1	K_2	K_l	
11x11 Blockage							
B11D1	yes	PH	0.00212 s	$\begin{bmatrix} 0.7226 & 0 \\ 0 & 0.2667 \end{bmatrix}$	$\begin{bmatrix} 0.1473 & 0 \\ 0 & 0.4958 \end{bmatrix}$	$\begin{bmatrix} 0.6703 & 0 \\ 0 & 0.2784 \end{bmatrix}$	☺
B11D2	yes	VA	1.4901 s	"	"	$\begin{bmatrix} 0.5943 & 0.0758 \\ 0.0758 & 0.2038 \end{bmatrix}$	😐
B11ND1	no	PH	0.00145 s	$\begin{bmatrix} 0.7226 & 0.4338 \\ 0.4338 & 0.2667 \end{bmatrix}$	$\begin{bmatrix} 0.1473 & 0.1253 \\ 0.1253 & 0.4958 \end{bmatrix}$	$\begin{bmatrix} 0.6537 & 0.4181 \\ 0.4181 & 0.2784 \end{bmatrix}$	☺
B11ND2	no	VA	2.0687 s	"	"	$\begin{bmatrix} 2.6477 & 2.7776 \\ 2.7776 & 2.9138 \end{bmatrix}$	☹
31x31 Blockage							
B31D1	yes	PH	0.00886 s	$\begin{bmatrix} 0.7226 & 0 \\ 0 & 0.2667 \end{bmatrix}$	$\begin{bmatrix} 0.1473 & 0 \\ 0 & 0.4958 \end{bmatrix}$	$\begin{bmatrix} 0.7040 & 0 \\ 0 & 0.2707 \end{bmatrix}$	☺
B31D2	yes	VA	1.6183 s	"	"	$\begin{bmatrix} 0.6328 & 0.0712 \\ 0.0712 & 0.1997 \end{bmatrix}$	😐
B31ND1	no	PH	0.00682 s	$\begin{bmatrix} 0.7226 & 0.4338 \\ 0.4338 & 0.2667 \end{bmatrix}$	$\begin{bmatrix} 0.1473 & 0.1253 \\ 0.1253 & 0.4958 \end{bmatrix}$	$\begin{bmatrix} 0.6980 & 0.4284 \\ 0.4284 & 0.2707 \end{bmatrix}$	☺
B31ND2	no	VA	7.6023 s	"	"	$\begin{bmatrix} 28.56 & 15.99 \\ 15.99 & 8.952 \end{bmatrix}$	☹
101x101 Blockage							
B101D1	yes	PH	0.07777 s	$\begin{bmatrix} 0.7226 & 0 \\ 0 & 0.2667 \end{bmatrix}$	$\begin{bmatrix} 0.1473 & 0 \\ 0 & 0.4958 \end{bmatrix}$	$\begin{bmatrix} 0.7169 & 0 \\ 0 & 0.2679 \end{bmatrix}$	☺
B101D2	yes	VA	2.1700 s	"	"	$\begin{bmatrix} 0.6487 & 0.0682 \\ 0.0682 & 0.1977 \end{bmatrix}$	😐
B101ND1	no	PH	0.06516 s	$\begin{bmatrix} 0.7226 & 0.4338 \\ 0.4338 & 0.2667 \end{bmatrix}$	$\begin{bmatrix} 0.1473 & 0.1253 \\ 0.1253 & 0.4958 \end{bmatrix}$	$\begin{bmatrix} 0.7150 & 0.4321 \\ 0.4321 & 0.2679 \end{bmatrix}$	☺
B101ND2	no	VA	timeout	"	"	N/A	☹
31x31 Laminated							
L31D1	yes	PH	0.00935 s	$\begin{bmatrix} 0.7226 & 0 \\ 0 & 0.2667 \end{bmatrix}$	$\begin{bmatrix} 0.1473 & 0 \\ 0 & 0.4958 \end{bmatrix}$	$\begin{bmatrix} 0.4256 & 0 \\ 0 & 0.3502 \end{bmatrix}$	☺
L31D2	yes	VA	1.440 s	"	"	$\begin{bmatrix} 0.3984 & 0.0635 \\ 0.0635 & 0.2908 \end{bmatrix}$	😐
L31ND1	no	PH	0.00681 s	$\begin{bmatrix} 0.7226 & 0.4338 \\ 0.4338 & 0.2667 \end{bmatrix}$	$\begin{bmatrix} 0.1473 & 0.1253 \\ 0.1253 & 0.4958 \end{bmatrix}$	$\begin{bmatrix} 0.3627 & 0.3212 \\ 0.3213 & 0.3502 \end{bmatrix}$	☺
L31ND2	no	VA	7.469 s	"	"	$\begin{bmatrix} 5.6020 & 7.2797 \\ 7.2797 & 17.312 \end{bmatrix}$	☹
31x31 Gaussian Random							
G31D1	yes	PH	0.01062 s	$\mu = 3,\ \sigma = 0.1$	N/A	$\begin{bmatrix} 3.0041 & -5e-5 \\ -5e-5 & 2.9997 \end{bmatrix}$	☺
G31D2	yes	VA	1.5473 s	$\mu = 3,\ \sigma = 0.1$	N/A	$\begin{bmatrix} 3.00438 & -1e-3 \\ -1e-3 & 2.9974 \end{bmatrix}$	☺
G31ND1	no	PH	0.17131 s	$\mu = 3,\ \sigma = 0.1$	N/A	$\begin{bmatrix} 2.8181 & 2.8187 \\ 2.8187 & 2.8340 \end{bmatrix}$	☺
G31ND2	no	VA	7.528 s	$\mu = 3,\ \sigma = 0.1$	N/A	$\begin{bmatrix} 209.5 & 279.66 \\ 279.66 & 373.29 \end{bmatrix}$	☹
31x31 Homogenous							
H31D1	yes	PH	0.00707 s	$\begin{bmatrix} 0.7226 & 0 \\ 0 & 0.2667 \end{bmatrix}$	N/A	$\begin{bmatrix} 0.7226 & 0 \\ 0 & 0.2667 \end{bmatrix}$	☺
H31D2	yes	VA	1.5966 s	"	"	$\begin{bmatrix} 0.6532 & 0.0694 \\ 0.0694 & 0.1974 \end{bmatrix}$	☹
H31ND1	no	PH	0.00798 s	$\begin{bmatrix} 0.7226 & 0.4338 \\ 0.4338 & 0.2667 \end{bmatrix}$	N/A	$\begin{bmatrix} 0.7226 & 0.4338 \\ 0.4338 & 0.2667 \end{bmatrix}$	☺
H31ND2	no	VA	7.637 s	"	"	$\begin{bmatrix} 15.615 & 10.065 \\ 10.065 & 6.4875 \end{bmatrix}$	☹

Table 5: Comparison of periodic homogenization (PH) and volume averaged gradient (VA) in their performance and calculation of homogenized results for different grid resolutions as well as diagonal and non-diagonal inputs. PH performs significantly faster and more accurately in both diagonal and non-diagonal cases while VA struggles to converge to analytical results results in diagonal cases and blows up in non-diagonal cases. The type of SGS layout is indicated by the first letter of the tag where B is blockage, L is laminated, G is Gaussian, H is homogeneous.

5.2.1 Special cases

Some special cases where accuracy is determined qualitatively are denoted by asterisks when shown in Figure 12. The homogenized permeability matrices for these cases are also given in

Table 5 under K_l.

Case H31D2, which uses the VA on a homogeneous and strictly diagonal SGS structure introduces unexpected off-diagonal components in the homogenized permeability matrix. This is likely an artifact of the method, as homogenizing an already homogeneous fine-scale permeability matrix must reflect exaclty that permeability matrix when upscaled with any any method and any upscaling ratio.

Cases G31D1 and G31D2 (Gaussian diagonal configuration), where SGS permeability matrices are diagonal and sampled from a normal distribution, generate negative off-diagonal components using both methods. In the context of a permeability-based friction applied to the porous SWE (discussed further in following sections), negative off-diagonal components leads to acceleration of the outflow in a perpendicular direction of the inflow. Whether this type of homogenization result is physical requires further testing in the context of a fluid model. We choose to classify these results as accurate, since the diagonal components reflect the mean they were generated from and off-diagonal components remain small. G31D1 (using PH) is classified as slightly more accurate than G31D2 (using VA) due to its components' tighter convergence to the expected average.

Case G31ND1, or the non-diagonal Gaussian substructure upscaled using PH, has homogenized permeability coefficients with values lower than the mean used to generate the subgrid scale. Studies of upscaling statistically isotropic media generally find that the resulting effective permeability tensor elements are smaller than their means due to local geometries that produce bottlenecks [12]. Generally, component values are expected to approach the geometric mean [17]. For this reason, both G31ND1 and the corresponding diagonal cases mentioned above are marked as "observationally" accurate.

Case B101ND2 is not included in Figure 12 since it did not converge in a reasonable time. We therefore conclude that VA is not usable for non-diagonal higher resolution cases.

6 Testing permeability homogenization for the porous shallow water equations

This section focuses on the implementation of PH for homogenizing the two-dimensional porous shallow water equation (pSWE) model with a specified fine-scale permeability configuration. A major benefit for upscaling through permeability is that the homogenization does not have to be done in dynamic time, but rather just once as a pre-processing step. Once the coarse-scale permeability configuration has been obtained, it can then be used in the time-dependent pSWE model. Single and multilayer SWE approximations are the standard for ocean modelling, and it is shown by [10, 11] for the CROCO ocean model that it is relatively easy to extend preexisting software to a pSWE model that includes an adaptation for the handling of permeability and porosity.

