

Computational analysis of transport in three-dimensional heterogeneous materials: An OpenFOAM®-based simulation framework

Original

Computational analysis of transport in three-dimensional heterogeneous materials: An OpenFOAM®-based simulation framework / Boccardo, G.; Crevacore, E.; Passalacqua, A.; Icardi, M.. - In: COMPUTING AND VISUALIZATION IN SCIENCE. - ISSN 1432-9360. - ELETTRONICO. - 23:1-4(2020). [10.1007/s00791-020-00321-6]

Availability:

This version is available at: 11583/2856580 since: 2020-12-11T08:22:22Z

Publisher:

Springer

Published

DOI:10.1007/s00791-020-00321-6

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

Springer postprint/Author's Accepted Manuscript

This version of the article has been accepted for publication, after peer review (when applicable) and is subject to Springer Nature's AM terms of use, but is not the Version of Record and does not reflect post-acceptance improvements, or any corrections. The Version of Record is available online at: <http://dx.doi.org/10.1007/s00791-020-00321-6>

(Article begins on next page)



Computational analysis of transport in three-dimensional heterogeneous materials

An OpenFOAM®-based simulation framework

Gianluca Boccardo¹ · Eleonora Crevacore² · Alberto Passalacqua⁴ · Matteo Icardi³

Received: 22 June 2017 / Accepted: 4 December 2019
© The Author(s) 2020

Abstract

Porous and heterogeneous materials are found in many applications from composites, membranes, chemical reactors, and other engineered materials to biological matter and natural subsurface structures. In this work we propose an integrated approach to generate, study and upscale transport equations in random and periodic porous structures. The geometry generation is based on random algorithms or ballistic deposition. In particular, a new algorithm is proposed to generate random packings of ellipsoids with random orientation and tunable porosity and connectivity. The porous structure is then meshed using locally refined Cartesian-based or unstructured strategies. Transport equations are thus solved in a finite-volume formulation with quasi-periodic boundary conditions to simplify the upscaling problem by solving simple closure problems consistent with the classical theory of homogenisation for linear advection–diffusion–reaction operators. Existing simulation codes are extended with novel developments and integrated to produce a fully open-source simulation pipeline. A showcase of a few interesting three-dimensional applications of these computational approaches is then presented. Firstly, convergence properties and the transport and dispersion properties of a periodic arrangement of spheres are studied. Then, heat transfer problems are considered in a pipe with layers of deposited particles of different heights, and in heterogeneous anisotropic materials.

Keywords Porous media · OpenFOAM® · Multi-scale simulation · Heterogeneous materials · CFD

1 Introduction

In recent years computer simulations are becoming an ever more important tool to study a large spectrum of physical problems, a trend supported by the increasing processor throughput and advances in the parallel computing research. Nonetheless, while the constraint due to the cost in computational time is gradually easing, other bottlenecks can remain, such as the limited or expensive access to licensed (i.e. commercial) or otherwise closed-source codes, that are hardly verifiable and rigorously validated. While several promising open-source libraries and codes (Deal.II, UG, dune, GetFEM, just to cite a few) are now available, there is still an urgent need of developing specific tools and workflow to approach particular classes of problems. This is particularly problematic in the case of very complex physics or inherently multi-scale problems such as heterogeneous media where, despite the current growing trend in non-physical data-driven science, computer modelling still remains the only practi-

Communicated by Gabriel Wittum.

✉ Matteo Icardi
matteo.icardi@nottingham.ac.uk

Gianluca Boccardo
gianluca.boccardo@polito.it

Eleonora Crevacore
eleonora.crevacore@polito.it

Alberto Passalacqua
albertop@iastate.edu

¹ Dipartimento di Scienza Applicata e Tecnologia, Politecnico di Torino, 10129 Turin, Italy

² Dipartimento di Scienze Matematiche “G. L. Lagrange”, Politecnico di Torino, 10129 Turin, Italy

³ School of Mathematical Sciences, University of Nottingham, Nottingham NG8 3LS, UK

⁴ Department of Mechanical Engineering, Iowa State University, Ames, IA 50011, USA

cal path to an improved knowledge and understanding and, eventually, design and optimisation. These problems, in fact, often require, not only efficient discretisation and solvers, but also specific features, and a wide range of pre- and post-processing tools. Porous media, in their wide interpretation, are a prominent category of multi-scale problems, and, as such, form the basis and the motivation of the model problems explored in this paper [23,35].

The purpose of this work is thus to present a series of tools and methods which can be used to approach a wide range of problems and applications in porous media, characterised by different geometrical descriptions and involved in several transport phenomena and dynamics: from transport of dilute suspensions to heat transfer in porous structures. Eventually, the aim of this paper is to propose a fully open-source simulation workflow, show the validity and feasibility of a fully *in-silico* investigation, and introducing some specific novel methods and formulation to streamline the modelling and upscaling step. We demonstrate this approach by solving three proof-of-concept problems. Although these are sometimes characterised by some ideal assumptions (e.g., geometrically simplified models, mostly linear and constant parameters, decoupled multi-physics), the tools presented in this work can be readily used for a number of important applications.

Within the computational workflow mentioned above, we thus present here specific developments and advancements to:

- generate realistic packings of solid grains of arbitrary shapes, suitable for the modelling of both natural formations (e.g. aquifers) and artificial packings (e.g. packed bed reactors), a class of unit operations central to process engineering [5];
- generate random arrangements of spheres or ellipsoids with a tunable porosity, which can be used to represent realistic solid dispersions with a variable amount of dispersed matter (as it will be shown in this work);
- improve on the classic boundary conditions used in unsteady linear transport problems, to strongly reduce computational times in the investigation of the upscaled dynamics of a quasi-periodic problem, bringing a definite improvement on the “naive” methods, especially for problems characterised by very large scales of variation of the quantity of interest (e.g.: advection-diffusion problems at very high Péclet numbers); this technique was already successfully used in a recent work of ours [6];
- define and solve efficiently closure cell problems for diffusion in heterogeneous anisotropic materials.

Other mathematical models and open-source simulation tools, based on the same computational platform, can be also profitably employed in conjunction to the methods explored




Fig. 1 Schematic representation of a representative volume element of a generic two-phase heterogeneous material

here, to overcome some of the limitations. These include a quadrature-based method of moments to deal with the solution of PDF transport equations [32] (e.g., evolution of particle populations or other additional internal coordinates such as transit time), and stochastic sampling techniques based on multi-level Monte Carlo methods to deal with uncertainty quantification and stochastic upscaling in heterogeneous porous media systems [19].

The structure of the paper is thus the following: in Sect. 2 the theoretical bases of the investigated problems are laid down, given that the real-world test cases explored share the same theoretical background, fully or in part. Section 3 summarises the numerical techniques, based on a finite-volume formulation. Eventually, the actual computational workflow is detailed in Sect. 4 for each case: from the generation of the geometrical model, to the creation of a suitable mesh, to the presentation of the actual results. Specifically, the first part will deal with transport of a solute in porous medium composed of a periodic arrangement of solid grains, with added attention to the meshing strategies employed in the simulation. The second part explores fluid flow and heat transfer in a pipe where varying amounts of solid matter has settled on its bottom, focusing on the resulting effect on the heat transfer coefficient of the system. Lastly, another heat transfer problem is presented, this time considering the case of a solid dispersion, where the continuous matrix and the dispersed inclusions are characterised by different thermal diffusivity. In this way, the reader will be able to more effectively peruse this work, focusing on the application of interest and on the specific computational pipeline for its development.

2 Model equations

Let us consider a generic two-phase heterogeneous material (cf. Fig. 1), denoted with $\Omega = \Omega_1 \cup \Omega_2$, Γ being the interface between them, and let us split the external boundary in the inlet $\partial\Omega_-$, outlet $\partial\Omega_+$, and lateral boundary $\partial\Omega_\ell$.

In the following, we will consider a few simplifications, the chief of which is to assume a saturated porous medium,

constituted by solid non-connected irregular grains immersed in a continuous phase, either liquid (i.e. Newtonian incompressible fluid at room temperature) or solid (with different characteristics from the grains).

