AN ABSTRACT OF THE DISSERTATION OF

Justin Richard Finn for the degree of Doctor of Philosophy in Mechanical Engineering presented on March 20, 2013.

Title: A Numerical Study of Inertial Flow Features in Moderate Reynolds Number Flow Through Packed Beds of Spheres

Abstract approved: _____

Sourabh V. Apte

In this work, flow through synthetic arrangements of contacting spheres is studied as a model problem for porous media and packed bed type flows. Direct numerical simulations are performed for moderate pore Reynolds numbers in the range, $10 \leq Re \leq 600$, where non-linear porescale flow features are known to contribute significantly to macroscale properties of engineering interest.

To first choose and validate appropriate computational models for this problem, the relative performance of two numerical approaches involving body conforming and non-conforming grids for simulating porescale flows is examined. In the first approach, an unstructured solver is used with tetrahedral meshes, which conform to the boundaries of the porespace. In the second approach, a fictitious domain formulation (Apte et al., 2009. J Comput Phys 228 (8), 2712-2738) is used, which employs non-body conforming Cartesian grids and enforces the no-slip conditions on the pore boundaries implicitly through a rigidity constraint force. Detailed grid convergence studies of both steady and unsteady flow through prototypical arrangements of spheres indicate that for a fixed level of uncertainty, significantly lower grid densities may be used with the fictitious domain approach, which also does not require complex grid generation techniques.

Next, flows through both random and structured arrangements of spheres are simulated at pore Reynolds numbers in the steady inertial ($10 \lesssim Re \lesssim 200$) and unsteady inertial ($Re \approx 600$) regimes, and used to analyze the characteristics of porescale vortical structures. Even at similar Reynolds numbers, the vortical structures observed in structured and random packings are remarkably different. The interior of the structured packings are dominated by multi-lobed vortex rings structures that align with the principal axes of the packing, but perpendicular to the mean flow. The random packing is dominated by helical vortices, elongated parallel to the mean flow direction. The unsteady dynamics observed in random and structured arrangements are also distinct, and are linked to the behavior of the porescale vortices.

Finally, to investigate the existence and behavior of transport barriers in packed beds, a numerical tool is developed to compute high resolution finite-time Lyapunov exponent (FTLE) fields on-the-fly during DNS of unsteady flows. Ridges in this field are known to correspond to *Lagrangian Coherent Structures* (LCS), which are invariant barriers to transport and form the skeleton of time dependent Lagrangian fluid motion. The algorithm and its implementation into a parallel DNS solver are described in detail and used to explore several flows, including unsteady inertial flow in a random sphere packing. The resulting FTLE fields unambiguously define the boundaries of dynamically distinct porescale features such as counter rotating helical vortices and jets, and capture time dependent phenomena including vortex shedding at the pore level.

A Numerical Study of Inertial Flow Features in Moderate Reynolds Number Flow Through Packed Beds of Spheres

by

Justin Richard Finn

A DISSERTATION

submitted to

Oregon State University

in partial fulfillment of
the requirements for the
degree of

Doctor of Philosophy

Presented March 20, 2013
Commencement June 2013

Doctor of Philosophy dissertation of Justin Richard Finn presented on March 20, 2013.

APPROVED:

_____

Major Professor, representing Mechanical Engineering

_____

Head of the School of Mechanical Industrial and Manufacturing Engineering

_____

Dean of the Graduate School

I understand that my dissertation will become part of the permanent collection of Oregon State University libraries. My signature below authorizes release of my dissertation to any reader upon request.

_____

Justin Richard Finn, Author

ACKNOWLEDGEMENTS

CONTRIBUTION OF AUTHORS

Dr. Sourabh V. Apte contributed to the development of the numerical methods used in Chapters 2, 3, and 4, as well as to the interpretation of the results throughout this dissertation.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF APPENDIX FIGURES

LIST OF APPENDIX TABLES

To Lillian

A Numerical Study of Inertial Flow Features in Moderate Reynolds Number Flow Through Packed Beds of Spheres

Chapter 1: Introduction

Though we may not always notice it, porous media are ubiquitous in the world around us. On the most basic level, porous media can be thought of as any complex matrix of solids with interconnected voids or "pores" through which a fluid may flow. The space and timescales of fluid transport encountered in both natural and engineered porous media flows vary widely, from atomic to geologic. However, at the geometric scale of individual pores (the *porescale*), such systems often look quite similar. This motivates fundamental studies of the hydrodynamics in simple, prototypical porous geometries. In this work, synthetic arrangements of contacting spheres are used as a simple model to understand the characteristics of porescale flow in more complex systems.

Despite the important connection to both natural and engineered systems, flows through packed beds of spheres are not, generally, well understood. This is particularly true at moderate to high Reynolds numbers where fluid inertia results in non-linear porescale flow features such as jets, vortices and backflow regions. Although these features are local *porescale* phenomena that occur at small scales, they are of great importance for their ability to affect *macroscale* properties of the flow. For example, non-uniform distribution of temperature and nutrients in stream beds depends crucially on on "dead zones", where closed streamlines create hydrodynamic transport barriers [5, 12]. In fixed bed chemical and nuclear reactors, relatively high flow rates are used to improve mixing, heat transfer and reaction rates [1], but can result in non-uniform flow, which may not be predictable a-priori during the reactor design process. To better understand these types of processes and others, it is therefore important to develop an improved characterization

of inertial flow features in porous media.

The inherent geometric complexity of porous media and packed beds has challenged theoretical, numerical, and experimental investigations alike. Relatively modern non intrusive experimental techniques such as magnetic resonance imaging [108, 57], ultrasonic profiling [51], and particle image velocimetry [81, 52] have provided important datasets from which the porescale behavior of inertial flows can be explored. At the same time, rapidly increasing computational power has made high resolution simulation [50, 48, 102, 78, 67, 64] of porous media flows increasingly feasible, and this work is an example of progress in that direction. As a complement to existing experimental data, numerical simulations have the advantage of being able to more easily capture continuous, time resolved, three dimensional data, which at present seems to challenge most experimental techniques.
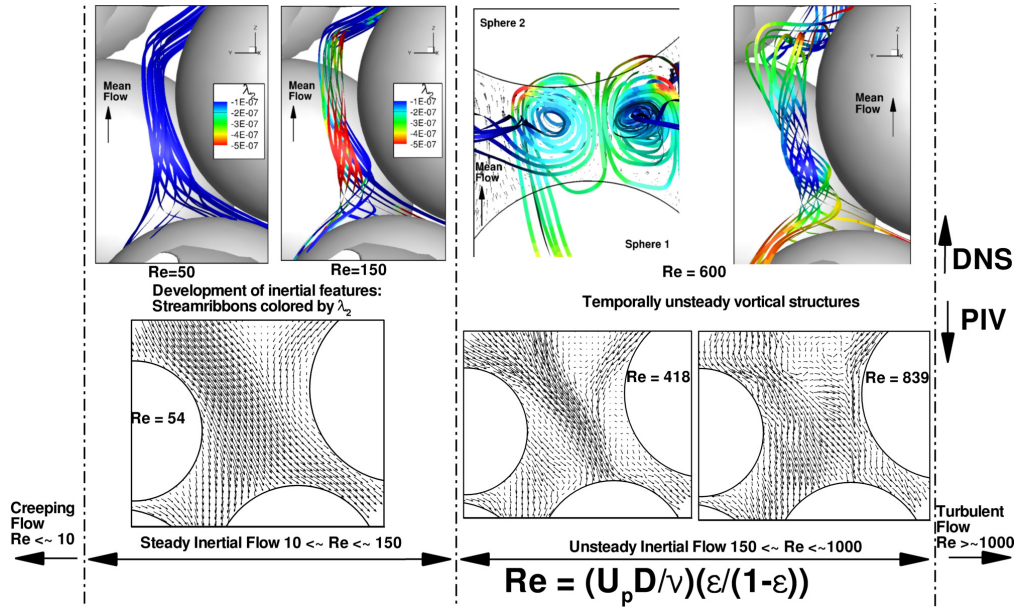
Figure 1.1: Overview of the porescale flow regimes studied in this work. Top and bottom rows correspond to DNS results from the present work and PIV results from [82] respectively.

Dybbs and Edwards [21] used MRI techniques to classify porous media flows

as a function of the pore Reynolds number, $Re_p = \frac{\rho U_p D}{\mu} \frac{\epsilon}{1-\epsilon}$, where ρ and μ are the fluid density and viscosity, U_p is the pore velocity scale, D is the pore length scale, and ϵ is the void fraction of the media. They determined four flow regimes based on the hydrodynamic behavior at the porescale, which are still widely accepted: (i) creeping flow ($Re_p \lesssim 10$), (ii) steady inertial flow ($10 \lesssim Re_p \lesssim 150$), (iii) unsteady inertial flow ($150 \lesssim Re_p \lesssim 1000$), and (iv) chaotic turbulent flow ($Re_p \gtrsim 1000$). As noted by [96, 51], the Reynolds numbers at which these transitions occur is very sensitive to the packing geometry, and values from one packing cannot necessarily be extrapolated to another. Figure 1 provides an overview of the porescale flow regimes covered by this work, and shows representative images of porescale flow features observed by direct numerical simulations (this study, Chapter 3) as well as in experiments by [82]. The figure emphasizes the development of pore scale vortices with increasing Reynolds number and illustrates the contribution of fluid inertia to porescale complexity. The attention of this thesis is focused on the inertial regimes (ii) and (iii), where steady and unsteady non-linear flow features are dominant in the porespace, but where the flow does not have all the symptoms of turbulence [110] such as broad length and time scale separation, and an energy cascade.

Even with powerful computational resources, porous media is challenging for resolved simulation techniques because of the complexity of the solid-fluid interface. In Chapter 2, the relative performance of two techniques capable of resolving the porescale flow field is examined. In the first approach, an unstructured solver is used with tetrahedral meshes that conform to the boundaries of the porespace. This is the more frequently used approach [34, 4, 18], but requires complex unstructured mesh generation techniques. In the second approach, a fictitious domain formulation [2] is used that employs non-body conforming Cartesian grids and enforces the no-slip conditions on the pore boundaries implicitly through a rigidity constraint force. Both approaches are found to be capable of capturing porescale inertial flow features, but it is shown that the fictitious domain approach does so with significantly less computational effort for similar levels of numerical uncer-

tainty and without the meshing overhead required by body fitted meshes. This chapter is in under review (after first revision) for publication in the International Journal of Multiphase Flow.

Vortical flow features are of particular interest in porous media because of the implications they have for mixing and dispersion in the porespace [55]. In Chapter 3, the development and behavior of vortical features are examined using DNS datasets of flow through both random and arranged packed beds of spheres. Even with similar porosities and Reynolds numbers, a number of striking differences exist between the porescale flow fields in the different arrangements studied, in both qualitative and statistical senses. This chapter is in preparation for submission to the journal Physics of Fluids

Although it is clear that there is an organizing structure to flows in porous media and packed beds, it is not obvious how to define the coherent porescale regions that are important to mixing and dispersion properties. The recently developed notion of Lagrangian coherent structures (LCS) [42, 38] as invariant transport barriers in steady and unsteady fluid flows has evolved from dynamical systems theory, and has proven its utility in understanding a number of mixing and transport problems [83, 94]. In Chapter 4 a computational approach is developed that allows the LCS theory to be applied to complex packed bed flows. The LCS are most often defined from ridges in the finite-time Lyapunov exponent (FTLE) field [98]. However, their practical use has been limited to this point because computing the FTLE field typically involves expensive post-processing of large fluid velocity datasets generated either from experiments or numerical simulations. The integrated approach proposed in this work is to compute evolving, transient, three dimensional FTLE fields on-the-fly during the direct numerical simulation. By integrating the computations in this way, tedious post-processing of velocity fields is no longer needed, and larger, more complex problems, such as flow through packed beds become accessible to the application of LCS theory. This chapter is in under final revisions for publication in the journal Chaos.

Chapter 2: Relative Performance of Body Fitted and Fictitious Domain Simulations of Flow Through Packed Beds

Justin R. Finn and Sourabh V. Apte

Abstract

The relative performance of two numerical approaches involving body conforming and non-conforming grids for simulating porescale flow in complex configurations of fixed packed beds of spheres at moderate pore Reynolds numbers ($12 \leq Re \leq 600$) is examined. In the first approach, an unstructured solver is used with tetrahedral meshes that conform to the boundaries of the porespace. In the second approach, a fictitious domain formulation (Apte et al., 2009. J Comput Phys 228 (8), 2712-2738) is used that employs non-body conforming Cartesian grids and enforces the no-slip conditions on the pore boundaries implicitly through a rigidity constraint force. Sphere to sphere contact points, where the fluid gap between solid boundaries becomes infinitesimal, are not resolved by either approach, but this is shown to have a negligible effect on the local flow field at the Reynolds numbers considered. Detailed grid convergence studies of both steady and unsteady flow through simple cubic packings indicate that for a fixed level of uncertainty, significantly lower grid densities may be used with the fictitious domain approach, which also does not require complex grid generation techniques. This translates into large savings for simulation of flow through realistic packed beds, which is shown by both analytic estimates and actual CPU timings. The applicability of the fictitious domain approach is demonstrated by simulating unsteady flow through a randomly packed bed of 51 spheres at a pore Reynolds number of 600. The results are used to examine the dominance of porescale helical vortices.

2.1 Introduction

Even at modest flow rates through porous media and packed beds, non-linear porescale flow features such as jets and vortices can have strong effects on macroscale properties of broader interest including pressure drop, heat transfer, and mass transfer. Many processes in porous media occur at lower flow rates where these inertial effects may be modeled or neglected (ie. many geologic processes). However,

higher flow rates are also of importance to both natural and engineered porous systems. For example, in river and stream beds, recirculating regions can result in non-uniform nutrient and temperature distributions [12], while higher flow rates are sometimes employed by design to increase heat transfer and reaction rates in packed bed chemical and nuclear reactors [1].

Flow through packed beds of spheres can be broadly classified into four regimes [21] based on the characteristics of the porescale flow: (i) Darcy flow, where viscous forces are dominant and fluid streamlines conform to the porespace boundaries, (ii) steady, inertial flow, where boundary layers and an "inertial core" develop in the porespace, and non-linear flow features begin to emerge, (iii) unsteady inertial flow where the flow becomes time dependent but remains laminar, and (iv) turbulent flow, with chaotic velocity fluctuations and a full energy cascade. The transition from one regime to another is Reynolds number dependent, but is also very sensitive to the packing configuration and boundary conditions. This makes developing a unified understanding of porescale flows challenging, and has motivated theoretical, experimental and numerical investigations. Modern experimental techniques such as magnetic resonance imaging (MRI) [108, 92], and particle image velocimetry (PIV) [81, 52] that allow for non-invasive measurement of the porescale velocity field have been applied to study these flows. At the same time, rapidly increasing processing power has enabled highly resolved simulation (DNS/LES) techniques as viable tools to investigate the porescale flow physics at play [48, 64]. In addition to providing physical insight into porescale flow behavior, DNS data sets are valuable in the process of developing lower order models for practical engineering calculations. Upscaling of the momentum equations for example, as done by [114] to predict macroscale dispersion properties, requires closures that can be obtained from three dimensional porescale velocity fields. Moving forward, it is important to continue developing efficient and accurate predictive simulation techniques capable of resolving porescale flow through packed beds.

Packed beds and porous media present unique challenges to resolved simulation techniques, mostly due to the complexity of the solid-fluid interface. Nonetheless,

successful simulations of these flows have been carried out for laboratory scale packings containing hundreds of spheres, most often using a body fitted grid combined with a finite volume approach [34, 4, 18]. While the body fitted approach has the ability to directly resolve the solid-fluid interface, generating and working with body conformal meshes for contacting sphere geometries can be extremely cumbersome. Most of the difficulty is due to the sphere-to-sphere contact points, which require careful treatment to avoid skewed, high aspect ratio control volumes. In our experience, the time required to generate a quality unstructured mesh for a large scale, complex geometry like a packed bed can often take longer than performing the simulation itself. Importantly, the significant meshing overhead and unstructured nature of the body fitted meshes for this problem makes systematic uncertainty quantification more difficult.

These factors leave the door open for alternative simulation approaches. The mesh related pitfalls encountered in body fitted simulations can be avoided with immersed boundary or fictitious domain approaches that use regular Cartesian grids over the entire fluid/solid domain, inclusive of the solid regions. In these types of approaches, a synthetic force is applied in or around the solid regions to satisfy the desired boundary condition (typically no-slip). The way in which this force is computed and applied is a field of active research, and a number of formulations have been proposed. A recent review on general techniques for resolved simulation of particle laden flow covering body fitted, immersed boundary, and fictitious domain approaches is given by Haeri & Shrimpton [37], and the interested reader is referred here for a more detailed discussion. In general, immersed boundary techniques add a corrective term to the momentum equation corresponding to a solid-fluid interaction that satisfies the no-slip condition. In the fictitious domain (or distributed Lagrange multiplier) approach, the entire domain is treated as a single fluid (inclusive of solid regions), and the rigid motion of immersed solid objects is projected in one or more steps onto the flow field using Lagrangian force points located at the solid-fluid interface.

These non body-conformal methods have evolved over the last several decades

from the original works of Peskin [85, 86], and have been used extensively in a number of fields including fluidized beds & suspensions [32, 111, 101], fluid structure interaction [116], and swimming/flying [19, 112]. Not surprisingly, their broadest application has been in problems where the motion of immersed solid boundaries would require adaptive re-meshing during simulation using a body fitted approach, and application to fixed bed type problems has been more limited. Some notable exceptions are the recent study of transitional flow through arrays of fixed two dimensional square cylinders by Malico & Ferriera de Sousa [67], and the study by Smolarkiewiscz & Larabee Winter [102] of Darcy flow through reconstructed porous media. Other alternatives to the body fitted approach exist that have been successfully applied to simulation of flows through fixed beds and porous media, including the Lattice-Boltzmann method [48] and methods based on smoothed particle hydrodynamics [78].

In this paper, the performance of fictitious domain simulations of flow through packed bed geometries based on the formulation of Apte et al. [2] is examined, relative to simulations using a well established body fitted unstructured grid approach developed from the work of Mahesh et al. [65] with modifications to account for mesh skewness proposed by Ham & Iaccarino [43]. The main goal is to demonstrate the capability and advantages of using the former for porous media and packed bed applications, especially at moderate Reynolds numbers when it is necessary to capture complex steady and unsteady porescale flow features. In addition, it is important to establish grid density requirements when simulating large scale, complex flow problems, so that solutions can be reported with a high level of confidence (or at least an understanding of their limitations). For example, in direct or large eddy simulations of turbulent channel flows, placing the grid cells within $Y^+ <= 1$ (where + denotes wall units), of the wall should produce reasonable confidence that boundary layers and small scale structures are resolved properly. This estimate takes advantage of the existence of the log-law in wall bounded flows. Likewise, for simulations of atmospheric turbulence or mixing layers, the grid spacing may be chosen to resolve the Kolmogorov scale. For inertial, transitional, and fully

turbulent flows through porous media, it is not completely clear on how to estimate grid resolution requirements apriori so that all important porescale features may be captured. It is possible to estimate the net average shear stress on the porous bed by relating the net pressure drop across the bed (through for example the Ergun [24] correlation) to the Reynolds number. Knowing the porosity, one could then compute the average shear stress and in turn estimate a boundary layer thickness. However, owing to complex bed geometries, these properties will vary substantially across the bed, and it is not clear what the grid resolution requirements are for accurate predictive simulations. Conducting grid refinement studies on very large scale simulations is also difficult due to the extensive computational cost. For this reason, our goal is to understand and establish grid resolution requirements DNS in the steady and unsteady inertial Reynolds number regimes by performing systematic grid refinement studies of flow through porous media on smaller domains with representative packings. While random packings require careful treatment, such estimates could be a valuable starting point for estimating the computational overhead of larger scale simulations.

The remainder of the paper is structured as follows. First the surface representation, and numerical solution procedure used by both methods is discussed in Section 2.2. Next we examine several test cases, which are designed to test the ability of each approach for complex flows in packed beds in Section 2.3. Emphasis is placed on estimating the uncertainty associated with the porescale flow fields produced by each method, and determining required grid densities for acceptable confidence levels. To this end, we employ the *Grid Convergence Index* (GCI), originally proposed by Roache [89]. In Section 2.4, expressions for the scalability of both methods for more general, randomly packed bed simulations are developed. Finally, in Section 2.5, the fictitious domain approach is used to simulate flow through a random arrangement of 51 spheres at a pore Reynolds number of 600. This case demonstrates the capability of the approach for more realistic packed bed problems of general interest and allows us to explore the porescale structure of unsteady inertial flow in a random arrangement of spheres.

2.2 Computational methods

Consider the Navier-Stokes equations for constant density and viscosity, incompressible fluid motion:

$$\nabla \cdot \mathbf{u} \;\; = \;\; 0 \tag{2.1}$$

$$\rho \left(\frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla) \, \mathbf{u} \right) \;\; = \;\; -\nabla p + \mu \nabla^2 \mathbf{u} + \rho \mathbf{g} + \mathbf{f} \tag{2.2}$$

where ρ is the density field, $\mathbf{u}$ the velocity vector, p the pressure, μ the fluid viscosity, $\mathbf{g}$ the gravitational acceleration, and $\mathbf{f}$ is an additional body force that is zero in the body fitted approach, and will be used to enforce rigidity within the solid phase in the fictitious domain approach. The two approaches to the solution of these equations are both implemented in a similar finite volume framework, and share several of the same basic techniques including a fractional step method [71], and an algebraic multigrid (AMG) solver for the pressure Poisson equation [25]. The codes are parallelized using Message Passing Interface (MPI), allowing for larger scale simulations by distributing the required memory over many processors.

Regardless of the approach used, precise representation of the solid-fluid interface is critical to obtain an accurate solution. Below, details are provided concerning the numerical representation of the porespace boundaries during simulation. For additional details regarding the numerical implementation, verification and validation of each method, the reader is referred to Moin & Apte [71] and, Ham & Iaccarino [43] for the body fitted approach, and to Apte et al. [2] for the fictitious domain approach.

2.2.1 *Body fitted approach*

In the body fitted approach, the pore space is first discretized into unstructured tetrahedral control volumes with a triangular surface mesh on the solid boundaries. Although arbitrarily shaped control volumes can be utilized by the solver, we have

found that it is easiest to generate quality meshes in packed beds with tetrahedral cells, as opposed to other shapes (hexahedral for instance). Unstructured mesh generation for complex geometries is a non-trivial procedure in general, and in packed beds the process is complicated by sphere-to-sphere contact points, near which elements can become unmanageably small, have high aspect ratio, and be skewed. Several methods have been proposed to mitigate this problem. Most commonly, the spheres are created at reduced diameter, typically 98 or 99 percent, eliminating all contact points [4, 11, 75], and creating a small gap between spheres. A related approach first pursued by Guardo et al. [34] is to create the spheres slightly larger than they actually are, so that they overlap, creating a continuous edge in the plane perpendicular to the contact line. Magnico [63, 64] has used a structured grid approach where the surface representation is stair stepped due to a voxelized treatment of the solid boundaries. While the meshing overhead is low with this approach, an artificial surface roughness is imposed, even with significantly refined grids. The approach that is pursued here takes advantage of the fact that the fluid very close to the solid contact points tends to be more or less stagnant even at moderate Reynolds numbers. In light of this, Kuroki et al. [56] have proposed a *bridge method*, wherein they do not change the diameter of the sphere, but rather unite two contacting spheres with a cylinder placed on the contact line as shown in Figure 2.1. This technique has the potential to significantly reduce overall mesh size because the regions where small element sizes are required have been eliminated. Nelson [74] extended this technique, and created a smooth fillet between the two contacting surfaces allowing for more continuity in the surface mesh for the application of prismatic surface layers.

The bulk behavior of flows in porous media and packed beds is strongly affected by porosity, so it is important that the bridges do not add significant solid volume to the porespace. The overall *increase in solid volume* due to a single bridge can be shown to be

$$\mathcal{V}_b = \frac{\pi}{8} D_b^2 \left(D - \sqrt{D^2 - D_b^2} \right) - \frac{\pi}{12} \left(D - \sqrt{D^2 - D_b^2} \right)^3 \qquad (2.3)$$

In this work, the bridge diameter is chosen to be $D_b = 0.25D$, so the volume of a single cylinder bridge is $\mathcal{V}_b = 0.0015\mathcal{V}_{sp}$, where $\mathcal{V}_{sp}$ is the volume of a single sphere. Thus, even for packings with moderate coordination numbers, we can expect the total solid volume to increase by less than 1% (and only in the mostly stagnant contact regions). By comparison, the more common technique of shrinking the spheres to 99% of their original size will decrease the solid volume (everywhere) by roughly 3%.
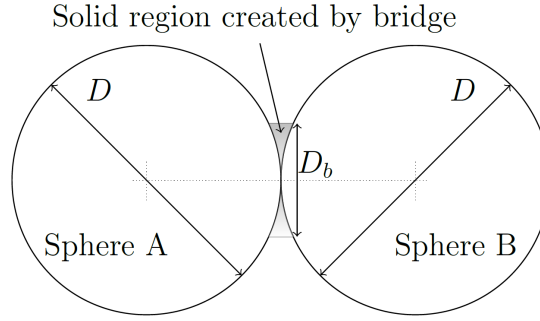
Figure 2.1: Schematic of the cylinder bridge created between two contacting spheres.

It quickly becomes challenging and time consuming to generate this type of geometry and mesh for more than a few spheres using the GUI of a typical mesh generation package because of the large number of geometric entities and high surface area to volume ratio. In light of this, a parameterized and automated approach has been developed that takes advantage of the commercial meshing package Pointwise's® full TCL-TK programmability. All sphere-sphere or sphere-boundary contact points are bridged, then trimmed and joined into a single watertight model. Once the solid geometry is assembled, a triangular surface mesh is generated on all solid surfaces. Because of the complex arrangements encountered in porous media, uniform mesh spacing, $\Delta = \mathcal{V}_{cv}^{1/3}$, where $\mathcal{V}_{cv}$ is the cell volume is used everywhere in the porespace. In geometries with entry/exit regions, mesh coarsening is used to reduce the total mesh size. An example of the cylinder bridge geometry and surface mesh is shown in the lower half of Figure 2.2 for a simple

two sphere system. This meshing tool is robust and can handle arbitrary random or arranged packings of spheres contained in box or tube geometries. It has been used to generate meshes with over 80 million control volumes and for geometries containing over 500 spheres. The interested reader is referred to [29] for examples.
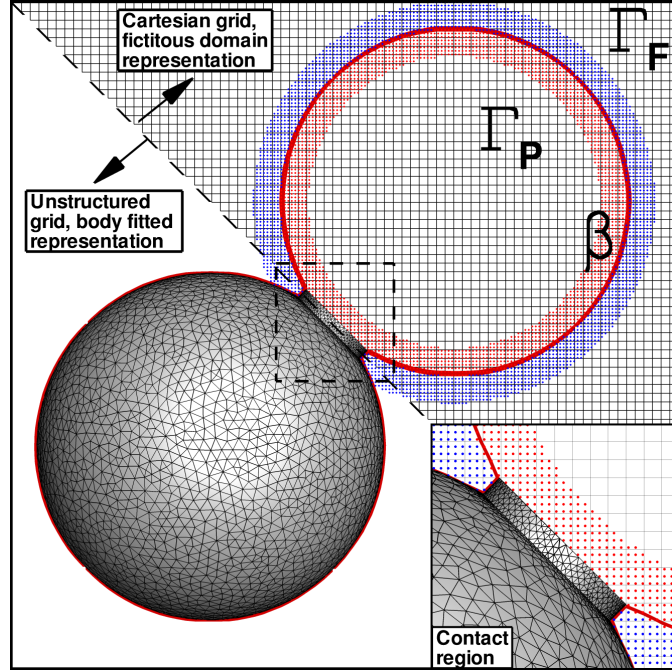
Figure 2.2: Comparison of discrete surface representations used by the body fitted (bottom) and fictitious domain (top) approaches.

2.2.2 *Fictitious domain approach*

The computations carried out with the fictitious domain approach utilize a hybrid Lagrangian-Eulerian (HLE) formulation for representation of arbitrarily shaped immersed solid objects and is not limited to the spherical objects used here. Although we consider only fixed beds, this fictitious domain approach also allows accurate representation of either forced or freely moving boundaries embedded in

the flow and is easily extensible to fluidized bed simulations. Let Γ be the the entire computational domain, which includes both the fluid (Γ_F) and the solid particle (Γ_P) domains shown in upper half of Figure 2.2. Let the fluid boundary not shared with the solid be denoted by β and have a Dirichlet condition (generalization of boundary conditions is possible). The basis of a fictitious domain approach is to extend the Navier-Stokes equations for fluid motion over the entire domain Γ inclusive of immersed solids [32]. The natural choice for these fixed bed problems is to assume that the immersed solid region, Γ_P, is filled with the same Newtonian *fluid* with density (ρ) and viscosity (μ) as Γ_F. Both the real and fictitious fluid regions will be assumed as incompressible and thus equations 2.1 and 2.2 apply everywhere in the domain. In addition, as the immersed solids are assumed rigid, the motion of the material inside Γ_P is constrained to rigid body motion. Several ways of obtaining the rigidity constraint have been proposed [32, 80, 100]. We follow the formulation developed by Sharma & Patankar [100] and described in detail by Apte et al [2]. A brief description is given here for completeness. The solid fluid interface, β is located on the Cartesian grid with subgrid scale marker points as shown in the top half of Figure. 2.2. These marker points carry a color function indicating the relative location of the interface, and remain fixed during the simulations.