6.1 Model configuration

For simplicity we assume that the flow is perfectly porous (i.e., porosity equals one), but with location-dependent permeability matrices. In other words, we allow all the penalization to come from the permeability term itself. Because we assume the fluid is perfectly porous, the mass conservation equation is unchanged from the non-porous SWE,

$$\frac{\partial \eta}{\partial t} + \boldsymbol{\nabla} \cdot hu = 0, \tag{6}$$

where u is the velocity vector and $h(x,t) = H + \eta(x,t)$ is the fluid column height. H is the constant mean depth and η is the free surface perturbation. The momentum equation for the pSWE has an extra Brinkman-type penalization term ($-\mathbb{1}K^{-1}(hu)$) compared to the SWE to include the effect of permeability:

$$\frac{\partial hu}{\partial t} = -\boldsymbol{\nabla} \cdot (hu \otimes u) - gh\boldsymbol{\nabla}\eta - f\hat{z} \times hu - F_{fric} + \tau + \boldsymbol{\nabla} \cdot (\nu h \boldsymbol{\nabla} u) - \mathbb{1}K^{-1}(hu), \tag{7}$$

where f is the Coriolis parameter, F_{fric} is an additional bottom friction parameter, $\tau = (\tau_x, 0)$ is the wind stress, and ν is the kinematic viscosity. $K(x, y)$ is the position-dependent 2×2 permeability matrix characterizing the material. The indicator function $\mathbb{1}(x, y)$ ensures that permeability-induced "friction" is not added where the material is a perfect fluid (i.e., where $K = K^{-1} = I$),

$$\mathbb{1}(x, y) = \begin{cases} 1, & \text{solid or semi-permeable} \\ 0, & \text{perfect fluid} \end{cases} .$$

The equations are time-stepped using an RK4 scheme and spatially discretized using a first order finite difference method. Since finite differences are used to determine velocity derivatives, the computational domain is spatially augmented by an extra grid point in all directions to allow for the calculation of gradients at the boundaries. Velocity vectors span this augmented domain with a correspondingly augmented amount of entries that represent measurements at cell edges. This discretization follows the staggered Arakawa C-grid layout commonly used in ocean modelling (Figure 13), where scalars are located at cell centres and fluxes at cell edges. If west/east periodic boundary conditions are used (which this study does), one ghost cell is added to each lateral boundary to ensure continuity. The north/south boundaries are closed and can either be set to free-slip or no-slip boundary conditions.

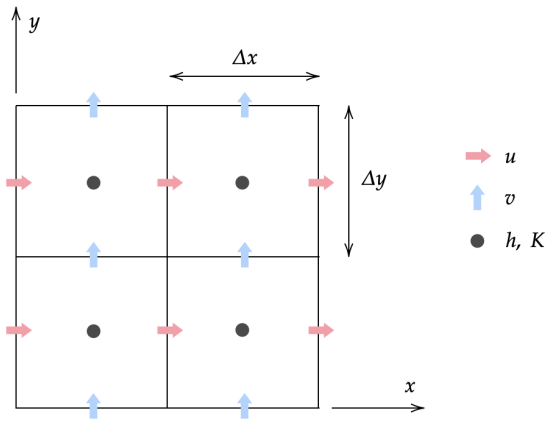

Figure 13: An Arakawa C-grid layout, where velocities are sampled at cell edges and free surface height and permeability is measured at cell centres.

A constant mean velocity u_∞ is imposed after each iteration of RK4 via a split-step approach,

$$u^*_{ij} = u^n_{ij} - \overline{u^n_{ij}} + u_\infty$$

where $\overline{u^n_{ij}}$ is the calculated spatial mean at time t_n and u_∞ is the $x-$directed imposed mean velocity. We take a computational domain that resembles a longer than wide channel with dimensions $[0, 4] \times [0, 1]$ to allow the streamwise flow to develop. The time step Δt is determined by the advective CFL criterion,

$$\Delta t = \frac{CFL_{adv}}{\left\|u_\infty + \sqrt{gH}\right\|\left(\frac{1}{\Delta x} + \frac{1}{\Delta y}\right)}, \quad CFL_{adv} \leq 1. \tag{8}$$

Along with enforcing the advective CFL criteron, we also ensure that the time step satisfies diffusive stability conditions,

$$\nu \Delta t \left(\frac{1}{\Delta x} + \frac{1}{\Delta y}\right) \leq \frac{1}{4}. \tag{9}$$

6.2 Testing model features

We begin by validating the SWE simulations (i.e., without permeability). We first test the case with zero mean velocity and two initial conditions: initial perturbations to the free surface

(standing wave and Gaussian droplet). The initial conditions diffuse as expected. Both periodic and closed boundary conditions were tested and perform as expected as well.

Adding the mean flow, free-slip boundary conditions generates a uniform flow while no-slip boundary conditions generate the expected Poiseuille flow. All tests confirm the numerical conservation of both mass and momentum.

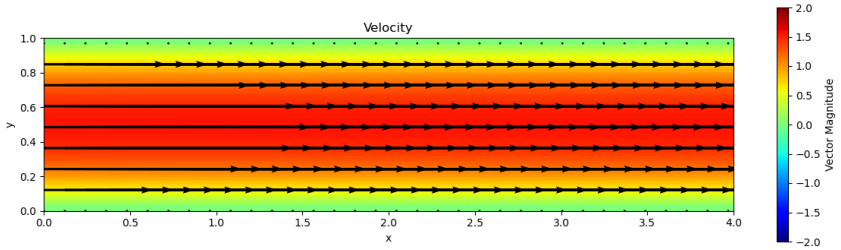

Figure 14: No-slip boundary conditions create a Poiseuille flow that displays a velocity gradient towards the centre of the channel with the typical bullet-shaped velocity structure.

6.3 Permeability distribution for the porous SWE

Permeability is included by supplementing the $n_x \times n_y$ data structure with an $n_x \times n_y \times 2 \times 2$ matrix that assigns a 2×2 permeability matrix to each cell in the computational domain. Since the code extends the domain via ghost cells, the permeability matrix is extended in the same way by mirroring edge permeabilities. Permeabilities are input as cell-centre quantities where, due to the fact that the differential equation (7) is evaluated at cell-edges, cell-edge permeabilities need to be found by averaging neighbouring matrices. This is a preliminary practice and we encourage that further developments on this method give more thought to how permeabilities should be sampled on staggered grids.

Permeability configurations are chosen using a fine-scale structure that has practical significance to real ocean models: a marshy domain and a narrow channel, which each respectively represent regions like the bayous of Louisiana and the Strait of Gibraltar. Note that both of these regions are extremely challenging to model using conventional methods.

Figure 15: Left: LaBranche Wetlands of Lousiana. Right: The Strait of Gibraltar, connecting the Atlantic Ocean to the Mediterranean Sea.

Lastly, in upscaled pSWE simulations we rely only on the PH method to homogenize fine-scale permeability structure given its far superior test results to VA from the previous section.

6.3.1 Fine-scale permeability configurations

The two permeability configurations considered (labelled “marsh" and a “tunnel") impose their fine-scale substructure by characterizing each coordinate as having one of two sample perme-

ability matrices K_1 or K_2 (Figure 16). These sample permeability matrices can either represent solid, fluid, or semi-permeable media depending on their values. In terms of practical application, the solid/fluid configuration better represents cases like bathymetry or coastline modelling while the semi-permeable configuration is a good model for marsh-like regions. Since PH can homogenize full matrices (i.e., with non-zero off-diagonal entries), both diagonal and non-diagonal permeability matrices are tested in order to investigate the method's potential.

The permeability matrix for a perfect fluid is described by the identity matrix, while a perfect solid is described either by a full or diagonal matrix of some small constant ϵ. Computationally, permeability $\epsilon = 0.01$ is used for solids, where the diagonal entries are set slightly larger than the off-diagonal entries to ensure the permeability matrix is non-singular. Semi-permeable matrix components are set to small values bounded between 0 and 1. An overview of the sample permeability matrices used is given in Table 6.

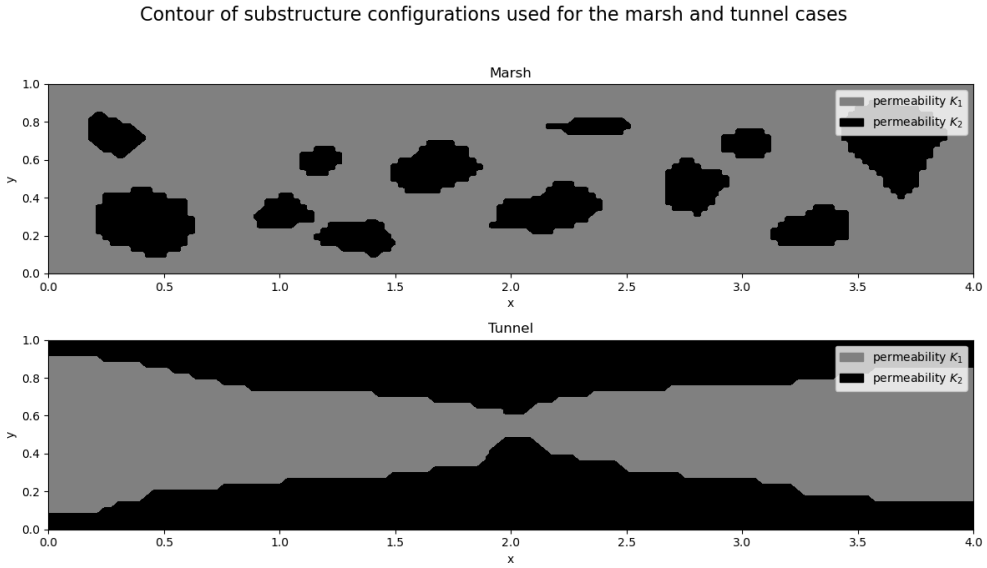

Figure 16: Masks used to generate the fine scale permeability in the marsh and channel cases. In the solid/fluid configurations, permeability matrices K_1 and K_2 represent fluid and solid matter respectively. In the semi-permeable configurations, K_1 and K_2 are non-equal intermediary permeability matrices. Exact permeability matrices can be found in Table 6.

Configuration type	K_1	K_2
fluid/solid non-diagonal	$\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$	$\begin{bmatrix} 0.0101 & 0.01 \\ 0.01 & 0.0101 \end{bmatrix}$
fluid/solid diagonal	$\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$	$\begin{bmatrix} 0.01 & 0 \\ 0 & 0.01 \end{bmatrix}$
semi-permeable non-diagonal	$\begin{bmatrix} 0.1473 & 0.1253 \\ 0.1253 & 0.4958 \end{bmatrix}$	$\begin{bmatrix} 0.7226 & 0.4338 \\ 0.4338 & 0.2667 \end{bmatrix}$
semi-permeable diagonal	$\begin{bmatrix} 0.1473 & 0 \\ 0 & 0.4958 \end{bmatrix}$	$\begin{bmatrix} 0.7226 & 0 \\ 0 & 0.2667 \end{bmatrix}$

Table 6: Permeability matrices for each configuration type, used for both the marsh and channel configuration.

6.3.2 Simulation parameters

The ultimate goal of this project is to compare coarse simulations with homogenized permeability coefficients to corresponding fine-scale simulations averaged to the coarse resolution. An ideal homogenization method would produce the same results as the averaged fine scale simulation.