2.1 Flow field

The equations governing the fluid phase are the stationary Navier-Stokes equation¹ and the mass balance law:

$$\rho(\nabla \mathbf{u})\mathbf{u} = -\nabla p + \mu \Delta \mathbf{u}, \quad (1)$$

$$\nabla \cdot \mathbf{u} = 0, \quad (2)$$

where ρ is the constant fluid density (kg m^{-3}), \mathbf{u} is the effective fluid velocity (m s^{-1}), p is the pressure ($\text{kg m}^{-1} \text{s}^{-2}$) and μ is the fluid dynamic viscosity ($\text{kg m}^{-1} \text{s}^{-1}$). When inertia is negligible, Eq. (1) reduces to the Stokes's equation:

$$-\nabla p + \mu \Delta \mathbf{u} = 0, \quad (3)$$

this case is usually referred to as “creeping flow”, typically encountered in several environmental applications: all our simulations are performed in this range, in laminar conditions.

2.2 Transport models

Coupled to the fluid phase, we generally solve a transport problem assuming it does not have a back-coupling with the flow. This is true for small dilute particles or for solutes. When no flow is present (i.e. the continuous phase is a solid), there is no need of solving the flow field, although in many applications advection can be due to electrostatic forces or other physics which we neglect here.

2.2.1 Eulerian equations

Working in an Eulerian framework, it is possible to model the concentration transport of solutes or of a diluted suspension of colloidal particles as:

$$\frac{\partial c}{\partial t} + \mathbf{u} \cdot \nabla c = \nabla \cdot (D \nabla c), \quad (4)$$

where c is the concentration (kg m^{-3}), D is the diffusion coefficient ($\text{m}^2 \text{s}^{-1}$); for the advective term we made use of the incompressibility condition Eq. (2). As it has been said, the interactions between solid and fluid phases are described with *one-way coupling*, stating that the fluid (and external forces) will affect the motion of the particles, but not the

vice-versa. The diffusion coefficient for small particles can be estimated with the Stokes-Einstein equation:

$$D = \frac{\kappa_B T}{3\pi\mu d_p}, \quad (5)$$

where κ_B is the Boltzmann constant, T is the temperature (K) and d_p is the particle diameter (m). When dealing with solid phases (inside the grains or in both domains), $\mathbf{u} = 0$ and D is a solid diffusion coefficient (possibly a tensor).

It is worth noting that an equation of the form of Eq. (4) can apply for heat transfer problems (studied in Sects. 4.2 and 4.3) as well. In this case, it is convenient to introduce the following notation:

$$\frac{\partial T}{\partial t} + \mathbf{u} \cdot \nabla T = \nabla \cdot (\alpha \nabla T), \quad (6)$$

where T is the fluid temperature (K) and the diffusive coefficient is now the thermal diffusivity (scalar or tensor) α ($\text{m}^2 \text{s}^{-1}$).

2.3 Scalar quasi-periodic boundary conditions

Eulerian transport problems, even in the case of periodic geometry and flow, cannot be solved with simple periodic conditions on a single periodic unit cell due to intrinsic evolution in space and time of the equations that requires larger domain to be solved, or due to the non-conservative nature of the equations. However, when the equation is linear and we are solely interested in the asymptotic (long-time, infinitely far, self-similar) behaviour, we can reformulate the problem to find a stationary *quasi-periodic solution* (up to a multiplicative or additive constant, depending on the nature of the problem). Dropping the time dependence, we can formulate a generic linear transport problem as

$$\nabla \cdot (\mathcal{L}c) = \mathfrak{R}c + \mathfrak{F} \quad \text{in } \Omega \quad (7)$$

where \mathcal{L} is a flux operator (e.g., advection-diffusion fluxes), \mathfrak{R} is a non-conservative operator (e.g., reactions, differential operators not in a divergence form), \mathfrak{F} is a source term. These operators are all assumed to be linear in c but possibly depending on space x . We consider the following internal boundary conditions²:

$$\mathcal{L}c = f(c) \quad \text{on } \Gamma \quad (8)$$

with f being a generic (space-dependent, possibly non-linear) flux at the wall.

² On the porous matrix walls, in the case of perforated domain. Otherwise, for dual domain problems, this has to be replaced with interface conditions and an additional equation for the other domain has to be considered.

¹ Body forces are neglected.

We aim to find a self-similar solution of Eq. (7), i.e., a solution in the smallest periodic cell with quasi-periodic external boundary conditions of the type³:

$$c|_{\partial\Omega^-} = \phi(c|_{\partial\Omega^+}), \quad \nabla_n c|_{\partial\Omega^-} = \psi(c|_{\partial\Omega^+}) \quad (9)$$

where ϕ and ψ are linear functions and $\partial\Omega^\pm$ are two (geometrically opposite) periodic boundaries⁴. It follows that, generally, $\psi(c) = \phi'(c)\nabla_n c$ to ensure the quasi-periodicity of the normal gradient. From simple compatibility conditions for the existence and uniqueness of solutions, we can distinguish the following cases:

- *Conservative transport*: $\mathfrak{R} = \mathfrak{F} = \mathfrak{f} = 0$. This setup is the one of interest for studying *passive asymptotic dispersion* (or effective diffusion when there is no advection) properties as what remains is a fully conservative equation. The quantity of interest to quantify the dispersion is the average of local gradients ∇c . However, for periodic conditions, only a trivial constant solution (with $\phi(c) = \psi(c) = c$) exists. Therefore, we have to reintroduce a fictitious non-conservative constant source $\mathfrak{f} = \mathfrak{f}_0 = \mathfrak{L}(\mathbf{p} \cdot \mathbf{x})$ where \mathbf{p} is the direction in which we want to study the transport (x in this case). This source term is zero for pure diffusion. The resulting equation has now a non-trivial solution up to an additive constant. We therefore select a quasi-periodicity of the type:

$$\phi(c) = c + \phi_0, \quad \phi_0 = \langle c \rangle_{\partial\Omega^+} - \langle c \rangle_{\partial\Omega^-} \quad (10)$$

where $\langle \cdot \rangle$ represents an averaging (surface- or flux-weighted) operator⁵, and the constant ϕ_0 is a variable of the problem and it has to be computed to counter-balance the volumetric source term. Without loss of generality we can assume then $\langle \phi(c|_{\partial\Omega^-}) \rangle = 0$. We can then compute the local gradients (whose norm gives the dispersion coefficient), produced by the transport operator \mathfrak{L} under the macroscopic gradient ϕ_0 , by subtracting the source term \mathfrak{f}_0 as

$$\frac{\nabla c - \nabla(\mathbf{p} \cdot \mathbf{x})}{\phi_0}$$

- *Homogeneous equation*: $\mathfrak{F} = 0$ and $\mathfrak{f}(c) = \mathfrak{f}_1 c$. This is the case of linear bulk and surface reaction, where, in general, the solution is given up to a multiplicative constant.

³ Depending on the form of the flux operator, the resulting PDE can be elliptic, parabolic or hyperbolic. We consider here conditions both for the solution and its normal gradient, generally needed for elliptic PDEs.

⁴ In three dimensional cubical elementary volumes, one can write separate equations for x , y , z periodicity.

⁵ In fact, considering a linear and periodic operator \mathfrak{L} , surface and flux-weighted averages are equivalent, i.e., $\langle \mathfrak{L}c \rangle = \langle c \rangle$.

This suggests us to look for quasi-periodic solutions with

$$\phi(c) = \phi_1 c, \quad \phi_1 = \frac{\langle \mathfrak{L}c \rangle_{\partial\Omega^-}}{\langle \mathfrak{L}c \rangle_{\partial\Omega^+}} = \frac{\langle c \rangle_{\partial\Omega^-}}{\langle c \rangle_{\partial\Omega^+}} \quad (11)$$

in such a way that the quasi-periodic boundary conditions balance the non-conservative terms and allows to identify ϕ_1 as an equivalent reaction rate (observing that it can be related to the volume-averaging of all non-conservative terms).