In order to enforce that the material inside the immersed solid remains rigid, a constraint is required that leads to a non-zero forcing function $\mathbf{f}$. Inside the solid region, the rigid body motion, $\mathbf{u}^{RBM}$ implies vanishing deformation rate tensor:

$$\left. \begin{array}{l} \frac{1}{2}\left(\nabla\mathbf{u} + (\nabla\mathbf{u})^T\right) = \mathbf{D}[\mathbf{u}] = 0, \\ \Rightarrow \mathbf{u} = \mathbf{u}^{RBM} = \mathbf{U} + \Omega \times \mathbf{r} \end{array} \right\} \text{ in } \Gamma_P, \tag{2.4}$$

where $\mathbf{U}$ and Ω are the translation and angular velocities of the object and $\mathbf{r}$ is the position vector of a point inside the object from its centroid. For the fixed bed problems considered here, $\mathbf{u}^{RBM}$ is always zero. The vanishing deformation rate tensor for rigidity constraint automatically ensures the divergence free, incompressibility constraint inside the solid region. The incompressibility constraint

gives rise to the scalar field (the pressure, p) in a fluid. Similarly, the tensor constraint $\mathbf{D}[\mathbf{u}] = 0$ for rigid motion gives rise to a tensor field inside the solid region. Distributed Lagrange multipliers based approaches have been proposed to solve for the rigid body motion and impose the rigidity constraint, which requires an iterative solution strategy. Sharma & Patankar [100] proposed an approach that provides the rigidity constraint explicitly, thus reducing the computational cost significantly. Noting that the tensorial rigidity constraint can be reformulated to give:

$$\nabla \cdot (\mathbf{D}[\mathbf{u}]) = 0 \text{ in } \Gamma_P \tag{2.5}$$

$$\mathbf{D}[\mathbf{u}] = 0 \text{ on solid/fluid interface, } \beta \tag{2.6}$$

A fractional-step algorithm can be devised to solve the fictitious domain problem [80, 100, 2]. Knowing the solution at time level t^n the goal is to find $\mathbf{u}$ at time t^{n+1}.

1. In the first step, the rigidity constraint force $\mathbf{f}$ in equation 2.1 is set to zero and the equation together with the incompressibility constraint (equation 2.4) is solved by standard fractional-step schemes over the entire domain. Accordingly, a pressure Poisson equation is derived and used to project the velocity field onto an incompressible solution. The obtained velocity field is denoted as $\mathbf{u}^{n+1}$ inside the fluid domain and $\hat{\mathbf{u}}$ inside the solid object.

2. To solve for $\mathbf{u}^{n+1}$ inside the solid region we require $\mathbf{f}$. The constraint on the deformation rate tensor given by equation 2.4, along with the no-slip specification at the solid-fluid interface can be reformulated to obtain:

$$\nabla \cdot \left(\mathbf{D}[\mathbf{u}^{n+1}]\right) = \nabla \cdot \left(\mathbf{D}\left[\hat{\mathbf{u}} + \frac{\mathbf{f}\Delta t}{\rho}\right]\right) = 0; \tag{2.7}$$

$$\mathbf{D}[\mathbf{u}^{n+1}] \cdot \mathbf{n} = \mathbf{D}\left[\hat{\mathbf{u}} + \frac{\mathbf{f}\Delta t}{\rho}\right] \cdot \mathbf{n} = 0. \tag{2.8}$$

The velocity field in the solid is zero for fixed beds. Thus $\hat{\mathbf{u}}$ is split into a

rigid body motion ($\mathbf{u}^{RBM} = \mathbf{U} + \Omega \times \mathbf{r} = 0$) and residual non-rigid motion ($\mathbf{u}'$). The above formulation can be easily generalized to solid bodies with *specified motion* by directly setting $\mathbf{u}^{RBM}$ to the specified velocity.

3. The rigidity constraint force is then simply obtained as $\mathbf{f} = \rho(\mathbf{u}^{RBM} - \hat{\mathbf{u}})/\Delta t$. This sets $\mathbf{u}^{n+1} = \mathbf{u}^{RBM}$ in the solid domain. Note that the rigidity constraint is non-zero only inside the solid domain and zero everywhere else. This constraint is then imposed in a third fractional step.

The utility of the Lagrangian marker points (see Figure 2.2) is in locating the solid/fluid interface and enforcing the boundary condition on β with subgrid scale resolution in the above steps. A quantity can be defined and calculated at the marker points (for example $\mathbf{f}$) and then be projected onto the Eulerian grid, or vice-versa, using accurate interpolation kernels [93]. In practice, because the no-slip condition at the boundary of the porespace is enforced indirectly, the precision of the boundary location is directly linked to the grid resolution. This is especially true near regions of sharp boundary curvature such as sphere to sphere contact points. This is illustrated in the top half of Figure 2.2 where the solid line denotes the fictitious domain solid-fluid interface. Even with subgrid marker points, the high curvature contact region appears as a *bridge* similar to the one obtained (intentionally) in the body fitted meshing procedure.

2.2.3   Uncertainty Estimation

In the absence of an analytic solution, it is important to be able to estimate and report the uncertainty associated with a CFD calculation. Furthermore, uncertainty should be computed in a consistent manner, which can be compared by future workers using new or different simulation approaches. For this reason, we choose the *grid convergence index* (GCI), originally proposed by Roache [89], to quantify the uncertainty associated with some of our simulation results. The GCI is based on the ideas of Richardson extrapolation and does not rely on the existence of an exact solution or the assumption that a very fine grid solution may be taken as

such. It is robust as a general post-processing tool for error estimation, and has proven its utility for numerical solutions of a variety of different flows [10].

To compute an uncertainty band for the flow variable, ϕ, the solution is first obtained on at least three grids with (not necessarily equal) refinement ratios $r_{21} = \Delta_2/\Delta_1$ and $r_{32} = \Delta_3/\Delta_2$. For the purposes of comparing results that utilize control volumes with various shapes, we define the grid spacing in this paper as $\Delta = \mathcal{V}_{cv}^{1/3}$, where $\mathcal{V}_{cv}$ is the average cell volume. For nominally two dimensional cases (such as Section 2.3.1), the analog to this definition is $\Delta = A_{cv}^{1/2}$, where A_{cv} is the cell area. The fine and medium grid solutions, ϕ_1 and ϕ_2, are then interpolated to the coarse grid, where the variations, $\epsilon_{32}(\mathbf{x}) = \phi_3(\mathbf{x}) - \phi_2(\mathbf{x})$, and $\epsilon_{21}(\mathbf{x}) = \phi_2(\mathbf{x}) - \phi_1(\mathbf{x})$, are computed. From $\epsilon_{32}(\mathbf{x})$ and $\epsilon_{21}(\mathbf{x})$, the *local* apparent order of accuracy, $\mathcal{P}(\mathbf{x})$, may be calculated, using the following equation [13].

$$\mathcal{P}(\mathbf{x}) = \frac{1}{\ln(r_{21})} \left| \ln |\epsilon_{32}(\mathbf{x})/\epsilon_{21}(\mathbf{x})| + \ln \left(\frac{r_{21}^{\mathcal{P}(\mathbf{x})} - sign(\epsilon_{32}/\epsilon_{21})}{r_{32}^{\mathcal{P}(\mathbf{x})} - sign(\epsilon_{32}/\epsilon_{21})} \right) \right| \tag{2.9}$$

In the event that $r_{21} \neq r_{32}$, a straightforward iteration of equation 2.9 can be used to determine $\mathcal{P}(\mathbf{x})$. The global order of convergence, $\mathcal{P}_G$, is then computed by averaging $\mathcal{P}(\mathbf{x})$ at nodes where monotone convergence is observed, indicated by $sign(\epsilon_{32}/\epsilon_{21}) > 0$. The percentage of nodes exhibiting monotone convergence is denoted $\%M_n$ and indicates the degree to which the results are in the asymptotic regime. Non-monotone grid convergence of CFD solutions is an unfortunate reality, and is not necessarily cause for excessive concern [22]. Using the global order of convergence, the GCI of the fine grid solution is then computed as

$$GCI(\mathbf{x}) = F_s \left| \frac{\phi_1(\mathbf{x}) - \phi_2(\mathbf{x})}{1 - r_{21}^{\mathcal{P}_G}} \right| \tag{2.10}$$

where $F_s = 1.25$ is a reasonably conservative factor of safety for a three grid refinement study [91]. In this form, the GCI has the same units as ϕ, and may be taken as a local uncertainty band. Alternatively, it may be multiplied by $100/\phi_{ref}$, where ϕ_{ref} is some meaningful reference value, to obtain a dimensionless percent relative

uncertainty. In this work we will report the global uncertainty, GCI_G, as the node averaged percent relative uncertainty with ϕ_{ref} equal to the spatial average of $RMS(\phi(\mathbf{x}))$. Boundary nodes, or nodes where ϕ is fixed (such as within rigid bodies in fictitious domain computations), are excluded from the global averages. Finally, the grid resolution required to obtain some target level of uncertainty, denoted GCI^*, may be estimated as [90]

$$\Delta^* = \Delta_1 \cdot \left(\frac{GCI^*}{GCI_G} \right)^{1/\mathcal{P}_G} \tag{2.11}$$

This is useful for determining simulation size requirements as is done in Section 2.4.

2.3 Performance assessment

The test cases examined here are designed to provide a basis for quantifying the performance of our two methods for fixed bed flow simulations. First, the influence of grid type and quality on solution error in the absence of solid boundaries is demonstrated using the case of decaying Taylor vortices. Second, as a basic hydrodynamic validation for packed bed and porous media flows, we examine Stokes flow through a dilute periodic array. Third, to test the effect of sphere to sphere contact points on both the body fitted and fictitious domain approaches, flow past a contacting pair of spheres is simulated. Finally flow through a simple cubic lattice is considered at both steady and unsteady Reynolds numbers. This serves two purposes. First, the predicted porescale velocity profiles can be compared to those measured experimentally by [108]. Second, the geometry is simple enough to allow a detailed grid refinement study from which we can generate error estimates, and determine the scalability of each approach for larger packed bed simulations over a broad range of pore Reynolds numbers.

2.3.1   Baseline accuracy and mesh quality: Decaying Taylor vortices

Based on the numerical implementations, in the absence of immersed solids the fictitious domain and body fitted approaches are identical. An important first step in understanding the solution behavior in complex packed bed flows is to assess the influence of grid type and quality on solution accuracy in the absence of solid boundaries. To accomplish this, we examine the case of decaying Taylor vortices [109]. This time dependent solution to the Navier-Stokes equations can be written as:

$$u_x = -\cos(\pi x)\sin(\pi y)\exp^{-\frac{2\pi^2 t}{Re}} \tag{2.12}$$

$$u_y = \sin(\pi x)\cos(\pi y)\exp^{-\frac{2\pi^2 t}{Re}} \tag{2.13}$$

$$p = -\frac{1}{4}\left(\cos(2\pi x)\cos(2\pi y)\right)\exp^{-\frac{4\pi^2 t}{Re}} \tag{2.14}$$

(a) Uniform Cartesian   (b) Uniform prism   (c) Non-uniform prism
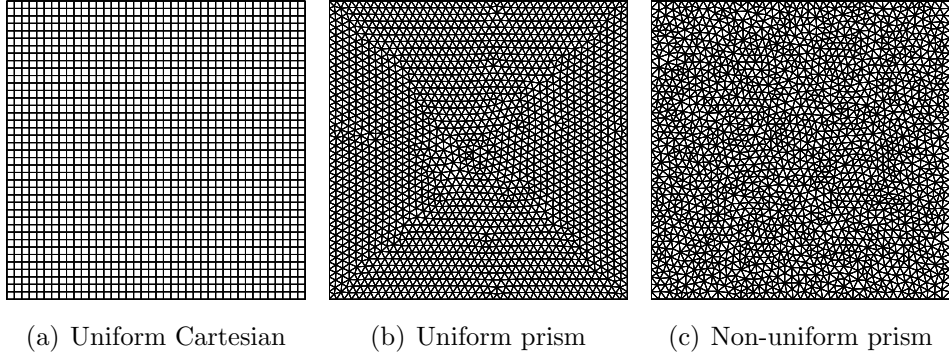
Figure 2.3: Three families of grids used in the Taylor vortex case.

We assign the initial condition at time $t = 0$ for $Re = 10$ in a periodic domain with $-1 < x < 1$, $-1 < y < 1$ on the three families of grids shown in Figure 2.3: (a) uniform, Cartesian grids with cubic control volumes, (b) prism grids with nearly uniform element size and aspect ratio near 1, (c) prism grids with non-uniform element size and aspect ratio. The vortex decay is simulated up to a time $t = 0.2$ on grids with average cell edge lengths of $l_e = 0.1$, 0.05, 0.025. This results in grid

spacing, $\Delta = A_{cv}^{1/2} = l_e$ for the Cartesian grids, and $\Delta \approx 0.66l_e$ for the prism grids. At the end of the simulation, the error in the u_x velocity component is computed at each cell center using equation 2.12. A constant time step, $\Delta_t = 2.5 \times 10^{-3}$, is used to isolate the spatial discretization errors, resulting in a maximum CFL number of $CFL = ||\mathbf{u}||\Delta_t/\Delta = 0.15$ for the initial condition on the finest prism grid. Figure 2.4 shows the L_1 and L_{max} error norms as a function of grid spacing for the three families of grids. As expected, the Cartesian grids show perfect second order spatial accuracy. Second order accuracy in the L_1 norm is retained by the uniform prism grids, but the L_{max} norm shows a significant increase in *local* error, which underscores the sensitivity of the methods to local grid quality. The convergence of error on the non-uniform prism grids is slightly better than first order in both the L_1 and L_{max} norms, and the magnitude of solution error is nearly 1 order of magnitude greater than the Cartesian grids. This demonstrates that unstructured meshes with lower quality cells compared to Cartesian grids can introduce significant local solution errors. This is important to consider for packed beds, where irregular geometries make generating uniform triangular/tetrahedral cells very difficult

(a) L_1 norm   (b) L_{max} norm
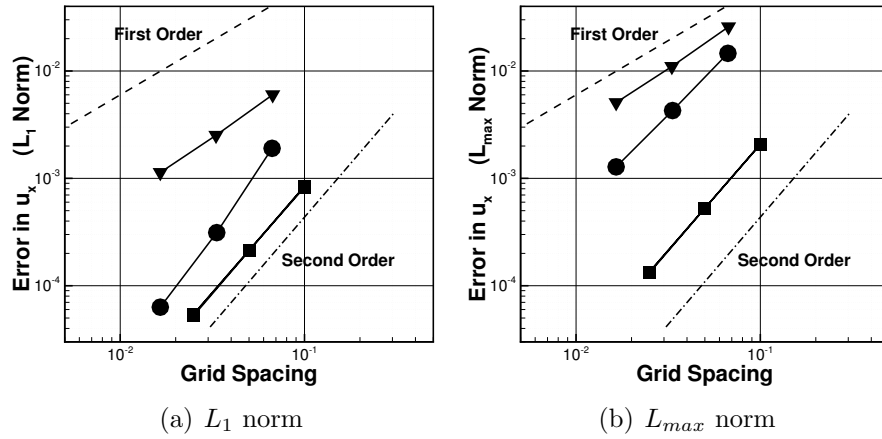
Figure 2.4: L_1 and L_{max} error norms for the u_x velocity component. Cartesian ■, uniform prism ●, non-uniform prism ▼.

2.3.2 Basic hydrodynamic verification: Stokes flow drag on a single sphere in a dilute array

To begin validating these approaches for packed bed type applications, we examine the case of Stokes flow through a dilute periodic array of spheres. The geometry we consider is a periodic box, with all sides of length L that contains a single sphere with diameter, D, centered at $\mathbf{x} = L/2$. The periodicity in all directions allows us to represent an infinite simple cubic arrangement with a single sphere. The ratio L/D is chosen to be 2.015 to allow for comparison with the integral equation solutions of [117] for the relatively dilute solid volume fraction of $(1 - \epsilon) = 0.064$, where ϵ is the fluid void fraction . Three Cartesian and tetrahedral grids are created with $D/\Delta = 12,\ 24,\ 48$, where $\Delta = \mathcal{V}_{cv}^{1/3}$.

The flow is started from rest, and accelerated by applying a small, uniform pressure gradient everywhere. When the flow reaches equilibrium, we confirm that the pore Reynolds number is small enough that inertial effects are negligible, and the total force exerted on the sphere is then computed. For this case, $Re_p = \frac{U_p D}{\nu} \approx 1 \times 10^{-3}$ where U_p is the average velocity in the fluid region of the porespace, or *pore velocity*. In the fictitious domain approach, the drag force, f_d, is taken to be the component of the rigid body force aligned with the applied pressure gradient. In the body fitted approach, the drag force is obtained by integrating the surface stress over the sphere boundary faces. The drag force is normalized by the Stokes drag on a single sphere in unbounded flow to obtain the drag coefficient, C_d.

$$C_d = \frac{f_d}{N_{sp} 3\pi \mu D U_p \epsilon} \tag{2.15}$$

Here, N_{sp} is the number of spheres and is equal to 1 for this case. The computed values are summarized in Table 2.1 for both approaches. The solutions converge to within about 2% of the analytic value, $C_d = 2.810$ obtained by Zick & Homsy. We also have computed the apparent convergence rate, $\mathcal{P}$ and the relative percent uncertainty using the GCI approach. The body fitted approach, wherein the no-slip condition is directly enforced at the sphere surface, has a very rapid apparent rate

of convergence, $\mathcal{P} = 3.19$, leading to very low uncertainty in C_d for the fine grid solution. The fictitious domain approach, which less precisely defines the surface via the sub-grid marker points, converges at a more modest rate of $\mathcal{P} = 1.65$, and achieves a relative uncertainty of 0.59% for $D/\Delta = 48$. The roughly 4% discrepancy between the converged body fitted and fictitious domain solutions may be due to differences in grid quality (body fitted) and interface representation (fictitious domain)

Table 2.1: Comparison of Stokes flow drag coefficient to analytic solution of [117] for a simple cubic array with solid concentration, $(1 - \epsilon) = 0.064$

D/Δ	Analytic	Body Fitted			Fictitious Domain		
	C_D	C_D	$\mathcal{P}$	GCI (%)	C_D	$\mathcal{P}$	GCI (%)
12	2.810	2.784	-	-	3.086	-	-
24	2.810	2.857	-	-	2.831	-	-
48	2.810	2.865	3.19	0.016%	2.750	1.65	0.59%

2.3.3   Effect of the contact point: *The contacting pair*

Flow past a contacting pair of spheres in a confined channel is chosen as a way to isolate the effect of sphere to sphere contact regions. The configuration of the connected body and flow domain is shown in Figure 2.5a. The spheres are located such that their contact point always lies at the origin, $(X, Y, Z) = 0$. Uniform flow enters from a square cross section, $3.5D \times 3.5D$ at $Z = -5D$. A convective outflow boundary is located downstream of the contact at $Z = +5D$. A no-slip condition is enforced on the channel walls. The spheres are rotated around the Y axis by the angle of incidence, θ, so that the leading and trailing sphere centers have the

coordinates

$$(X_l,\ Y_l,\ Z_l)\ =\ \left(0,\ -\frac{D}{2}\sin(\theta),\ -\frac{D}{2}\cos(\theta)\right)$$

$$(X_t,\ Y_t,\ Z_t)\ =\ \left(0,\ \frac{D}{2}\sin(\theta),\ \frac{D}{2}\cos(\theta)\right)$$

The value of θ is varied from 0^o to 90^o in increments of 15^o in order to fully sample the range of orientations found in realistic packed beds. The uniform inflow velocity is assigned so the Reynolds number, defined $Re = U_{in}D/\nu$ is between 50 and 175, in intervals of 25. The flow is allowed to develop for $tU_{in}/D = 30$ non-dimensional time units. To accelerate the development of the flow, solutions from consecutive Reynolds numbers are used as initial conditions (ie. the $Re = 100$ solution is used as an I.C. for the $Re = 125$ simulation). The total hydrodynamic force on the connected body is monitored and plotted in Figure 2.6 for $\theta = 0^0$ and $Re = 50,\ 175$, along with similar results from Section 2.3.4, to confirm that 30 non dimensional time units is sufficient to achieve a steady, developed flow condition. Mean grid spacing on the surface and in the vicinity of the pair is $D/\Delta = 40$ for both approaches. Since the goal of this study is to determine the effect of near sphere grid resolution, the grids are allowed to coarsen far away from the pair near the inflow and outflow regions. At a distance of $1.5D$ away from the contact line, the grid spacing is never greater than twice the surface spacing.

We first confirm that our flow is steady by monitoring the drag force signal on the surface of the spheres and verifying that no oscillations are present for the entire range of present conditions. For the confined cases studied here, we see the development of complex, yet steady flow features as shown in Figure 2.5b-e. Here, we show select stream-ribbons for $\theta = 0^o,\ 30^o,\ 60^o$,and $\theta = 90^o$ along with pressure contours in the $X = 0$ plane. At $\theta = 0^o$, the flow stagnates evenly on both spheres and the contact point, generating a large symmetric recirculation bubble. As the angle of incidence is increased to 30^o and 60^o, the recirculation bubble bends away from the leading sphere and towards the trailing sphere and the flow is no longer symmetric about the contact line. At $\theta = 60^o$ there is a noticeable

(a) Connected solid   (b) $\theta = 0$ (c) $\theta = 30$ (d) $\theta = 60$ (e) $\theta = 90$
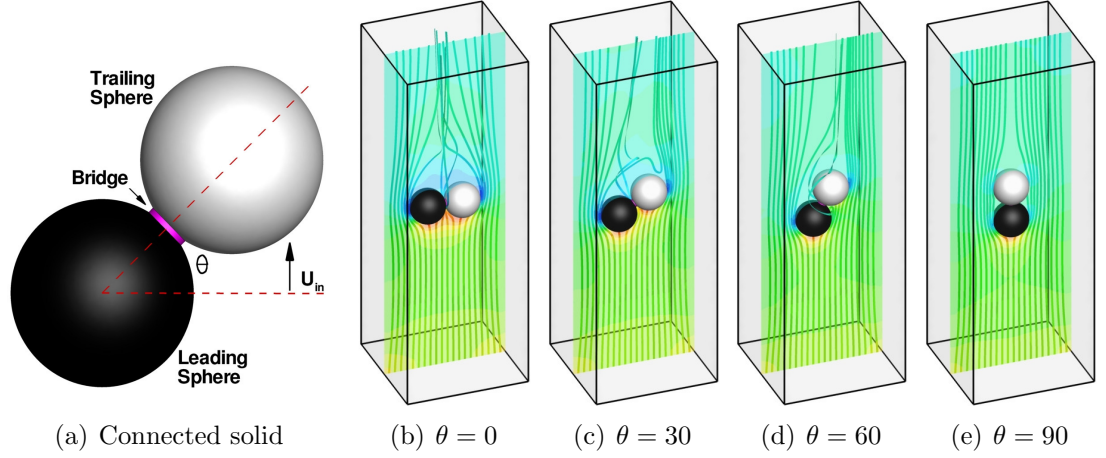
Figure 2.5: Configuration and flow visualization for the contacting pair case. (a) Shows the configuration of the connected body. (b-e) are stream ribbon visualization of flow around the pair of spheres at $Re = 175$ and select angles of incidence. Pressure contours are shown in the $X = 0$ plane; red indicates high pressure, blue indicates low pressure.

decrease in stagnation pressure on the trailing sphere, suggesting a drafting effect is present. At $\theta = 90^0$, symmetry is regained as the body becomes streamlined, and the contact point is completely hidden from the oncoming flow by the leading sphere.

Using the body fitted approach, we can isolate the force on the leading sphere, F_z^l, the trailing sphere, F_z^t, and the bridge F_z^b by computing directly the surface integral of viscous and pressure forces on each of the solid boundary zones separately. In Figure 2.7 the ratio of drag force on the trailing and leading sphere (F_z^t / F_z^l), as well as the ratio of drag force on the bridge to the leading sphere (F_z^b / F_z^l) is plotted as a function of Re. The first ratio demonstrates the increased drafting ability of the trailing sphere at increased θ or increased Re as would be expected. The second ratio demonstrates that for all combinations of Re and θ, the contribution of the bridge to the total drag of the pair is less than 1% of the leading sphere's contribution. This ratio is maximized at all Re for $\theta = 0^o$, and decreases as the bridge surface is rotated out of the stagnation region. Roughly
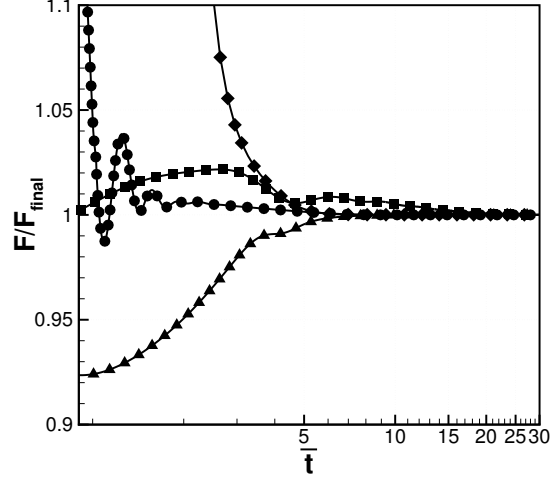
Figure 2.6: Semi-logarithmic plot of the time history of hydrodynamic force exerted on the spherical bodies for steady flow conditions in the contacting pair and simple cubic lattice test cases. The force in each case is normalized by the force at $tU/D = 30$. ▲ contacting pair at $\theta = 0^o$, $Re = 50$, ■ contacting pair at $\theta = 0^o$, $Re = 175$, ● simple cubic lattice at $Re = 12$, ♦ simple cubic lattice at $Re = 204$.

extrapolating, the cylinder bridge drag will be 1% of the leading sphere drag at $\theta = 0^o$ for $Re \approx 600$, the largest Reynolds number considered in this study.

To verify that the fictitious domain representation of the surface is consistent with the body fitted representation, especially near the contact points, the velocity profile in the near wake behind the cylinder bridge is plotted in Figure 2.8 for $Re = 175$, $\theta = 0$. Despite the close proximity of the probes to the solid boundaries, the wake profiles predicted by each approach are in very good agreement. These results demonstrate that the inability of both methods to completely resolve the contact point between two spheres should not significantly affect global or local flow properties in packed beds.
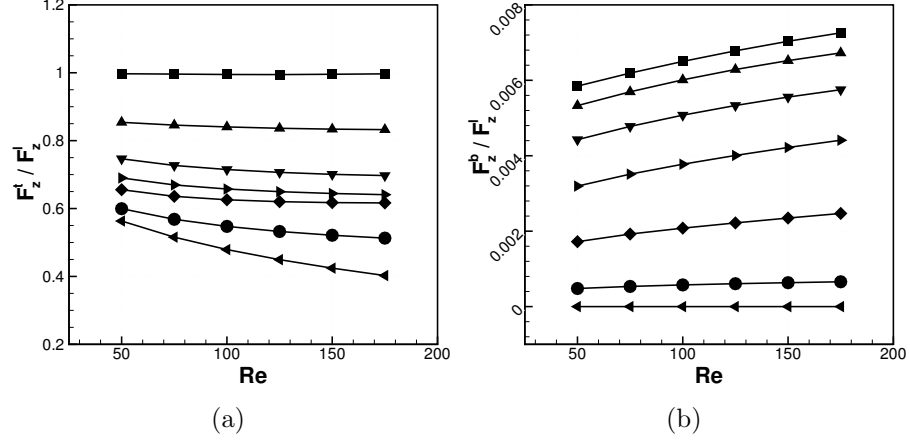
Figure 2.7: Drag force on portions of the connected body relative to the leading sphere drag as a function of Re for all values of θ tested. (a) The trailing sphere drag, F_z^t/F_z^l. (b) The bridge drag, F_z^b/F_z^l. ■ $\theta = 0^o$, ▲ $\theta = 15^o$, ▼ $\theta = 30^o$,▶ $\theta = 45^o$, ♦ $\theta = 60^o$,● $\theta = 75^o$,◀ $\theta = 90^o$.

(a) Probe location   (b) Streamwise velocity
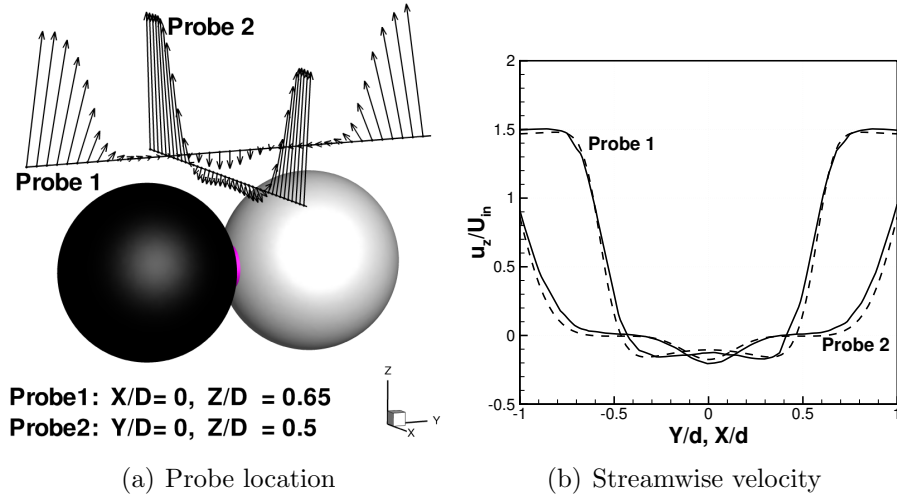
Figure 2.8: Velocity profiles in the near wake behind the contacting pair at $Re = 175$, $\theta = 0^o$. (a) Probe locations. (b) Streamwise velocity profile (—)body fitted, (- - -)fictitious domain.

2.3.4 Performance in porous geometries: Flow through a simple cubic lattice

In this test case, pore scale flow through square channels filled with a simple cubic lattice of spheres is examined. The configurations are chosen to be similar to the experiments of Suekane et al. [108] who used MRI techniques to make detailed three dimensional measurements of the porescale velocity field. To our knowledge, it is one of the only experimental measurements of its kind, and has served as a validation for several other numerical studies including [35] and [78]. We consider two slightly different computational domains shown in Figure 2.9a. The first configuration, intended to match the experiments of [108] as closely as possible, consists of six layers of diameter $D = 28mm$ quarter spheres located in the corners of the channel, with no-slip boundaries enforced on the exterior channel walls. This domain is used to simulate Reynolds numbers based on the pore velocity, U_p, of $Re = U_p D/\nu = 12.17$, 105.57, and 204.74, allowing for direct comparison with the experiments. Flow in this configuration is driven by a constant inflow boundary, located $5D$ upstream of the first sphere (not shown in the Figure), where $U_{in} = \epsilon U_p$, and the fluid fraction is $\epsilon = 0.476$ for this arrangement. A convective outlet is located $5D$ downstream of the last sphere. The flow is started from rest and simulated for $tU_p/D = 10$ non dimensional time units for each Reynolds number, which is sufficient to obtain a fully developed steady flow field as confirmed in Figure 2.6.