Therefore, the computational domain for both the fine and coarse-scale models is the same, but discretized with different grid resolutions.

The coarse scale consists of grid cells that represent 11×11 blocks in the fine-scale simulation (i.e., a upscaling ratio of 11). Each 11×11 block of fine-scale substructure has a corresponding $11 \times 11 \times 2 \times 2$ permeability matrix, which is upscaled to find an effective 2×2 coarse-scale permeability matrix via PH.

Feature	Code object	Value
Domain and time		
Discretization (fine-scale)	`(nx,ny)`	$(264, 66)$
Discretization (coarse-scale)	`(nx,ny)`	$(24, 6)$
Length-scale	`(Lx,Ly)`	$(4, 1)$
Integration time	`total_time`	10
Physical parameters		
Coriolis initializer	`f0`	7×10^{-7}
Coriolis multiplier	`beta`	2×10^{-11}
Gravity	`g`	1
Mean free surface height	`H`	1
Constant wind forcing	`tau0`	0.015
Bottom friction	`cb`	5×10^{-7}
Density of seawater	`rho0`	1000
Numerical viscosity	`mu`	0.001
Mean horizontal velocity	`mean_u`	0.1
Numerical parameters		
Advective CFL condition	`cfl`	0.99
Boundary conditions		
Periodic West/East boundary	`bc`	1
Free-slip	`noslip`	0
Land parameters		
No mask	`masking`	`false`
Permeability on	`permeability`	`true`

Table 7: Necessary parameter values to reproduce the simulations in the following sections.

These simulations also ensure that the time step is chosen to satisfy numerical stability. When the mean velocity is set to $u_\infty = 0.1$, the Reynolds number $Re = \frac{uL}{\nu}$ corresponding to the domain width is $Re = 100$. Note that this number obviously varies depending on the local flow velocity and size of substructure features like the blockages in the marsh case or the narrowest point of the tunnel case. The following test cases are all performed using closed, free-slip north/south boundary conditions and periodic west/east boundary conditions.

6.4 Understanding how flow interacts with permeability-defined substructure

Typically, solids in fluid models are defined either by masking or via the boundary geometry. This allows for no-slip or free-slip conditions to be set explicitly at edges where solid and fluid meet. In addition to the application of the porous shallow water equations (pSWE) to the typically-tested semi-permeable medium, we want to explore if this method extends to allow for the physically realistic modelling of flow around completely solid blockages, and whether diagonal or full solid-representing matrices accomplish this task better. This is obviously a necessary intermediary for the ultimate use of homogenization on pure solid/fluid domains. Therefore, in order to establish permeability as a reasonable means of defining bluff bodies, we aim to use it to reproduce telltale dynamic features that specific boundary geometries will impose on flow.

The type of flow expected within any domain is dependent on the Reynolds number. Larger objects (high L), larger speeds, and smaller viscosities can all take the flow from a laminar to a

turbulent regime. Low Reynolds numbers around each blockage should produce a creeping flow, where higher Reynolds numbers should exhibit vortex shedding. Though the resolutions tested are not high enough to handle the conditions that produce vortex shedding, low to moderate Reynolds numbers (∼10-50) are still large enough to generate adverse pressure gradients on the downstream side of a solid body that result in the steady recirculation of two symmetrically placed vortices on each side of the wake [21].

A good first test is to insert a single disk-shaped bluff body at the beginning of a channel domain that is otherwise filled with pure fluid. The matrices used for each permeability are those from the fluid/solid non-diagonal configuration type given by Table 6. A mean flow carries fluid from left to right in the domain, where it interacts with the bluff body and splits around it. The low to moderate Reynolds numbers $Re = 97.34$ corresponding to the domain width (spanwise) and $Re = 32.45$ corresponding to the diameter of the blockage indeed generate opposing vortices on each side of the blockage (Figure 17). The flow splits into regions of opposing vorticity on either side of the body that expand in the streamwise direction with time.

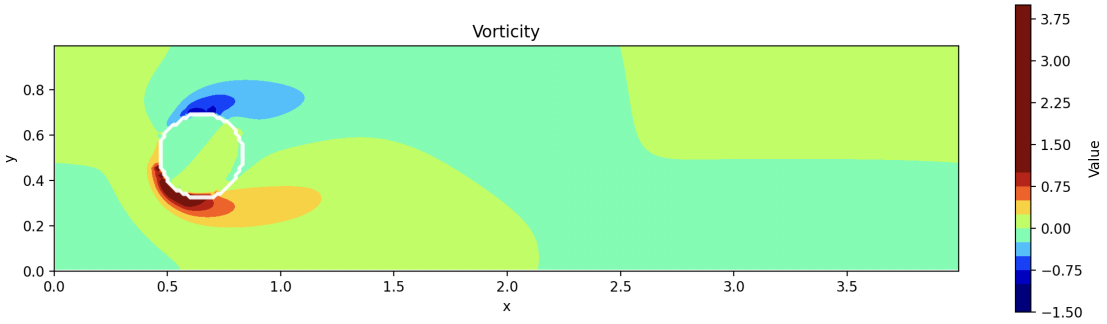

Figure 17: Vorticity plot of a mean flow interacting with a bluff body defined by solely permeability (no boundary or mask) with low to moderate Reynolds number Re = 97.34 and resulting generated opposing vortices (time = 26.23).

6.4.1 Marsh substructure tests

Using the marsh substructure from Figure 16, the fluid/solid configuration from Table 6, and the simulation parameters from Table 7, a constant mean flow is applied to a channel with solid blockages. The solid obstacles divert flow to open regions in both diagonal (Figure 18) and non-diagonal cases (Figure 19). In this configuration, a steady state is reached well before the integral time-scale (the time-scale for the mean flow to traverse the entire channel). There are

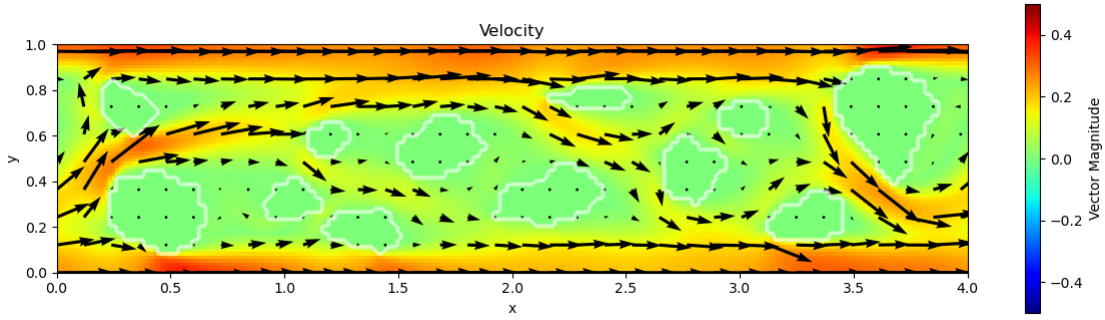

Figure 18: Diagonal permeability matrix case: vector velocity plot for the solid/fluid diagonal marsh configuration with a mean flow at time = 15. The Reynolds number associated with the spanwise length-scale is Re = 395.90, just slightly lower than the Reynolds number of its non-diagonal counterpart below. Since the open pathways and blockages are smaller than the lateral length-scale, this Reynolds number indicates laminar flow where viscosity dominates.

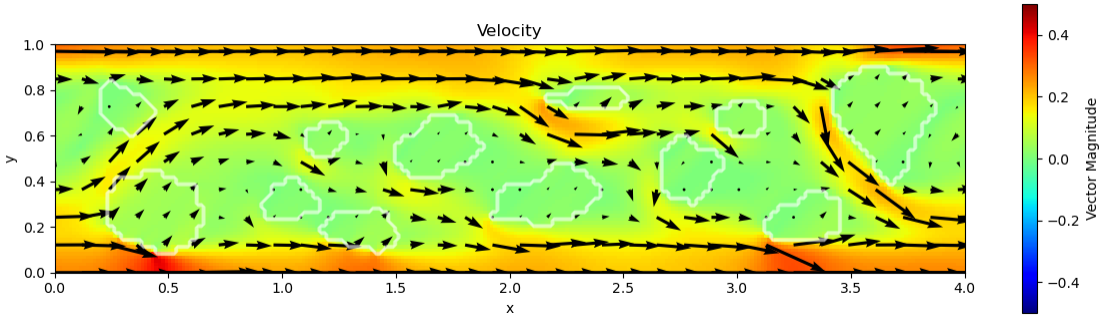

Figure 19: Non-diagonal permeability matrix case: vector velocity plot for the solid/fluid marsh configuration with a mean flow at time = 15. The flow velocity approaches a steady state around permeability-defined obstacles. The Reynolds number associated with the spanwise length-scale is Re = 403.62, i.e. viscosity dominates and the flow is laminar.

few differences in the flow between the diagonal and non-diagonal cases. The main one being that the non-diagonal seems to suffer to keep zero velocities within all solid regions. Seen most clearly in the lower-left corner of the velocity plots Figures 18 and 19, the allowance of non-zero velocities to move into these regions essentially reduces how strongly a solid body will divert flow. In turn, the velocity gradients that are created by the flow "squeezing" between blockages are generally stronger in the diagonal case, even though the non-diagonal case generates stronger gradients between solid blockages and the lower boundary. Otherwise, both diagonal and non-diagonal cases form the expected low velocity regions downstream of all blockages. There are no characteristics of these voids (e.g. length, circulation around) that are consistent enough between either of the two cases to be of note.

In the context of the momentum equation (7), diagonally and non-diagonally defined permeability matrices produce qualitatively different effects, much due to the structure of their inverses (Table 8), which apply a friction to the flow via the term $-K^{-1}(hu)$. A scalar matrix permeability (diagonal with equal components) will decelerate both streamwise inflow u in the $x-$direction and spanwise inflow v in the $y-$direction equally and independently (e.g., $\partial_t hv = \ldots - K_{22}^{-1}(hv)$). A non-diagonal permeability will also decelerate the velocity components u and v in their respective parallel directions, but the off-diagonal components will accelerate flow in their respective perpendicular directions (e.g., $\partial_t hv = \ldots - K_{21}^{-1}(hu)$, where $K_{21}^{-1} = K_{12}^{-1} < 0$).