- *General linear case*: Integrating Eq. (7) over the volume, and considering a standard advection diffusion operator $\mathfrak{L} = \mathbf{u} + D\nabla$, linear bulk reaction \mathfrak{R} , linear surface reaction

$$\mathfrak{f}(c) = \mathfrak{f}_1 c + \mathfrak{f}_0$$

and a linear transformation $\phi(c) = \phi_1 c + \phi_0$, the following conditions hold for ϕ_1

$$\begin{aligned} \phi_1 &= \frac{\langle \mathfrak{L}c \rangle_{\partial\Omega^-}}{\langle \mathfrak{L}c \rangle_{\partial\Omega^+}} = 1 + \mathfrak{R} \frac{\int_{\Omega} c}{\int_{\partial\Omega^+} \mathfrak{L}c} + \mathfrak{f}_1 \frac{\int_{\Gamma} c}{\int_{\partial\Omega^+} \mathfrak{L}c} \\ &= 1 + \mathfrak{R} \frac{\langle c \rangle}{\langle \mathfrak{L}c \rangle_{\partial\Omega^+}} \frac{|\Omega|}{|\partial\Omega^+|} + \mathfrak{f}_1 \frac{\langle c \rangle_{\Gamma}}{\langle \mathfrak{L}c \rangle_{\partial\Omega^+}} \frac{|\Gamma|}{|\partial\Omega^+|} \end{aligned} \quad (12)$$

and ϕ_0

$$\begin{aligned} \phi_0 &= \mathfrak{F} \frac{\int_{\Omega} 1}{\int_{\partial\Omega^-} \mathfrak{L}1} + \mathfrak{f}_0 \frac{\int_{\Gamma} 1}{\int_{\partial\Omega^-} \mathfrak{L}1} \\ &= \mathfrak{F} \frac{1}{\langle \mathbf{u} \cdot \mathbf{n} \rangle_{\partial\Omega^-}} \frac{|\Omega|}{|\partial\Omega^-|} + \mathfrak{f}_0 \frac{1}{\langle \mathbf{u} \cdot \mathbf{n} \rangle_{\partial\Omega^-}} \frac{|\Gamma|}{|\partial\Omega^-|} \end{aligned} \quad (13)$$

As it can be seen from the last equality, the integrals can be rewritten highlighting the geometric factors $\frac{|\Omega|}{|\partial\Omega^-|}$ and $\frac{|\Gamma|}{|\partial\Omega^-|}$, and the mean concentration values $\langle c \rangle$ in the volume and on the surfaces. If the flow field is incompressible, $\langle \mathbf{u} \cdot \mathbf{n} \rangle_{\partial\Omega^-}$ is equivalent to the mean (volumetric) velocity.

- *Non-linear case*: For general non-linear operators, one has to find a function $\phi(c)$ such that

$$\int_{\partial\Omega^+} \mathfrak{L}\phi(c) = \int_{\partial\Omega^+} \mathfrak{L}c + \int_{\Omega} \mathfrak{R}c + \int_{\Gamma} \mathfrak{f}(c)$$

This problem is not generally easily solvable. One possibility is to assume a “macroscopically” linear reaction, i.e., ϕ to be linear. If $\phi(c) = \phi_1 c$, the problem of finding ϕ_1 becomes equivalent to a generalised non-linear eigenvalue problem. There is no guarantee, however, that the problem admits a solution. More in general, in our future works, we will investigate a more general algorithm to find linear and non-linear functions ϕ .

This proposed approach can be proven to be equivalent to the several cell problems formally derived by two-scale expansions [17], and it is particularly convenient in the cases where the phenomenon of interest would require very long physical time or especially, a very large spatial domain. This is true in most upscaling problems where, to derive upscaled equations, the asymptotic regime is often sought. Simulations on a smaller periodic domain with these sets of *quasi-periodic boundary conditions* can significantly reduce the computational effort. It is interesting to notice that, in two-scale asymptotic homogenisation, this is usually overcome with the formal derivation of cell problems with periodic conditions. Our approach, despite being more phenomenological, is indeed equivalent and can also be applied to problems where an homogenisation limit does not exist. From the computational point of view, Eq. (10), Eq. (11), and Eq. (12), have to be implemented via outer iterations, with ϕ_1 and ϕ_0 computed from the field c . In our implementation, outer iterations are always present also to allow for non-linear terms, non-linear discretisation schemes, and explicit corrections for non-orthogonal cells. At each iteration, the boundary conditions are computed with the equations above. The iterations stop when the residuals of the equations fall below a certain threshold and, at the same time, the estimated pseudo-periodic boundary conditions converges to a constant value.

3 Numerical discretisation

3.1 Numerical schemes

OpenFOAM® implements the finite volume method (FVM) [14] with co-located grid arrangement. Internal values are stored at the cell centre, while boundary values are stored at the face centroid on the corresponding boundary cell faces. The Rhie-Chow approach [36,38] is used to address the pressure-velocity decoupling observed when this grid arrangement is adopted. When computing fluxes, values of each variable are computed at the face centroid of each interior cell through a reconstruction technique, following [11] (more details on the implementation can be found in [21]), as shown below.

The implementation of numerical schemes for convection and diffusion schemes follows the standard FVM approach [14,30].

The convective term is discretised by applying Gauss' theorem. If we consider a scalar c , advected with a velocity \mathbf{u} , and we indicate the volume of a computational cell with V , while ∂V is its surface area, we can write

$$\int_V \nabla \cdot (\mathbf{u}c) dV = \int_{\partial V} (\mathbf{u}c) \cdot \mathbf{n} dS \approx \sum_f c_f \mathbf{u}_f \cdot \mathbf{n}_f S_f. \quad (14)$$




Fig. 2 Schematic representation of two computational cells. In this picture, the cell on the right (in blue) represents the upwind cell

The value of both \mathbf{u}_f and of c_f need to be evaluated at cell faces; S_f is the surface area of the considered faces and \mathbf{n}_f is the associated outward unit vector. If a second-order TVD scheme is used, this is achieved by means of the reconstruction procedure illustrated in [11] (Sec. 2), in which the face value of a variable is obtained as

$$c_i = c_P + \frac{1}{2} \psi(r_f)(c_W - c_P), \quad (15)$$

where c_P is the value of c at the cell centre (labelled P , as illustrated in Fig. 2), c_W the one at the downwind node, and $\psi(r_f)$ is the limiter function. The quantity r_f is defined as [11]

$$r_f = \frac{c_P - c_E}{c_W - c_P}, \quad (16)$$

where c_E is the upwind value of c .

The advantage of this procedure consists in its generality, which allows the first-order upwind scheme and the central differencing scheme to be recovered by setting $\psi(r_f) = 0$ and $\psi(r_f) = 1$ respectively, [11].

The diffusion term is discretised considering the Laplacian operator [15]

$$\begin{aligned} \int_V \nabla \cdot (\Gamma \nabla c) dV \\ = \int_{\partial V} (\Gamma \nabla c) \cdot \mathbf{n} dS \approx \sum_f \Gamma_f (\nabla c)_f \cdot \mathbf{n}_f S_f, \end{aligned} \quad (17)$$

with

$$(\nabla c)_f \cdot \mathbf{n}_f = \frac{c_E - c_P}{|\mathbf{d}|} S_f \cos \theta + C_g, \quad (18)$$

where \mathbf{d} is the vector joining the cell centre of the considered adjacent cells. The term C_g is an explicit correction for non-orthogonality [30].

Gradients are computed either with Gauss' integration

$$\int_V \nabla c \, dV \approx \sum_f c_f S_f, \quad (19)$$

where the face values of c are found using a linear reconstruction of the cell values on cell faces, or using the least-squares approach. Several implementations of the latter are available in OpenFOAM®: the second-order least-squares approach, available in different flavours, either using the cell-centred values or the centre and nodal values for the calculation, and a fourth-order least-square approach.

The adoption of Gauss' gradient approach (Eq. (19)) is only appropriate on regular meshes, where it has second-order accuracy (see, among others [14]). Additionally, this approach is sensitive to mesh anisotropy. Consequently, on meshes with varying cell size, mesh refinement and arbitrarily shaped cells, the least-squares approach is the recommended choice to ensure second-order accuracy. Thus, we adopt it when dealing with irregular meshes, as described in the following cases.