In order to test the performance of the solvers for unsteady flows with vortex shedding, a second configuration is considered to simulate $Re = 450$ and 600 as there is no experimental data in this flow regime. This configuration consists of six layers of complete spheres stacked in the center of the channel. It has periodic boundaries on all exterior channel walls that allows for faster symmetry breaking and transition to unsteady flow. The flow is initialized with large wavelength sinusoidal fluctuations and the flow is simulated for at least $tU_p/D = 60$ time units to allow the unsteady dynamics to reach a stationary state. During the simulation, a body force is continuously adjusted so that the the target Reynolds

numbers are achieved. At the end of all simulations, the flow is probed along the line $Y = 0$, $Z = 0$, as shown in Figure 2.9a, which traverses the center of the fifth pore in both configurations.

(a) Computational Domains
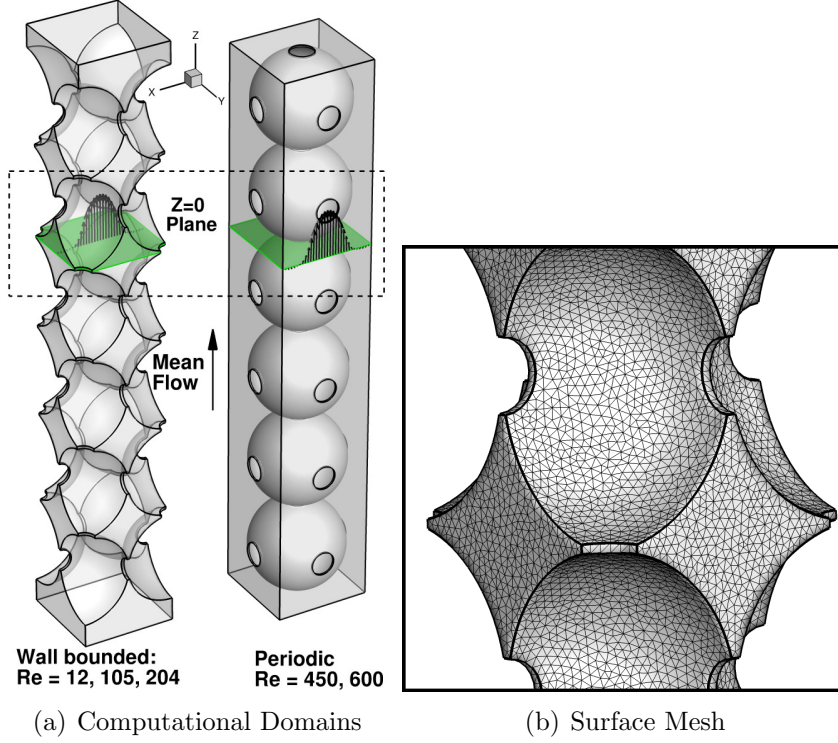
(b) Surface Mesh

Figure 2.9: Setup of flow through a simple cubic lattice. (a) shows the two computational domains used. The orientation of the line probe used is indicated by the vector velocity profiles. (b) shows the $D/\Delta = 64$ surface mesh used in the body fitted simulations.

For the body fitted approach, the domains are discretized using unstructured, tetrahedral meshes as described earlier. Each full or quarter sphere was created at its exact diameter inside the channel. Then a small cylindrical region of fluid, with diameter $D_b = 0.25D$, is removed from the near the contact points. The entire connected solid is subtracted from the channel resulting in a watertight solid boundary. Uniform size tetrahedral elements with $\Delta \approx \mathcal{V}_{cv}^{1/3}$ are created throughout the channel. A closeup view of the surface of the body fitted mesh

Table 2.2: Grids used to simulate flow through the simple cubic lattice.

	$D/\Delta = 128$		$D/\Delta = 64$		$D/\Delta = 32$		$D/\Delta = 16$	
	Tet.	Cart.	Tet.	Cart.	Tet.	Cart.	Tet.	Cart.
N_{cv} In/Out	7.62M	17.9M	1.13M	2.245M	190k	281k	54k	35k
N_{cv} Periodic	6.33M	12.6M	867k	1.57M	119k	196k	n/a	n/a

with $D/\Delta = 64$ is shown in Figure 2.9b. The tetrahedral grids are high quality; the maximum cell aspect ratio is 4 or less for all meshes. Regular Cartesian grids are used for the fictitious domain approach that include both the fluid and solid portions of the channel. Four meshes were generated for each method with mean spacing, $D/\Delta = 16, 32, 64, 128$. The total cell count for each mesh is summarized in table 2.2.

Fine, medium, and coarse grid results are used to estimate the uncertainty in the fine grid solution obtained with each approach by applying the GCI method outlined in Section 2.2.3. For $Re = 12$, grids with $D/\Delta = 64, 32$, and 16 are used, while the $Re = 105, 204, 450$, and 600 simulations use grids with $D/\Delta = 128, 64$, and 32. Since grid independence of one flow property does not necessarily imply grid independence of another, it is important to use multiple measures to fully understand the convergence behavior. For the steady flow rates, the grid convergence properties are computed for the streamwise velocity component, $\phi = u_z$, as well as the non-dimensional macroscale pressure gradient, $\phi = \Psi = \frac{\Delta P}{L} \frac{D}{\rho U_p^2} \frac{\epsilon}{1-\epsilon}$. For $\phi = u_z$ we are directly assessing the convergence of the porescale velocity field, while for $\phi = \Psi$, we are assessing the convergence of a macroscale integrated property. For the unsteady flow rates, the analysis is done for the time averaged streamwise velocity, $\phi = \overline{u_z}$, as well as the time averaged turbulent kinetic energy, $\overline{TKE} = \overline{\sum u_i'^2}$. These two measures capture the mean and fluctuating character of the porescale flow respectively. Key results of the analysis are reported in table 2.3. For all porescale properties the max and RMS value of ϕ predicted by the two approaches are in reasonably good agreement, with the exceptions discussed further below. The monotone convergence rate is significantly higher for the fictitious domain approach in all cases, indicating that the solutions on the three

Cartesian grids are more fully in the asymptotic regime relative to the body fitted tetrahedral grid results. The global apparent rate of convergence, $\mathcal{P}_G$, is bound by the formal second order accuracy of the spatial discretization schemes (with a few exceptions discussed below), and is observed to be between roughly 1.3 and 2. In general, it is higher for the body fitted results. Even though the convergence rate is higher, the difference between the fine and medium grid solutions ($\phi_1 - \phi_2$) is in general much more significant with the body fitted approach, leading to a much higher global uncertainty, GCI_G, of the porescale quantities in most cases. The convergence of the integral pressure drop at the steady flow rates behaves differently and is discussed in more detail below. From the uncertainty estimates associated with the fine grid solutions, the grid density required to achieve a 5% uncertainty level, $D/\Delta_{5\%}$ is estimated from equation 2.11. In many cases, the grid requirements based on the porescale quantities are significantly stricter for the the body fitted approach.

It is helpful to take stock of the results in the table in more detail while at the same time examining the grid convergence behavior graphically. The fine, medium, and coarse grid predictions of streamwise velocity along the $Y = 0$, $Z = 0$ line are plotted in Figure 2.10a-c for the steady flow rates. In this Figure as well as Figures 2.12 and 2.14, the body fitted solutions are shown in the left hand column and the fictitious domain solutions are shown in the right hand column. In the main portion of the sub-figures, only the fine grid result with error bars corresponding to $GCI(\mathbf{x})$ is plotted for clarity alongside the experimental data of [108]. In the insets, the medium and coarse grid solutions are shown as well, for the region near $|X/D| \lesssim 0.15$, to illustrate the convergence with grid refinement. In both methods, refinement results in convergence of the solution towards the experimental data. The maximum and RMS values of u_z are similar for all cases on the fine grids. In all cases there is good agreement between the experimental data points, and the medium to fine grid solutions. It is evident that as Reynolds number is increased, the significant flow inertia gives rise to a strong jet through the center of the pore ($|X/D| \lesssim 0.15$), and symmetric backflow regions close to the

Table 2.3: Grid convergence results for flow through the simple cubic lattice. In all cases D/Δ_1, D/Δ_2, $D/\Delta_3 = 128$, 64, 32 except $Re = 12.12$ where D/Δ_1, D/Δ_2, $D/\Delta_3 = 64$, 32, 16. Values are reported in mm/s for velocity, and mm^2/s^2 for TKE.

Re	ϕ	$MAX(\phi)$ BF	FD	$RMS(\phi)$ BF	FD	$M_n[\%]$ BF	FD	$\mathcal{P}_G$ BF	FD	$GCI_G[\%]$ BF	FD	$D/\Delta_{5\%}$ BF	FD
12	u_z	1.98	1.71	0.53	0.50	62	95	1.73	1.35	3.61	1.60	53	28
12	Ψ	n/a	n/a	13.9	13.4	n/a	n/a	1.43	1.67	2.94	3.69	44	53
105	u_z	16.0	15.9	4.61	4.53	65	82	1.76	1.65	1.89	0.54	74	33
105	Ψ	n/a	n/a	2.17	2.18	n/a	n/a	2.72	1.30	0.10	1.98	30	62
204	u_z	30.7	30.9	8.99	8.85	66	88	1.82	1.62	1.99	0.90	77	44
204	Ψ	n/a	n/a	1.29	1.33	n/a	n/a	3.34	1.41	0.03	2.91	28	87
450	$\overline{u_z}$	61.4	56.2	24.2	24.2	41	84	1.84	1.56	4.77	2.20	125	76
450	$\overline{TKE}$	230	284	64.8	82.8	52	66	1.76	1.71	16.1	3.93	249	111
600	$\overline{u_z}$	75.5	72.3	32.1	32.1	47	80	2.08	1.29	2.75	3.79	96	103
600	$\overline{TKE}$	436	458	143	153	47	67	1.84	1.45	10.4	5.42	191	135

walls ($|X/D| \approx 0.35$) for $Re = 105$ and 204. For these steady results, the global apparent rate of convergence, $\mathcal{P}_G$, is roughly 1.75 for the body fitted results and varies from 1.35 to 1.65 for the fictitious domain results. This overall lower rate of convergence may be due to the higher percentage of monotone nodes, $\%M_n$, obtained with the fictitious domain approach, 82% or more. Despite enjoying a slightly better convergence rate, the global relative uncertainty of the fine grid body fitted solution, GCI_G, is roughly two to three times more than the fictitious domain solution in each case. This indicates that local errors are introduced by the unstructured meshes, similar to the observations made in the Taylor vortex case. Using equation 2.11, the grid density required to obtain a (relatively arbitrary) target global uncertainty level of 5% was computed and found to be roughly twice as much for the body fitted approach. For this fixed level of uncertainty, these steady Reynolds numbers will require between 53 and 77 cv/D with the body fitted tetrahedral grids, but only 28 to 44 cv/D with a Cartesian grid and the fictitious domain approach.

The non dimensional pressure drop for the steady flow rates predicted by both methods using the fine grid is plotted in a Figure 2.11 alongside the correlation of Ergun [24] for random packed beds. To produce the data points, the pressure is spatially averaged in two planes located $L = 2D$ apart in the streamwise direction once the flow is fully developed. The figure and the results in table 2.3 show that the results from both types of simulations are in agreement in the fine grid solutions. However, the GCI analysis indicates that there is significantly less uncertainty in the body fitted result (especially for $Re = 105$ and 204), which is surprising given the observations of the interstitial velocity profiles. Inspection of the convergence trends shows that the body fitted pressure drop converges at a rate higher than the formal second order of accuracy of the code. This raises questions about the reliability of using such a result, and shows the danger of using integral measurements alone to assess grid convergence of the pore scale flow field. Deviation of the results from the Ergun correlation at increased Reynolds numbers is most likely a result of the simple cubic arrangement, and has also been observed

(a) Re = 12.17

(b) Re = 105.57

(c) Re = 204.74
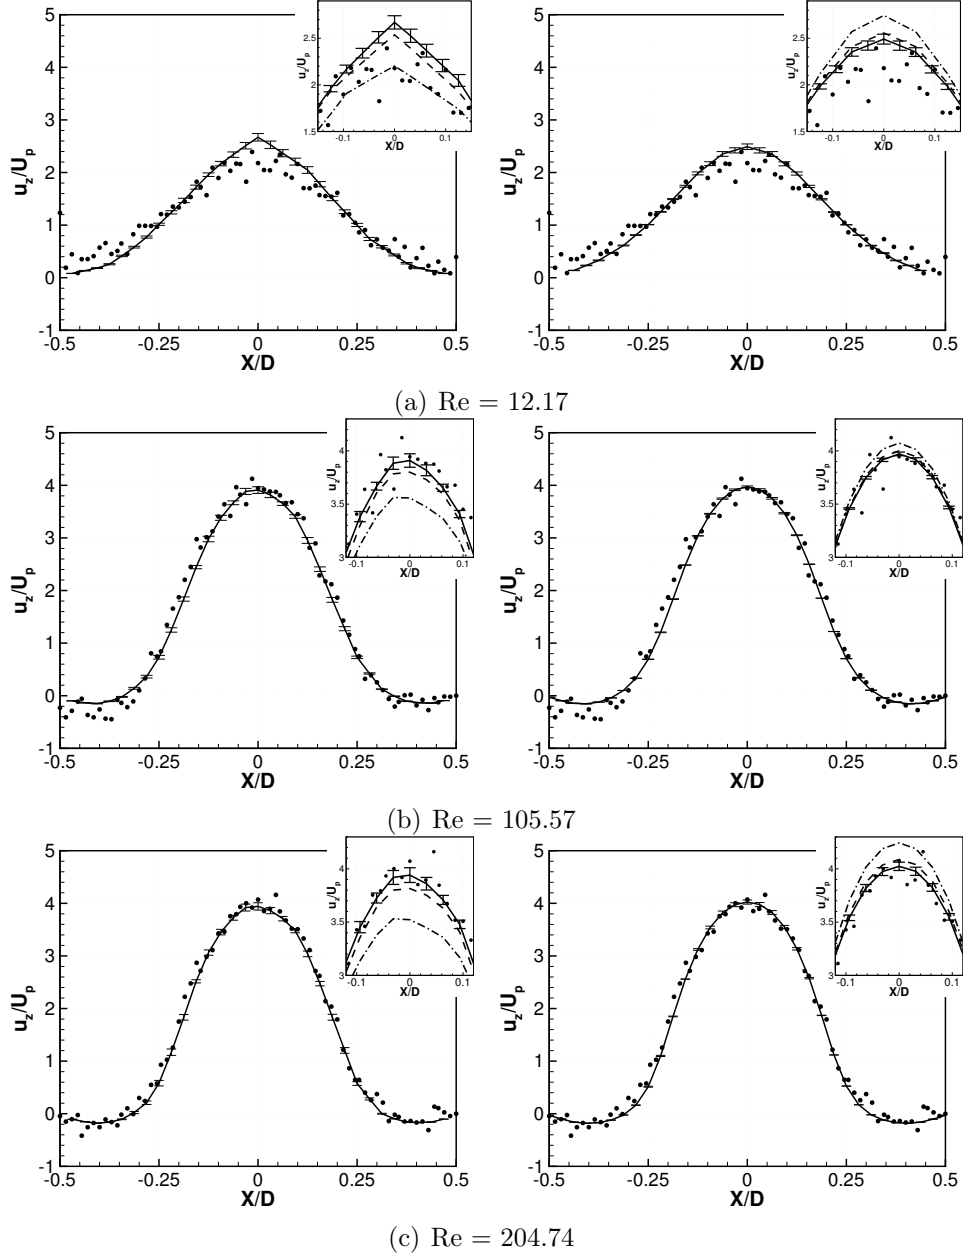
Figure 2.10: Comparison of interstitial velocity profiles with the data of [108] Along the $Y = 0$, $Z = 0$ line probe shown in Figure 2.9a. (—) Fine Grid Δ_1 shown with error bars, (- - -) Medium Grid Δ_2, (- · -) Coarse Grid Δ_3, ● Experiment. Left column shows body fitted solution. Right column shows fictitious domain solution.

by [35]. This particular configuration allows a high inertia jet to develop in the center of the channel that is not typically seen in the large random configurations the correlation was developed from.
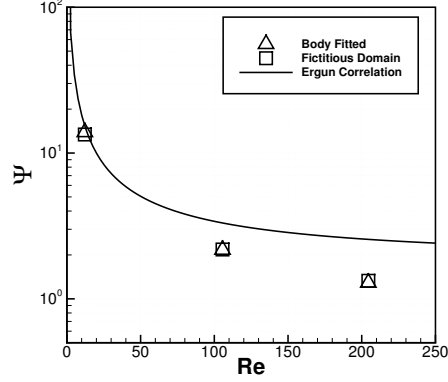
Figure 2.11: Non-dimensional pressure drop as a function of Reynolds number for steady flow through the simple cubic lattice.

We now consider results from the two unsteady Reynolds numbers in the periodic domain. In Figure 2.12, the interstitial profiles of the time averaged streamwise velocity $\overline{u_z}$ are shown along the $Y = 0$, $Z = 0$ line probe. The global relative uncertainty for the fine grid solution is good for both methods, less than 5% in each case. At $Re = 600$ the fine grid mean velocity profiles (Figure 2.12b) are in good agreement. Examination of the interstitial velocity profiles at $Re = 450$ (Figure 2.12a), however, reveals a noticeable discrepancy between the body fitted and fictitious domain predictions. In the fictitious domain solution, the average velocity profile is flattened out somewhat in the pore center, while the body fitted profile retains a larger maximum velocity and the parabolic character observed at Re=105 and 204. This suggests that the porescale unsteady dynamics predicted by the two approaches may be different. This suspicion is confirmed by Figure 2.13, which is shows the time history and frequency spectra of the transverse (X) velocity component at the point $X/D = -0.36$, $Y/D = 0$, $Z/D = 0$ computed on the fine grid using both approaches. At $Re = 450$, a single dominant frequency with a Strouhal number of $St = fD/U_p \approx 0.75$ is evident in the body fitted solution,

(a) Re = 450

(b) Re = 600
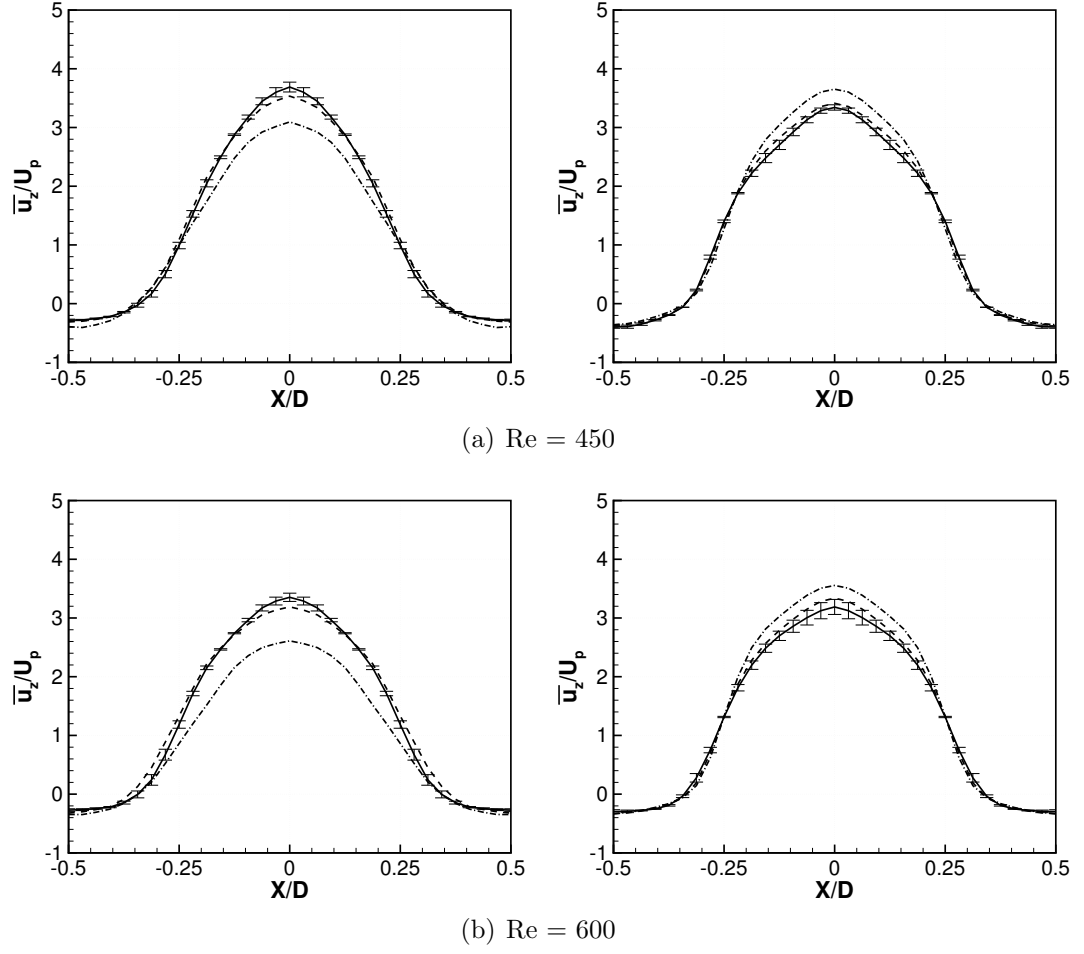
Figure 2.12: Time averaged streamwise velocity component along the line $Y/D = 0$, $Z/D = 0$ for unsteady flows in the periodic simple cubic lattice. Left column shows body fitted solution. Right column shows fictitious domain solution. (—) Fine Grid Δ_1 shown with error bars, (- - -) Medium Grid Δ_2, (- · -) Coarse Grid Δ_3.

while the fictitious domain solution shows energetic fluctuations over a broader range of frequencies. Interestingly, the fluctuations observed by both methods at $Re = 600$ seem to be in better agreement. The transition to unsteadiness and turbulence in packed beds is a very sensitive process, and it is possible that the dynamics of weakly unsteady flows could be altered by the type of grid being used. Tetrahedral grids cells will, in general, be more susceptible to numerical diffusion than uniform hexahedral cells. We believe that especially for flows in this transitional (sub turbulent) regime, modest amounts of dissipation due to the grid type could lead to noticeable changes in the dynamics of the mean flow.

A more appropriate level of confidence for these Reynolds numbers can be obtained by examining the uncertainty in the $\overline{TKE}$ distribution, shown in Figure 2.14. The fictitious domain approach has relatively low global uncertainty of around 4-5% and the main features seem to be captured along the interstitial profiles, even on coarse grids. The body fitted approach on the other hand has much larger global uncertainty, 16% and 10% for $Re = 450$ and 600 respectively, and the peaks in the interstitial profile change significantly on the coarse, medium and fine grids. At $Re = 450$, both the mean and RMS values of TKE are significantly lower in the body fitted solution relative to the fictitious domain solution, indicating more dissipation may be occurring due to the unstructured meshes.

These unsteady flow results underscore a troubling aspect of uncertainty assessment for CFD results in general. Quantities of interest in numerical solutions to complex flow problems need not converge in the same way, and false positives may indicate convergence in one variable when another is completely unresolved. In our case we may very reasonably judge that when using the body fitted approach the medium grid should provide a high quality solution for $Re = 450$ if we only consider convergence of the mean flow, when in fact the unsteady velocity fluctuations are completely unresolved and appear to degrade the accuracy of the time averaged solution. This emphasizes the need for thorough and quantitative grid convergence studies in CFD, especially for the complex configurations presented by packed bed and porous media flows.

(a) $Re = 450$

(b) $Re = 600$
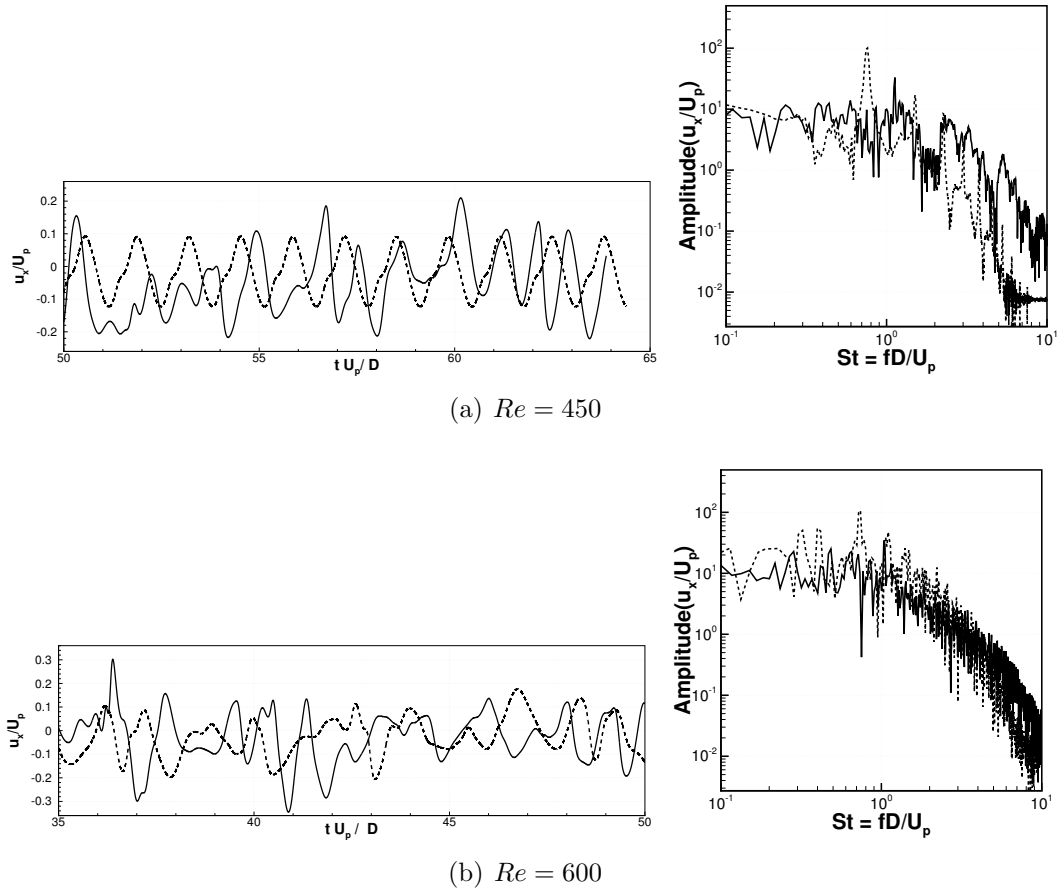
Figure 2.13: Time history of the transverse velocity component, u_x, and it's non-dimensional frequency spectra at a point in the main recirculation zone ($X/D = -0.36$, $Y/D = 0$, $Z/D = 0$) for the two unsteady Reynolds numbers simulated. (- - -) Body fitted approach, (—) fictitious domain approach.

(a) Re = 450

(b) Re = 600
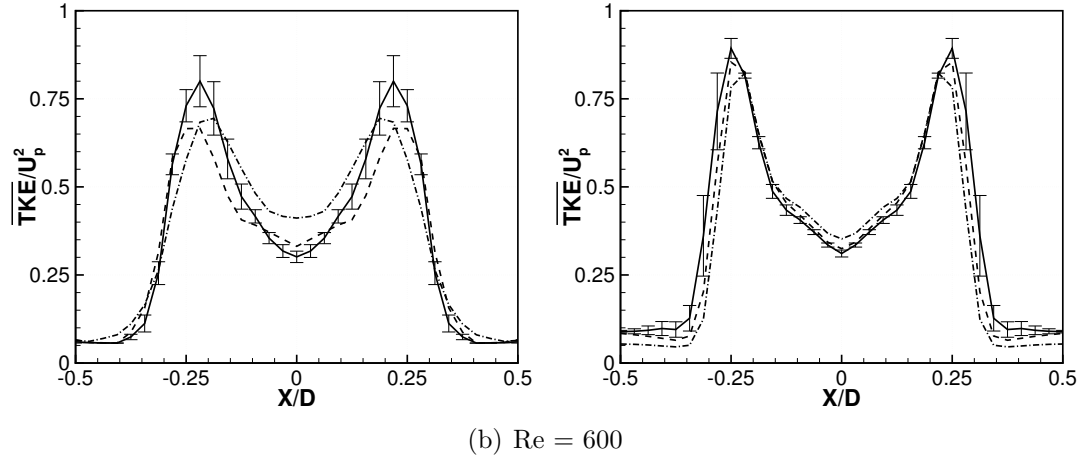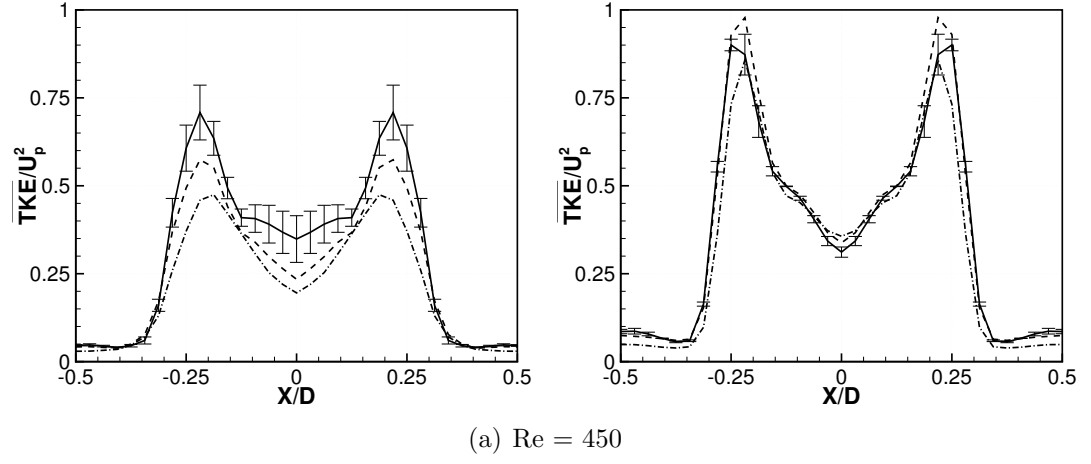
Figure 2.14: Time averaged turbulent kinetic energy along the line $Y/D = 0$, $Z/D = 0$ for unsteady flows through the simple cubic lattice. Left column shows body fitted solution. Right column shows fictitious domain solution. (—) Fine Grid Δ_1 shown with error bars, (- - -) Medium Grid Δ_2, (- · -) Coarse Grid Δ_3.