	K	K^{-1}
non-diagonal solid	$\begin{bmatrix} 0.0101 & 0.01 \\ 0.01 & 0.0101 \end{bmatrix}$	$\begin{bmatrix} 5024.87 & -4975.12 \\ -4975.12 & 5024.87 \end{bmatrix}$
diagonal solid	$\begin{bmatrix} 0.01 & 0 \\ 0 & 0.01 \end{bmatrix}$	$\begin{bmatrix} 100 & 0 \\ 0 & 100 \end{bmatrix}$

Table 8: Non-diagonal and diagonal choices of permeability matrices for solid material and their inverses. Applied as a Brinkman friction $-K^{-1}(hu)$, the diagonal case decelerates both streamwise and spanwise velocities without any changing of direction. The non-diagonal case accelerates outflows that are perpendicular to the incumbent inflow but removes slightly more velocity from parallel outflow.

The deceleration of the velocity in the parallel direction causes a mass build-up and therefore a pressure gradient on the upstream edge of the solid body, leading fluid flow towards the lower pressure areas around the blockage (Figure 20). In the non-diagonal case, the acceleration in the perpendicular direction should theoretically aid in the turning of the flow as well. These turning effects are well present in the marsh tests (Figures 18, 19).

In SWE models that represent solid regions by boundary geometry or masking, the velocity gradient resulting from no-slip conditions applied to these regions creates a shear flow at the

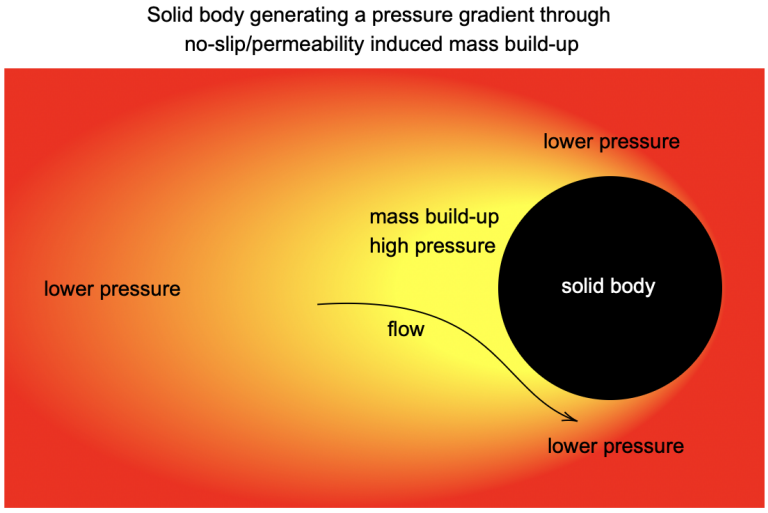

Figure 20: However flow is slowed at the edge of a solid region, either by a penalizing permeability or a no-slip condition, will cause a mass build-up and consequently, a pressure gradient that redirects flow around the bluff body.

viscous boundary layer that generates vorticity, especially at higher Reynolds numbers. This vorticity not only contributes to turning the flow around a blockage, but is also the reason vortices are formed behind bluff bodies. Since, in this simulation, the no-slip boundary condition is enforced implicitly using a Brinkman friction penalization, the permeability matrix is solely responsible for replicating this behaviour.

Permeability does the job as expected, as shown by the vorticity plot of the marsh configuration (Figure 21). Similar vortices are generated for both the diagonal and non-diagonal case, due to the Reynolds numbers for all cases being within a transitional regime. Even when subjected to higher mean velocities, vortices do not shed but form a steady state.

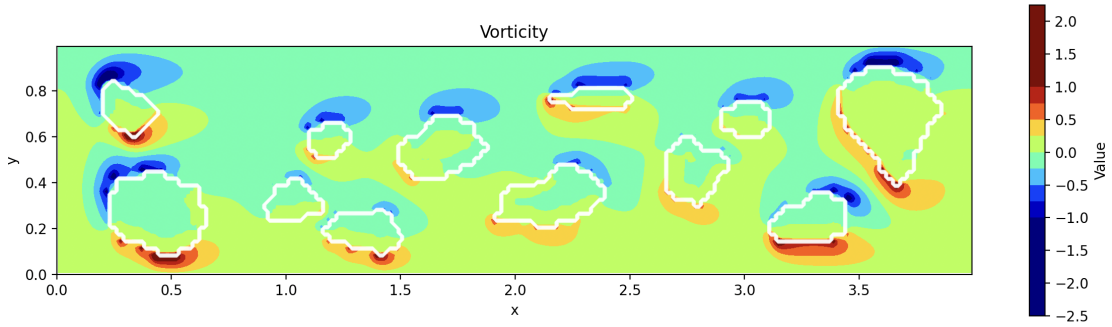

Figure 21: Vorticity plot of the solid/fluid diagonal marsh configuration with a gentle mean flow ($u_\infty = 0.02$). Vortices form on top and bottom edges of solid permeability-defined blockages. The Reynolds number corresponding to the spanwise length-scale is Re = 76.47 (time = 22.5). The non-diagonal case displays a similar distribution of vortices and an equal Reynolds number.

In addition to solid/fluid permeabilities, semi-permeable input matrices were also considered. The flow was diverted through blockages similarly to the solid/fluid case, but in the semi-permeable case a small amount of flow being is able (and expected) to traverse the blockage. Vorticity followed a similar structure to the solid/fluid cases but at a much weaker level. The non-diagonal semi-permeable test cases suffered from stability issues when the same mean velocity as the other cases was used and therefore the semi-permeable runs documented in Table 9 were performed with half the mean velocity (and consequently half the starting Reynolds

number) in order to allow each system to comparably reach a steady state. We find that models initialized with semi-permeable substructure produce flows with lower Reynolds numbers than their solid/fluid counterparts.

	solid/fluid (homogenized)	solid/fluid (fine-scale)	semi-permeable (homogenized)	semi-permeable (fine-scale)
diagonal	238.76	403.33	57.06*	85.92*
non-diagonal	366.96	404.36	109.72*	235.53*

Table 9: Steady-state (t = 30) Reynolds numbers of simulations with different permeability types within the marsh substructure configuration. Each simulation starts with Re = 100 (except for starred cases that start off with Re = 50 via a reduced mean velocity for stability). Increasing resolution or adding off-diagonal elements to the permeability structure starts to transition flows towards a turbulent regime.

6.4.2 Aside: addressing the instability of non-diagonal semi-permeable cases

Observing a velocity flow plot of the unstable case before its blow-up indicates an intersection of large velocity build-ups in the lesser permeable region, which could potentially be catapulted into instability due to the specific choice of permeability matrix or the blockage layout. Table 10 compiles the stiffness ratios of the permeability matrices in Table 6, where we see that the choice of non-diagonal semi-permeable and solid/fluid matrices will actually generate a stiff system, which will cause instability if the time step is not small enough. Since the permeability matrices themselves do not influence the CFL-determined time step, this is a likely outcome, especially in the stiffer semi-permeable case.

We can even get more specific and look at the absolute stability limit for the time stepping scheme (RK4), which is given by

$$z \in [-2.785, 0]$$

where $z = \lambda \Delta t$ for largest eigenvalue λ and time step Δt [8]. However, the largest eigenvalue for the semi-permeable non-diagonal case is $\lambda = -217$, meaning that any time step smaller than $\Delta t = 0.013$ should be stable. Since the time stepping in fine-scale simulations takes $\Delta t = 0.0075 < 0.013$, this configuration lies within the stability limit for RK4. Interestingly, the largest eigenvalue for the solid/fluid non-diagonal configuration is $\lambda = -10000$, which should fail the stability criterion yet does not crash. In the future, more exploration should be done on how the permeability-induced evolution time-scales affect the stability of the scheme since this level of analysis is not enough to determine a clear conclusion.

Configuration type	K_1 approx. stiffness	K_2 approx. stiffness
fluid/solid non-diagonal	1	201
fluid/solid diagonal	1	1
semi-permeable non-diagonal	5	214
semi-permeable diagonal	3	3

Table 10: Approximate stiffness ratios of the permeability matrices in Table 6. The non-diagonal "less permeable" matrices introduce stiffness which may contribute to instability, especially since this is not taken into account for the time-stepping criterion (which is only based on CFL conditions).

6.4.3 Tunnel set-up

The tunnel set-up (Figure 16) was specifically defined to mimic a land/water system and therefore we test only the solid/fluid configuration for this substructure. In this case, we expect a mean flow to reflect no-slip boundary conditions on the top and bottom walls of the channel where the streamwise velocity exhibits a gradient towards the centre of the channel, not unlike the Poiseiulle flow in Figure 14. As the flow squeezes through the narrowest point in the channel,

it should speed up. Furthermore, since the channel tightens more steeply on the bottom wall upstream and opens wider on the top wall downstream, the flow should accelerate in the positive $xy-$direction. The solid regions of the substructure should display zero velocity.

Both cases exhibit the expected speeding up directly downstream of the narrowed region of the tunnel as well as the general gradient towards a centrally dominant velocity. The diagonal case (Figure 22) ubiquitously keeps solid areas at a zero velocity, where the non-diagonal case (Figure 23) displays two small problem regions doing so: directly at the channel narrowing and slightly downstream. In the diagonal case, the flow maintains a smoother, mostly central pattern. This promotes two clear regions of opposing recirculation at the widest area of the channel (which are actually next to each other due to periodic boundary conditions). The non-diagonal case behaves differently, where flow seems to bounce between the top and bottom channel walls as it evolves. Both simulations have much larger Reynolds numbers than the marsh cases which indicate more turbulent flow (though not quite past the laminar regime yet). The non-diagonal case has a comparable Reynolds number to the diagonal case, unlike the marsh configurations. Like the marsh case, however, homogenized simulations display much lower Reynolds numbers (Table 11).