Another source of potential difficulties is the presence of the explicit non-orthogonal correction in the calculation of the face-normal gradient (Eq. (18)), which may lead to unbounded results, depending on the mesh quality and on the problem being solved; the difficulty with the correction term is that, being it explicit, it may lead to unboundedness when the non-orthogonality is too large, with consequent loss of robustness of the scheme. The approach used in OpenFOAM® is presented in [21].

To avoid unboundedness, one can either apply of the correction only when mesh non-orthogonality is high, or exclude the correction altogether. Direct consequence of this is the reduction of the accuracy of the numerical scheme, which deteriorates when the non-orthogonality problem is particularly serious.

3.2 Meshing

Aside from the issue concerning the generation of the geometry, the inherent randomness which characterises most porous systems also affects the generation of the computational mesh. Far from being a marginal part of the workflow, a fundamental trade-off in the setup of the computational model has to be evaluated, weighing a numerically satisfactory mesh with the increasing costs in machine time for a finer discretisation, both for the mesh generation itself and the consequent flow/transport problem.

Studying a satisfactory portion of a porous medium (i.e.: a *representative elementary volume*) means having to deal (in most cases) with a random system, which immediately excludes the possibility of employing a structured (or block-structured) mesh, usually favoured for their efficiency.

Moreover, one often has to deal with grids composed of a number of elements up to tens of millions [18] or much more [28], making any kind of “naked-eye” inspection unfeasible.

As a result, it is essential to develop a robust pipeline for both the generation of the mesh, which can be controlled by an a-priori discretisation strategy (not reliant on qualitative analysis of the resulting grid), and an efficient process of testing for error convergence based on the simulation results. Here, we have tested two different unstructured mesh generators, both present in the OpenFOAM® suite: *snappy-HexMesh* and *foamyHexMesh* (hence respectively *sHM* and *fHM* for brevity). More details on these tools are available in the appendix while a test-case with the results of the meshing algorithms and their effects on the numerical solution are reported in Sect. 4.1.1. We note, for the interested reader, that some valid meshing alternatives can be found in *Cubit* or *Gmsh* [16], the latter being a freely available and consolidated tool.

4 Results

4.1 Transport in periodic sphere packings

One of the fields which most commonly deals with the study of transport in porous media is filtration theory. For example, this has been applied for the investigation of aquifer contamination and remediation, [9,25,27,37]. Correspondingly, a lot of effort has been expended in the past decades to approach this problem, both experimentally and via theoretical investigation. These kind of problems are basically equivalent, in their simplest formulation, to linear advection-diffusion equations. Upscaling techniques are therefore robust and well-known [2,17]. However, despite the simplicity of the problem, some issues are still unsolved, such as the quantitative understanding of the role of geometrical parameters in transport processes or the effect of non-linearities. In recent years, computational studies have begun to complement the available analytical and empirical results, especially when the physical problem presented further difficulties (e.g.: mixing controlled reactions, heterogeneous materials, etc.) which rendered an a-priori upscaling process exceedingly difficult, if not outright impossible [1].

As mentioned elsewhere, the pre-processing step constitutes a big part of a CFD case of transport in porous media. While in the next sections we explored the process of the generation of the geometrical model, here we will explore the issue of the choice of the optimal meshing strategy and its consequences in the convergence of the finite volume scheme. To this end, we will consider here the Stokes flow over a unit cell of a periodic arrangement of spheres. Then, in Sect. 4.1.3, we turn to look to the upscaling step, and specifically on how to conveniently solve unsteady trans-




Fig. 3 Comparison of four meshing strategy: sHM_U (up-left), sHM_R (up-right), fHM_U (low-left), fHM_R (low-right). Each mesh appearing in this picture has a total number of cells that is a little less than 5×10^5

port Eulerian equations and on an alternative Lagrangian approach to extract residence-time distribution curves for a solute transport problem.

4.1.1 Unstructured meshing generation

Tests were performed meshing a body-centred cubic structure, where grains were spheres of equal size [10], and the medium porosity was held fixed at 0.366 (i.e.: close to the minimum porosity, but sufficiently high to avoid issues due to the meshing of contact points between adjacent spheres). Four different meshing strategy were then employed:

- snappyHexMesh with uniform cell size (sHM_U);
- snappyHexMesh with one level of refinement (sHM_R);
- foamyHexMesh with uniform cell size (fHM_U);
- foamyHexMesh with one level of refinement (fHM_R).

A representation of these different meshing strategies can be found in Fig. 3. When present, the refinement level is located around the grain surface, in order to allow a better discretisation of the boundary layers (of the momentum and concentration/temperature gradients). For each meshing strategy, meshes of different cell size were built, paying attention to both the mean cells size and the total number of cells. This latter parameter, see Table 1, was used to analyse the grid convergence results. Technical details on the meshing generation are reported in the Appendix. Equations (1) and (2) were solved imposing a pressure drop between two opposite

Table 1 Total number of cells for ten increasingly fine meshes

| | sHM_U | sHM_R | fHM_U | fHM_R |
|----|---------|---------|---------|---------|
| 1 | 1296 | 1296 | 1278 | 914 |
| 2 | 3744 | – | 3768 | 2799 |
| 3 | 5952 | 4176 | 5991 | 4944 |
| 4 | 11,376 | 8016 | 11,509 | 10,133 |
| 5 | 22,800 | 21,232 | 22,792 | 22,919 |
| 6 | 44,976 | 48,288 | 44,987 | 46,638 |
| 7 | 86,352 | 84,280 | 87,856 | 90,763 |
| 8 | 172,944 | 168,688 | 182,910 | 169,449 |
| 9 | 302,880 | 309,360 | 333,254 | 280,978 |
| 10 | 601,352 | 606,760 | 698,629 | 501,542 |

The four columns report cell number for an uniform and a wall-refined meshing strategy for snappyHexMesh and foamyHexMesh respectively

faces of the domain (along the x -axis). The resulting (mean) velocity is thus comparable with environmental applications ($Re \approx 3$).

4.1.2 Fluid flow convergence study

Using the meshes described above, we solve the Stokes⁶ problem Eqs.(1) and (2), and analyse the numerical performance of the finite-volume schemes. Technical details on the meshing generation are reported in the Appendix.

In Fig. 4 we can observe the effect of the various meshing strategies, combined with different schemes. The mean velocity (or equivalently the drag coefficient) is compared against a finer resolution result (and reference results in literature [24]). First of all, it is interesting to notice how in the first part of the plot (i.e.: coarse grids) the main error is always approximately second-order. This is due to the fact that, in our meshing strategy, the surface of the sphere is not exactly preserved but it is itself discretised. This can be easily estimated from simple geometrical arguments and the theory of elliptic operators. This “geometrical error” can instead behave very differently in presence of significant boundary-layer effects (Navier-Stokes or boundary reactions). After this first regime, we can observe that, in general, the uniform meshes generated with snappyHexMesh perform significantly better reaching a second- and third-order convergence, according to the finite-volume scheme chosen. On the other hand, the refined grid sHM_R loses the second-order accuracy due to the introduced non-orthogonality and skewness of the cells. As explained earlier, this can be corrected only with explicit

⁶ We remind the reader that although in the OpenFOAM® implementation the equation actually solved is Navier-Stokes, given that the simulations are performed in laminar flow, this is equivalent to the simple Stokes problem.

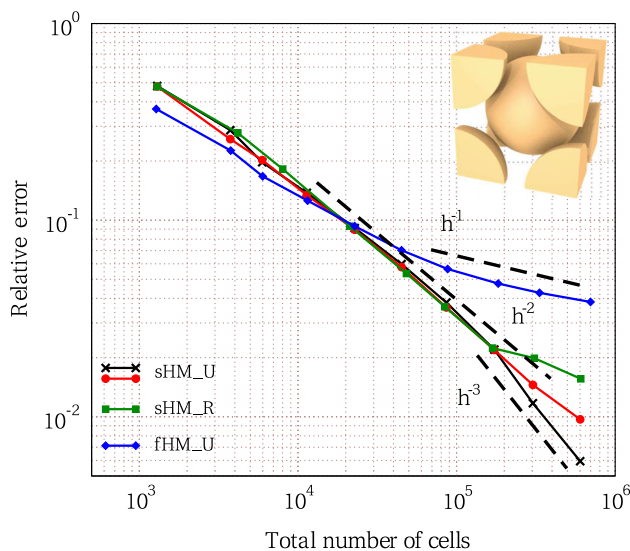


Fig. 4 Numerical test to verify the convergence properties of the solver for Stokes flow around the body-centred cubic packing (shown in the right upper corner). Relative error in the mean velocity with respect to the total degrees of freedom of the mesh: uniform mesh with third order schemes (black \times), uniform mesh with second order scheme (red \circ), refined mesh (green \square), unstructured Voronoi mesh (blue \diamond). Power law curves are reported according to an equivalent cell size h

terms that could cause unboundedness. Therefore we limit the corrections, and obtain a sub-optimal convergence rate.