2.4   Scale up to realistic packed bed simulations

With the results presented up to this point in mind, it is possible to estimate the computational expense required for simulation of more general packed bed problems. Consider a rigid container that contains a matrix of randomly arranged spheres. The volume of the computational domain that must be meshed can be written in terms of the sphere diameter, D, as $\mathcal{V}_c = \epsilon_m N_c D^3$, where N_c relates the container volume to the sphere diameter (ie. a $3D \times 3D \times 6D$ container has $N_c = 54$), and ϵ_m is the *effective meshing porosity*. It is equal to 1 in the fictitious domain approach, where the entire fluid/solid region is meshed, and equal to ϵ in the body fitted approach where only the fluid domain must be discretized. For $\Delta = \mathcal{V}_{cv}^{1/3}$, the total number of control volumes, N_{cv}, required to mesh the volume, $\mathcal{V}_c$, will be:

$$N_{cv} = \epsilon_m N_c \left(\frac{D}{\Delta} \right)^3 \tag{2.16}$$

For the time accurate solutions obtained in this work, the time step, Δ_t, is restricted for temporal accuracy by the CFL number, which we define following [54] for an arbitrary cell as,

$$CFL = \frac{1}{2} \frac{1}{\mathcal{V}_{cv}} \sum_{faces} |U_f A_f| \Delta_t \tag{2.17}$$

Where A_f is the area of a single cell's face. If we assume that the length $A_f / \mathcal{V}_{cv}$ is constant for all faces of a given cell (true for all regular polyhedra), and we assume the face velocity to be proportional to the pore velocity, we can re-write the CFL number as,

$$CFL = \frac{\mathcal{C}_s |U_p| \Delta_t}{\Delta} \tag{2.18}$$

Where cell shape factor, $\mathcal{C}_s = \frac{A_{fa}}{\mathcal{V}_{cv}^{2/3}}$, has been introduced. This ratio of areas is equal to 1 for cubic cells, and 1.81 for *perfect* (regular) tetrahedral cells. In practice, it is

difficult to mesh a complex three dimensional domain with perfect tetrahedral cells everywhere and $\mathcal{C}_s$ can be much higher. On average for our tetrahedral meshes, we see that the time step must be 5 times lower when compared to the Cartesian mesh of the same resolution to obtain the same CFL, implying that $\mathcal{C}_s \approx 5$. With this in mind, we can express the total number of timesteps, N_t, required to simulate one non-dimensional time unit, $t = \frac{D}{U_p}$, as

$$N_t = \frac{t}{\Delta_t} = \frac{\mathcal{C}_s}{CFL} \cdot \frac{D}{\Delta} \qquad (2.19)$$

For perfect parallel scalability the computational expense (CPU-Hrs), of a time accurate, finite volume simulation is proportional to N_{cv} times N_t, implying that,

$$CPU \propto (\epsilon_m N_c) \cdot \left(\frac{\mathcal{C}_s}{CFL}\right) \cdot \left(\frac{D}{\Delta}\right)^4 \qquad (2.20)$$

The present finite volume solver shows good (nearly perfect) scalability for up to 500 processors [44]. The constant of proportionality in equation 2.20 could in general be dependent on Re, CFL, D/Δ, as well as the hardware being used to run the simulation, and is difficult to measure. Nonetheless, it is important to note that the computational expense should scale like $\left(\frac{D}{\Delta}\right)^4$. The speedup that is obtainable by moving from a body fitted to fictitious domain simulation of flow through a packed bed can be estimated from the GCI analysis presented in Section 2.3.4. If we require a fixed level of uncertainty in our solution variable, ϕ, say 5%, we can set the grid resolution, D/Δ in equation 2.20 to the required grid resolutions from table 2.3. We find that the theoretical speedup factor, $\mathcal{S} = CPU_{BF}/CPU_{FD}$, listed in table 2.4 for each case is between 1.9 and 65.

These estimates may seem unreasonably large until looking at actual CPU timings in as close of a 1:1 comparison as possible, for example, flow through the periodic simple cubic lattice. For this case, the CPU times required to simulate one non-dimensional time unit are summarized in table 2.5. This case was run with both methods at $Re = 450$ and 600 using $CFL \approx 0.5$. Simulations using the

Table 2.4: Estimated simulation speedup for obtaining a solution with 5% uncertainty. $\mathcal{S} = CPU_{BF}/CPU_{FD}$ computed using equation 2.20.

Re	ϕ	$D/\Delta_{5\%}$		$\mathcal{S}$
		BF	FD	
12	u_z	53	28	36
105	u_z	74	33	63
204	u_z	77	44	24
450	$\overline{u_z}$	125	76	19
450	$\overline{TKE}$	249	111	65
600	$\overline{u_z}$	96	103	1.9
600	$\overline{TKE}$	191	135	10

$D/\Delta = 64$ and 128 grids were performed on the Lonestar supercomputer at the Texas advanced supercomputing center. For the same Re and D/Δ, the fictitious domain simulation requires significantly less CPU time, between 63% and 84% as much as the body fitted simulation. The GCI analysis of the results has shown that if less than 5% uncertainty is required in the solution, the body fitted simulation could require roughly double the grid density of the fictitious domain simulation. Combining these two observations suggests that an actual speedup of about 5 times is likely for solutions of equivalent uncertainty. In practice the desired uncertainty level may be more or less strict than the 5% level used for the estimates here, leading to higher or lower grid density requirements. However, these scalability and speedup estimates should retain roughly the same character.

2.5   Flow through a randomly packed bed

We now examine the case of flow through a random sphere pack using the fictitious domain approach. This is intended to demonstrate the capability of the fictitious domain approach for the random sphere arrangements typical of engineering applications. The packing we consider is a relatively loose arrangement of 51 spheres in a box with sides $L_x = L_y = L_z = 4D$, resulting in the void fraction $\epsilon = 0.58$. It was

Table 2.5: Observed computational expense using both approaches for simulating the simple cubic lattice case in a periodic domain with $CFL \approx 0.5$

		CPU-Hr/(D/U)	
D/Δ	Re	Body Fitted	Fictitious Domain
64	450	111	80
128	450	340	284
64	600	118	87
128	600	470	298

generated using a *ballistic deposition* algorithm, similar to the method employed by [4]. The flow is driven in the positive Z direction by a constant inflow velocity boundary located $3D$ upstream of the packing so that the Reynolds number for the flow is $Re = \frac{U_p * D}{\nu} = 600$, where $U_{in} = U_p \epsilon$ is the assigned inflow velocity. A convective outlet condition is located $3D$ downstream of the packing. The flow is started from rest and allowed to develop to a non-dimensional time $tU_p/D = 130$.

The fictitious domain approach was chosen for this case because of the superior scalability demonstrated in the previous section as well as the guarantee of good mesh quality, even for random packings. Three Cartesian grids with uniform spacing $\Delta = D/80,\ D/56,\ D/40$ in the porespace are used to quantify the grid convergence. The grids are stretched toward the inlet and outlet faces so that the grids contain a total of 47 million, 16 million and 5 million control volumes, respectively. These grids provide a roughly constant refinement ratio of $r_{21} \approx r_{32} = 1.4$.

A qualitative visualization of the simulation domain and instantaneous flow field is provided in Figure 2.15. For this Reynolds number and packing geometry, the porescale flow is unsteady, but not turbulent, and is dominated by helical vortices, elongated in the mean flow direction. In their simulations of transitional flows in close packed ordered arrays, [48] observed that regions of strong helicity became more pronounced with increasing Reynolds number, up to the limit of unsteadiness. Following their definition, we compute the cosine of the angle between
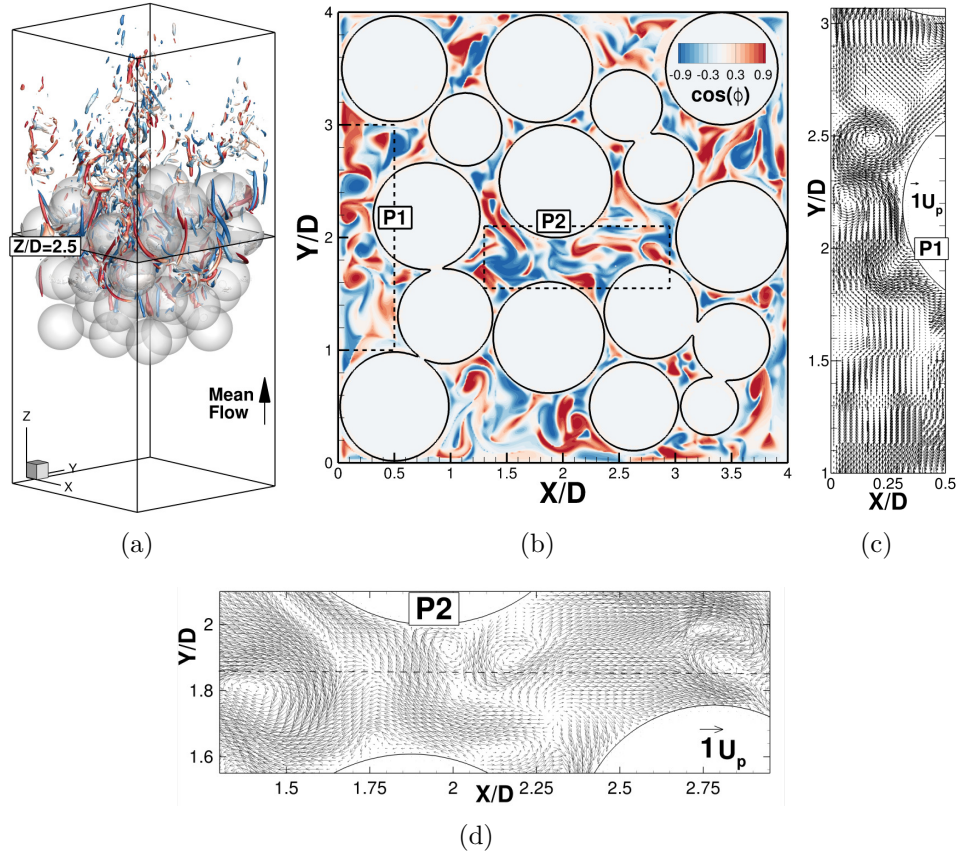
Figure 2.15: Visualization of the instantaneous porescale flow through the randomly packed bed of $N_{sp} = 51$ spheres. (a) shows the simulation domain and sphere surfaces (transparent) along with instantaneous vortical structures corresponding to isosurfaces of $\lambda_{ci}/\lambda_{ci}^{max} = 0.25$. The isosurfaces are colored by $\cos(\phi)$, indicating relative helicity. (b) shows contours of $\cos(\phi)$ on the $Z/D = 2.5$ cross stream slice and the location of two pores, $P1$ and $P2$ used in convergence analysis. (c) and (d) show the instantaneous, cross stream velocity vectors in the pores labeled $P1$ and $P2$.

the vorticity and the velocity vectors,

$$\cos(\phi) = \frac{\boldsymbol{\omega} \cdot \boldsymbol{u}}{\|\boldsymbol{\omega}\|\|\boldsymbol{u}\|}. \tag{2.21}$$

Regions with $\|\cos(\phi)\| \approx 1$ indicate strong helical or corkscrew like trajectories. Instantaneous vortex cores are detected using the swirling strength criteria, λ_{ci} [115], and are visualized as isosurfaces of $\lambda_{ci}/\lambda_{ci}^{max} = 0.25$ in Figure 2.15a, where λ_{ci}^{max} is the maximum value of swirling strength in the domain. The surfaces are colored by $\cos(\phi)$, indicating the relative orientation and magnitude of helical vortex motion. In Figure 2.15b, the contours of the $\cos(\phi)$ field are shown in the cross-stream, $Z/D = 2.5$ plane. Near the surface of the spheres, and at the exterior wall boundaries, the flow is nearly two dimensional, and $\cos(\phi)$ is small. In the pores however, regions of $\cos(\phi) \approx \pm 1$ are dominant. Assuming the velocity and vorticity are directed nominally in the $+Z$ direction (out of the page), blue regions indicate clockwise helical trajectories $(\cos(\phi) < 0)$ and counter-clockwise helical trajectories are indicated by red regions $(\cos(\phi) > 0)$. Instantaneous snapshots of the cross-stream vector field in two pores, labeled $P1$ and $P2$ are shown in Figures 2.15c and 2.15d. The length scale of these vortical structures appears to be roughly $D/10$ to $D/20$, and the magnitude of the cross-stream velocity is on the order of $1U_p$. In these vector plots, we also show the location of the line probes used to assess grid convergence of the pore velocity.

The grid convergence of the solution is demonstrated in Figure 2.16. First, the total hydrodynamic force on the spheres is non-dimensionalized according to equation 2.15, and plotted as a function of time in Figure 2.16a for the three different grid resolutions. Both the time average, $\overline{C_D}$, as well as the fluctuations in C_D are sensitive to the grid spacing. Performing a convergence analysis on $\overline{C_D}$, using equations 2.9 and 2.10, we estimate the order of convergence to be $\mathcal{P} = 1.30$, and determine the uncertainty in the fine grid solution to be $GCI = 8.24\%$ or $\overline{C_D} = 245.1 \pm 20.2$.

The convergence of the time averaged streamwise velocity, $\overline{u_z}$, in the $Z/D = 2.5$

slice is also examined. Averaging over all nodes with monotone convergence (44%
of all nodes in the post processing grid), the globally averaged observed convergence
rate for this quantity is $\mathcal{P}_G = 3.3$. This is above the formal order of accuracy of
the code, and could be due to the $\Delta = D/40$ solution being out of the asymptotic
convergence range in some regions. Using this convergence rate results in a globally
averaged uncertainty of $GCI_G = 5\%$ of $RMS(\overline{u_z})$.

Because of the randomness of the bed, certain regions of the porespace are more
grid converged than others. This is demonstrated by the velocity profiles extracted
from pores $P1$ and $P2$, which are plotted in Figure 2.16b and Figure 2.16c. The
line probe through pore $P1$, along the line $X/D = 0.16$, $Z/D = 2.5$ demonstrates
good convergence and low uncertainty for $\Delta = D/80$. In fact, the only slight
differences in the solution on all three grids seems to be near the counter-rotating
helical features between $Y/D = 2$ and $Y/D = 2.5$. The probe in the $P2$ pore along
the line $Y/D = 1.85$, $Z/D = 2.5$ indicates larger uncertainties and worse overall
convergence. We believe that the randomness of the bed could be the cause of
this behavior. Slight shifts in mean flow upstream of this plane for example could
propagate non-linearly downstream. This was not an issue when examining the
results for structured packings (Section 2.3.4), but should certainly be a considera-
tion when performing resolved simulations of random arrangements. Nonetheless,
the major features of the velocity profile along the line, including a maximum near
$X/D = 2.1$ and minima near $X/D = 1.7$ and $X/D = 2.6$, are consistent across the
three grids.

The statistical nature of the porescale flow has been examined by plotting the
PDFs of the instantaneous velocity components as well as $\cos(\phi)$ for one instant
in time in the porespace $(0.5 < Z/D < 3.5)$. In Figure 2.17 the PDF of each
of the velocity components is shown. The PDFs have been normalized to have a
standard deviation of one, for comparison with the Gaussian distribution, plotted
as a solid line, and with the lattice Boltzmann simulations of [48] for structured
close-packings. Similar to their study, there is a peak in the PDF of each velocity
component for $u_i = 0$, corresponding to the large amount of fluid influenced by the

(a)

(b)   (c)
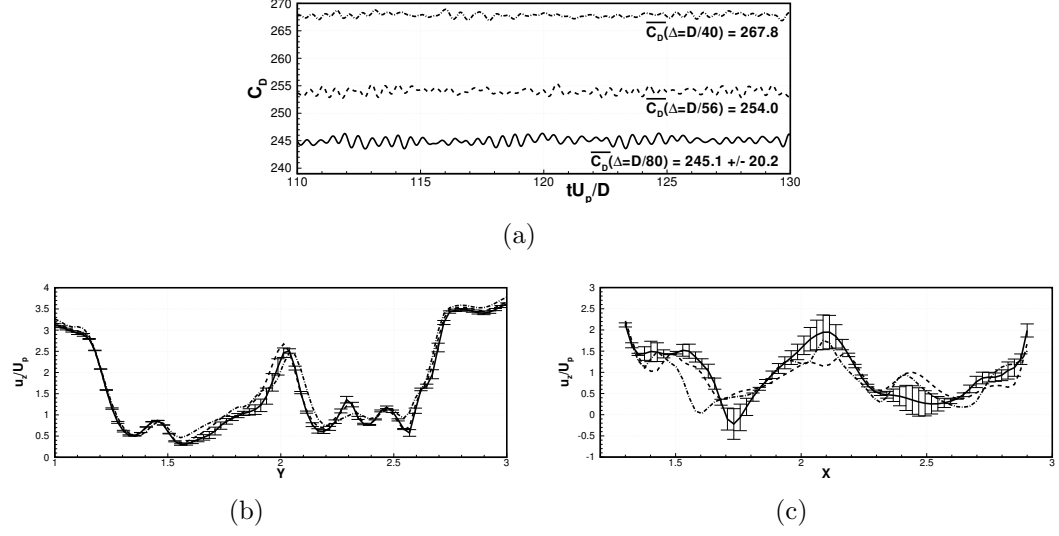
Figure 2.16: Grid convergence analysis for the randomly packed bed. (a) Time history of the drag coefficient. (b) Time averaged streamwise velocity, $\overline{u_z}$ non-dimensionalized by U_P in pore $P1$, along the line $X/D = 0.16$, $Z/D = 2.5$. (c) Same as (b) but along the line $Y/D = 1.85$, $Z/D = 2.5$ in pore $P2$.

solid boundaries. The cross stream components have a nearly Gaussian spread, but with wider tails at high velocity. In comparison, Hill and Koch observed tails with less spread than the normal distribution. This difference is likely due the high void fraction of the present sphere packing, which can accommodate larger cross-stream fluid motion relative to the structured close packing. The PDF of the streamwise component, u_z indicates that a significant portion of fluid, roughly 7%, is directed upstream, also similar Hill and Koch's observations. The PDF of of $\cos(\phi)$, shown in Figure 2.18, corresponds well with their results. The peak at 0 is due to the strong influence of the walls, near which the flow is two dimensional, even at this relatively high Reynolds number. The symmetric tails of the PDF, with peaks at ± 1 show the high degree of strong helical motion in the porespace.

(a)
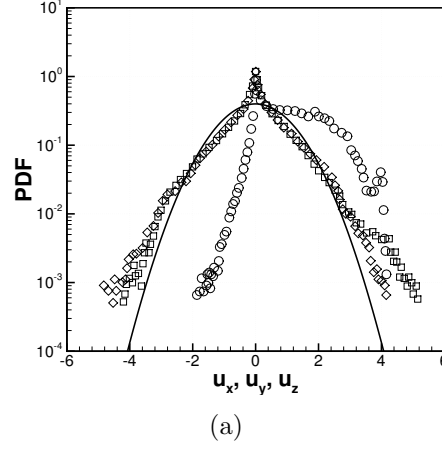
Figure 2.17: Normalized probability distribution function of the velocity components in the porespace for flow through the random packing. $\Diamond$ u_x, $\Box$ u_y, $\bigcirc$ u_z, — Gaussian distribution with zero mean and standard deviation of one.

(a)
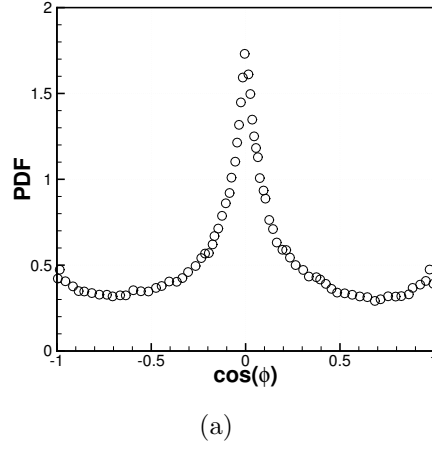
Figure 2.18: Probability distribution function of the cosine of the angle between the velocity and vorticity vectors in the porespace for flow through the random packing.

2.6 Conclusions

The performance of a fictitious domain approach relative to a body fitted approach has been assessed for simulation of flow through packed beds of spheres. The body-fitted approach utilizes a cylinder bridge model during mesh generation to avoid troublesome sphere to sphere contact points. This strategy avoids major modification of the solid geometry and removes areas where small elements are required. The main advantage of the fictitious domain approach for fixed bed and porous media problems is that it can use regular Cartesian grids, and avoids unstructured mesh generation all together. In this case, the solid-fluid interface is accurately represented using Lagrangian marker points with subgrid resolution, and a rigidity constraint within the solid bodies is imposed to enforce the no-slip boundary condition.

Several test cases have been examined to address particular concerns associated with simulating flow in packed beds. It was first shown using decaying Taylor vortices that local grid quality issues can result in a significant increase in solution error, even for only mildly skewed cells. This is an important point to consider for packed bed simulations because in large scale random packings, some skewed elements are inevitable when creating a body fitted mesh. The case of Stokes flow in a dilute periodic array was used as a basic hydrodynamic validation case relative to packed beds, without having to consider sphere to sphere contact points. Using the case of two contacting spheres, it was shown that for moderate Reynolds numbers the region of fluid very close to the contact point has little effect on the bulk flow, and may be safely removed as is deliberately done in the bridge meshing technique. This is also reassuring for the fictitious domain technique, which can only resolve sphere to sphere contact points at the scale of the Cartesian mesh. The final test case considered, flow through a section of a simple cubic lattice, allowed us to perform a detailed comparison of both methods in a prototypical sphere packing. In this case, a systematic grid refinement study has been performed, with the grid convergence index (GCI) used to assess uncertainty of the fine grid solutions. Both methods perform well at steady flow rates ($Re = 12, \ 105, \ 204$)

in the sense that they predict interstitial velocity profiles in agreement with the experimental data of [108]. The difference is that fictitious domain approach is able to obtain a solution with low uncertainty using a much lower grid density than the body fitted approach. Similar trends are observed for the unsteady flow rates ($Re = 450, 600$). In these cases, the body fitted approach is observed to have large uncertainty associated with the TKE distribution, making it undesirable for use at higher Reynolds numbers.

The scalability of both approaches for general, random packed bed flow simulations has been estimated. For a predetermined target level of uncertainty the computational expense, estimated to be proportional to the number of control volumes times the number of timesteps, is found to be between 1.9 and 65 times more for the body fitted approach compared to the fictitious domain approach for the cases considered. Actual CPU timings for the periodic SCP channel indicate that speedup may be more modest, perhaps 5 times. Finally, the fictitious domain approach has been applied to a more practical case of unsteady flow at $Re = 600$ through a random packing of 51 spheres. The results show grid convergence in global and local time averaged quantities, although the randomness of the packing results in significant variation in the local uncertainties. Nonetheless, results capture complex instantaneous porescale flow features and are used to further explore helical motions in random arrangements. Good statistical agreement is obtained with prior Lattice Boltzmann results of [48].

2.7 Acknowledgments

Chapter 3: Characteristics of Pore-Scale Vortical Structures in Random and Arranged Packed Beds of Spheres

Justin R. Finn and Sourabh V. Apte

Abstract

The characteristics of pore scale vortical structures observed in moderate Reynolds number flow through mono-disperse packed beds of spheres are examined. Our results come from direct numerical simulations of flow through (i) a periodic, simple cubic arrangement of 54 spheres, (ii) a wall bounded, close packed arrangement of 216 spheres, and (iii) a realistic randomly packed tube containing 326 spheres with a tube diameter to sphere diameter ratio of 5.96. Pore Reynolds numbers in the steady inertial ($10 \lesssim Re \lesssim 200$) and unsteady inertial ($Re \approx 600$) regimes are considered. Even at similar Reynolds numbers, the vortical structures observed in flows through these three packings are remarkably different. The interior of the arranged packings are dominated by multi-lobed vortex ring structures, which align with the principal axes of the packing. The random packing and the near wall region of the close packed arrangement are dominated by helical vortices, elongated in the mean flow direction. In the simple cubic packing, unsteady flow is marked by periodic vortex shedding that occurs at a single frequency. Conversely, at a similar Reynolds number, the vortical structures in unsteady flow through the random packing oscillate with many characteristic frequencies.

3.1 Introduction

The flow of fluid through a packed bed of spheres is a fundamental problem, rich in its applicability to a variety of disciplines. For example, groundwater hydrologists [114, 88], chemical engineers [17] and nuclear reactor designers [45] often consider flow through strikingly similar arrangements of fixed, contacting spheres as a prototypical problem for more complex systems. Despite such broad importance, the existing body of knowledge related to the flow physics of such problems is limited, particularly at moderate to large flow rates. For Stokesian, creeping flow, fluid streamlines conform naturally to the boundary of the porespace. In this case, the general behavior of the flow may be adequately described using geometric
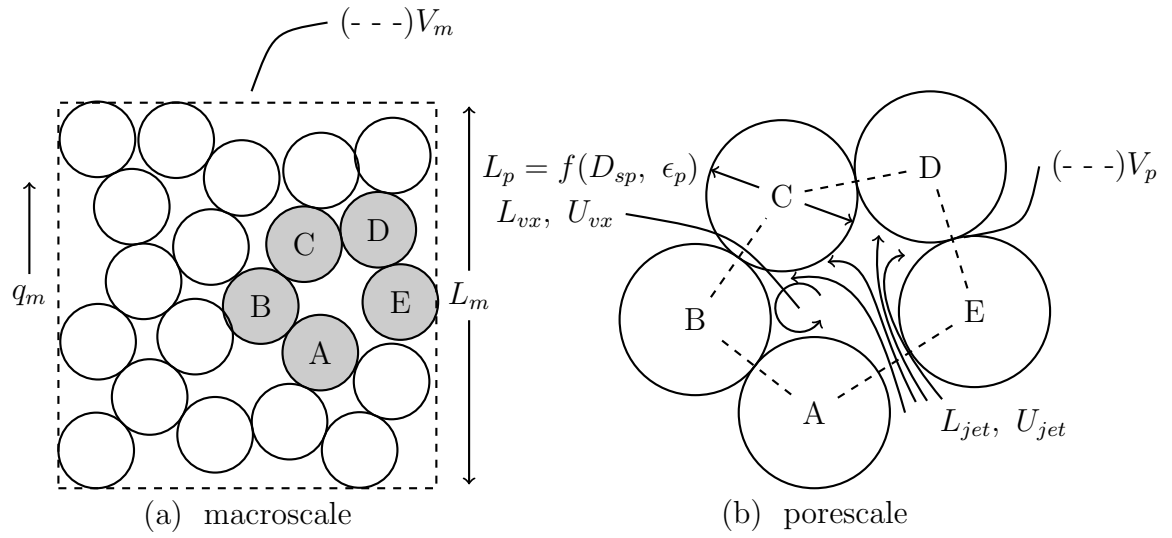
Figure 3.1: Illustration of the multiscale nature of porous media flows. (a) shows a representative *macroscale* volume. (b) shows the *porescale* where complex hydrodynamic interactions occur.

or network models [8], and simple relationships are available for macroscale properties such as *Darcy's law* for permeability. However, at larger flow rates, porous media and packed bed flows become highly non-linear and multiscale in nature due to the contribution of flow inertia. This makes a-priori determination of the flow characteristics difficult or impossible even for simple, homogeneous sphere packings.

This multiscale nature of inertial flows through porous media and packed beds is illustrated in FIG. 3.1. The *macroscale* is often used to describe the largest representative scale encountered in the field or laboratory. The macroscale flow can be simply characterized with a length scale L_m, volume V_m, porosity, ϵ_m and flux q_m. Typically it is macroscale properties including permeability, dispersion coefficients, or reaction rates that are of interest to practical engineering applications. Such properties derive from hydrodynamic interactions at the *porescale* or *Darcy scale*, shown as a single pore subset of the macroscale volume in FIG. 3.1b. At the porescale the length and volume scales are functions of the local pore geometry, L_p, $V_p = f(D_{sp}, \epsilon_p)$. These scales, along with the Reynolds number, govern the hydrodynamic character of the flow. We define the Reynolds number, following Ergun [24] and Dybbs & Edwards [21], as $Re = \frac{\rho U_p D_{sp}}{\mu} \cdot \frac{\epsilon}{1-\epsilon}$, where D_{sp} is the sphere diameter, $U_p = q_m/\epsilon_m$ is the pore averaged flow velocity. As Reynolds number is increased from creeping flow rates, it has been observed both experimentally [21, 108], and computationally [50, 48] that the porespace becomes increasingly dominated by inertial flow features such as steady or unsteady jets, vortices, and stagnation regions.

Vortical flow features are of particular interest in porous media because of the implications they have for mixing and dispersion in the porespace. Vortices can enhance mixing by collecting streamlines from distant sources and ejecting them along new trajectories after exchange of fluid. Helical vortices with corkscrew trajectories are good examples of this, and are common features in inertial flows through porous media [48]. At unsteady Reynolds numbers, flow oscillations provide an effective means for enhancing fluid transport, and vortical features are often

the backbone of these oscillations. Alternatively, vortices can act as *dead zones*, such as stationary recirculation bubbles in bluff body wakes, where entrained fluid is bounded by closed streamlines and residence time is high.

The case for detailed and accurate simulation to investigate these types of features is particularly compelling considering the challenges associated with making detailed three dimensional flow measurements in the interior of a solid sphere matrix, although new experimental techniques such as index of refraction matched PIV, and MRI are helping [3, 52, 106, 51]. Broadly, resolved simulation methods used for flow through porous media can be classified by the way in which they represent the fluid/solid interface. To obtain a detailed description of the pore-scale flow field, appropriate boundary conditions (typically no-slip) need to be suitably enforced on the solid boundaries of the porespace, and the porespace needs to be sufficiently resolved by the computational grid. Body fitted grid methods, where the Navier Stokes equations are solved on unstructured grids that conform to the solid boundaries of the porespace, have probably been the most popular type of approach [18, 4, 34, 74], and is the method that is pursued here. However, body fitted grid generation is complicated for porous media and packed bed geometries, and a care must be taken to avoid poor mesh quality, as described in the next section. Alternatives to the body fitted approach exist, including the *lattice-Boltzmann method*, *immersed boundary methods* and *fictitious domain methods*. The Lattice-Boltzmann method solves the Boltzmann equation of particle motion on a regular grid or *lattice*, which approximates the statistical distribution of fluid motion in space and time. This method has been used by Hill, Koch and coworkers [50, 49, 48, 47] to investigate inertial flows in packed beds and by Pilotti [87] in a synthetically generated porous medium. Immersed boundary and fictitious domain methods, which have been applied widely to problems with complex or moving boundaries, are also beginning to see application to porous media flows [78, 30]. These methods use Cartesian, non body-conformal grids, and enforce surface boundary conditions with the addition of a local forcing term in the momentum equation. Recently Finn & Apte [30] have made comparisons of the

traditional body fitted approach to a Cartesian grid fictitious-domain approach, and shown the latter to be favorable in terms of computational overhead required for accurate simulations of flow through packed beds.