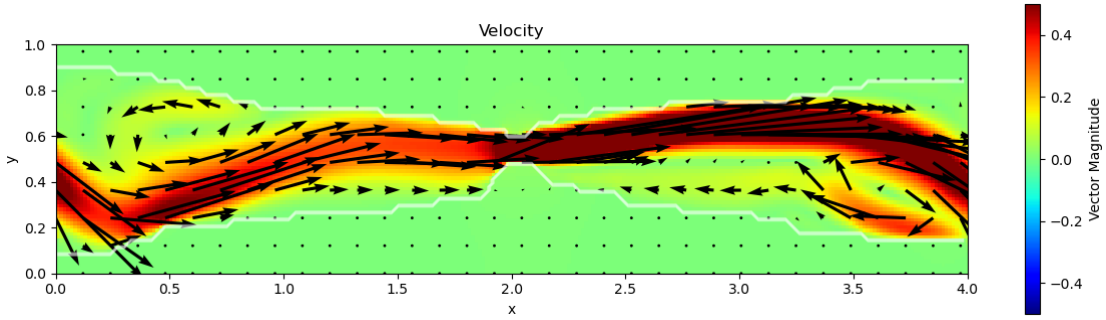

Figure 22: Diagonal permeability matrix case: vector velocity plot for the solid/fluid tunnel configuration with a mean flow at time = 22.5. The flow velocity speeds up on downstream side of a narrowing gap and does not approach a steady state when left to evolve. The Reynolds number associated with the spanwise length-scale is Re = 1227.11, which is technically still laminar but more turbulent than the marsh substructure.

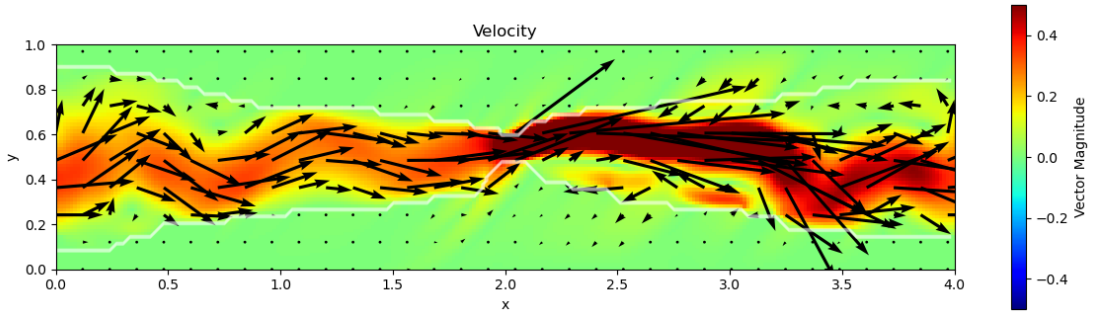

Figure 23: Non-diagonal permeability matrix case: vector velocity plot for the solid/fluid tunnel configuration with a mean flow at time = 22.5. The flow velocity exhibits a similar overall movement as the diagonal case above but seems to bounce between top and bottom walls rather than conforming more smoothly to a central flow. The Reynolds number associated with the spanwise length-scale is Re = 1179.99, slightly higher than the diagonal case.

	solid/fluid (homogenized)	solid/fluid (fine-scale)
diagonal	377.07	1223.97
non-diagonal	298.42	1227.77

Table 11: Reynolds numbers of homogenized and fine-scale simulations at t = 30 within the tunnel substructure configuration. Each simulation starts with Re = 100. Diagonal cases have higher Reynolds numbers than non-diagonal cases, which is opposite in behaviour to the marsh case.

6.4.4 Should non-diagonal permeability matrices be ruled out

Analytically, we can better understand the difference between the contribution of a diagonal and a non-diagonal permeability matrix by the simplified differential equation

$$\partial_t \begin{bmatrix} u \\ v \end{bmatrix} = -K_s^{-1} \begin{bmatrix} u \\ v \end{bmatrix}$$

which penalizes velocity through the solid representing permeability matrix K_s. In the case of a diagonal K_s, K_s^{-1} would just be a diagonal matrix populated with the inverses of the original components. If those components are equal, like they are in the solid diagonal cases in this paper, this penalization becomes equivalent to a bottom friction, where the scalar friction coefficient applied is tuned based on the location within the domain. If an initial velocity is applied solely in one direction (e.g. setting $u(0) = u_0$ and $v(0) = 0$) with the diagonal components of K_s being some small constant α, the solution to the differential equation is

$$u(t) = u_0 e^{-t/\alpha}$$
$$v(t) = 0$$

for which $u(t)$ decays exponentially and v is unchanged at 0, approximating a no-slip boundary condition. In the diagonal case, the decay of the velocity components are proportional only to themselves.

In order to understand the additional contribution off-diagonal components add to a solution, we can also solve for the flow in a sample non-diagonal case subject to the same initial conditions. Starting with the full permeability matrix K_s populated with small constant entries $\beta \ll \alpha \ll 1$ and eigenvalues $\lambda_{1,2}$

$$K_s = \begin{bmatrix} \alpha & \beta \\ \beta & \alpha \end{bmatrix}, \quad \lambda_{1,2} = \alpha \pm \beta$$

we see quickly how this system might become stiff. The more comparable the constants α and β, the larger the ratio between the eigenvalues and the stiffer the system. Continuing to solve, the eigenvectors can be assembled into a similarity transform matrix

$$P = P^{-1} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}$$

to diagonalize the system such that

$$K_s^{-1} = P D^{-1} P^{-1},$$

where D^{-1} is the inverse of the diagonal matrix of eigenvalues $\lambda_{1,2}$. If a coordinate transform to vector velocity z is done in the appropriate basis with

$$z = P^{-1} \begin{bmatrix} u \\ v \end{bmatrix},$$

the solution to the transformed differential equation

$$\partial_t z = D^{-1} z$$

is given by the flow

$$z(t) = \begin{bmatrix} z_1(0)e^{-\lambda_1 t} \\ z_2(0)e^{-\lambda_2 t} \end{bmatrix}.$$

Returned to the original coordinates and subject to the same initial conditions $u(0) = u_0, v(0) = 0$, the final solution is becomes

$$\begin{bmatrix} u \\ v \end{bmatrix} = \frac{1}{2} u_0 \begin{bmatrix} e^{-(\alpha+\beta)^{-1}t} + e^{-(\alpha-\beta)^{-1}t} \\ e^{-(\alpha+\beta)^{-1}t} - e^{-(\alpha-\beta)^{-1}t} \end{bmatrix}.$$

In the limit $\beta \ll \alpha$, the solution for u is clear and decaying

$$u(t) \approx u_0 e^{-t/\alpha},$$

whereas the limit of $v(t)$ requires a Taylor expansion:

$$\begin{aligned} v(t) &= \frac{1}{2} u_0 e^{-t/\alpha} \left(e^{\left(1+\frac{\beta}{\alpha}\right)^{-1}} - e^{\left(1-\frac{\beta}{\alpha}\right)^{-1}} \right) \\ &\approx u_0 e^{-t/\alpha} \left(-\frac{\beta t}{\alpha^2} \right). \end{aligned}$$

In the limit $t \to \infty$, both the diagonal and non-diagonal permeability set-ups lead to u and v decaying to zero and similarly approximate a no-slip boundary condition. During intermediate times, the non-diagonal solution generates a negative and non-zero $v(t)$ through the coupling of velocity components - which is qualitatively very different from the solution $v(t) = 0$ of the diagonal case, for which the rate of change of each component depends only on itself.

Even though both penalizations analytically approach a no-slip boundary condition in the limit $||K^{-1}|| \to \infty$, it is unclear at this point whether diagonal or non-diagonal permeability matrices better reproduce physical results. The instability in the semi-permeable marsh case and non-zero velocities in solid regions that non-diagonal solid permeability matrices seem to generate are obviously not ideal effects. On the other hand, the results generated by the non-diagonal case in the tunnel configuration are qualitatively different from those generated by the diagonal case, indicative of the fact that non-diagonal permeability matrices posses unique flow-turning capabilities. Provided that there may be some cases which benefit from these capabilities, there may be ways to constrain the non-physical effects and instability by testing different choices of non-diagonal matrices and exploring configuration options like "coating" diagonally-defined solid regions in a non-diagonal layer. An ideal next step would be to determine test cases that turn the flow in a known manner (e.g. an impinging jet) or reproduce experimentally recorded trials (e.g. layered soils) with the specific purpose of validating different samples of diagonal and non-diagonal solid-representing permeability matrices and the qualities of their diverted flow.

The fact that this question is not properly resolved, especially in the context of accurate upscaling, means that we will be proceeding with the use of both diagonal and non-diagonal solid-representing permeability matrices. It is possible that non-diagonal permeability is needed to properly capture the unresolved fine-scale structures in an upscaled simulation. It is also not impossible that homogenization produce small, off-diagonal components from certain blocking effects of diagonal fine-scale structure. Therefore, at this point, it would be prudent that any permeability based upscaling method be prepared to handle both diagonal and non-diagonal permeability matrices. We find PH a reliable method of doing so.

6.5 Use of PH to upscale fine-scale pSWE simulations

The process of "homogenizing" a simulation starts with the fine-scale permeability matrix being fed into a method that performs PH over the entire domain. A visualization of the computational domain Ω_L and its discretization is given in Figure 6, where the global scale Ω_L represents the computational domain, Ω_l represents a grid-scale cell, and Ω_ϵ represents a subgrid-scale cell. The process of PH will take in one grid-scale cell, filled with heterogeneous SGS permeabilities, in order to assign it one homogenized permeability matrix through the process described in

§3.1.3. Once this process has been applied to every grid-scale cell, the permeability of the computational domain can be described using a coarsened grid, where each grid cell hosts an effective permeability matrix that is representative of its (now upscaled) SGS structure.

It is important to emphasize that when applying this method to an ocean model, the coarsened length-scale represents the model's computational resolution - which in a global setting could correspond to a length-scale of as large as 10 km. This physical limit obviously leaves many significant coastal features unresolved, that PH could potentially provide a way of capturing through its representation of SGS structure, until resolution capabilities are improved upon. Currently, the effect of unresolved SGS bathymetry is represented by an additional bottom friction term, where a scalar friction coefficient is tuned based on the location, therefore nothing is to be lost by replacing this method with PH, where equal component diagonal cases replicate this exactly.

This is a preliminary study of permeability-based Brinkman penalization and we are not yet at the point where homogenization can be tested in real ocean model. Nevertheless, we can evaluate how closely coarse-scale simulations with homogenized permeability match averaged fine-scale simulations. This section primarily focuses on features generated by fine and coarse-scale simulations and how they compare.

6.5.1 Steady state comparison of averaged fine-scale and homogenized simulations

The homogenized simulations tend to lose detail in the velocity structure as compared to the fine-scale simulations and their averages. However, in order to be classified as productive, one must remember that homogenization doesn't need to fully reproduce fine-scale features, it just has to capture them better (or at least not worse) than a bottom friction. At a minimum, a homogenized flow should preserve key qualitative features of system such as if, when, and how a steady state is reached, as well as a general similarity in spacial velocity distribution. The questions of whether diagonal or non-diagonal matrices are better suited to a domain or produce more realistic flow become redundant in this section, because a successful homogenized flow should approximate its corresponding fine-scale behaviour at a reasonable level of error regardless of permeability choice.