For this linear problem, the advantage of refining near the boundary does not pay off, as the computational savings are not enough to justify this deterioration of the numerical convergence. Similarly, the unstructured meshing strategies based on Voronoi tessellation (fHM both uniform and refined), despite being attractive for their robustness and simplicity, are characterised only by a linear decay of the error; please note that fHM_R is not reported in Fig. 4 for sake of simplicity, as it practically behaves like the fHM_U. However they seem to perform better for very coarse meshes where the geometrical features can be better represented with a small number of Voronoi tiles.

4.1.3 Residence time distribution and breakthrough curves

The solute transport problem can be solved, either in a Lagrangian or Eulerian way, given the computation of the flow field (from now on, we will refer only to the results obtained with the finest uniform mesh with second-order schemes). For infinite Péclet numbers (purely hyperbolic equation), the transport can be fully characterised by the streamlines. This is particularly convenient as, within an Eulerian formulation, specific advection schemes and highly refined meshes would be needed to avoid significant numerical artificial diffusion or instabilities [3]. Streamlines are integrated with the streamLine post-processing tool available

in OpenFOAM® 4, and they are generated seeding a cloud of points that can be located wherever within the domain. In particular, we provided points seeded on the inlet face of the domain, perpendicular to the flow direction. Points can be placed either randomly or on a regular grid; the former strategy should be preferred to provide statistically reliable samples.

In presence of finite (possibly very small) diffusion, a linear drift diffusion Ito's Stochastic Differential Equation should instead be solved. This requires more sophisticated and expensive random walk algorithms. These are usually coupled to the fluid flow equations, allowing a possible coupling between the solute concentration and the fluid. Despite the higher accuracy of Lagrangian methods for advection-dominated transport, the Lagrangian discretisation of more complex non-linear advection diffusion reaction and the coupling with the fluid flow can be non trivial and give rise to significant statistical errors due to the finite number of trajectories.

Here we follow 1000 particle trajectories from the inlet to the outlet and compute the cumulative distribution of arrival (or residence) times $f(t)$ in the pure advection regime, and compare them with Eulerian simulations at increasing Péclet numbers. The cumulative arrival times of a random ensemble of particles initially positioned at the inlet is equivalent to the time integral of the concentration of particles at the outlet with inlet concentration delta-distributed in time (the “response” function of an impulse at time $t = 0$). This is therefore equivalent to the time-derivative of the so-called “breakthrough curve”, i.e. the response to an Heaviside function.

Having established a connection between Eulerian and Lagrangian simulations in the non-reacting case, what we want to study is the appearance of “anomalous transport” characteristics when the diffusion is small enough not to allow for a significant mixing, therefore resulting in particles following preferential directions and staying for long times in the same streamline. This, apart from very special cases (e.g. Gaussian distributed flow field), generates non-Fickian transport behaviour. This is a well-known problem in porous media who has received significant attention (see, for example, [12] and [13]). It can be easily understood if one tries to fit the three-dimensional results with an equivalent one-dimensional linear advection diffusion equation (with constant velocity, equal to the mean velocity, and unknown diffusion⁷).

For the purpose of this section, simulation were performed on a face-centred body structure (made of spherical grains of

⁷ In this case, it would be more correct to talk about hydro-dynamical dispersion as the effective diffusion is caused both by molecular diffusion and by mechanical dispersion due to the heterogeneous velocity field.




Fig. 5 Breakthrough curves $f(t) = \langle c \rangle_{u_x}$ for a single face-centered cubic cell computed from Eulerian simulations at finite Péclet numbers (respectively in red, green, and blue for $Pe = 10, 100$, and 1000) compared with the Lagrangian residence time distribution (in black). For the Eulerian cases, the continuous lines result from employing the classic boundary conditions, while the dashed lines are the results employing the scalar quasi-periodic boundary conditions described in Sect. 2.3

equal size) with porosity 0.40. In Fig. 5 the residual concentration $1 - f(t)$ is shown for Lagrangian streamlines (equivalent to infinite Péclet number) and for Péclet numbers ranging from 10 to 1000, against the dimensionless advective time t/τ , where τ is the arrival time for the advective front (i.e., considering the mean velocity). At first, a uniform (across the inlet plane) boundary condition has been tested (ref. Fig. 5 continuous lines). Here it can be seen that for $Pe = 10$ the transport is still dominated by diffusion, with a first arrival time lower than the front arrival time. Then, for higher Péclet numbers, the curves start to show a significant power-law “tailing”, while the solution of a 1D advection diffusion equation would predict an exponential decay. This is due to the incomplete mixing and preferential trajectories. Although it is not in the objectives of the paper to analyse the anomalous transport characteristics, it is interesting to notice here that a significantly different behaviour is observed when the technique and boundary conditions described in 2.3 are used (dashed lines). In fact, when the non reactive stationary quasi-periodic problem is solved within the porous space, an “asymptotic” self-similar profile is obtained. When using this asymptotic profile as the boundary condition at the inlet, the resulting transport becomes clearly “Fickian”, with exponential tails. This is because the profile found by the quasi-periodic “closure” represents the profile for which the stationary outlet concentration is uniform. Therefore, it can be considered as an inverse problem of finding the profile that cancels out the preferential trajectories. The black solid line finally represents the arrival (or transit or residence) times of purely advective Lagrangian trajectories. As it can be seen, the first arrival times match well with the high Péclet case

but the absence of diffusion cause a very significant tailing with a power-law behaviour.

4.2 Flow over porous surfaces

In this section, we will explore the case of flow and heat transfer in a pipe containing, for a small portion of its volume, a number of spherical solid grains settled on its lower side. The problem of studying heat transfer in pipes, and consequently evaluating the heat transfer coefficient in pipes has been extensively studied [4,8,20,26,29]; in this work, we want to show how the presence of randomly arranged solid matter affects fluid flow and heat transfer. We compare the results with the analogous “clean” cases where no deposited matter is present. In particular, this section will serve as a proof-of-concept for the methodology we employed for both the generation of the geometry comprising the settled matter (as described in the following paragraphs) and the easiness with which it is possible to explore geometrical parameter space and operating conditions. Then, an overview of the numerical details of the performed simulations, as well as the results obtained, are presented.

4.2.1 Geometry generation

As it has been mentioned, all the geometrical models used in this work were built *in-silico*, avoiding the costly step of obtaining experimental samples and adapting them for use in CFD codes. Specifically, the case study presented in this section was created with Blender, a free and open-source software used in computer graphics applications, which can be used to perform rigid-body simulations.

The settled matter is represented by uniformly-sized spheres, which were chosen to be of a diameter equal to one-twentieth of the pipe diameter. In the Blender simulation setup, grains are first located above the container (which is now represented by just the bottom half of the pipe, to let the granular matter settle into it), taking care of adding randomness to their initial position in order to avoid unwanted “arranged” structures in the final packing. Then, the solid grains are free to fall due to the effect of gravity, which is the only outside force. The final state of the system is given by the solution of the rigid-body problem, where the solid grains are treated as non-deformable and impenetrable entities⁸, and the only interactions considered are instantaneous collisions and sliding. After reaching the final state, identified by the absence of relative motion between all the solid grains and the pipe, the geometric model is exported for its use in the CFD code (substituting the half-pipe with the complete pipe for the simulations). This kind of methodology has been suc-

⁸ More precisely, the definition for the *rigid bodies* is given as solids possessing an infinite repulsive potential present on their surface.