The remainder of this paper will focus on the characterization of porescale vortical features in packed beds of spheres. These features are investigated using a parallel Navier-Stokes solver with unstructured, body fitted meshes. To understand the effects of the both flow inertia and solid geometry on these flow features, three different sphere packings are investigated at Reynolds numbers spanning the steady inertial ($10 \lesssim Re \lesssim 200$), and unsteady inertial flow $200 \lesssim Re \lesssim 1000$ regimes.

3.2 Computational Approach

Consider the incompressible Navier-Stokes equations for the entire fluid domain:

$$\rho \left(\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} \right) = -\nabla p + \mu \nabla^2 \mathbf{u} \tag{3.1}$$

$$\nabla \cdot \mathbf{u} = 0 \tag{3.2}$$

where ρ is the density, $\mathbf{u}$ the velocity vector, p the pressure, μ the fluid viscosity. The equations are discretized on unstructured, tetrahedral meshes, and solved using a parallel, second order accurate fractional step solver. The code is parallelized using Message Passing Interface (MPI), allowing for larger scale simulations by distributing the required memory over many processors. For additional details regarding the numerical implementation, the reader is referred to [71].

Unstructured mesh generation for complex geometries is a non-trivial procedure in general, and in packed beds the process is complicated by sphere-to-sphere contact points, near which elements can become unmanageably small and have high aspect ratios. We employ a cylinder bridge model [56], whereby every contact point is bridged by a small cylinder, removing a small amount of typically stagnant fluid as is illustrated in FIG. 3.2. This process has been parameterized and automated using the commercial package Pointwise®. Details regarding the mesh generation
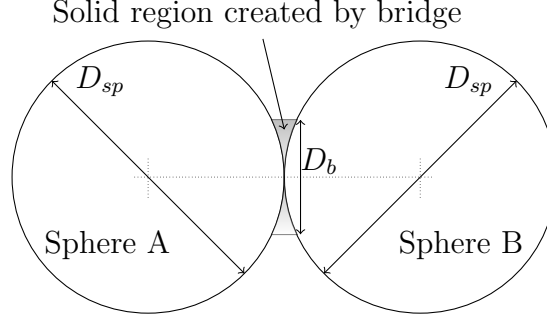
Figure 3.2: Schematic of the fillet bridge created between two contacting spheres.

Table 3.1: Packed bed simulations

	Packing	Dimension			Re	D_{sp}/Δ	N_{cv} total
Steady	SCP	$3 \times 3 \times 3$			$10 < Re < 250$	29	$1.2m$
	HCP	$6 \times 6 \times 6$			$10 < Re < 150$	29	$12.3m$
	Random	D_{tube}/D_{sp}	$=$	$5.96,$	$10 < Re < 150$	29	$16.8m$
		326 sp. total					
Unsteady	SCP	$3 \times 3 \times 6$			529	56	$7.3m$
	Random	D_{tube}/D_{sp}	$=$	$5.96,$	600	56	$32m$
		326 sp. total					

procedure as well as detailed validation of the flow solver for packed bed type problems can be found in [30].

3.3 Flow Through Packed Beds of Spheres

We now present results for flow through three different packed beds of spheres at both steady and unsteady Reynolds numbers. The sphere packings considered are shown in FIG. 3.3. Their characteristics, and the relevant parameters used in the simulations of flow through each are summarized in TAB. 3.1. The three packings were chosen to complement each other, as their porescale geometries vary significantly. The simple cubic packing (SCP) shown in FIG. 3.3a is the loosest possible packing arrangement with $\epsilon_{SCP} = 0.47$. It contains periodic boundaries on

(a) SCP(periodic)

(b) HCP
(wall bounded)

(c) Random
(wall bounded)
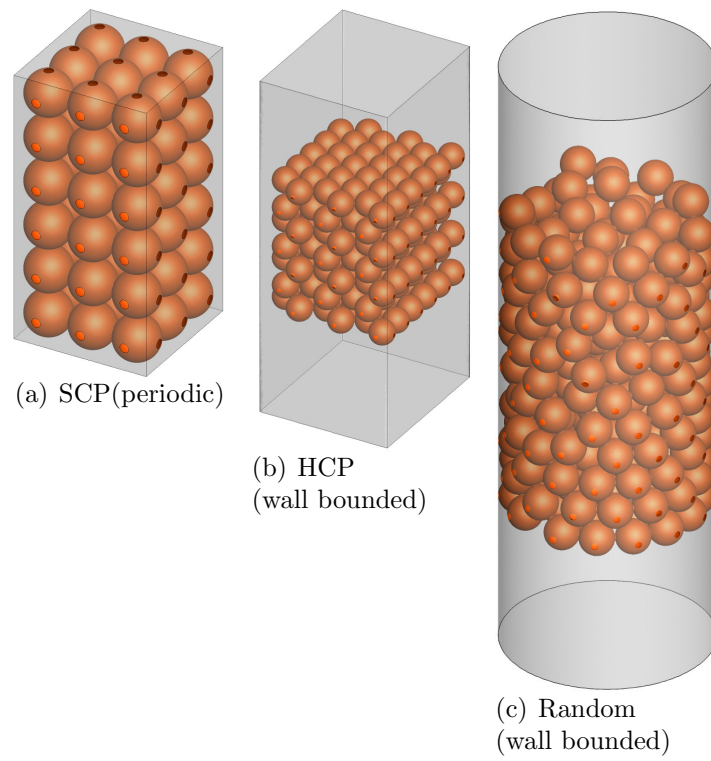
Figure 3.3: The three sphere packings considered. Flow is from bottom to top in each case.
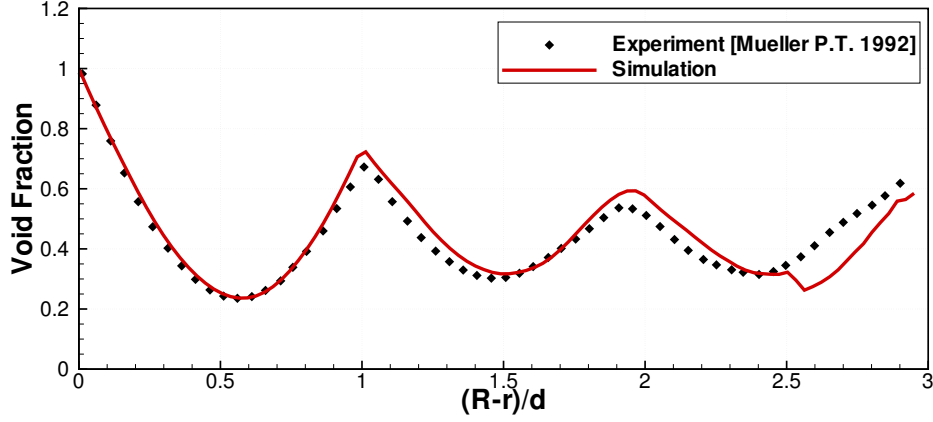
Figure 3.4: Radial void fraction variation of the random packing. Comparison of the present simulation to the experimentally measured void fraction of Mueller [72].

all sides, which allows for the generation of a flow field unaffected by the presence of external walls. Also, the repeating nature of the packing provides a single geometric length scale in the pore. A $3 \times 3 \times 3$ arrangement is used for the steady flows and a $3 \times 3 \times 6$ arrangement for the unsteady flows. The second packing is a $6 \times 6 \times 6$ hexagonal close packing containing (216 spheres) shown in FIG. 3.3b. It is one of two arrangements (face centered cubic being the other) with lowest possible void fraction of $\epsilon = 0.26$. However, this is only in the interior of the packing, and the solid, external boundaries result in a significantly higher mean void fraction, $\epsilon_{HCP} = 0.40$. This *wall effect* increases the Reynolds number near the wall and can lead to significant flow channeling. Finally, the packing shown in FIG. 3.3c is a tube packed randomly with 326 spheres. The tube diameter to sphere diameter ratio is $D_{tube}/D_{sp} = 5.96$. The packing was generated following the parameterized, sequential procedure of Mueller [73] and has a mean void fraction of $\epsilon_{random} = 0.47$ and a total packing height of $h/D_{sp} \approx 12$. Because of the exterior tube boundary, there is a significant wall effect in the radial porosity variation, $\epsilon(r)$, which is shown in FIG. 3.4 alongside the void fraction variation of a laboratory generated packed column [72]. The trend of $\epsilon(r)$ obtained with the current packing do a good job of matching the measured trend.

The pore volume of each geometry is meshed using the parameterized method described in [30]. All meshes used in the steady flow simulations have a grid spacing of $D_{sp}/\delta_{cv} \approx 29$, which was shown to provide grid convergence and good experimental agreement in prototypical test cases. The unsteady simulations use a more refined mesh with $D_{sp}/\delta_{cv} \approx 56$. The ratio of fillet diameter to sphere diameter is $D_f/D_{sp} = 0.2$ for all simulations. The total number of tetrahedral control volumes, N_{cv} used in the body fitted meshes of each geometry varies from $1.2m$ to $16.8m$. The flow is forced in the positive Z direction by a constant pressure gradient in the SCP cases, and by a constant flux inflow condition in the HCP and random cases. Each simulation is started with the flow at rest, and the flow solution is advanced in a time accurate manner until the solution has reached a stationary state. In the steady flow cases, we use the total kinetic energy, $\sum_{cv} u_i u_i$, in the porespace as an indicator. When this quantity does not change significantly, the inflow flux or pressure gradient is increased to the next target Reynolds number. In the unsteady flow cases, the simulation is run for several non-dimensional flow through times, $\overline{T} = tU_p/D_{sp}$, so that the velocity fluctuations in the porespace reach a statistically stationary state.

3.3.1   Steady Flow

Steady Reynolds numbers were chose to thoroughly sample the inertial flow regime range between creeping flow ($Re \lesssim 10$) and inertially dominated flow near the threshold of unsteadiness ($Re \gtrsim 200$) [21]. The non-dimensional pressure drop, $\Psi = -\frac{\Delta P}{L_m} \frac{D_{sp}}{\rho U_p^2} \frac{\epsilon_m}{1-\epsilon_m}$ measures the relative resistance to flow through the packed bed. Here, ΔP is the difference in pressure measured at the outlet and inlet of the packing. This is plotted in FIG. 3.5 vs Reynolds number for each of the three packings. Also plotted are the correlations of Ergun [24] for infinite, random porous media, and Eisfeld & Schnitzlein's [23] correlation for randomly packed tubes with finite tube to sphere diameter ratios. The periodic SCP arrangement offers the least flow resistance at all Re, followed by the HCP arrangement. This is due in part to the preferential flow channeling between rows of spheres in the SCP

arrangement and along the walls in the HCP arrangement. The random packing, which lacks long open channels, shows the highest resistance for all Re, and agrees reasonably well with the correlation from [23].

Next, the porescale vortical structures are examined in each arrangement. Due to the large amounts of shear in the porespace, using vorticity alone for vortex detection is generally noisy and unfruitful. Instead, the λ_2 criteria of Jeong & Hussain [53] is used. This criteria identifies vortices as regions with a local pressure minima and can therefore be identified as regions where

$$\lambda_2(\mathbf{S}^2 + \Omega^2) < 0 \qquad (3.3)$$

Here, $\mathbf{S}$ and Ω are the symmetric and anti-symmetric components of the velocity gradient tensor respectively, and the operator λ_2 returns the intermediate (second) eigenvalue of a symmetric matrix. In FIG.'s 3.6, 3.7 and 3.8, several visualizations are provided of the characteristic vortical structures seen in SCP, HCP, and random arrangements respectively. The SCP arrangement, shown in FIG. 3.6 is dominated by two distinct regions that can be identified easily with a simple streamline visualization (FIG. 3.6a): (i) high velocity channels through which the flow passes undisturbed, explaining the low values of Ψ, and (ii) large recirculation regions in the gaps between rows of spheres. These recirculation regions are actually complex, multi-lobed vortex ring structures, shown with an isosurface of λ_2 in FIG. 3.6b. These *vortex rings* are actually more like *vortex cubes*, with four distinct sides each aligned along a cross-stream axis. This is further illustrated in FIG. 3.6c, where the tangent vectors are plotted on a plane passing through the center of one of these vortex cubes, perpendicular to the mean flow. The vectors show the complex three-dimensionality of the flow feature, and the close relationship between the pore geometry and the behavior of the vortical structure.

The characteristic vortical structures observed in the HCP arrangement are shown in FIG. 3.7 as isosurfaces of λ_2. Similar to the SCP arrangement, the dominant porescale vortical feature in the interior of the packing is a multi-lobed vortex ring (FIG. 3.7a-b). These rings are located upstream of the forward stagnation
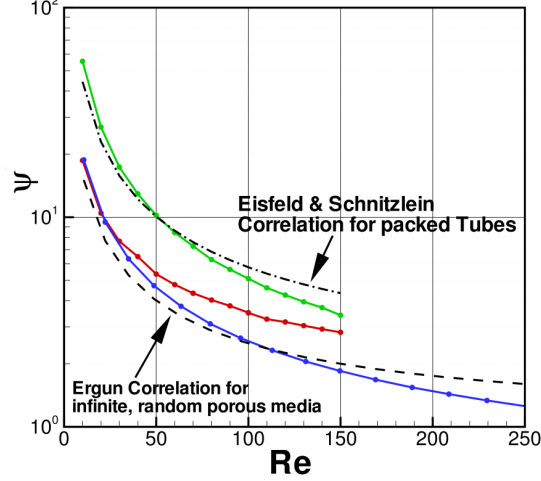
Figure 3.5: Non-dimensional pressure drop as a function of Reynolds number in the steady flow regime for the three sphere packings. SCP (blue) HCP (red) Random (green) The correlations of Ergun [24] and Eisfeld & Schnitzlein [23] are plotted for comparison.

points on the spheres in the interior of the packing. Due to the geometric orientation of the packing, the ring has 3 lobes that align with the principal axes of the packing as well as the major flow directions away from the stagnation point. Near the bounding walls, the flow is accelerated significantly due to the increased porosity and less tortuous channels compared to the interior. These factors give rise to elongated helical vortices shown in FIG. 3.7c. These helical structures are regions where the velocity and vorticity vectors are roughly aligned, and the flow spirals through the porespace with a roughly corkscrew motion.

This type of helical vortex is the dominant inertial feature everywhere in the random packing. In FIG. 3.8a-d, one typical feature is visualized using stream ribbons colored by λ_2, (red colors indicating the strongest vortical region) for $Re = 10$, 50, 100, and 150, showing its development from near creeping flow to strongly inertial flow. The stream ribbons that enter the feature at all Re originate from two sides of the bottom sphere and collect in a channel where the streamwise flow is accelerated. As Re is increased from the creeping regime, flow inertia pushes

the streamlines away from the right hand sphere, and the two streams entering the feature begin to twist around one another as they travel upward in a corkscrew like motion. This radial motion is non-existent at $Re = 10$ (FIG. 3.8a), first appears at $Re \approx 50$ (FIG. 3.8b), and strengthens with Re up to $Re = 150$ (FIG. 3.8c-d). Despite the variety of pores that exist in the random packing, this single feature is very representative of the majority of vortical structures seen for these steady flow rates. There are few, if any, of the vortex-ring like structures that were observed in the interior of the SCP and HCP arrangements, likely due to the lack of geometric symmetry and organization.
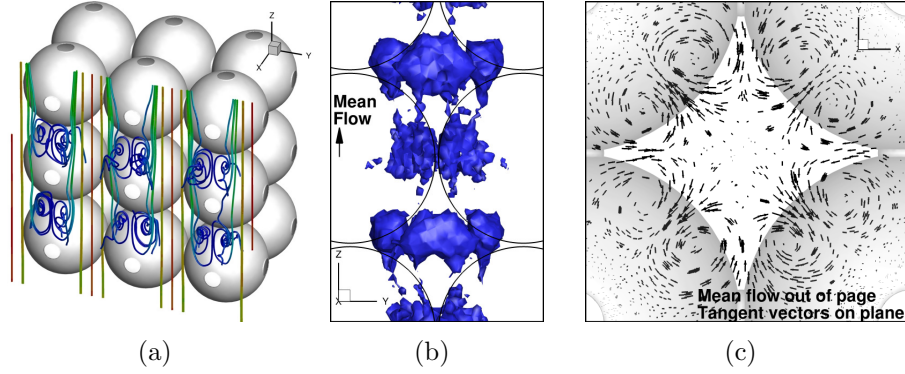
Figure 3.6: SCP flow at $Re = 230$. (a) Streamlines colored by streamwise velocity (red = fast), (b) isosurfaces of λ_2, (c) velocity field on a cross-stream plane.

3.3.2 Velocity, Vorticity & Helicity statistics

In this section, the porescale flow fields are examined in a statistical sense to quantify the influence of pore geometry and Reynolds number on the development of the various vortical flow features visualized in the previous section. The probability distribution functions are computed for the streamwise (z) and cross-stream (x) components of both velocity, $\boldsymbol{u}$, and vorticity, $\boldsymbol{\omega} = \nabla \times \mathbf{u}$, as well as the λ_2 vortex detection criteria, and the cosine of the angle between the velocity and vorticity
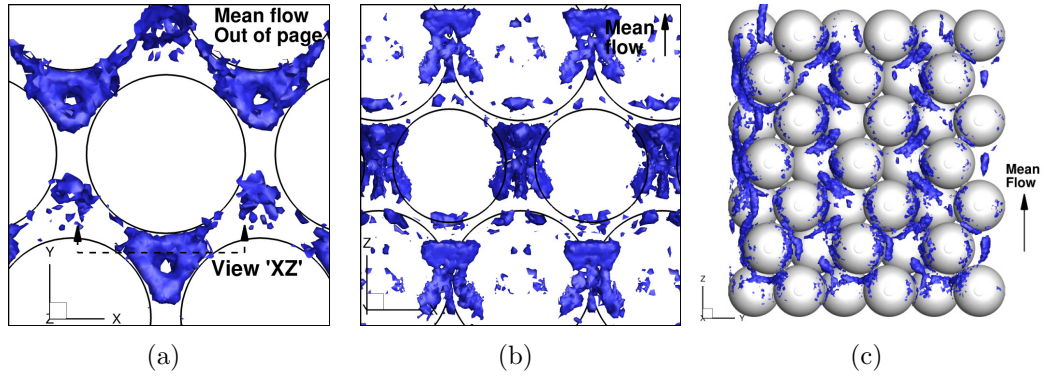
Figure 3.7: HCP flow at $Re = 150$. Isosurfaces of λ_2 showing the interior vortex ring structures (a,b), and the elongated helical vortices in the wall region (c).
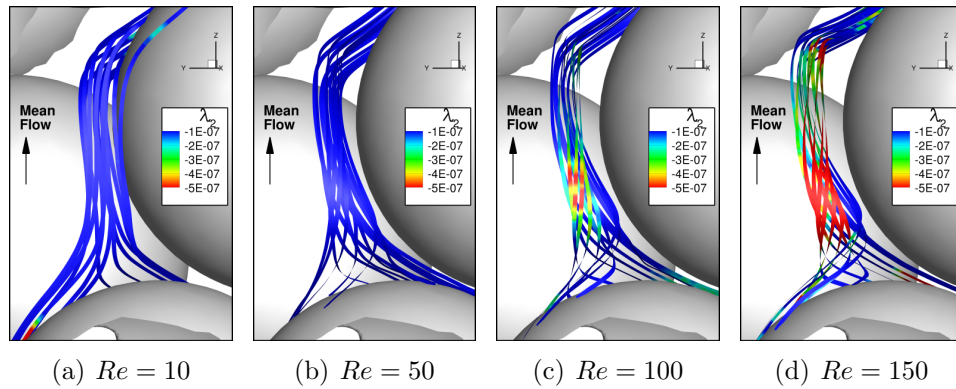
Figure 3.8: Development of a porescale helical vortex in the random arrangement. Stream ribbons colored by λ_2 (red indicates strong swirl).

vectors,

$$\cos(\phi) = \frac{\boldsymbol{\omega} \cdot \boldsymbol{u}}{\|\boldsymbol{\omega}\| \|\boldsymbol{u}\|}. \tag{3.4}$$

Regions with $\|\cos(\phi)\| \approx 1$ indicate strong helical or corkscrew like trajectories. In their simulations of transitional flows in close packed ordered arrays, Hill & Koch [48] observed that regions of strong helicity became more pronounced with increasing Reynolds number, up to the limit of unsteadiness. To facilitate comparison with their study, all PDFs with the exception of $\cos(\phi)$ are normalized so that they both integrate to and have a standard deviation of unity. Where appropriate, the standard normal distribution or the double exponential distribution are also plotted for comparison. The PDFs have been computed in the interior of the porespace only, away from the influence of the inflow and outflow regions in the random and HCP arrangements.

Figure 3.9 shows the PDFs of the cross stream (a-c) and streamwise (d-f) components of velocity for the three different arrangements. The four sets of symbols indicate the results for $Re = 10$, 50, 100, 150, which spans the range of development of the inertial core in these flows. The standard normal curve is shown for comparison to the cross stream PDFs. When normalized, the cross-stream PDFs (Fig. 3.9a-c) obtained for each sphere arrangement collapse nicely for all Reynolds numbers, similar to the observations of [48, 57]. The SCP and random arrangements have PDFs of u_x that agree well with the standard normal distribution, but with slightly wider tails. The HCP arrangement shows a much wider distribution of cross stream velocity, most likely due to the influence of high porosity, high velocity wall region. In the streamwise velocity, a more noticeable change occurs with increasing Reynolds number, especially in the SCP arrangement (Fig. 3.9d). In this case, the $Re = 10$ PDF declines gradually with increasing velocity, while higher Reynolds numbers have a high velocity "inertial knee" in the PDF. This is due to the development of the inertial core, which is especially prevalent in the channel-like SCP geometry. For all geometries, the amount of fluid directed in the negative z (upstream) direction increases with Reynolds number, indicating the increasing influence of vortical regions aligned perpendicular to the
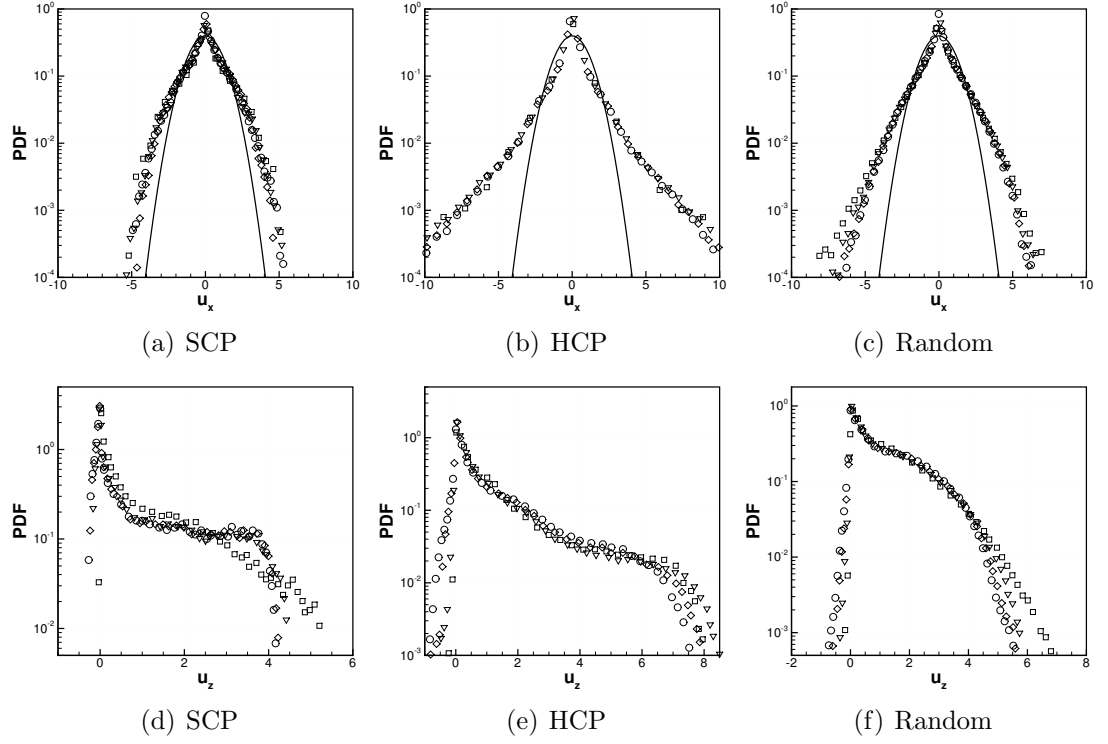
Figure 3.9: Normalized velocity PDFs for steady flow rates through the three different packed beds. Top row (a-c) shows the cross stream, u_x component. Bottom row (d-f) shows the streamwise, u_z component. Different symbols correspond to pore Reynolds numbers of roughly $Re = 10$ □, $Re = 50$ ▽, $Re = 100$ ◇, $Re = 150$ ○

mean flow.
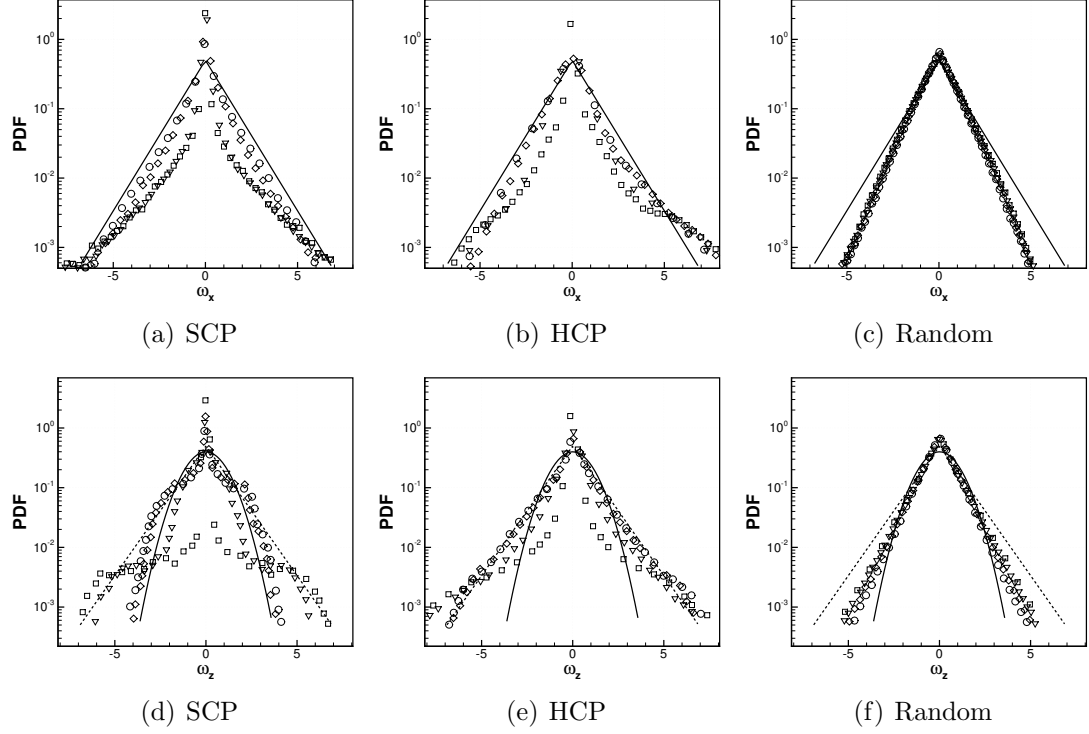
Figure 3.10: Normalized vorticity PDFs for steady flow rates through the three different packed beds. Top row (a-c) shows the cross stream, u_x component. Bottom row (d-f) shows the streamwise, u_z component. Different symbols correspond to pore Reynolds numbers of roughly $Re = 10$ □, $Re = 50$ ▽, $Re = 100$ ◇, $Re = 150$ ○

In Figure 3.10, similar PDFs are shown for the streamwise and cross stream vorticity components. Here, the normalized distributions of vorticity collapse with Reynolds number in the random arrangement but not the structured arrangements, especially the SCP arrangement. This represents a major difference between flow in the structured and random packings. It indicates the development of local regions with disproportionately strong ω_x and ω_z in the structured arrangements, while at the same time showing that the porescale vorticity components scale nearly the same everywhere in the random packing. This can be seen from the qualitative

visualizations presented earlier . In the arranged packings, the multi-lobed vortex ring structures appear in very specific locations in the structured packings and are the dominant vortical feature. On the other hand, in the random arrangement, there is still significant vortical motion, but it is distributed to different types of pores and strengthens more uniformly with Reynolds number.

Finally, Figure 3.11 shows the PDFs of two vortex detection criteria, $\cos\phi$ (top row) and λ_2 (bottom row). These PDFs demonstrate the difference in the type of vortical structures observed in the structured and arranged packings. The SCP arrangement shows that as Reynolds number is increased, the amount of helicity decreases. At the same time, the proportion of fluid with negative λ_2, indicating swirling motion, increases (Fig. 3.11d). This reinforces the fact that the vortical structures in the SCP arrangement are vortex rings with vorticity aligned perpendicular to the mean flow. The random PDFs show that the probability of strong helicity increases significantly with Reynolds number, as does the probability of negative λ_2. This confirms that the main type of vortical structure emerging at these Reynolds numbers in the random arrangement is the helical vortex. Similar trends are observed in the HCP arrangement, where there are strong helical vortices near the high porosity solid walls.

3.3.3   Unsteady Flow

The random and SCP arrangements were chosen for unsteady flow simulations because of their distinctly different geometries and vortical structures observed in the steady flow regime. The simple cubic and random meshes used for these unsteady flows are composed of 7 million and 32 million CVs respectively. As in the steady flow cases, the mean flow is provided by a constant inflow condition on the bottom boundary in the random packing, and a constant pressure gradient in the triple periodic SCP arrangement. The resulting, time averaged Reynolds numbers are $Re = 600$ and $Re = 529$ respectively.