To be able to state that coarse-scale simulations using homogenized permeability are reasonable models for unresolved SGS structures, it is necessary to compare them with suitably averaged corresponding fine-scale simulations. Initially, we will evaluate this for flows that have reached a stationary state. Finding this stationary state (either deterministically or statistically) generally entails allowing the system to evolve to multiple times its integral timescale. However, in our chosen configurations, the mean velocity and measure of Δt imposed are small and the flow converges to a deterministic steady state well before its integral timescale (5333 time-steps at the fine scale or $t = 40$).

A good diagnostic for a steady state is the total kinetic energy

$$KE(t) = \sum_{i,j} \frac{1}{2} u_{ij}^T u_{ij} \Delta x \Delta y,$$

where $u_{ij}(t)$ is the vector velocity at location (i, j). It should be noted that $\Delta A = \Delta x \Delta y$ is a factor that should be treated with care depending on whether velocities summed over are sampled from cell edges or cell centres, whether ghost cells are present, and what resolution samples are taken at. When this is done over both the averaged fine-scale simulation as well as the homogenized coarse-scale simulation, a steady state is seen to be achieved at around $t = 22.5$ (Figure 24). This steady-state is realized at more or less the same time throughout all permeability configurations tested.

In the solid/fluid non-diagonal marsh case, although the kinetic energies of both the homogenized and the fine-scale simulations reach a steady state, the homogenized simulation achieves a slightly lower kinetic energy (Figure 24). Other marsh cases, such as the diagonal solid/fluid, as well as both semi-permeable cases also reach an equilibrium where the homogenized kinetic energy tends to be smaller than its fine-scale counterpart. The more variance there is in velocity measurements in a certain area, the larger the kinetic energy will measure. As such, it is natural

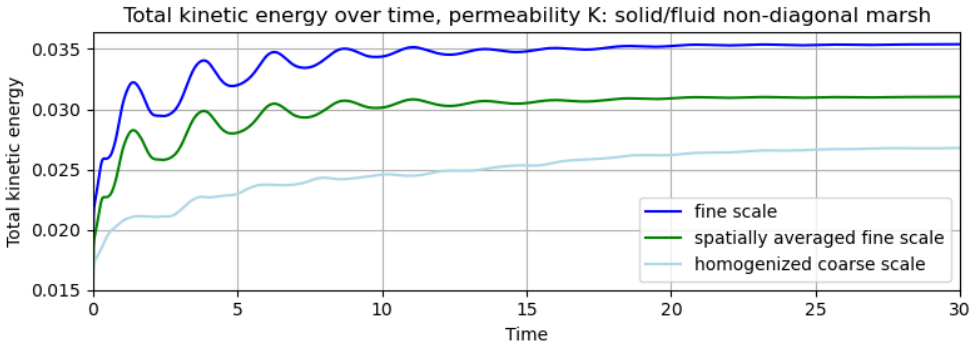

Figure 24: Total kinetic energy measured over the computational domain for the solid/fluid non-diagonal marsh configuration. The system appears to approach a steady state at a time of about 22.5. The spatially averaged fine scale calculates total kinetic energy from velocities averaged over blocks with the same resolution as the coarse scale and naturally compares more closely to coarse-scale measurements.

to see that when fine-scale velocity measurements are spatially averaged to the resolution of the coarse-scale simulation and then squared and summed, the resulting kinetic energy is lower and better compares to that of the coarse scale.

We see this effect again very clearly when we plot kernel density estimates (KDEs) for the steady state velocities averaged over time steps 3000 to 4000 ($t \in [22.5, 30]$). KDEs are essentially a probability density estimate of a normal variable where probabilities are weighted by a distribution (in this case a Gaussian) in order to smooth a (potentially discrete) density plot. We choose to use KDEs to visualize due to their clarity over histograms given the continuous nature of the velocity variable.

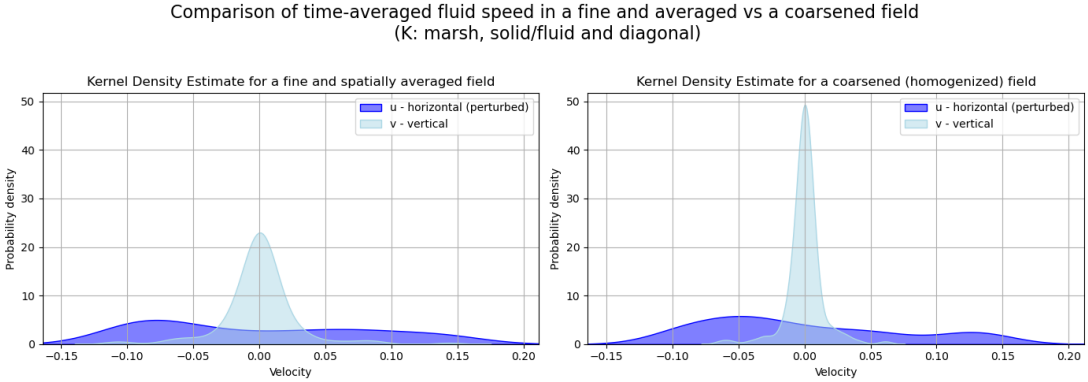

Figure 25: Kernel density estimates for time-averaged velocities within a steady state: averaged fine-scale velocities exhibit more spread than homogenized ones but both tend to centre around the same mean.

A KDE for the time-averaged steady state velocities that compares the distribution from a fine and spacially averaged field to a coarsened field shows that though both spanwise and perturbed streamwise velocities pass the sanity check of centring around their expected means, they exhibit different statistics (Figure 25). First of all, the variance of the perturbed streamwise velocity u is much larger than that of the spanwise velocity v. However, the streamwise velocity is much more similar between the fine and averaged field and the coarsened field than the spanwise is. In that, the spanwise contribution dominates the total velocity variance . This trend is consistent over all substructure and permeability configurations other than the solid/fluid non-diagonal marsh case, where higher variance is displayed in the opposite sense.

The connection between this observed trend and the comparison of kinetic energy between homogenized and fine-scale averaged cases is actually quite intuitive, as the formula for kinetic energy, which can also be written as $\frac{1}{2}\sum(u_{ij}^2 + v_{ij}^2)\Delta x \Delta y$, is directly related to the variance of u and v. The variance of some distribution of normal variable x with number of samples n and mean $\bar{x}$ is

$$var(x) = s_x^2 = \sum_i \frac{x_i^2}{n} - \bar{x}^2,$$

from which total kinetic energy can be written as

$$KE = \frac{A}{2}(var(u) + \bar{u}^2 + var(v) + \bar{v}^2)$$

where A is the area of the domain sampled over. Together with the variance trends seen in the KDE plot, we see clearly that this velocity variance is the reason as to why kinetic energy in homogenized simulations is generally lower than that of its fine-scale counterpart. Plotting the variance of both velocity components in each configuration shows that this trend is indeed realized, with all simulations displaying higher kinetic energies at fine scaling than at homogenized coarse scaling (Figure 26). The marsh cases have comparable kinetic energies between scales, but coarsened tunnel cases only retain about 60% of the kinetic energy of their fine-scale counterparts.

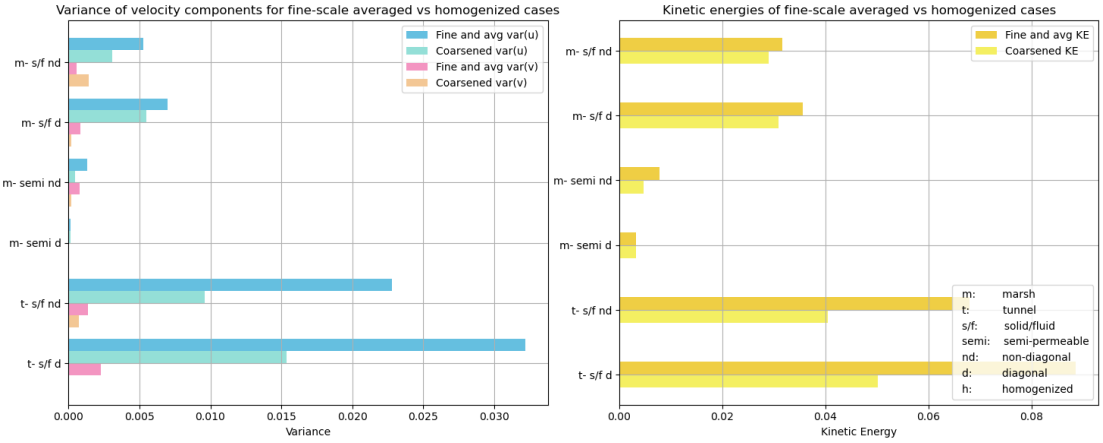

Figure 26: A figure relating the variance of time-averaged velocity components to total kinetic energy values for all solid/fluid, semi-permeable, diagonal, and non-diagonal cases. Larger variance implies a larger kinetic energy. Fine-scale cases ubiquitously display a larger combined variance in velocity components u and v which consequently implies higher total kinetic energy values. Solid/fluid cases suffer from larger discrepancies in their coarsened versus fine-scale total kinetic energies, where semi-permeable systems tend to keep a tighter relationship.

After having found a time period that corresponds to steady state behaviour, the spatial distribution of the velocities time-averaged over this period can be analyzed to gain more detailed understanding of the difference between coarse and averaged fine-scale cases than analyzing the kinetic energy brings. Figures 27 and 28 display the spatial distribution of relative error in the time-averaged velocities in the solid/fluid diagonal marsh and tunnel cases respectively. Relative error is measured at the resolution of the coarsened simulation, where homogenized values are compared to fine-scale spatially averaged values. The marsh case constrains relative error within 10×10^{-2} throughout all marsh-type configurations in both the streamwise and spanwise velocities, though the spanwise direction generally stays within an order of magnitude lower. The tunnel cases also perform similarly, constrained mostly within a relative error 10×10^{-2}, with lower error (expectedly) within regions of solid and variable error within the channel interior. The only exception to this is in small the region that has a high velocity gradient just downstream of the channel narrowing, which has a larger relative error in the streamwise velocity

only. In general, relative error is quite low and consistent throughout all configuration types, leading us to believe that periodic homogenization actually does a good job at approximating fine-scale averaged behaviour.