Fig. 6 Velocity contour plot inside the pipe and fluid streamlines (coloured by fluid velocity) inside the granular portion of the domain, showing the increasing distortion of the fluid paths moving from the bulk towards the bottom of the pipe

cessfully used to represent both packings of catalyst particles in chemical reactors [5,7] and in digital rock physics applications [19]: the numerical details of the rigid-body simulations can be found therein.

With the process just described, we created three different geometries, each characterised by the (increasing) volume of settled matter, roughly corresponding to an average height, H_g , of the settled matter packings equal to two, three, and four grain diameters respectively. In all cases, the pipe length is equal to three times the pipe diameter. The spatial domain was discretised with a sHM_R strategy. Also, an in-depth description of the use of this mesh generator in the case of random packings and an analysis of mesh convergence for the study of advection-transport problems can be found in our earlier work [18].

4.2.2 Numerical details

The first step, as for the case described in the previous section, has been to solve the equation for the motion of fluid inside the pipe. We thus solved Eqs.(1) and (2) in laminar regime by setting a constant pressure drop between the inlet and the outlet of the domain, corresponding to a Reynolds number $Re \approx 1$ (we used again second-order schemes). A snapshot of the velocity contour plot inside the pipe can be found in Fig. 6, together with a representation of the slower fluid streamlines inside the volume occupied by the settled granular matter.

A simulation of heat transfer was then performed solving Eq. (6), setting a constant fluid temperature at inlet, equal to $T=293$ K, and a zero flux at the outlet. A Dirichlet condition of constant temperature $T=343$ K was set on the pipe wall, while a zero thermal flux (corresponding to $\nabla T = 0$) was set on the grains. This choice of boundary conditions is equivalent to assume a situation of perfect thermal equilibrium between the flowing fluid and the solid grains: this is justified in the limiting cases of infinitely fast heat transport




Fig. 7 Temperature contour plots on an axial plane of the pipe for cases with average granular matter height from the bottom of the pipe equal to roughly four, three, and two grain diameters (from top to bottom), and for the “clean” pipe (bottom)

in the solid, if one considers the long-time, steady-state, solution of the problem. A snapshot of the contour plots of fluid temperature on the median length-wise section of the pipe⁹ is found in Fig. 7, which shows the amount of solid matter in each case and the slight difference in the thermal boundary layer at the bottom.

In the next section, we will instead explore the problem of pure heat transfer in solids, between a continuous matrix and a disperse solid phase, characterised by two different thermal diffusivity values: these two cases are meant to represent, in a way, two opposite situations in the study of heat transfer in heterogeneous materials.

4.2.3 Heat transfer coefficient

When the steady-state solution of the problem of advection-diffusion has been obtained, the heat fluxes along the main flow direction are extracted in order to calculate the local heat transfer coefficient h_{loc} [4] evaluated at a number of successive planes orthogonal to the main flow direction (i.e. the pipe axis):

⁹ More precisely, taken on a plane parallel to both the main flow direction and the direction of gravity, positioned on the pipe axis.




Fig. 8 Heat transfer coefficient h_{loc} as a function of the distance from the inlet (in meters): values for the “clean” pipe (black \times , line at the top) and for granular matter height from the bottom of the pipe equal to roughly two, three, and four grain diameters (respectively red \circ , green \square , and blue \diamond curves, in descending order in the graph)

$$h_{loc} = \frac{q\rho C_p}{\pi D} \frac{dT_{bulk}(x)}{dx} \frac{1}{T_{wall} - T_{bulk}} \quad (20)$$

where q (m s^{-1}) is the average flow velocity, ρ (kg m^{-3}) its density, C_p ($\text{J kg}^{-1} \text{K}^{-1}$) its specific heat capacity, x (m) the distance of the averaging plane from the inlet boundary, and T_{wall} and T_{bulk} (K) respectively the temperature on the wall of the pipe and the local average of the fluid bulk temperature.

The h_{loc} results for each case are shown in Fig. 8, compared with the heat transfer coefficient of the “clean” pipe devoid of settled matter. As it can be seen, the presence of granular solids at the bottom of the pipe cause two effects. The first is that of reducing the total heat transfer coefficient (and consequently Nusselt number), which is smaller the bigger the volume of settled particles: this is due to the reduction of the contribution of advective heat transport with respect to conductive transport, caused by the much lower fluid velocities in the inter-granular portion of the domain. Then, it can also be noticed that the presence of randomly arranged grains have an effect on the dynamics of reaching an asymptotic regime: while in the “clean” pipe case the heat transfer coefficient reaches a plateau not far from the inlet (as it is expected), for the cases with granular matter it is clear that a decreasing trend is still present even close to the domain outlet.

4.3 Heat transfer in heterogeneous materials

In the preceding section, we only considered the effect of the presence of a granular solid on flow and heat transfer in pipes, neglecting the heat transfer between the bulk of the fluid (or the pipe wall) and the grains themselves. In this section, we will instead explore a case where no fluid flow is present, investigating the case of heat transfer between two

solid phases with different thermal diffusivities α_1 and α_2 , the first phase constituting a continuous matrix and the second phase dispersed (as solid inclusions) in the first; moreover, the diffusivity in the continuous matrix is anisotropic, having different values for its longitudinal and transversal components.

4.3.1 Extended Jodrey–Tory algorithm (EJT)

For this application, a periodic extension of the classical Jodrey–Tory algorithm [22] has been implemented to deal with randomly oriented ellipsoids. Despite not being computationally optimal for low-porosity (high-density) materials, it can easily generate random packings with a specific porosity and a specific overlapping between grains, as well as clustering. The new algorithm has been implemented in the open-source repository *mlmc-porescale*[19], and can be sketched as follows:

1. The user specifies a domain size (including periodicity information), a target porosity, the minimum degree of overlapping θ ($\theta < 1$ allows overlapping, $\theta > 1$ prescribe a certain distance from grain to grain), a maximum distance between grains $\Theta > \theta$ (this is enforced on clusters of N_Θ particles), and a maximum limit of particle numbers and displacement iterations. Furthermore, geometrical statistics about the ellipticity of the grains and their size have to be provided. In the simplest (isotropic) case, a single random variable is selected (log-normal, truncated Gaussian or uniform) to describe the ellipses canonical axes length.
2. Random ellipsoids are generated until they reach the desired porosity (not considering overlapping). For each ellipsoid, a random unitary matrix C is sampled [34] according to the axes lengths and orientation statistics. This is then decomposed $C = C_1 C_2$ into a rigid rotation part $C_1 \in SO(3)$ and a diagonal part C_2 (responsible for scaling).
3. Once all ellipsoids have been randomly placed in the box, a greedy-type algorithm is executed: at each iteration, the pair-wise distances are generated and the first N_{moves} pairs, characterised by the largest overlapping ratio, are detached along the line of the pair-wise distance axis. The displacement length controls the convergence properties of the algorithm but its optimal value strongly depends on the porosity of the packing. Usually a relative displacement of 0.2–0.5 times the particle size is chosen.
4. If a maximum distance is specified, the same type of moves are applied to shorten the pair-wise distance. Here, however, this can be only applied to a specific cluster of particles selected as the closest ones.
5. The moves are applied iteratively until all the pair-wise distances satisfy the criteria. Periodicity is ensured by




Fig. 9 Example of random ellipsoids with random (statistically uniform in $SO(3)$) orientation, generated with the Extended Jodrey-Tory algorithm (EJT)

implementing the appropriate toric metric to compute the distances.

It has to be noted that there is not ensured stopping point for the algorithm and a total number of moves of the greedy algorithm described in point 3. is fixed as an upper bound. For a given available volume and number of ellipsoids, in the current implementation of the code, one needs to choose the overlapping layer appropriately and large enough for a packing solution to exist. In the future, however, simple adaptivity could be implemented to override the settings for overlaps or for ellipsoid size, to converge faster to a final arrangement.

An example of 200 touching ($\theta = 1$, $\Theta = \infty$) ellipsoids in a periodic arrangement with 0.45 porosity is shown in Fig. 9.