At $Re = 529$, the flow in the SCP arrangement is dominated by spatially and temporally periodic porescale vortex shedding. The characteristics of the

(a) SCP   (b) HCP   (c) Random

(d) SCP   (e) HCP   (f) Random
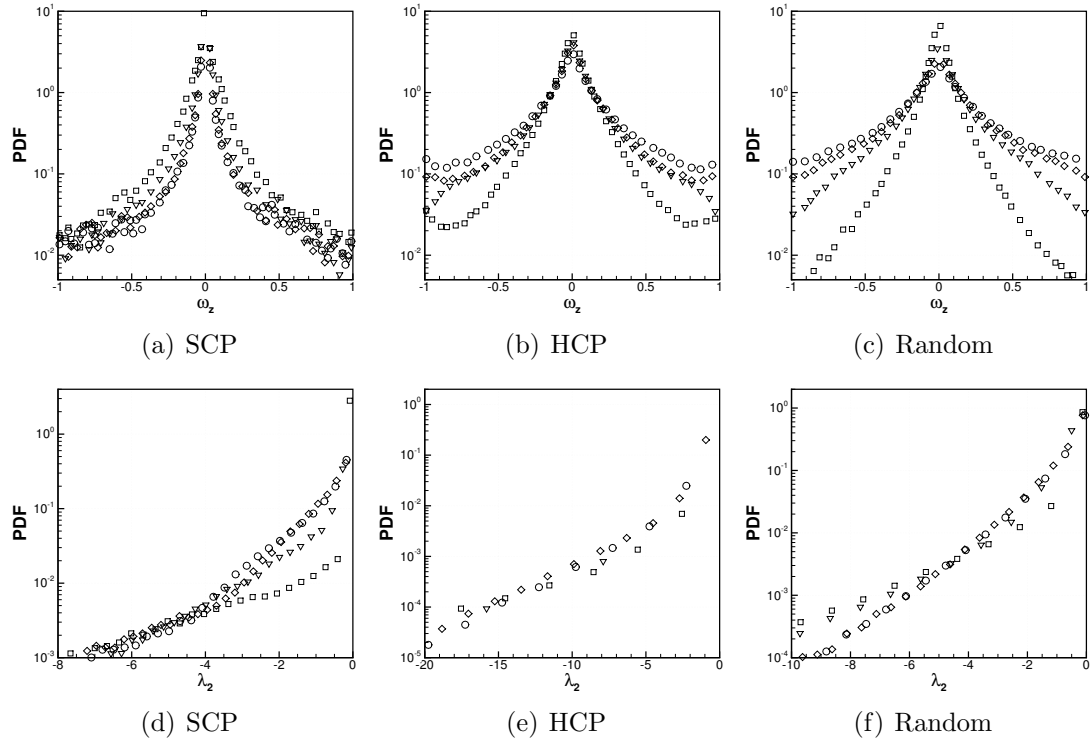
Figure 3.11: PDFs of vortex detection criteria, $\cos(\phi)$ (top row) and λ_2 (bottom row) for steady flow rates through the three different packed beds. Symbols correspond to pore Reynolds numbers of roughly $Re = 10$ $\square$, $Re = 50$ $\triangledown$, $Re = 100$ $\diamondsuit$, $Re = 150$ $\circ$
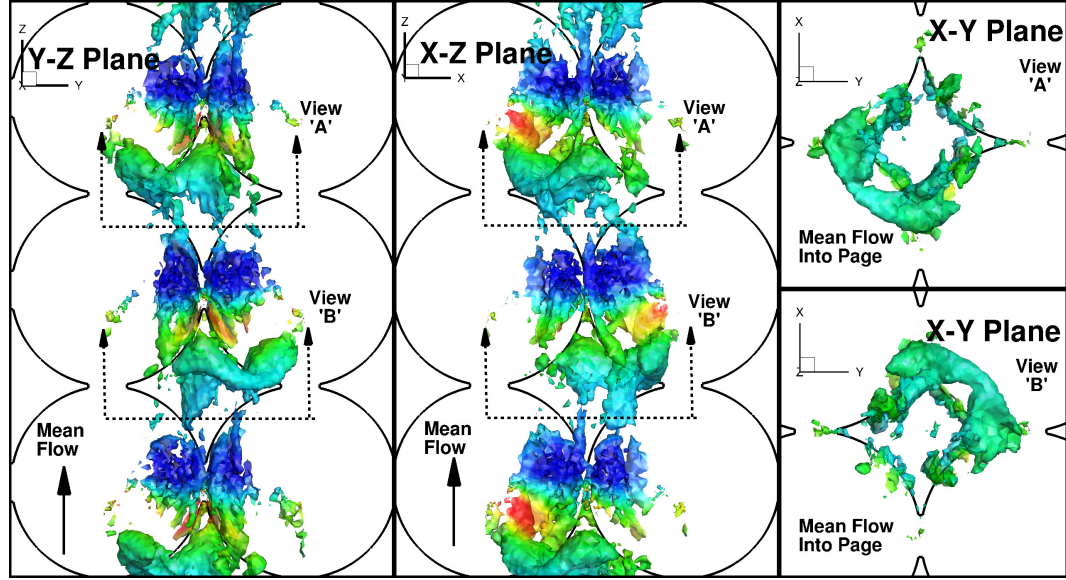
Figure 3.12: Unsteady porescale Kármán like vortex shedding observed in the simple cubic packing. isosurface of λ_2 shown, colored by pressure (blue indicates low pressure region).

three dimensional vortical structures are shown in FIG. 3.12, where isosurfaces of λ_2 colored by local pressure are shown for a single snapshot in time. Several orthogonal views of consecutive pores, with sphere boundaries denoted with solid black lines, are provided to give three dimensional context to the vortical structure relative to the sphere packing. The main vortical feature in *every pore* is a half vortex ring that is bent upward in the direction of the mean flow. These half rings have a larger lobe at each end. These lobes are similar to the lobes observed in the steady flows through the SCP arrangement, in that they are aligned with the major cross-stream axes of the sphere packing (X and Y in this case). The difference however is that only two lobes exist at any given time in a single pore; the vortical structure never forms a complete, four lobed vortex ring as was seen for steady flow rates. These structures alternate spatially along the mean flow direction (Z), such that the half ring is located in opposite sides of consecutive pores. This is best seen in the X-Y views on the right hand side of FIG. 3.12. In the cross stream

directions, these vortical features are *in phase* with one another. That is, across pores in the same X-Y plane, they will be located in the same quadrant of every pore (not shown).

The transient character of these structures strongly resembles the periodic Kármán vortex shedding observed in the wake of a single cylinder or sphere. The half vortex ring structures described above alternate between the $(+X, +Y)$ corner of the pore shown in View 'B' of FIG. 3.12 and the $(-X, -Y)$ corner shown in View 'A'. This 180^o phase shift occurs in every pore in the packing. In the middle of this cycle, the half ring structures are flushed into the high velocity channels by the mean flow, and are stretched into thinner tail-like regions. As a new half ring forms, the old structure is being flushed through the center of channel. These tail regions can be seen in the pore throats in the snapshot of FIG. 3.12.

Despite similar Reynolds numbers, the flow through the random packing contains remarkably different vortical structures. Due the the variety of geometric pore configurations that exist in this packing, the local value of Re can vary considerably from the bed average value of $Re = 600$. This results in not just one, but a variety of coherent porescale vortical structures. We will concentrate on the two main types of feature types seen, which are shown in FIG. 3.13. To illustrate these features, instantaneous stream-ribbons colored by the local pressure are shown in three different pores. The most common feature is the helical vortex, which was also seen at reduced flow rates in the steady flow simulations. In fact, the same pore is shown in FIG. 3.13a, as in FIG. 3.6g-k, and a helical vortex can be seen in a similar location as the one that develops over the range $50 < Re < 150$. At this higher, unsteady flow rate, the streamlines of the feature have much stronger radial motion as they spiral through the pore compared to the steady flow feature, and the length of the vortex is stretched further in the streamwise direction. These elongated helical vortices persist throughout the packing, often with multiple vortices in a single pore. Another of these features is shown in FIG. 3.13b, located near the solid outer tube wall. This feature results from the high momentum streamlines in the high void fraction wall region (see FIG. 3.4) mixing as

(a) (b)
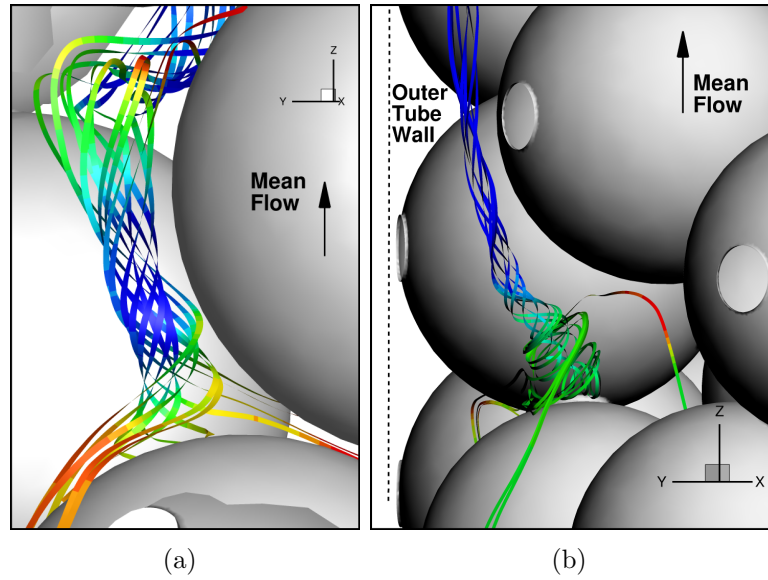
Internal and external helical vortices
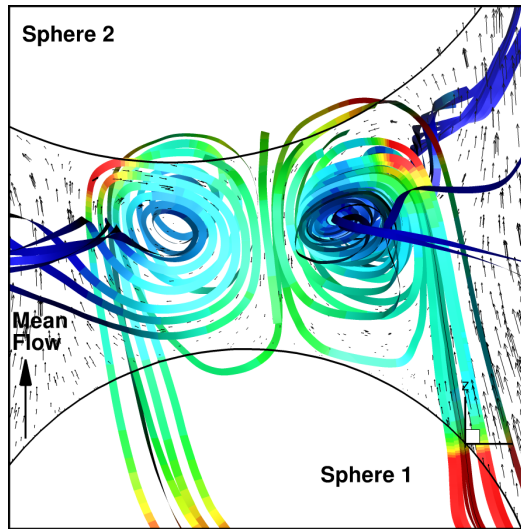
(c) Counter rotating vortex tubes

Figure 3.13: Unsteady porescale vortical features observed in the random packing. Stream ribbons colored by pressure are shown in three different pores (blue indicates low pressure region)

they enter the pore from below, and swirling around each other during their exit upward. Because of the locally higher Re at the wall, and the other features that exist in this pore, this vortical feature is more *active* in it's unsteady motions than the internal feature of FIG. 3.13a, although it remains in this pore at all times. Future work will focus on quantifying the relationship between the vortical length and time scales and Re over the entire steady and unsteady flow regime.

The other main type of feature that is observed in the random packing at $Re = 529$ is a vortex tube whose axis is roughly aligned perpendicular to the flow. A counter-rotating pair of these features are shown in FIG. 3.13c. The instantaneous flow vectors in a perpendicular plane are shown in the background. Such features tend to exist in the near wake behind a sphere, or in the gaps between two spheres as shown. Despite the slow turnover time of these eddies exhibited by the vector magnitudes in the circulating regions, they are not complete dead zones. They expand and contract in tandem, meaning one grows while the other shrinks in a periodic fashion, exchanging fluid amongst themselves and with the mean flow.

None of the vortical features observed in the unsteady flow through the random packing exhibit any vortex shedding and there does not appear to be vortex destruction or creation at this Reynolds number. Rather, they tend to oscillate in their position without much net change of shape. The helical vortex cores translate periodically, mostly in the in the cross stream plane, but their size and mean location remains nearly constant for the length of the simulation.

In order to quantify the timescales of these unsteady flows, the fluctuating velocity field, $\mathbf{u}'(x,t)$, is computed via a Reynolds decomposition for a time series of instantaneous velocity fields, $\mathbf{u}(\mathbf{x},t)$.

$$\mathbf{u}(\mathbf{x},t)' = \mathbf{u}(\mathbf{x},t) - \overline{\mathbf{u}}(\mathbf{x}) \tag{3.5}$$

Here, $\overline{\mathbf{u}}$ is the time averaged velocity field. With the fluctuating velocity field, the time autocorrelation function is computed for the porespace region of the SCP and

(a) SCP   (b) Random
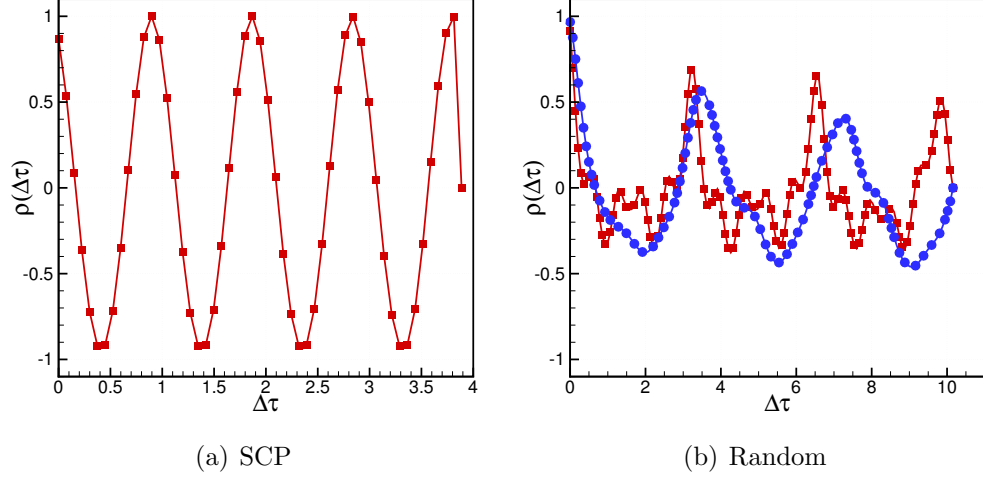
Figure 3.14: Time autocorrelation functions of streamwise velocity fluctuations, ρ_{xx}. ($-\blacksquare-$) scp & random wall region (outermost $0.5d_{sp}$), ($\bullet$) random inner core (innermost $0.5d_{sp}$).

random arrangements:

$$\rho_{ii}(\Delta\tau) \;=\; \sum_{\mathbf{x}}\sum_{\tau} \frac{u_i'(\mathbf{x},\tau)\,u_i'(\mathbf{x},\tau+\Delta\tau)}{u_i'(\mathbf{x},\tau)^2} \qquad (3.6)$$

$$(3.7)$$

Here, $\Delta\tau = \frac{tU}{D_{sp}}$ is the non-dimensional time separation, between two velocity samples at position $\mathbf{x}$ in the porespace. The streamwise autocorrelation function, ρ_{xx} is plotted for both arrangements in FIG. 3.14. The cross-stream correlation functions, ρ_{yy} and ρ_{zz}, show similar trends in both arrangements and are omitted for clarity. The SCP time correlation function is periodic with a single dominant frequency. Recently, Horton & Pokrajac [51] have experimentally investigated the flow structure in SCP arrangements of spheres driven by constant pressure gradients with ultrasonic velocity profilers (UVP), and observed similar periodic correlation functions at transitional Reynolds numbers. These resolved simulation results, which show the shedding and motion of coherent vortical structures at the

porescale, provide corroborating evidence that these periodic velocity fluctuations are linked to porescale vortical motion and vortex shedding .

Because of the radial porosity variations present in the random packing (FIG. 3.4), and the locally higher Re near the wall, the correlation function is computed *separately* for different annular regions of the flow to show the effect of the wall on the unsteady timescales. In FIG. 3.14b we have plotted the function for the region within $0.5D_{sp}$ of the tube wall and the region within $0.5D_{sp}$ of the tube centerline. These two correlations both show time periodic behavior, and the existence of several modes of velocity fluctuations in the porespace. The higher Reynolds number wall region shows exhibits a more complex function with more modes than does the inner region. Both correlation functions decay with long time separation, similar to the results of [36] for converging/diverging channel flows, but shows a strongly periodic behavior for the short time separations investigated here. This shows that despite the randomness of the bed, there are very specific dominant modes of velocity fluctuation. The increased number of modes compared to the SCP arrangement signals that the flow through the random packing may be closer to a turbulent chaotic breakdown.

By performing a Fourier transform on the autocorrelation functions, the velocity frequency spectra can be obtained. This is shown in FIG. 3.15 for both sphere arrangements, where the frequency is non-dimensionality to give the Strouhal number, $St = (\Delta\tau)^{-1}$. The SCP arrangement contains a dominant Strouhal number of nearly unity, showing the vortex shedding is strongly linked to the length scale of the sphere packing. The random packing contains several different Strouhal numbers between 0.29 and 1.57. In the inner region, the high frequency modes ($St > 1$ have been eliminated. This shows that the larger Reynolds number wall region contains a larger spectrum of energetic time (and most likely spatial) scales,
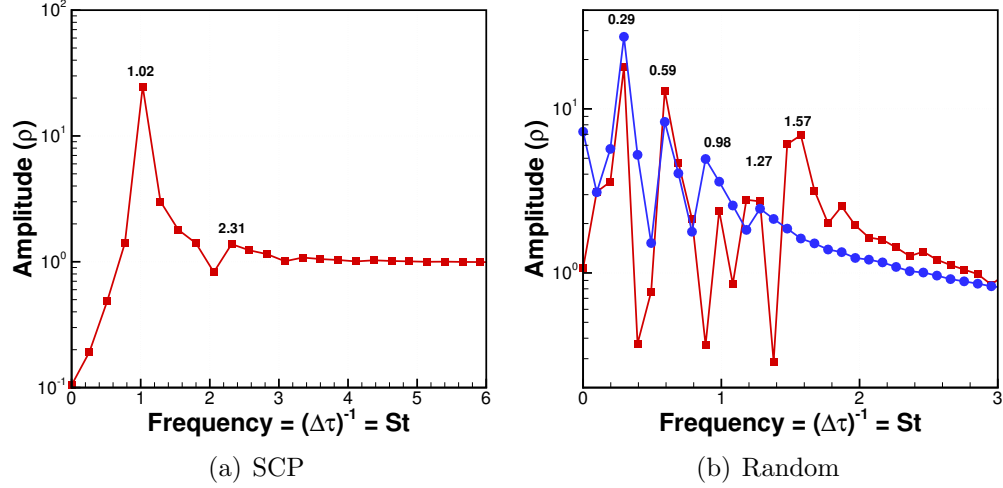
Figure 3.15: Velocity frequency spectrum obtained from the Fourier transform of ρ_{xx}. ($-\blacksquare-$) scp & random wall region (innermost $0.5d_{sp}$), ($\bullet$) random inner core (innermost $0.5d_{sp}$)

3.4   Conclusions

Resolved and time accurate simulations of flows through three different types of packed bed geometries have been performed to identify and characterize pore scale vortical structures. These simulations considered the steady and unsteady inertial regimes ($10 < 600$), where the effect of porescale flow features on macroscale properties is not well understood. Vortical features are detected by regions of swirl and pressure minima using the λ_2 criteria. The results show that there is a clear relationship between pore geometry and the types of pore-scale vortical structures observed. In the SCP and HCP arrangements, a multi-lobed vortex ring structure is present in the gaps between spheres at steady flow rates. The lobes of these structures align with the principal geometric axes of the packing, and are highly three dimensional. At steady flow rates through the random packing, and in the wall region of the HCP arrangement, elongated helical vortices are the dominant porescale flow feature. Such features appear to be located near accelerating streamlines in regions of high curvature and low geometric symmetry.

The emergence of these different types of vortical features are confirmed through statistical analysis of velocity, vorticity and helicity in the porespace.

At higher flow rates near $Re = 600$, the unsteady dynamics of flow through the SCP and random arrangements are significantly different. In the simple cubic arrangement, spatially and temporally periodic Kármán like vortex shedding is observed at a single Strouhal number near 1. Flow through the random arrangement exhibits a wider range of time scales with Strouhal numbers between 0.29 and 1.57 but without any observed vortex shedding. The higher frequency modes seem to be active mostly in the near wall region, where the Reynolds number is locally higher due to the wall effect. At this Reynolds number, helical vortices are still the dominant porescale feature, along with vortex tubes aligned perpendicular to the mean flow.

Future work on this topic will focus on quantifying the length and time scales of the porescale vortical features identified here, and linking their characteristics to macroscale properties. Also, direct comparison with experimental measurements in similar packed beds will be used as more detailed validation of the computational methods.

3.5   Acknowledgments

Chapter 4: Integrated Computation of Finite-Time Lyapunov Exponent Fields During Direct Numerical Simulation of Unsteady Flows

Justin R. Finn and Sourabh V. Apte

Abstract

The computation of Lagrangian coherent structures typically involves post-processing of experimentally or numerically obtained fluid velocity fields to obtain the largest finite-time Lyapunov exponent (FTLE) field. However, this procedure can be tedious for large-scale complex flows of general interest. In this work, an alternative approach involving computation of the FTLE on-the-fly during direct numerical simulation of the full three dimensional Navier-Stokes equations is developed. The implementation relies on Lagrangian particle tracking to compose forward time flow maps, and an Eulerian treatment of the backward time flow map [S. Leung, Journal of Computational Physics 230, 2011] coupled with a semi-Lagrangian advection scheme. The flow maps are accurately constructed from a sequence of smaller sub-steps stored on disk [S. Brunton and C. Rowley, Chaos 20, 2010], resulting in low CPU and memory requirements to compute evolving FTLE fields. Several examples are presented to demonstrate the capability and parallel scalability of the approach for a variety of two and three dimensional flows.

The notion of Lagrangian coherent structures (LCS) as invariant transport barriers in steady and unsteady fluid flows has evolved from dynamical systems theory, and has proven its utility in understanding a number of mixing and transport problems. These structures are most often defined from ridges in the finite-time Lyapunov exponent field. However, their practical use has been limited to this point because computing the FTLE field typically involves expensive post-processing of large fluid velocity datasets generated either from experiments or numerical simulations. Using some recently developed tools [61, 9], we have designed and implemented an integrated approach to compute evolving, transient, three dimensional FTLE fields during a CFD simulation. By integrating the computations in this way, tedious post-processing of velocity fields is no longer needed, and larger, more complex problems become accessible to the application of LCS theory.

4.1   Introduction

That fluid flows are organized by an underlying structure is not a new idea; da Vinci described the similarity of hair-like curls (eddies) generated by a narrow jet issuing into a pool as early as the year 1500, and much of turbulence theory has evolved from notions of organized scales of motion. Modern advances in experimental techniques and the advent of direct numerical simulation (DNS) of the Navier-Stokes equations have provided databases with incredible microscopic detail for a variety of natural and engineered flows. However, a structural description of the coherent motions exhibited by unsteady flows remains elusive in most cases. As available databases continue to grow rapidly in size and complexity, a number of tools have been developed to help extract coherence from chaotic fluid motions. For example, variants of the Eulerian velocity gradient tensor have aided in the detection of instantaneous vortical regions [14], and proper orthogonal decomposition has been helpful in extracting energetic modes of complex turbulent flows [6]. More recently, a class of methods to detect *Lagrangian coherent structures* (LCS) has emerged from dynamical systems theory and the works of Haller [42, 38].

LCS are the codimension one manifolds[1] that form the skeleton of tracer trajectories [68] and separate time dependent flows into regions of dynamically distinct behavior [98]. When properly defined, they act as the locally most attracting or repelling material surfaces [39, 27] and therefore have important consequences for mixing and transport. In addition, as boundaries of dynamically distinct regions they provide an exceptional tool with which to visualize and understand coherent fluid motion. The identification of LCS in experimentally measured or numerically simulated flows has most often relied on the relation of these special material surfaces to ridges of the finite-time Lyapunov exponent (FTLE) field. Ridges in the forward time FTLE field can be thought of as candidates for *repelling* LCS about which there is large fluid stretching, while ridges in the backward time FTLE field are candidates for *attracting* LCS about which folding occurs [42]. An im-

---

[1]This implies that LCS are curves in two dimensional flows and surfaces in three dimensional flows.

pressively broad range of flows have now been studied using the FTLE ridges as a proxy for LCS (see [83, 94] for recent reviews). For example, new insight has been provided in areas such as fundamental vortex dynamics [76, 97], aerodynamics [33], biological feeding [84], ocean and atmospheric transport [7], and granular flows [15]. Significant work has been undertaken to explore the link between LCS and ridges of the FTLE field from a theoretical standpoint [38, 98, 60, 39, 27, 40]. Unfortunately, it appears that there is not necessarily a one to one mapping between FTLE ridges and LCS, and both false LCS positives and negatives can be obtained without sufficient additional restrictions [39, 27] on the ridges. [42] provided the counter example of two-dimensional shear flow, $u = -x, \ v = y$, where the line $x = 0$ is a positive time FTLE ridge, but is not an unstable material surface. Recently, the problem of defining LCS has been revisited, and a variational theory [39, 27] and geodesic theory [40] have been developed that more rigorously distinguish Lagrangian transport barriers. Practical implementation of these more general LCS theories, as done by [26] and [40], is presently limited to two dimensional flows, but algorithmically the preliminary steps are similar to a standard FTLE computation (requiring the computation of the Cauchy-Green deformation tensor). In this sense computation of the FTLE, or similar field, seems to be an important step to determining *LCS candidates* in general time dependent flows.

Typically, FTLE fields are computed in a post-processing procedure using either experimentally or numerically obtained velocity fields. First a flow map over the finite-time interval $[t_0, \ t_1]$ is computed by numerical integration of a grid of Lagrangian particles through snapshots of the time dependent velocity field. The flow map Jacobian is then obtained by finite differencing on the initial tracer grid, and used to build the right Cauchy-Green deformation tensor. Finally, the maximum eigenvalue of this tensor is used to compute the maximum FTLE field, which corresponds to the rate of Lagrangian separation of initially adjacent trajectories. As just described, this procedure can be resource intensive for a number of reasons. Most significantly, to obtain sharp, well defined ridges in the FTLE field, the initial grid of tracers must be refined relative to the smallest spatial scales of fluid motion,

and the integration time must be long relative to the dominant time scales of the flow. For non-trivial flow fields, especially in three dimensions, this can result in an enormous number of tracer advections. Additionally, for every time the FTLE is needed, for example to visualize its evolution, the tracer advections must be re-computed in forward and backward time. Finally, for acceptable accuracy, the spatial and temporal resolution of the flow field snapshots must be fine enough for accurate interpolation of the velocity field to the Lagrangian tracers. For large datasets, this implies careful memory management during post-processing.

A number of strategies have been employed to address these difficulties. [62] noted that the LCS are themselves Lagrangian objects that are advected with the flow and devised an efficient ridge tracking algorithm so that the FTLE field need only be computed in the vicinity of the LCS. A few authors have enlisted unstructured meshes and automated mesh refinement (AMR) schemes to enhance the resolution of FTLE field near the LCS [59, 70] while reducing the number of tracer advections far from the LCS. In a similar spirit, [26] used a staggered auxiliary grid of tracers to increase the accuracy of the flow map gradient while maintaining a relatively coarse tracer background grid. Recent developments in computer hardware have also been leveraged to speedup the computation. Due to their streaming capabilities, graphical processing units (GPUs) are naturally suited for the problem of particle advection that is the bottleneck of the FTLE computation and impressive speedups have been obtained relative to standard CPU computations [31, 16]. From a theoretical perspective, a recent observation that the largest FTLE in forward time is related to the smallest FTLE in backward time (and vice-versa) [41] could allow for the computation of both attracting and repelling LCS candidates with only a single set of tracer integrations. However, practical implementation issues arise for long integration times and three dimensional flows due to the need to interpolate from a highly deformed grid of particles. Moving forward, it is important to continue to develop efficient computational strategies so that LCS theory can be applied to more complex problems of general interest.

In this paper, we describe an integrated procedure for computing both forward

and backward time FTLE fields *on-the-fly* during DNS of the Navier-Stokes equations. Two recent observations have provided the necessary building blocks for our current implementation. The first, by [61], is that the backward time flow map may be treated as a collection of Eulerian scalar fields that can be evolved forward in time by solving three level set equations on a fixed grid. This provides a natural way to compute the backward time flow map during a simulation that evolves only in forward time. The second important observation, made by [9], exploits the property [77] that a time T flow map, where $T = |t_1 - t_0|$, may be constructed from a sequence of N smaller time h flow maps, where $h = T/N$. This removes the need to do redundant tracer integrations when many concurrently evolving FTLE fields must be computed, for example to animate their evolution, meaning the simulation only needs the resources to evolve one forward and one backward flow map at a time.

Integrating the FTLE computations into a CFD simulation has several potential advantages over the post-processing approach. Perhaps the most compelling is that the full temporal and spatial resolution of the simulation become available when computing the flow maps. This is not generally true in the post-processing approach, where the temporal and spatial resolution of the velocity field, and by consequence the flow map uncertainty, could be limited by available hard disk or memory resources. This makes applying LCS theory to study flows with broad length and timescale separations, such as multiphase and turbulent flows, as well as flow in complex geometries, such as packed beds and porous media, much more feasible. Second, by performing the computations during the simulation, we can harness the parallelism of the simulation code as well as tools already built into many flow solvers such as accurate gradient calculations, interpolations, particle tracking schemes, and passive scalar solvers for the advection-diffusion equation. Finally, integration of the FTLE computations directly into a CFD simulation framework makes it possible to extend LCS theory to study a number of new applications. For example, new explorations of active feedback control could be pursued based on characteristics of the LCS, or measurement of Lagrangian length

scales could be made using the FTLE field in evolving three dimensional turbulent flows. Some of these benefits have been suggested previously (for example by [61]), but to our knowledge, this work represents the first fully integrated computation.

Our main objective is to describe a relatively simple algorithm that can be used to compute evolving forward and backward time FTLE fields during a computational fluid dynamics (CFD) simulation, with minimal additional overhead. The remainder of the paper is organized as follows. We first give a brief mathematical background for the computation of FTLE fields in Section 4.2. Next, we present numerical details related to the integration of the computations in a direct numerical simulation framework in Section 4.3.3. In Section 4.4, we demonstrate the performance of the implementation and discuss the computational overhead for several two and three dimensional flows. Finally, in Section 4.5 we make some general conclusions and proposals for future work.