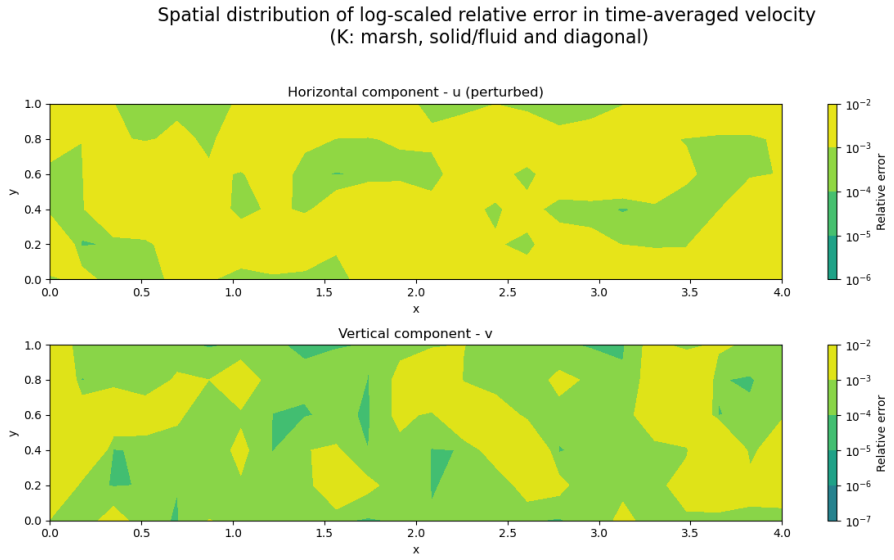

Figure 27: Spatial distribution of relative velocity error averaged over steady state times 22.5-30 for the solid/fluid diagonal marsh case. The relative error of homogenized velocities when compared to averaged fine-scale velocities is generally around 10×10^{-2} or lower. This result is similar in other permeability configurations.

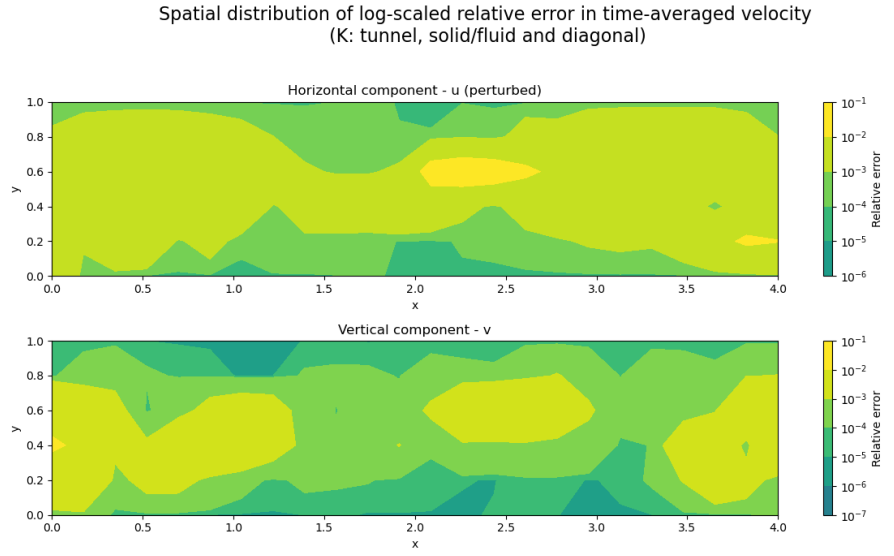

Figure 28: Spatial distribution of relative velocity error averaged over steady state times 22.5-30 for the solid/fluid diagonal tunnel case. The relative error of homogenized velocities when compared to averaged fine-scale velocities is generally around 10×10^{-2} or lower, with the exception of the streamwise velocity just downstream of the channel narrowing which has higher error by an order of magnitude. Regions of solid are (expectedly) well approximated, especially in the spanwise direction. Similar results hold in other permeability configurations.

6.5.2 Decomposition of kinetic energy into its components

Another interesting metric to look at is the breakdown of the kinetic energy rate into components, which is naturally done by multiplying Equation 7 by vector velocity u^T. This results in a measure of the rate of change of kinetic energy on the left hand side and the breakdown of components that contribute to it on the right hand side. These terms include pressure, advection, Coriolis, wind, bottom friction, diffusion, and permeability, but are missing the contribution to the kinetic energy rate that imposing a mean flow in a split step adds to the picture. Given that kinetic energy tends to a steady state, the sum of all components that contribute to its rate should tend to zero - which they do, shown in Figure 29.

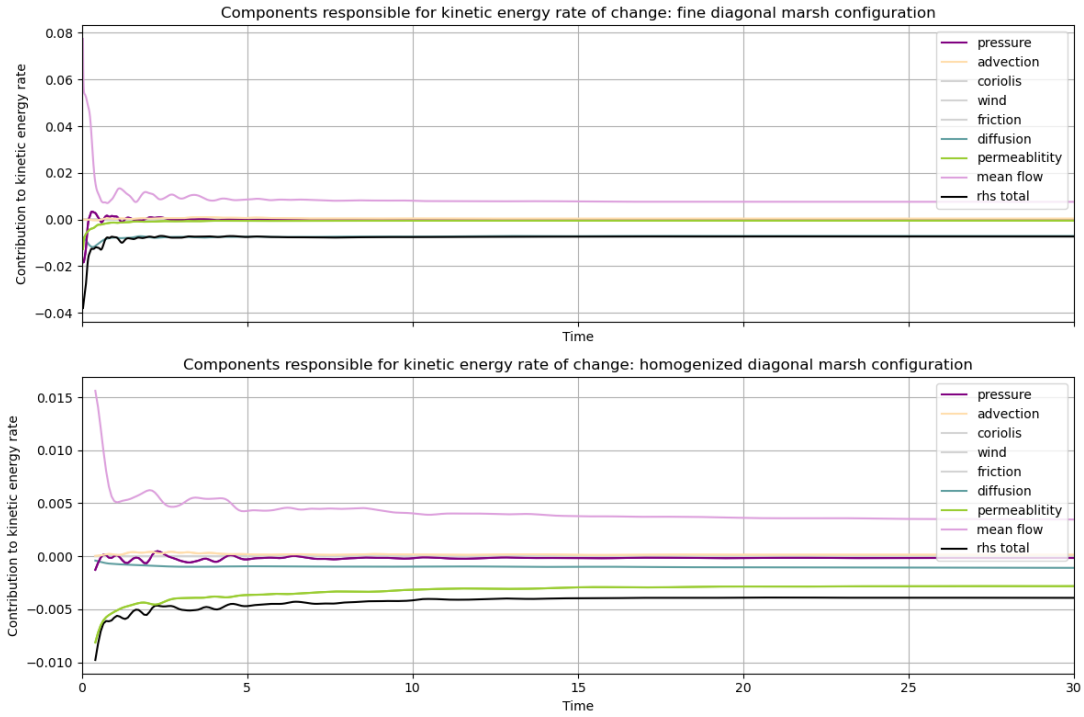

Figure 29: A breakdown of contributions to the rate of change of kinetic energy by term in Equation 7 (summed to "rhs total") as well as the imposition of a mean speed for the diagonal marsh configuration. The above components comprehensively cover all sinks and sources of kinetic energy in a simulation given that their sum tends to 0 while kinetic energy tends to a steady state. Permeability contributes significantly more to the kinetic energy rate of change in the averaged fine-scale case than the homogenized case.

Laplacian diffusion plays the dominant role (supported by pressure and advection) in decreasing the kinetic energy in the homogenized simulation, while both diffusion and permeability contribute similarly in the averaged fine-scale case. The addition of a mean flow provides a forcing to the system that counteracts the general decrease in kinetic energy imposed by the right hand side contributions, allowing the flow to reach a steady state rather than decaying to zero.

The breakdown of components for the complete set of configurations run is compiled in Figure 30, time-averaged over a steady state period. Percent contribution to total change in kinetic energy is plotted on the $y-$axis, giving an idea of the ratios between all contributing terms from both the momentum equation as well as the mean flow. Generally, all components contribute at about an order of magnitude less in the homogenized cases than the fine-scale averaged cases. An overview of the residual rate of contribution for each configuration type is available in Table 12. Most configurations show a general agreement between homogenized and fine-scale cases that an equilibrium will be reached, except the diagonal solid/fluid tunnel case, in which the fine-scale does not reach an equilibrium but the homogenized version does.

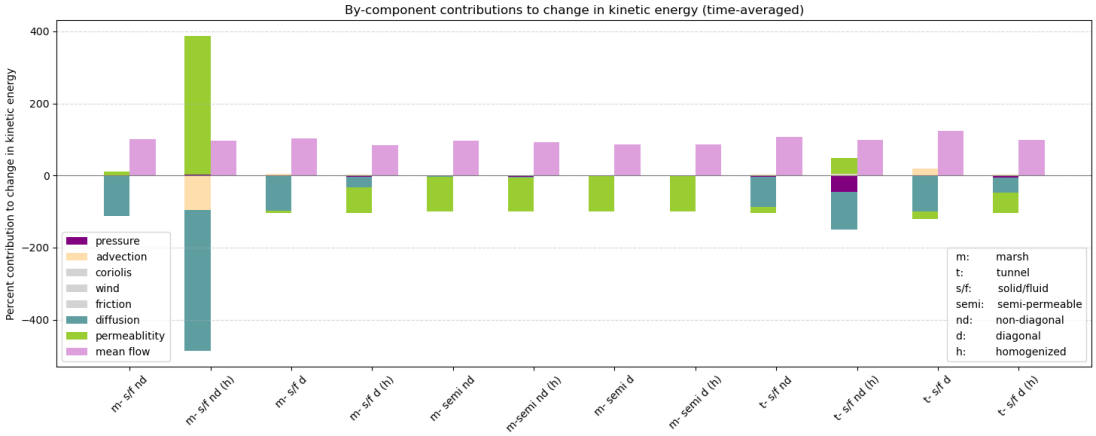

Figure 30: Percent contribution breakdown of components that influence rate of change of kinetic energy. Equilibrium rate of change contributions consistently decrease by an order of magnitude in homogenized cases and most cases feature permeability as the predominant sink of kinetic energy rate of change.