4.3.2 Effective diffusion

The micro-scale heat transfer problem to be solved is Eq. (6) with $\mathbf{u} = 0$ and

$$\alpha(\mathbf{x}) = \begin{cases} \text{diag}(0.1, 10, 10) & \text{for } \mathbf{x} \in \Omega_1, \\ \text{diag}(1, 1, 1) & \text{for } \mathbf{x} \in \Omega_2 \end{cases}$$

Where diag is the diagonal matrix constructed from the vector. This represents the case of a continuous phase made of highly conductive layers (e.g., graphene, therefore with large conductivity in only two directions), reinforced with ellipsoidal inclusions with isotropic conductive properties.




Fig. 10 Representation of the complete geometry, showing the mesh of two distinct domains, with the dispersed solid inclusions generated with the Extended Jodrey-Tory algorithm highlighted as transparent red volumes

In Fig. 10, the periodic structure (with the same properties of the one in Fig. 9 but with up to a 20% overlapping allowed) is shown with the underlying locally refined mesh (sHM_R). We have shown above that the local refinements destroys the second-order convergence of the numerical schemes. However, due to the discontinuous diffusion coefficient, a second-order convergence would not be possible with our numerical schemes.

To compute the effective (homogenised) diffusion coefficient, one may follow a volume averaging strategy, applying the quasi-periodic BCs in Eq. (10) (results shown in Fig. 11). This is more convenient than solving the homogenisation cell problem where the derivative of the discontinuous conductivity appears as a source term of the equation [17]. The resulting effective diffusion is, in this case:

$$\alpha_{eff} = \begin{pmatrix} 4.955 & \epsilon & \epsilon \\ \epsilon & 39.86 & \epsilon \\ \epsilon & \epsilon & 38.41 \end{pmatrix}$$

where ϵ represents small coefficients $\approx 10^{-3}$. As it can be seen, the resulting effective diffusion, as expected, is again anisotropic.

5 Conclusions

In this work, we have shown the computational challenges involved in the simulation of transport problems in the field of porous media or, more generally, in problems which feature multiple *spatial* scales, or whose evolution takes place over multiple *temporal* scales: examples of which we have




Fig. 11 Local temperature fluctuations $c - \mathbf{p} \cdot \mathbf{x}$ for the quasi-periodic cell problem. The volume average of their gradient gives the effective diffusion coefficient

respectively provided in the form of the case of a pipe with small-size granular sedimentation, and the study of solute transport at very high Péclet numbers. Moreover, we have also presented applications of a number of third-party open-source codes and showcased our freely available extensions. These range from tools used to quickly and robustly generate a wide variety of three-dimensional realistic porous media models, to innovative computational setups providing for a great decrease of the computational expense in quasi-periodic cases [10]. In the first case, for example, an extended Jodrey–Tory algorithm has been developed, to generate random packings of ellipsoids, while, in the latter, a simple steady-state closure problem, fully consistent with volume averaging and homogenisation, is derived for time-dependent transport with linear (and constant in time) transport parameters. Then, the semi-periodic setting allows to solve one single periodic cell, while a Representative Elementary Volume is, in the case of advection-diffusion, dependent on the Péclet numbers of the system, and possibly including tens or hundreds of periodic cells.

Given our focus on free and open-source software and having primarily chosen, due to its robustness and flexibility, the OpenFOAM® platform to discretise our models, we have illustrated the main features of the meshing and the discretisation schemes underlying its classical finite-volume approach. We have explored the alternatives provided in the OpenFOAM® suite regarding the choice of the meshing preprocessors, providing an in-depth analysis about the order of relative error convergence of the two approaches, while also exploring different meshing strategies. Again, this is

done with the objective of providing guidance and an initial benchmark with respect to the pre-processing part of the numerical simulation, which probably has the greatest impact on the accuracy of the final results of the model.

The simulations presented, although representing three-dimensional, realistic, and physically meaningful systems, do not represent real large-scale structures and therefore do not require any sophisticated adaptivity and parallelism.

In this work therefore, we have chosen to limit our investigation to showing the difference between two radically different body-fitted approaches, a Cartesian-based cut-cell approach and a Voronoi unstructured approach, presenting the differences in accuracy between the two, with the focus on building a general-purpose package for medium-scale simulation. OpenFOAM® has, however, a very flexible parallel structure and, as such, most of our solvers are inherently parallel, with some extra functionality (particularly for I/O and averaging operations) implemented in a parallel way. Depending on the overall size of the geometry, this can easily be parallelisable, with good scaling properties, up to a few hundreds cores. It has to be noted that, while OpenFOAM® can easily be parallelised on many thousands cores (although, depending on the specific solver, losing the optimal scaling), in our experience, a multi-scale micro-macro offline coupling could be a more viable and robust option to deal with porous media.

Finally, the aggregate results of this work demonstrate how it is possible to build a completely computational pipeline for the study of a variety of physical problems, from the generation of a suitable and realistic geometric model, to the actual simulation run, to the extraction of the relevant data. This already found many applications in the field of environmental and chemical engineering and will be extended for other applications, for example, in energy storage. This emphasis is meant and presented both as an alternative between free and open-source software as opposed to licensed and “closed” software, and between *in-silico* methods of reconstruction of the geometry of the porous medium as opposed to more costly (and often less available) methods such as micro-computed-tomography or x-ray imaging. In both cases, the clear purpose is to work towards the development of open scientific frameworks providing easily reproducible and verifiable results that links rigorous mathematical models and numerics with physics and engineering applications.

Acknowledgements MI acknowledge the financial support of AVL and the computational resources provided by the Warwick Centre for Scientific Computing.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material

in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

A OpenFOAM® meshing details

In this appendix, we explain the features of `foamyHexMesh` (fHM) and `snappyHexMesh` (sHM) used to generate the meshes of all our numerical results. fHM was introduced in the 2.3.0 version; it is based on a Delaunay tessellation, and the resulting Voronoi dual is typically a hex-dominated mesh [31]. This mesher tool was thought to provide a good feature conformation, avoiding high aspect ratio cells, [31]. To the best of our knowledge, detailed guidelines for the usage of fHM for porous materials are not available. Thus in the following of this section, prior to provide the parameters used to generate the presented meshes, we present a quick description of the key parameters of fHM, while making a parallelism with sHM, highlighting advantages and disadvantages of these two meshing tools. Both sHM and fHM deals with in-built analytical surfaces and .stl files. Analytical surfaces (spheres, boxes, cylinders) are readily implementable in sHM, while this is not true in fHM, where the user must explicitly manage the interception of the surfaces with the (domain) bounding box (or among the surfaces themselves when necessary); thus the usage of externally generated geometries is of easier implementation. In this work, when needed, .stl geometries have been generated with openSCAD [33]. With fHM the use of .stl allows a straightforward implementation of the surface refinement, as it is sufficient to include an extra +.stl file (or a .stl composed of different regions/surfaces) to address those surface on which a greater refinement (moving from the surface to the bulk region) is required.

With sHM the dimension and number of cells, initialised in the starting Cartesian mesh properties, encoded in the dictionary `blockMeshDict`, can be manipulated with `refinementRegions`. In fHM, the way cell dimensions can be controlled depends instead on the `initialPointsMethod` selected. In particular, we tested both the `uniformGrid` and the `autoDensity` functions. In the former case, the cell size (and thus the total number of cells) is controlled fixing the `initialCellSize` parameter, while in the `autoDensity` case it is possible to define `defaultCellSize`, in the `motionControl` section¹⁰. In both cases, the size variable makes reference to the mean cell side length (keeping in mind that the mesher tries to

produce hexaedral cells). Notice that, since sHM relies on a Cartesian approach, while fHM does not, this implicitly implies that sHM usually allow a better control on the final number of cells per grains diameters. Another parameter of interest in fHM is the `maxSmoothingIterations` number (motionControl section). It is intuitive that the higher is the number of iterations, the higher is the mesh quality, and this is particularly true when going from single core to parallel run of fHM, as this last approach requires a greater number of iterations to achieve meshes of comparable quality.