4.2 The Finite Time Lyapunov Exponent

We now briefly develop the theory and notation related to the FTLE calculation. More thorough expositions describing the FTLE and its relation to LCS are available [98, 60]. Assume that some velocity field, $\mathbf{u}(\mathbf{x}, t)$, is defined on the three dimensional domain, $\mathbf{x} \subseteq \mathbb{R}^3$, over the time interval, $t \in (t_0, t_1)$. The flow map, $\mathbf{\Phi}_{t_0}^{t_1}(\mathbf{x}_0, t_0)$ integrates passive tracers from their initial position, $\mathbf{x}_0$, at time t_0 along pathlines to their "advected" position, $\mathbf{x}$, at time t_1,

$$\mathbf{\Phi}_{t_0}^{t_1}(\mathbf{x}_0, t_0) = \mathbf{x}_0 + \int_{t_0}^{t_1} \mathbf{u}(\mathbf{x}(\tau), \tau) d\tau. \tag{4.1}$$

The flow map may be computed in *forward time* $(t_1 > t_0)$, or *backward time* $(t_1 < t_0)$. Two tracers initially displaced by a small perturbation, $\delta\mathbf{x}_0$, will find themselves separated by distance $\delta\mathbf{x}$ at time t_1. To a leading order, this separation

can be written in terms of the flow map as [77]

$$\delta\mathbf{x}(t_1) = \mathbf{\Phi}_{t_0}^{t_1}(\mathbf{x}_0 + \delta\mathbf{x}_0, t_0) - \mathbf{\Phi}_{t_0}^{t_1}(\mathbf{x}_0, t_0) \tag{4.2}$$

$$= \mathbf{D\Phi}_{t_0}^{t_1}(\mathbf{x}_0, t_0)\delta\mathbf{x}_0 + \text{H.O.T..} \tag{4.3}$$

Where $\mathbf{D\Phi}_{t_0}^{t_1}(\mathbf{x}_0, t_0)$ is the Jacobian of the flow map evaluated at the initial tracer coordinate, $\mathbf{x}_0$. The magnitude of the separation, $||\delta\mathbf{x}(t_1)||$, is found from the L^2 matrix norm,

$$||\delta\mathbf{x}(t_1)|| = \sqrt{\langle \delta\mathbf{x}_0, \left[\mathbf{D\Phi}_{t_0}^{t_1}(\mathbf{x}_0, t_0)\right]^* \left[\mathbf{D\Phi}_{t_0}^{t_1}(\mathbf{x}_0, t_0)\right] \delta\mathbf{x}_0 \rangle}, \tag{4.4}$$

where, $\langle \cdot, \cdot \rangle$ denotes the Euclidean inner product, and * denotes transposition. We introduce the right Cauchy-Green deformation tensor,

$$\mathbf{C}_{t_0}^{t_1}(\mathbf{x}_0, t_0) = \left[\mathbf{D\Phi}_{t_0}^{t_1}(\mathbf{x}_0, t_0)\right]^* \left[\mathbf{D\Phi}_{t_0}^{t_1}(\mathbf{x}_0, t_0)\right]. \tag{4.5}$$

From equation 4.4, the maximum stretching over the interval (t_0, t_1) occurs when $\delta\mathbf{x}_0$ aligns with the eigenvector associated with the maximum eigenvalue, λ_{max}, of $\mathbf{C}_{t_0}^{t_1}(\mathbf{x}_0, t_0)$. This leads to the definition of the finite-time Lyapunov exponent, which characterizes the maximal rate of stretching over the finite-time interval (t_0, t_1).

$$\sigma_{t_0}^{t_1}(\mathbf{x}_0, t_0) = \frac{1}{|t_1 - t_0|} \log \sqrt{\lambda_{max}(\mathbf{C}_{t_0}^{t_1}(\mathbf{x}_0, t_0))} \tag{4.6}$$

For simplicity, the standard convention is to notationally drop the dependence of the flow map, deformation tensor, and FTLE fields on $\mathbf{x}_0$ and t_0, and to introduce the integration time, $T = |t_1 - t_0|$. This way for example, the forward and backward time T flow maps can be written in shorthand as $\mathbf{\Phi}_{t_0}^{t_0+T}$, and $\mathbf{\Phi}_{t_0}^{t_0-T}$ respectively.

Under sufficient restrictions [39, 27], ridges in the FTLE field can mark attracting ($t_1 < t_0$) or repelling ($t_1 > t_0$) LCS. However, as mentioned earlier, it is possible to identify false positive and negative LCS using primitive ridge definitions alone.

At the present, we choose to use FTLE ridges as indicators of *LCS candidates*, knowing that more rigorous analysis is needed to truly confirm them as having all the properties of LCS.

4.3 Numerical implementation

In this section, we discuss the tools needed to integrate the FTLE computations within the framework of a CFD solver. We will discuss the details relevant to our present implementation, however, it is intended that the general algorithm that is presented could be followed by anyone wishing to add a similar capability to their own CFD code. The present solver was originally developed to perform large-eddy and direct numerical simulations of flows in complex geometries on potentially very large grids [71]. The code solves the Navier-Stokes equations on arbitrary shaped, unstructured grids using a co-located finite volume discretization, a fractional step method for time advancement, and an algebraic multigrid (AMG) solver for the pressure Poisson equation [25]. The code is parallelized using Message Passing Interface (MPI), allowing for larger scale simulations by distributing the required memory over many processors. The flow solver has a number of extended capabilities particularly relevant to multiphase flow simulations including Euler-Lagrange models for a subgrid scale dispersed phase [99], and a fictitious domain approach for fully resolved simulation of immersed rigid bodies [2].

4.3.1 Lagrangian and Eulerian flow map computations

In the post-processing approach to computing the FTLE fields, the forward and backward time flow maps are typically composed by seeding the flow with Lagrangian tracers at time $t = t_0$ and advecting them in positive and negative time over the intervals $[t_0 + T]$ or $[t_0 - T]$ respectively. This is the methodology that we employ for the calculation of the forward time flow maps. The three important steps of this process are illustrated schematically in Figure 4.1a. First, at time t_0, we *launch* a Lagrangian tracer from each cell, by setting the initial tracer

coordinates, $\mathbf{x}_0 = \mathbf{x}_{cv}$, where the subscript cv denotes the cell center coordinates of each control volume. The tracers are then advected passively with the flow by integrating the equation,

$$\frac{d\mathbf{x}}{dt} = \mathbf{u}(\mathbf{x}, t) \tag{4.7}$$

To do this, the position of the Lagrangian tracers is updated after each time step of the flow solver using a simple, explicit, trapezoidal approximation of equation 4.7.

$$\mathbf{x}^{n+1} = \mathbf{x}^n + \mathbf{u}(\mathbf{x}^n, t^{n+1/2})\Delta t \tag{4.8}$$

Where the time step, Δt, is the time step used by the flow solver. The fluid velocity used for the tracer position update, $\mathbf{u}^{n+1/2}$, is staggered in time with respect to the fluid velocity known on the fixed grid. Since we perform the tracer update after the fluid solve, this is approximated as simply, $\mathbf{u}^{n+1/2} = 0.5\,(\mathbf{u}^n + \mathbf{u}^{n+1})$. Higher order tracer integrations such as RK4 are possible. However, since our time step is small to assure good accuracy of the flow solver, we have found that time advancement according to equation 4.8 performs well and is computationally efficient. During the integration, the tracers will cross over cell and processor boundaries. To handle this in our unstructured solver we use a particle in cell approach and known vicinity search algorithms [99] to efficiently locate the tracers as they evolve with the flow field. At time $t_0 + T$, all processors gather the Lagrangian tracers currently located on their own grid partition. Each tracer carries two integer tags corresponding to the processor and cv index that it was launched from. Each processor sorts the tracers currently belonging to it by these tags and communicates their current position, equal to the forward time flow map $\mathbf{\Phi}_{t_0}^{t_0+T}$, back to its launch processor and cv.

During the simulation, if a tracer crosses over a solid boundary at any point during the integration, it is relocated to the cell center of the last control volume that contained it. In practice, this is a rare occurrence because the no-slip condition at solid boundaries and low simulation timesteps prevent these types of erroneous tracer trajectories. If a tracer crosses over an outlet boundary, its

position is fixed at the boundary for the remainder of the simulation. This condition is somewhat unphysical, but is not an issue if the simulation outlet is placed significantly downstream of the region of FTLE interest.

Unfortunately, there is not a natural way to compose the backward time flow map during the CFD simulation with the Lagrangian approach as just described. Doing so would involve saving several snapshots of the velocity field between t_0 and $t_0 + T$, either in memory or hard disk, then subsequently using them to integrate tracers backward in time once the simulation reached $t_0 + T$. This is both cumbersome and resource intensive. A clever alternative, proposed by [61], which integrates almost seamlessly with a forward evolving CFD simulation is to treat each component (x, y, and z) of the backward time flow map as a scalar field. These scalar fields represent the *"takeoff coordinates"* of tracers located on the fixed Eulerian grid, and are evolved forward in time by solving three level set equations at each timestep:

$$\frac{\partial \Phi_{t_0,i}^{t_0-T}}{\partial t} + (\mathbf{u} \cdot \nabla)\Phi_{t_0,i}^{t_0-T} = 0 \tag{4.9}$$

This equation states that the i^{th} component of $\boldsymbol{\Phi}_{t_0}^{t_0-T}$ remains constant along tracer trajectories, just as it would in the Lagrangian approach if we integrated backward in time. To embed the *takeoff coordinates* at time $t_0 - T$ in the Eulerian representation, we simply initialize the backward flow map at the cell centers to the grid coordinates, $\boldsymbol{\Phi}_{t_0}^{t_0-T}(\mathbf{x}_0, t_0 - T) = \mathbf{x}_{cv}$. The three scalars are then evolved according to equation 4.9 up to time t_0, at which point their value on the fixed grid is equivalent to the i^{th} component of the backward flow map for integration time T. This process is shown for the y component of the backward flow map being evolved by a simple flow field in Figure 4.1b.

Equation 4.9 can be solved in a number of ways, and [61] chose a higher order weighted essentially non-oscillatory (WENO) scheme. In our present implementation, we have chosen a semi-Lagrangian approach. Compared to many Eulerian solvers, these schemes are simple, have relatively low overhead, and are easy to

(a) Forward time
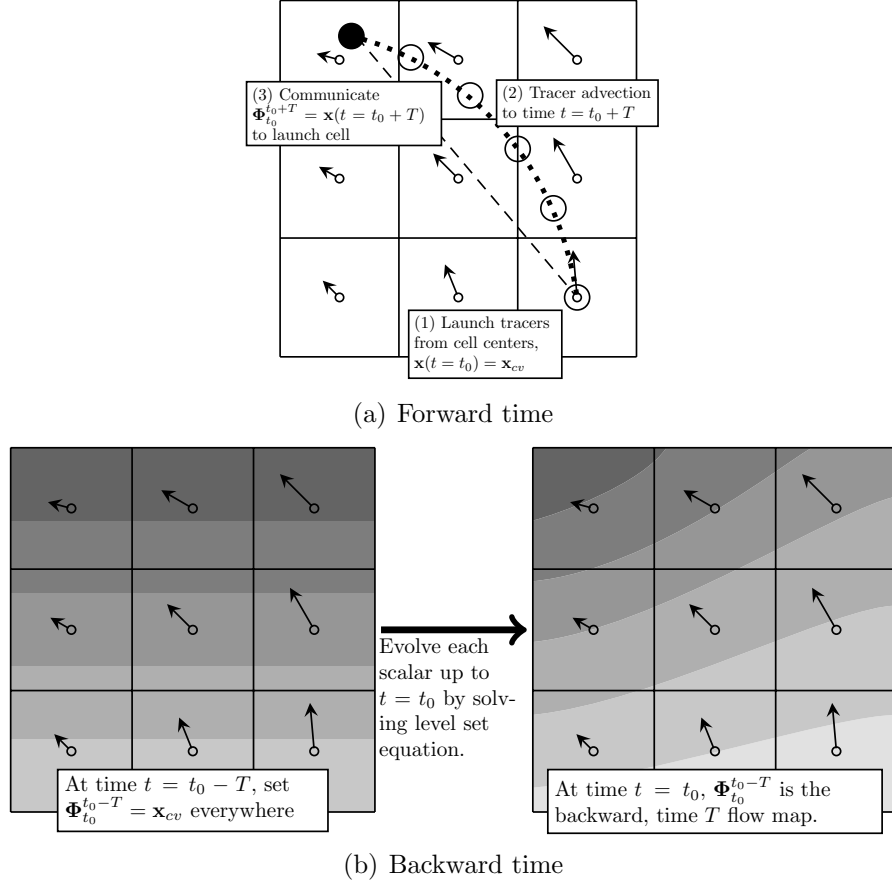
(b) Backward time

Figure 4.1: Lagrangian and Eulerian computations of the flow maps for time $t = t_0$. (a) The Lagrangian computation of the forward time flow map, $\mathbf{\Phi}_{t_0}^{t_0+T}$. (b) The Eulerian computation of the backward time flow map, $\mathbf{\Phi}_{t_0}^{t_0-T}$. The scalar field shows the evolution of the y component of the backward time flow map in a simple velocity field.

implement. They have gained popularity in atmospheric and multiphase flow simulations for their efficiency and stability with large timesteps [107]. The idea of semi-Lagrangian advection is to approximately integrate equation 4.9 along tracer trajectories:

$$\frac{\Phi_{t_0,i}^{t_0-T}(\mathbf{x}, t^{n+1}) - \Phi_{t_0,i}^{t_0-T}(\mathbf{x} - 2\mathbf{m}, t^n)}{\Delta t} = 0 \tag{4.10}$$

Here, the vector $\mathbf{m}$ is the distance traveled by a tracer in the time $\Delta t/2$, ie. the midpoint of the trajectory between its takeoff point at t^n and the fixed grid point $\mathbf{x}$. In order to solve this equation, we use the second order, two time level scheme described by [105].

1. The vector $\mathbf{m}$ may be approximated to the second order through the implicit relation, $\mathbf{m} = \mathbf{u}(\mathbf{x} - \mathbf{m}, t^{n+1/2})\frac{\Delta t}{2}$, which can be solved iteratively, provided some initial guess, $\mathbf{m}_0$

$$\mathbf{m}_{k+1} = \mathbf{u}(\mathbf{x} - \mathbf{m}_k, t^{n+1/2})\frac{\Delta t}{2}, \tag{4.11}$$

   where subscript k denotes the iteration index. On the first time step, we initialize $\mathbf{m}_0$ to zero everywhere. On subsequent timesteps, we simply use the $\mathbf{m}$ from the previous time step. Because major temporal changes to the flow field occur on much longer timescales than the simulation time step, we observe that typically very few iterations are required after the initial time step.

2. Evaluate $\Phi_{t_0,i}^{t_0-T}(\mathbf{x} - 2\mathbf{m}, t^n)$. This value is then projected forward in time along the trajectory according to equation 4.10.

At the boundaries of the simulation domain, we enforce the boundary conditions for $\mathbf{\Phi}_{t_0}^{t_0-T}$ suggested by [61]. At inflow, and no-slip boundaries, we enforce $\mathbf{\Phi}_{t_0}^{t_0-T} = \mathbf{x}|_{boundary}$. At all other boundaries, we enforce $\mathbf{n} \cdot \nabla \Phi_{t_0,i}^{t_0-T} = 0$, where $\mathbf{n}$ is the outward normal vector associated with the boundary.

Our use of a semi-Lagrangian scheme to solve the level set equation is somewhat different from most multiphase flow problems where tracking of an interface

is required. In such applications, particle level set methods [46] that employ tracer particles that carry a marker function, can be used to accurately capture the evolution of a deforming interface. As the interface deforms, so does the distribution of its marker particles, and in practice the tracer grid must be *re-meshed* or *reinitialized* to obtain a high quality representation of the evolving interface. Similar issues arise in our problem, where the goal is to track all values of the scalar fields (backward flow map), everywhere on our grid. As the scalar fields evolve from their initial states (eg. Figure 4.1b), stretching and folding over long times will result in large gradients that could lead to loss of accuracy or instability. One way to handle this would be to employ the forward time Lagrangian tracers in a way similar to the particle level set approach; Every few timesteps, the tracer *takeoff coordinates* could be interpolated fixed grid to correct the backward flow map representation. However, in many of our cases the initial tracers distribution is non-uniform (because it is associated with the fixed grid, see for example section 4.4.2), meaning such an interpolation could be of dubious quality. The approach we pursue instead, is to only compute the flow maps for short integration times over which the scalar field does not distort significantly. The long integration time flow maps are then composed approximately using the method of [9], explained in more detail in section 4.3.2.

The discrete equations for evolving both the forward and backward time flow maps involve variables that may be located at non-mesh locations in space. In general, this requires an interpolation of some variable, ϕ, known at the cell centers of the fixed grid to these points. In an orthogonal, structured computational grid, some form of trilinear interpolation [58, 69] can be applied to interpolate the values to the non-mesh points. Since our flow solver is more general and designed to handle complex geometries and unstructured grids, we use an interpolation, based on a two term Taylor series about the cell center that contains the non-mesh point. Let the non-mesh point, $\mathbf{x}_p$ be contained in the control volume with cell center coordinate, $\mathbf{x}_{\text{cv of p}}$, where $\phi|_{\text{cv of p}}$ and $\nabla \phi|_{\text{cv of p}}$ are known. The interpolant, $\mathcal{I}\phi$

provides an estimate for $\phi|_{\mathbf{x}_p}$.

$$\phi|_{\mathbf{x}_p} = \mathcal{I}\phi = \phi|_{\mathbf{x}_{\text{cv of p}}} + \nabla\phi|_{\mathbf{x}_{\text{cv of p}}}(\mathbf{x}_p - \mathbf{x}_{\text{cv of p}}) \tag{4.12}$$

To compute $\nabla\phi|_{\mathbf{x}_{\text{cv of p}}}$, a least squares gradient estimation for unstructured meshes [65, 79] is used that is robust in the presence of skewed mesh elements, and reduces to second order central differencing for regular Cartesian grids.

4.3.2 Efficient composition of flow maps

In practice, the desired temporal resolution of the FTLE fields is typically much finer than the integration time, T, over which the flow maps are computed. If this is the case, then several sets of tracer particles and scalar fields need to be evolved concurrently in the simulation in order to compute the exact forward and backward time flow maps. This problem has been addressed by [9], who used the property [77] that a time T flow map can be decomposed into a sequence of N sub-steps of length, $h = T/N$. Adopting their notation, we can write

$$\mathbf{\Phi}_{t_0}^{t_0+T} = \mathbf{\Phi}_{t_0+(N-1)h}^{t_0+Nh} \circ \cdots \circ \mathbf{\Phi}_{t_0+h}^{t_0+2h} \circ \mathbf{\Phi}_{t_0}^{t_0+h} \tag{4.13}$$

Because the flow maps are obtained discretely, the flow map from one sub-step will not necessarily point to the fixed grid points of the next sub-step, and implementing this reconstruction involves interpolation between sub-steps. Using the interpolation operator, $\mathcal{I}$, we have

$$\mathbf{\Phi}_{t_0}^{t_0+T} = \mathcal{I}\mathbf{\Phi}_{t_0+(N-1)h}^{t_0+Nh} \circ \cdots \circ \mathcal{I}\mathbf{\Phi}_{t_0+h}^{t_0+2h} \circ \mathbf{\Phi}_{t_0}^{t_0+h} \tag{4.14}$$

If the sub-step, h, corresponds to the frequency that FTLE fields are to be computed, then redundant tracer integrations and scalar evolutions can be eliminated. This means only one set of Lagrangian tracers and one set of Eulerian scalars, corresponding to the forward and backward time h flow maps, need to be evolved at any time during the simulation to compute the time T flow maps at all times

of interest. The integer N was referred to as the "speedup factor" by [9] because it represented the potential speedup of their post-processing approach relative to computing the "exact" time T flow map every time the FTLE field is needed.
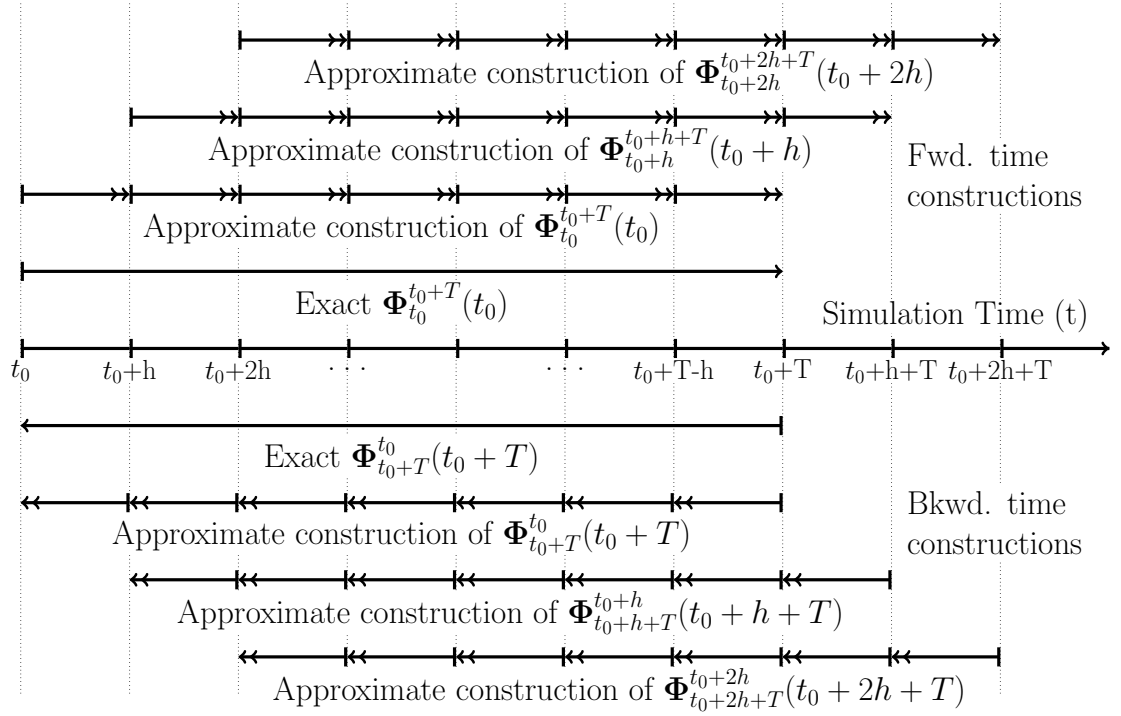
Figure 4.2: Integration of the flow map composition scheme with the flow solver forward time advancement. At any simulation time t that is a multiple of h, the forward time flow map, $\mathbf{\Phi}_{t-T}^{t}(t - T)$, and the backward time flow map, $\mathbf{\Phi}_{t}^{t-T}(t)$, are constructed from the N time h sub-steps.

A schematic of this flow map composition scheme, which corresponds to the "single-tiered unidirectional" approach of [9], and the way that it fits into the CFD simulation is shown in Figure 4.2. The simulation begins at $t = t_0$ and always advances in forward time. To compute the first time T flow maps, $\mathbf{\Phi}_{t_0}^{t_0+T}$ and $\mathbf{\Phi}_{t_0+T}^{t_0}$, exactly we could evolve a single set of Lagrangian tracers and Eulerian scalars over the entire interval $[t_0, t_0 + T]$, but this becomes computationally inefficient if we need to compute the FTLE field at intervals smaller than T. Instead the total

integration time, T is broken up into the N equal sub-steps of length $h = T/N$ shown in Figure 4.2. As the simulation proceeds, time h flow maps are recorded at these regular intervals and then re-initialized. Rather than storing all time h flow maps in memory, we write them to structured binary files on the hard disk. If sufficient memory is available to the computation, it is possible to avoid this step. However, we have found that performing parallel binary read/write operations accounts for a very small amount of the total simulation time (see section 4.4.6), and hard disk space is more readily available than memory on most computing platforms. Once the simulation reaches time t_0+T, the time T flow maps, $\boldsymbol{\Phi}_{t_0}^{t_0+T}(t_0)$ and $\boldsymbol{\Phi}_{t_0+T}^{t_0}(t_0 + T)$, are constructed using equation 4.14 and the time h flow maps that have been stored on disk. Because the flow solver uses unstructured grids and distributed memory, efficient interpolation of one flow map to the next requires careful memory and parallel task management. Details on the parallelization of this step are given in the Appendix. At subsequent simulation times, t, that are multiples of the flow map sub-step, h, the approximate reconstruction procedure is repeated to construct the forward time flow map, $\boldsymbol{\Phi}_{t-T}^{t}(t - T)$, and the backward time flow map, $\boldsymbol{\Phi}_{t}^{t-T}(t)$. For example, when the simulation time reaches $t = t_0 + 2h + T$, the time h flow maps, $\boldsymbol{\Phi}_{t_0+h+T}^{t_0+2h+T}(t_0 + h + T)$ and $\boldsymbol{\Phi}_{t_0+2h+T}^{t_0+h+T}(t_0 + 2h + T)$ are first written to disk. Then, the forward time T flow map, $\boldsymbol{\Phi}_{t_0+2h}^{t_0+2h+T}(t_0 + 2h)$, and the backward time T flow map, $\boldsymbol{\Phi}_{t_0+2h+T}^{t_0+2h}(t_0+2h+T)$, are constructed as shown in Figure 4.2.

4.3.3   Integrated Algorithm

Before discussing how the complete algorithm has been integrated into the flow solver, it is beneficial to review the relevant temporal scales, summarized in table 4.1, and how they are chosen in our integrated approach. In each problem considered in section 4.4, we can identify a reference time, t_{ref}, which characterizes the dominant hydrodynamic timescale of the flow. This could be an eddy turnover time, oscillation period, or any other meaningful reference time. To retain good accuracy of our flow solver, the simulation time, t, is advanced using

the timestep, Δt. The latter is set so that the maximum Courant-Friedrich-Lewy number is roughly $CFL = u_{cv}\Delta t/\Delta x = 0.3$, where u_{cv} is the velocity in a control volume with characteristic dimension Δx. In general, this results in a simulation timestep several orders of magnitude smaller than t_{ref}. The FTLE integration time, T, should be chosen to be long enough that all LCS candidates are detected. In practice, T is typically chosen to be larger than the longest hydrodynamic time scale of the flow. The final timescale, h, is the flow map sub-step. This is chosen to match the desired temporal resolution of FTLE fields. Typically, to create an animation of an FTLE field's evolution, one would desire h to be significantly smaller than t_{ref}, and therefore significantly smaller than T. Our experience with the several different flows presented in section 4.4 suggests that high quality animations with sharp FTLE features, and good temporal resolution between frames can be obtained with h between $T/20$ and $T/10$.

Table 4.1: Timescales involved in the integrated computations

Timescale	Meaning
t_{ref}	Hydrodynamic reference time
t	Current simulation time
Δt	Simulation timestep
T	FTLE integration time
h	Flow map sub-step

Figure 4.3 is a block diagram that shows how the combined simulation and FTLE computation proceeds. After the flow solver advances the velocity field per the usual procedure from t^n to t^{n+1}, the intermediate velocity field, $\mathbf{u}^{n+1/2} = 0.5\left(\mathbf{u}^n + \mathbf{u}^{n+1}\right)$ is used to update the forward and backward time h flow maps. The explicit trapezoidal scheme (equation 4.8) is used to update the Lagrangian tracer positions, and the two time level semi-Lagrangian scheme is used to update the value of the current time h backward flow maps on the fixed grid (according to equation 4.10). If Lagrangian tracers cross boundaries during the integration step, they are relocated back into the in-bounds cell, except in the case of an outflow, to

which they remain fixed for the duration of their integration. The backward time flow map boundary conditions suggested by [61] and described earlier are imposed at the boundary cells of the fixed grid.

Update time h flow maps
- Advance time h Lagrangian tracers (fwd)
- Advance time h Eulerian fields (bkwd)
- Handle flow map B.C.

$\text{mod}\,(t^{n+1}, h) = 0$

$\mathbf{u}^{n+1/2}$

$\text{mod}\,(t^{n+1}, h) \neq 0$

FTLE Computation
- Map current tracer positions to launch cells to obtain time h forward flow map
- Write time h flow maps to disk
- Construct time T flow maps from sequence of N time h flow maps
- Compute FTLE fields
- Re-initialize time h flow maps.

Main Flow Solver
- $t^{n+1} = t^n + dt$
- Advance $\mathbf{u}^n$ to $\mathbf{u}^{n+1}$

RETURN

Figure 4.3: Block diagram showing the flow of the CFD simulation with integrated FTLE computations.

At this point, if the current simulation time is not a multiple of the flow map sub-step ($\text{mod}\,(t^{n+1}, h) \neq 0$), we return to the main flow solver loop. If $\text{mod}\,(t^{n+1}, h) = 0$ is satisfied, then the following steps are initiated to compose the time T flow maps and compute the FTLE fields. First, the current position of all Lagrangian tracers is communicated back to the processor and cell that launched them time h ago, giving the time forward time h flow map, $\mathbf{\Phi}_{t-h}^{t}(t - h)$, on the fixed grid. The backward time h flow map, $\mathbf{\Phi}_{t}^{t-h}(t)$, is known because it has been evolved up to time t in forward time on the fixed grid. Next, these time h forward and backward flow maps are written to binary files on the hard disk. Then, the time T flow maps are constructed according to equation 4.14 from the sequence of time h flow maps stored on the hard disk. The Jacobian of each time T flow map is computed at the cell centers of the fixed grid using the least squares gradient

operator, and the right Cauchy-Green deformation tensors are computed according to equation 4.5. Finally, the largest eigenvalue of the forward and backward time tensors are determined, and the corresponding FTLE fields are computed using equation 4.6. As a last step before returning to the flow solver, the time h flow maps are re-initialized on the fixed grid. This amounts to positioning a Lagrangian tracer at each cell center, and setting value of the backward time flow map to the cell center coordinates.

4.4 Numerical Test Cases

We now present several examples designed to show the performance and capability of the integrated computations.