An interesting behaviour arises in several of the configurations, most notable in the solid/fluid non-diagonal marsh configuration, where permeability actually acts as a source for kinetic energy. Given the positive-definite requirement of a permeability matrix, this only possible when off-diagonal components are present. Indeed, we see no positive KE contributions from diagonal configurations. The off-diagonal entries can accelerate flow in the perpendicular direction, especially at the boundaries between fluid and solid. Since permeability matrices are averaged in order to find cell-centre tensors, there is a possibility of non-positive-definite matrices being generated - which could ultimately "create energy". However, this doesn't seem to prevent configurations that ultimately should evolve to a steady state from settling into one.

Figure 30 shows that diffusion plays a large role in removing kinetic energy in cases where permeability is less influential. Advective and pressure effects do contribute to the tunnel cases and the homogenized non-diagonal marsh cases more than other configurations, while the rest of the terms are negligible for all runs - mostly due to the fact that they were initialized with very small parameters to focus on permeability.

Overall, given these systems' comparable general tendencies, kinetic energies, and direct velocity distributions, we find that simulations coarsened through periodic homogenized do approach fine-scale simulations averaged to the same resolution, especially in cases that reach a steady state. We find this to be true for both solid/fluid permeability choices as well as semi-permeable ones.

Configuration type	Pos/neg KE rate contributions	Residual contribution
marsh		
solid/fluid non-diagonal	1.02	2×10^{-5}
solid/fluid non-diagonal (h)	0.96	-6×10^{-6}
solid/fluid diagonal	1.03	6×10^{-5}
solid/fluid diagonal (h)	0.84	-1.35×10^{-4}
semi-permeable non-diagonal	0.96	-1.6×10^{-3}
semi-permeable non-diagonal (h)	0.92	-3.4×10^{-4}
semi-permeable diagonal	0.87	-8.6×10^{-3}
semi-permeable diagonal (h)	0.85	-7.1×10^{-4}
tunnel		
solid/fluid non-diagonal	1.07	5.9×10^{-4}
solid/fluid non-diagonal (h)	0.99	-7×10^{-6}
solid/fluid diagonal	1.24	2.31×10^{-2}
solid/fluid diagonal (h)	0.99	-1×10^{-5}

Table 12: Ratios of positive to negative contributions to kinetic energy rate of change for all configurations shown in Figure 30. Ratios less than 1 indicate residual negative rate of change. Most cases seem to generally balance positive and negative contributions, with small residuals indicated in the table as well, with the exception of the diagonal tunnel configuration, of which the fine-scale case does not approach an equilibrium but its homogenized counterpart does.

7 Conclusions

We implemented two different homogenization methods for porous media defined through permeability matrices: volume averaging and periodic homogenization, to find that periodic homogenization is is not only more robust mathematically, but offers superior computational performance, increased accuracy, and the ability to upscale non-diagonally defined permeabilities. It accurately reproduces cases with empirically known results and approximates analytical solutions to machine precision, in both two and three dimensions.

We then apply periodic homogenization to a fluid model governed by the two-dimensional porous shallow water equations, which are derived by adding a Brinkman-type friction term that has been modified to include a tensorial permeability, to the original shallow water model. We do so in the context of presenting an extension to the existing literature on the use of Brinkman penalization, which relies on representing subgrid-scale features in the solid/fluid interface of a complex coastline as a porous medium to obtain a flow model along the coastline with numerous stability, accuracy, and computational efficiency benefits ([11]). Where the Brinkman method has previously only been applied to scalar porosity penalizations, we augment the flexibility of this penalization by imposing it through a permeabiliy-porosity matrix structure, which has the additional capability of characterizing a directional dependence in its associated friction. Since the equal-component diagonal representation of a permeability-defined medium is exactly equivalent to a scalar porosity, this extension only provides a scaffolding for improvement with little added cost. This added cost is mostly constrained to a preprocessing step that runs independently from the shallow water model, meaning that a given medium with a subgrid-scale structure only needs to go through periodic homogenization once to be used in any number of dynamical simulations. Furthermore, we have shown that periodic homogenization is extremely quick and reliable, and can be extended to high resolutions with a very small requirement on computational resources. Because of this and its ease of implementation in existing ocean models, we find that adding a matrix-defined permeability is a natural next step to the already successful Brinkman penalization method.

We also test an actual porous shallow water model subject to a permeability penalization on subgrid-scale structure that mimics a marsh and a tunnel environment. We define this substructure through a location-based choice of permeability matrix, where we test both solid and fluid representing matrices as well as semi-permeable ones. We find that solid-representing perme-

ability matrices accurately reproduce no-slip boundary conditions along blockages to reproduce expected results in all choices of substructure configuration. As an exploration, we test non-diagonal permeability-defined solid regions and find that they produce additional flow turning, which may be useful for some types of media or substructure, but further tests are required to understand their effects more comprehensively.

Overall, we find that when porous shallow water models with permeability-defined substructure are coarsened, they approach their corresponding fine-scale case, averaged to the same resolution. A spatial distribution of relative error between the time-averaged steady-state velocities of coarse simulations and fine, averaged simulations shows minimal discontinuity between these two cases (largely constrained within 1% relative error). We find that in models that reach a steady state, the kinetic energies of the coarse and fine, averaged simulations tend towards the same level, with coarsened simulations generally displaying slightly less total energy. This can be explained by homogenization's inherent smoothing of features, which tightens the variance in the coarsened flow velocities. We do not find this to be a failing, since the current prevailing method having to average scalar porosity penalizations experiences the same effect. All the permeability-defined method has to do is not perform worse than this method, which it fundamentally cannot do given that one makes appropriate choices for subtructure matrices.

Therefore, we advocate for the use of permeability-defined subgrid-structure in porous shallow water models employing a Brinkman penalization, if only to build in the capability of a directionally dependent friction. Furthermore, we think that the coarsening of subgrid-structure should be done by the method of periodic homogenization, which has proven to be quick, rigorous, and reliable for both diagonal and non-diagonal input matrices.

8 Acknowledgements

LP thanks Nicholas Kevlahan for his support and expertise throughout the development of this thesis and Blaise Bourdin for his (life and technical) advice as well as the introduction to the method of periodic homogenization. This work would not be where it is without either of you.

References

[1] M. K. Alam and A. Farid. "Hydraulic conductivity tensor of anisotropic soils: The impact on seepage flow". In: *Acta Scientiarum Polonorum. Architectura* 23 (2024).

[2] G. Allaire. "Homogenization". In: *Shape Optimization by the Homogenization Method* (2002), pp. 1–89.

[3] G. Allaire. "Homogenization and Two-Scale Convergence". In: *SIAM Journal on Mathematical Analysis* 23.6 (1992), pp. 1482–1518. DOI: 10.1137/0523084.

[4] P. Angot. "Analysis of singular perturbations on the Brinkman problem for fictitious domain models of viscous flows". In: *Mathematical Methods in the Applied Science* 22 (1999), pp. 1395–1412.

[5] P. Angot, C.-H. Bruneau, and P. Fabrie. "A penalization method to take into account obstacles in viscous flows". In: *Numerische Mathematik* 81 (1999), pp. 497–520.

[6] J. Bear. *Dynamics of fluid in porous media*. Courier Corporation, 1988, pp. 137–151.

[7] S. Brisard and F. Legoll. "Periodic homogenization using the Lippmann–Schwinger formalism". In: *arXiv* (2014). arXiv: 1411.0330 [math.NA].

[8] J. C. Butcher. *Numerical methods for ordinary differential equations*. John Wiley & Sons, 2016, pp. 448–460.

[9] D. Coulson. *Aerials by Dave Coulson*. Accessed: 2024-11-06. n.d. URL: https://www.davecoulson.ca/Portfolio/Aerials/8.

[10] L. Debreu, N.K.-R. Kevlahan, and P. Marchesiello. "Brinkman volume penalization for bathymetry in three-dimensional ocean models". In: *Ocean Modelling* 145 (2020), p. 101530. ISSN: 1463-5003. DOI: 10.1016/j.ocemod.2019.101530.

[11] L. Debreu, N.K.-R. Kevlahan, and P. Marchesiello. "Improved Gulf Stream separation through Brinkman penalization". In: *Ocean Model.* 179 (2022), p. 102121. DOI: 10.1016/j.ocemod.2022.102121.

[12] L.W. Gelhar and C.L. Axness. "Three-dimensional stochastic analysis of macrodispersion in aquifers". In: *Water Resources Research* 19.1 (1983), pp. 161–180.

[13] T.W.J. de Geus et al. "Finite strain FFT-based non-linear solvers made simple". In: *Computer Methods in Applied Mechanics and Engineering* 318 (May 2017), pp. 412–430. ISSN: 0045-7825. DOI: 10.1016/j.cma.2016.12.032.

[14] P. R. King. "The use of field theoretic methods for the study of flow in a heterogeneous porous medium". In: *Journal of Physics A: Mathematical and General* 20.12 (1987), p. 3935.

[15] Q. Liao et al. "Analytical solution for upscaling hydraulic conductivity in anisotropic heterogeneous formations". In: *Advances in Water Resources* 128 (2019), pp. 97–116.

[16] B. A. Lippmann and J. Schwinger. "Variational Principles for Scattering Processes. I". In: *Phys. Rev.* 79 (3 Aug. 1950), pp. 469–480. DOI: 10.1103/PhysRev.79.469.

[17] M. Masihi, P. A. Gago, and P. R. King. "Estimation of the effective permeability of heterogeneous porous media by using percolation concepts". In: *Transport in Porous Media* 114 (2016), pp. 169–199.

[18] H. Moulinec and P. Suquet. "A fast numerical method for computing the linear and nonlinear mechanical properties of composites". In: *Comptes Rendus de l'Académie des sciences. Série II. Mécanique, physique, chimie, astronomie* (Apr. 1994). URL: https://hal.science/hal-03019226.

[19] NOAA. *What is the Global Ocean Conveyor Belt?* Accessed: 2024-11-06. n.d. URL: https://oceanservice.noaa.gov/facts/conveyor.html.

[20] P. Renard and R. Ababou. "Equivalent Permeability Tensor of Heterogeneous Media: Upscaling Methods and Criteria (Review and Analyses)". In: *Geosciences* 12.7 (2022). ISSN: 2076-3263. DOI: 10.3390/geosciences12070269.

[21] C.H.K. Williamson. *Vortex dynamics in the wake of a cylinder*. Springer, 1995, pp. 155–234.

[22] S. Żak. *Developments in Hydraulic Conductivity Research*. IntechOpen, 2011.