The most important parameters used to generate the meshes in Sect. 4.1 are reported below:

sHM_U

- `blockMeshDict` → `blocks`: number of cells.
- `snappyHexMeshDict` → `refinementSurfaces`: level (0 0).

sHM_R

- `blockMeshDict` → `blocks`: number of cells.
- `snappyHexMeshDict` → `refinementSurfaces`: level (2 2).

fHM_U

- `initialPoints`: `initialPointsMethod` `autoDensity`.
- `motionControl`: `defaultCellSize`.
- `motionControl`: `minimumCellSizeCoeff` 10% `defaultCellSize` value.
- `motionControl`: `maxSmoothingIterations` 500.
- `motionControl`: `maxSmoothingIterations` 0.

fHM_R

- `initialPoints`: `initialPointsMethod` `autoDensity`.
- `motionControl`: `defaultCellSize`.
- `motionControl`: `minimumCellSizeCoeff` 10% `defaultCellSize` value.
- `motionControl`: `maxSmoothingIterations` 500.
- `motionControl`: `maxSmoothingIterations` 0.
- `motionControl` → `shapeControlFunctions` → `surfaceToBeRefined` → `uniformValueCoeffs`: `surfaceCellSizeCoeff` 0.5;

¹⁰ Notice that when the `uniformGrid` method is used, `motionControl` fields are practically ignored, particularly if `initialCellSize` \gg `defaultCellSize`.

References

- Battiato, I., Tartakovsky, D.: Applicability regimes for macroscopic models of reactive transport in porous media. *J. Contam. Hydrol.* **120**, 18–26 (2011)
- Bear, J.: *Dynamics of Fluids in Porous Media*. Dover, New York (1988)
- Benson, D.A., Aquino, T., Bolster, D., Engdahl, N., Henri, C.V., Fernández-García, D.: A comparison of Eulerian and Lagrangian transport and non-linear reaction algorithms. *Adv. Water Resour.* **99**, 15–37 (2017)
- Bird, R., Stewart, W., Lightfoot, E.: *Transport Phenomena*. Wiley, New York (1960)
- Boccardo, G., Augier, F., Haroun, Y., Ferré, D., Marchisio, D.L.: Validation of a novel open-source work-flow for the simulation of packed-bed reactors. *Chem. Eng. J.* **279**, 809–820 (2015)
- Boccardo, G., Crevacore, E., Sethi, R., Icardi, M.: A robust upscaling of the effective particle deposition rate in porous media. *J. Contam. Hydrol.* **212**, 3–13 (2018)
- Boccardo, G., Sethi, R., Marchisio, D.L.: Fine and ultrafine particle deposition in packed-bed catalytic reactors. *Chem. Eng. Sci.* **198**, 290–304 (2019)
- Choi, C.Y., Waller, P.M.: Momentum transport mechanism for water flow over porous media. *J. Environ. Eng.* **123**(8), 792–799 (1997)
- Crevacore, E., Boccardo, G., Marchisio, D., Sethi, R.: Microscale colloidal transport simulations for groundwater remediation. *Chem. Eng. Trans.* **47**, 271–276 (2016)
- Crevacore, E., Tosco, T., Sethi, R., Boccardo, G., Marchisio, D.: Recirculation zones induce non-Fickian transport in three-dimensional periodic porous media. *Phys. Rev. E* **94**(5), 053118 (2016)
- Darwish, M.S., Moukalled, F.: TVD schemes for unstructured grids. *Int. J. Heat Mass Transf.* **46**(4), 599–611 (2003)
- Dentz, M., Carrera, J.: Mixing and spreading in stratified flow. *Phys. Fluids* **19**, 017,107 (2007)
- Dentz, M., Gouze, P., Russian, A., Dweik, J., Delay, F.: Diffusion and trapping in heterogeneous media: An inhomogeneous continuous time random walk approach. *Adv. Water Resour.* **49**, 13–22 (2012)
- Ferziger, J.H., Peric, M.: *Computational Methods for Fluid Dynamics*. Springer, Berlin (2002)
- Foundation, T.O.: *OpenFOAM—The Open Source CFD Toolbox—Programmer’s Guide*. United Kingdom, 3.0.1 edn (2015)
- Geuzaine, C., Remacle, J.F.: Gmsh: A 3-d finite element mesh generator with built-in pre- and post-processing facilities. *Int. J. Numer. Methods Eng.* **79**(11), 1309–1331 (2009)
- Hornung, U.: *Homogenization and Porous Media*, vol. 6. Springer, Berlin (2012)
- Icardi, M., Boccardo, G., Marchisio, D.L., Tosco, T., Sethi, R.: Pore-scale simulation of fluid flow and solute dispersion in three-dimensional porous media. *Phys. Rev. E* **90**(1), 013,032 (2014)
- Icardi, M., Boccardo, G., Tempone, R.: On the predictivity of pore-scale simulations: estimating uncertainties with multilevel Monte Carlo. *Adv. Water Resour.* **95**, 46–60 (2016)
- Incropera, F.P., Lavine, A.S., Bergman, T.L., DeWitt, D.P.: *Fundamentals of Heat and Mass Transfer*. Wiley, New York (2007)
- Jasak, H.: *Error analysis and estimation for the finite volume method with applications to fluid flows* (1996)
- Jodrey, W., Tory, E.: Computer simulation of isotropic, homogeneous, dense random packing of equal spheres. *Powder Technol.* **30**(2), 111–118 (1981)
- Joekar-Niasar, V.: Pore-scale modelling techniques: balancing efficiency, performance, and robustness. *Comput. Geosci.* **20**(4), 773 (2016)
- Khirevich, S., Ginzburg, I., Tallarek, U.: Coarse- and fine-grid numerical behavior of MRT/TRT lattice-Boltzmann schemes in regular and random sphere packings. *J. Comput. Phys.* **281**, 708–742 (2015)
- Krol, M.M., Oleniuk, A.J., Kocur, C.M., Sleep, B.E., Bennett, P., Xiong, Z., O’Carroll, D.M.: A field-validated model for in situ transport of polymer-stabilized nZVI and implications for subsurface injection. *Environ. Sci. Technol.* **47**(13), 7332–7340 (2013)
- Manes, C., Pokrajac, D., McEwan, I., Nikora, V.: Turbulence structure of open channel flows over permeable and impermeable beds: a comparative study. *Phys. Fluids* **21**(12), 125,109 (2009)
- Messina, F., Marchisio, D.L., Sethi, R.: An extended and total flux normalized correlation equation for predicting single-collector efficiency. *J. Colloid Interface Sci.* **446**, 185–193 (2015)
- Mosby, M., Matouš, K.: Computational homogenization at extreme scales. *Extreme Mech. Lett.* **6**, 68–74 (2016)
- Möbner, M., Radespiel, R.: Modelling of turbulent flow over porous media using a volume averaging approach and a reynolds stress model. *Comput. Fluids* **108**, 25–42 (2015)
- Moukalled, F., Mangani, L., Darwish, M.: *The Finite Volume Method in Computational Fluid Dynamics*. Springer, Berlin (2016)
- OpenFOAM: <http://openfoam.org/> (2016)
- OpenQBMM: <https://www.openqbmm.org>, <https://doi.org/10.5281/zenodo.161617>
- OpenSCAD: <http://www.openscad.org/>
- Ozols, M.: How to generate a random unitary matrix (2009)
- Perović, N., Frisch, J., Salama, A., Sun, S., Rank, E., Mundani, R.P.: Multi-scale high-performance fluid flow: simulations through porous media. *Adv. Eng. Softw.* **103**, 85–98 (2017)
- Rhie, C.M., Chow, W.L.: A numerical study of the turbulent flow past an isolated airfoil with trailing edge separation. *AIAA J.* **21**, 1525–1532 (1983)
- Rolle, M., Chiogna, G., Hochstetler, D.L., Kitanidis, P.K.: On the importance of diffusion and compound-specific mixing for groundwater transport: an investigation from pore to field scale. *J. Contam. Hydrol.* **153**, 51–68 (2013)
- Shen, W.Z., Michelsen, J.A., Nørkaer Sørensen, J.: Improved Rhie–Chow interpolation for unsteady flow computations. *AIAA J.* **39**(12), 2406–2409 (2001)

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