4.4.1 Time Dependent Double Gyre

As the first test, we examine the analytic model of time periodic Raleigh-Bénard convection of [103]. This two dimensional stream-function model, sometimes referred to as "double gyre flow," consists of two counter rotating vortices that expand and contract in tandem. The time dependent velocity in the xy plane is described by

$$
\begin{aligned}
u_x &= -\pi A \sin(\pi f(x,t)) \cos(\pi y) \\
u_y &= \pi A \cos(\pi f(x,t)) \sin(\pi y) \frac{\partial f}{\partial y}.
\end{aligned}
\tag{4.15}
$$

where,

$$
\begin{aligned}
f(x,t) &= a(t)x^2 + b(t)x \\
a(t) &= \epsilon \sin(\omega t) \\
b(t) &= 1 - 2\epsilon \sin(\omega t)
\end{aligned}
\tag{4.16}
\tag{4.17}
\tag{4.18}
$$

To facilitate comparison with other previous studies of the FTLE field in this system [62, 61], we choose parameter values of $A = 0.1$, $\epsilon = 0.1$, $\omega = 2\pi/10$. The domain is discretized with a uniform, 512×256 Cartesian mesh in the region $[0, 2] \times [0, 1]$. The "simulation" is started at at time $t = 0$, and at advanced to $t = 1.5t_{ref}$, where the reference time is $t_{ref} = 2\pi/\omega$. A constant time step of $\Delta t = 2.5 \times 10^{-4}$ is used. The computational pipeline shown in Figure 4.3 is followed exactly, with the exception that in the main flow solver we force the velocity field to obey system (4.15) at every time step.
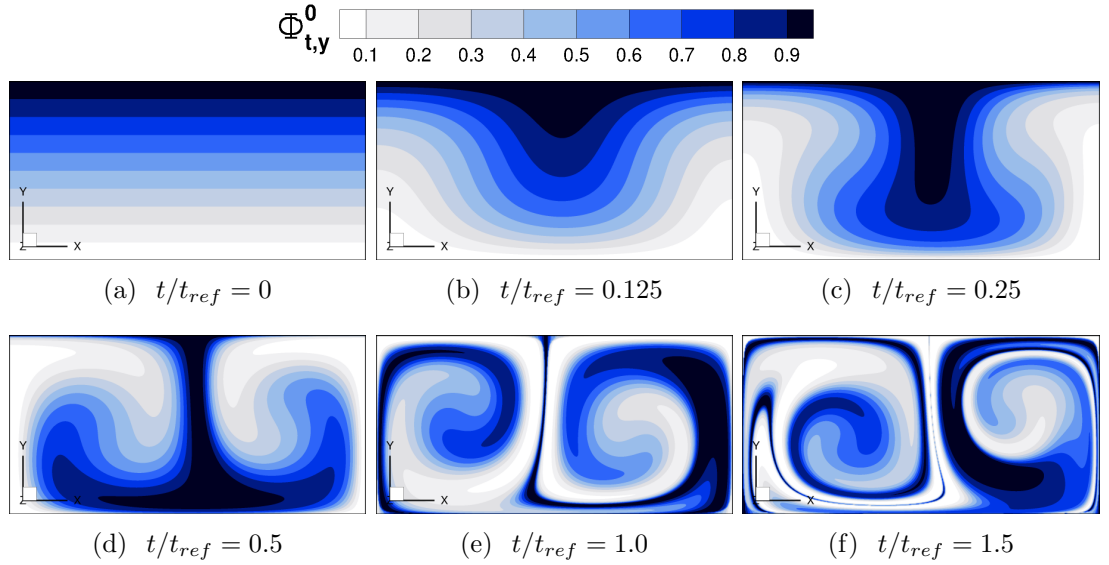
Figure 4.4: Evolution of the Eulerian representation of the backward time flow map on the fixed grid during simulation of the double gyre flow. The scalar field, $\Phi^0_{t,y}$, represents the y coordinate of the Lagrangian particle at time, $t = 0$, that arrives at the fixed grid points at each later time.

To illustrate the Eulerian treatment of the backward time flow map, we plot the evolution of its y component, $\Phi^0_{t,y}$, at several times in Figure 4.4. The Eulerian field is initialized to $\Phi^0_{t,y} = y_{cv}$ at $t = 0$. These takeoff coordinates are evolved in with flow in forward time using the semi-Lagrangian scheme. The flow maps are reinitialized at regular intervals, $h = 1.25t_{ref}$, and the longer time flow maps are constructed using the composition scheme described in section 4.3.2. For short

times $(t/t_{ref} \lesssim 0.5)$, it is evident that the flow transports the high $\Phi^0_{t,y}$ values in the negative y direction in the center of the cell, and transports the low $\Phi^0_{t,y}$ in the positive y directions along the left and right edges. At later times, the x direction oscillation results in stretching and folding of fluid transported between the two gyres, and a significantly more complex flow map. Despite the Eulerian treatment of $\Phi^0_{t,y}$, its sharp gradients are well preserved by the semi-Lagrangian advection scheme.

We can also use this flow to show that the Lagrangian treatment of the forward time flow map is equivalent to the Eulerian treatment of the backward time flow map. Because the system is symmetric in time about $t = 0$, forward trajectories in system (4.15) are equivalent to backward trajectories in the negative of system (4.15). This allows us to make direct comparisons of FTLE fields computed using the Lagrangian and Eulerian representations. To do this, we first apply the flow field of system (4.15) and simulate the flow for an integration time of $T = 1.5t_{ref}$. We then simulate the negative of system (4.15) up to the same time. In Figure 4.5a-d the forward and backward time FTLE fields are shown. It is nearly impossible to distinguish any differences by eye in the FTLE obtained by the Lagrangian and Eulerian representations of the flow map. To obtain a quantitative measure of the difference, we also plot the normalized relative error between the Lagrangian and Eulerian results, Relative Error $= |FTLE^L - FTLE^E|/max(FTLE^L)$ for forward and backward time in Figure 4.5e-f. Because this metric is very sensitive to slight shifts in ridges of the FTLE field, we are generally satisfied that for most of the domain, the difference between the Lagrangian and Eulerian FTLE field is less than 10% of the max FTLE. Qualitatively, the FTLE fields for the model parameters chosen are in good agreement with prior studies [62, 61] with the same parameters. This gives us confidence that our forward and backward time computations are consistent and accurate.

(a) σ_T^0 Lagrangian

(b) σ_0^T Lagrangian

(c) σ_T^0 Eulerian

(d) σ_0^T Eulerian

(e) σ_T^0 relative error

(f) σ_0^T relative error
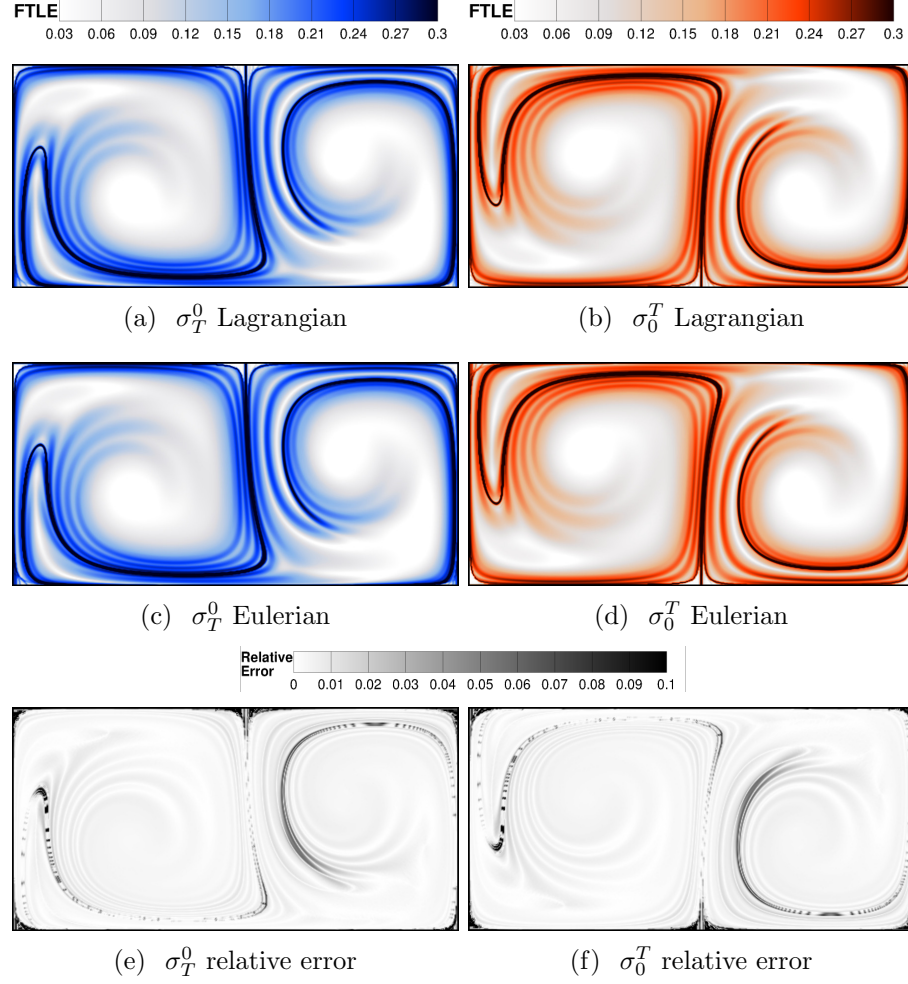
Figure 4.5: Comparison of the Lagrangian and Eulerian approaches to calculating the FTLE field. (a) through (d) show the FTLE fields computed in (a) backward time with the Lagrangian approach, (b) forward time with the Lagrangian approach, (c) backward time with the Eulerian approach (d) forward time with the Eulerian approach. (e) and (f) show the relative error between the Lagrangian and Eulerian approaches for backward and forward time respectively.

4.4.2   Flow Over a Fixed Cylinder at $Re = 300$

We now turn to flow over a fixed cylinder, a typical test case for an unsteady flow solver. A body fitted C-grid containing a total of 1.3 million hexahedral control volumes is used to discretize the computational domain, shown in Figure 4.6. At the no-slip cylinder boundary, the grid spacing is $D/100$, where D is the cylinder diameter. In the wake region the grid spacing is coarsened to $D/40$. A slip condition is imposed on the top and bottom boundaries in the y direction, and a convective outlet condition ($\partial \mathbf{u}/\partial \mathbf{n} = 0$) is set $25D$ downstream of the cylinder. The grid is two cells thick in the z direction, where periodicity is imposed. At the C boundary $20D$ upstream of the cylinder, a uniform inflow velocity, U, is specified so that the Reynolds number, $Re = UD/\nu$, is equal to 300 where ν is the kinematic viscosity.
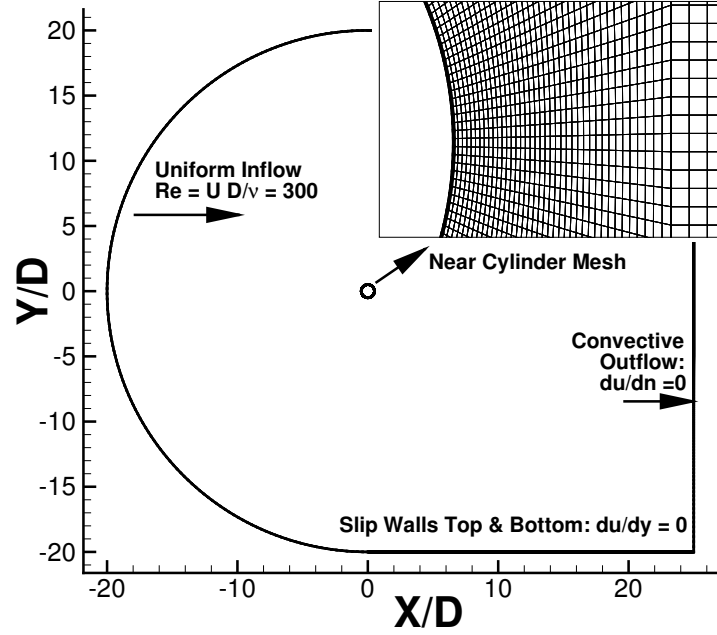

Figure 4.6: Domain and mesh used for simulating flow over a fixed cylinder.

The flow is started from rest and simulated for $t = 200t_{ref}$, where $t_{ref} = D/U$, using a constant timestep $\Delta t = 1.7 \times 10^{-3} t_{ref}$. The FTLE fields are computed for

an integration time of $T = 7.8t_{ref}$, and with a flow map sub-step, $h = T/13 = 0.6t_{ref}$. When updating the backward time flow map, the boundary conditions discussed earlier are used. That is, on the solid cylinder boundary and at the inflow we impose $\Phi_{t,i}^{t-T} = x_i$, while at the slip boundaries and outflow, we impose $\mathbf{n} \cdot \nabla \Phi_{t,i}^{t-T} = 0$. The influence of the non-physical fixing of Lagrangian particles that encounter the outflow does not travel upstream far enough to affect the forward time FTLE fields in the near wake. At this Reynolds number, the cylinder wake becomes unsteady, and develops into an unsteady Kármán vortex street. This phenomena is nicely captured by the FTLE fields shown for several instants in time in Figure 4.7. For $t \lesssim 50D/U$, the boundaries of the symmetric recirculation bubble are clearly marked as LCS candidates by the FTLE ridges. The bubble becomes unstable and the FTLE ridges begin to wander from the line of symmetry at $t = 60D/U$. By $t = 75D/U$, alternating vortex shedding has set in, and the FTLE ridges show the skeleton of the periodic vortex chain in the wake of the cylinder.

4.4.3   In-Line Oscillation of a Circular Cylinder

In this example, we show that our integrated approach to computing FTLE fields can be easily applied to simulations of flows with *moving boundaries*, which are encountered in multiphase flows, and fluid-structure interaction problems. Our flow solver has the extended capability to simulate resolved rigid particle-flow interactions on non-body conformal Cartesian grids using a fictitious domain approach [2], and we demonstrate the simultaneous FTLE computations here with a simple test case. For our example, we consider the flow generated by a circular cylinder performing linear oscillations. This problem can be described in terms of the maximum cylinder Reynolds number and the Keulegan-Carpenter number,

(a) $t = 40.2D/U$

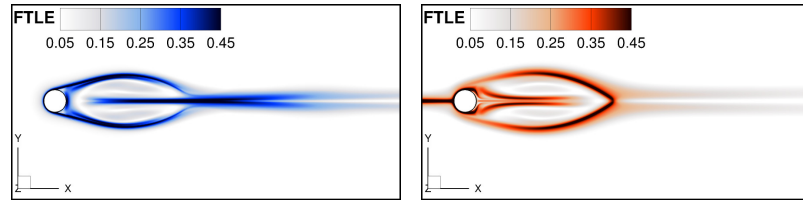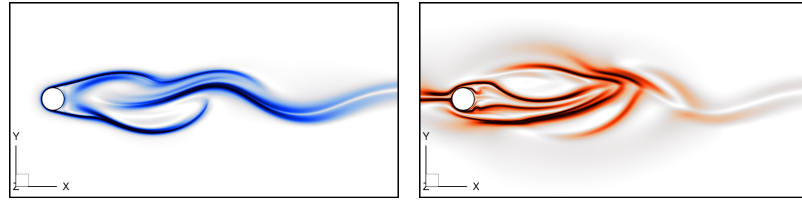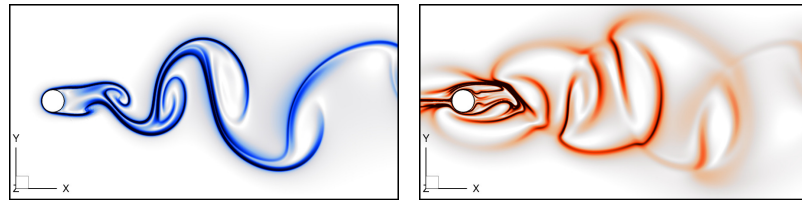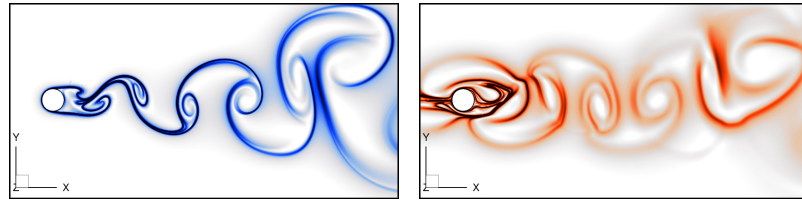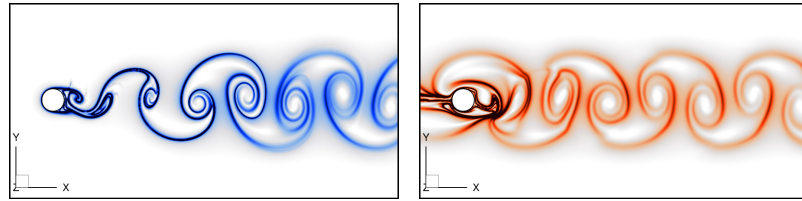(b) $t = 60.6D/U$

(c) $t = 69.0D/U$

(d) $t = 75.6D/U$

(e) $t = 168.0D/U$

Figure 4.7: Snapshot of FTLE fields for flow over a cylinder at $Re = 300$ for $T = 7.8D/U$. Left column shows backward time FTLE, right column shows forward time FTLE.

which characterizes the oscillation frequency of the cylinder:

$$Re = \frac{U_m D}{\nu} \tag{4.19}$$

$$KC = \frac{U_m}{fD} = \frac{2\pi A}{D} \tag{4.20}$$

where, U_m is the maximum cylinder velocity, D is the cylinder diameter, ν is the kinematic viscosity of the fluid, f is the frequency of cylinder oscillations, and A is the oscillation amplitude. The cylinder and fluid both start at rest. At $t > 0$, the cylinder position and velocity are described by the sinusoidal functions:

$$x_c(t) = -A\sin(2\pi ft) \tag{4.21}$$

$$U_c(t) = -2\pi Af\cos(2\pi ft) \tag{4.22}$$

The cylinder diameter, viscosity, maximum velocity, and oscillation amplitude and are chosen so that $Re = 100$, and $KC = 5$, corresponding with the experimental study of [20]. A block-type Cartesian grid is used, with a uniform patch in the region of cylinder motion, and periodicity assumed in the spanwise direction. The domain size is $50D \times 50D$ in the x and y directions. Near the cylinder, a uniform grid spacing of $D/100$ is used. The grid uses $750 \times 600 \times 2$ (9×10^5 total) cells in the x, y, and z directions.

The flow is allowed to develop for 9 complete oscillations before computing any FTLE fields. The integration time, $T = 1.5t_{ref}$, where $t_{ref} = 1/f$ is the oscillation period. The FTLE fields are constructed at a flow map sub-step, $h = T/30$. Figure 4.8 shows the forward and backward FTLE field at four different phase angles, $\theta = 2\pi ft$, of the oscillation. At this Reynolds number and KC number, the flow is characterized by a pair of counter rotating vortices being shed from the top and bottom of the cylinder every half cycle. Upon reversing direction at $\theta = 90^o$ and 180^o, the cylinder destroys the previously formed pair while creating a new pair in its wake. The attracting LCS candidates compare very well with experimental dye visualizations of the same flow (see FIG 5 of [20]). For this relatively long

(a) $\theta = 0^o$

(b) $\theta = 54^o$
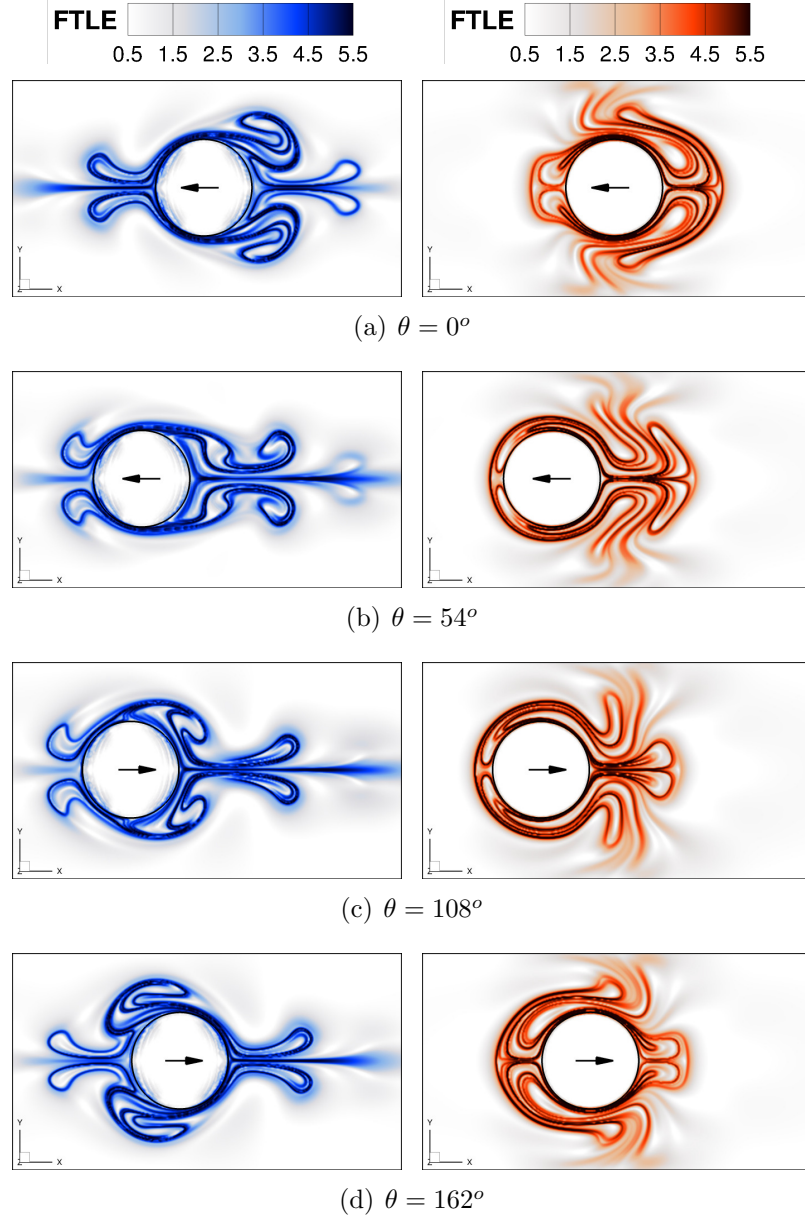
(c) $\theta = 108^o$

(d) $\theta = 162^o$

Figure 4.8: FTLE fields for selected phase angles in one half cycle of the oscillating cylinder flow. The arrow denotes the direction of cylinder motion. Left column shows backward time FTLE, right column shows forward time FTLE.

integration time the LCS candidates, which can be visually identified as ridges of the FTLE field, show how the cylinder entrains fluid into its wake, then stretches and folds the fluid as it oscillates.

During this simulation, the fluid inside the cylinder, and as a result the flow map, are constrained to the rigid body motion specified by equations 4.21 and 4.22 at all times. This results in a sharp FTLE ridge along the solid-fluid interface as would be expected. Despite the immersed representation of the moving rigid body on the fixed grid, the FTLE fields obtained contain sharp well defined features, which provides confidence that the integrated computations are easily extensible to problems with moving boundaries.

4.4.4   Three Dimensional Turbulent Vortex Ring

To show the capability of the proposed approach for three dimensional flows, we now examine the case of a traveling vortex ring. LCS theory has previously been applied to vortex ring flows as a way to to understand the lobe dynamics of entrainment, detrainment, and transport within the ring [97], as well as the identification of vortex pinch-off during formation of the flow structure [76]. This vortex ring is generated by a quick pulse of fluid from a circular inlet located at the $x = 0$ wall of an initially stagnant rectangular domain. The domain geometry, shown in Figure 4.9a, is $30cm \times 30cm$ in the y and z directions, and extends $80cm$ in the x direction. The velocity of the inflow jet as a function of time is similar to the studies of bubble vortex interaction by [104] and [28], but is scaled so that the initial circulation is equal to $79.5cm^2/s$. The shear layer of the jet rolls up into a vortex ring and travels downstream at a speed of roughly 14% of the maximum jet velocity. Also in Figure 4.9a, we visualize the three dimensional ring structure by showing an isosurface of swirling strength, λ_{ci}. This Eulerian vortex detection criteria corresponds to the imaginary part of the complex eigenvalue of the instantaneous velocity gradient tensor [115]. The isosurface shows the turbulent nature of the ring structure, particularly in its wake as it travels downstream. The domain is meshed with a $640 \times 241 \times 241$ (37×10^6 total points) Cartesian grid. The cells

are coarsened in the corners of the tank, away from the evolution of the vortex ring.

The forward and backward time FTLE fields are computed during the simulation for $T = 1.5s$, about the time it takes the ring to travel $10cm$ downstream. A flow map sub-step of $h = T/15 = 0.1s$ is used to compose the time T flow maps. Two dimensional slices in the $z = 0$ plane of the three dimensional FTLE fields are shown in Figure 4.9b and c for $t = 2s$, $t = 7s$, and $t = 12s$. Despite the three dimensionality of the flow, FTLE ridges can be clearly seen in the xy plane, particularly at $t = 2s$ when the candidate LCS surfaces lie perpendicular to the $z = 0$ plane. At later times, as the ring travels downstream and begins to break down, the strong FTLE ridges are less evident in this plane, implying the LCS candidate surfaces are no longer perpendicular to this slice. Three dimensional ridge extraction for cases such as this is an important but difficult task (see [31] and [68] for some unique approaches relevant to LCS) and is reserved for a future work.

4.4.5   Flow through a random sphere pack

The final case studied, unsteady flow through a random sphere pack, is designed to showcase the benefits of integrating the FTLE computations with the CFD simulation. At moderate Reynolds numbers, flow through packed beds of spheres can develop a variety of coherent porescale flow features including jets and helical vortices, which operate over a wide range of characteristic space and times scales. For packed bed reactors and other complex flow configurations, LCS theory could provide crucial new understanding of how geometric and hydrodynamic features affect transport and mixing, and how slight geometric changes can produce desirable or undesirable to be valuable design tools. However, the geometric scale and complexity of these types of configurations precludes computing the FTLE using the typical post-processing approach.

The geometry that we consider is a channel packed with 51 spheres with constant diameter D, shown in Figure 4.10a. A *ballistic deposition* algorithm, similar
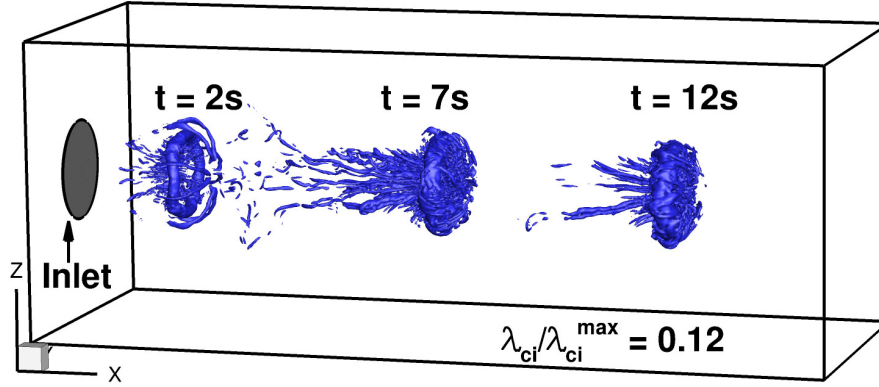
(a) Eulerian swirling strength
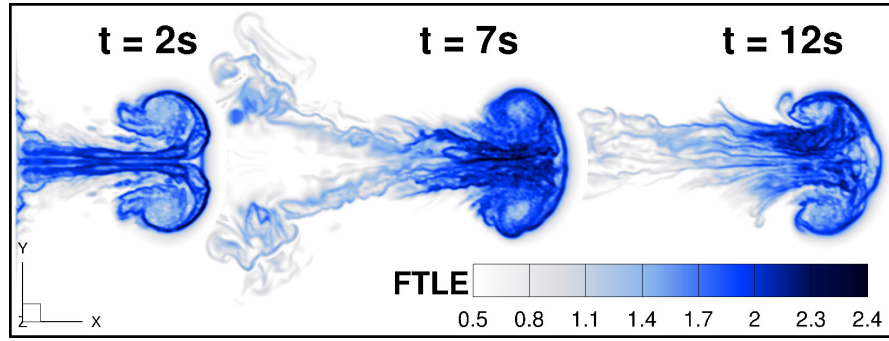
(b) Backward time FLTE
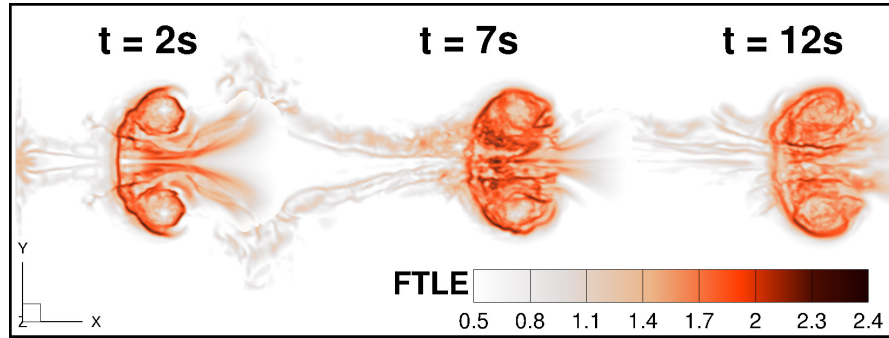
(c) Forward time FTLE

Figure 4.9: Snapshots of the three dimensional vortex ring at $t = 2s$ $7s$, $12s$. (a) Isosurface of swirling strength, $\lambda_{ci}/\lambda_{ci}^{max} = 0.12$. (b-c) FTLE fields for the $z = 0$ cross sectional slice through the center of the 3D vortex ring

to the method employed by [4], was used to pack the spheres into a $4D \times 4D \times 4D$ box resulting in a mean void fraction, $\epsilon = 0.58$. The channel was then formed by extending the box boundaries 3D upstream and downstream of the spheres. The flow is driven in the positive z direction by a constant inflow velocity at the upstream boundary so that the pore Reynolds number for the flow is $Re = \frac{UD}{\epsilon \nu} = 600$, where U is the assigned inflow velocity and ν is the viscosity. A convective outlet condition is located $3D$ downstream of the packing.

The same fictitious domain approach used for the oscillating cylinder case is used to represent both the fixed solid and fluid region on a non body conformal Cartesian grid. This grid has a uniform spacing of $D/80$ everywhere in the porespace, and is stretched toward the inflow and outflow boundaries. In total, it contains 47×10^6 control volumes. Validation of the fictitious domain approach for flow through packed beds of spheres , including justification for the chosen grid spacing in this case may be found elsewhere [30]. The flow was started from rest and allowed to develop to for $t = 60t_{ref}$ where $t_{ref} = D/U$. The flow becomes unsteady right away, and contains a number of interesting porescale features. The most notable of which is helical vortices, elongated in the mean flow direction. These are shown in Figure 4.10a by plotting isosurfaces of the Eulerian swirling strength field, $\lambda_{ci}/\lambda_{ci}^{max} = 0.25$.

The FTLE fields are computed for an integration time of $T = t_{ref}$ using a flow map sub-step of $h = T/7$. Because of the complex geometry, the flow is multiscale in nature. After some experimentation, these values of T and h were determined to provide sharp FTLE features (long enough T) and good temporal resolution of all timescales (small enough h). The fields are visualized in Figure 4.10b and c on the two dimensional $z/D = 2.5$ slice. This slice is oriented perpendicular to the mean flow, at roughly two thirds of the way through the packed bed. The fine grid resolution used for the simulation provides excellent resolution in the FTLE field, and a number of porescale LCS candidates can be easily located with the eye. We will discuss a few of the features in the plane here. (A) is a thin region sandwiched between a sphere and the $x = 0$ wall occupied by a pair of

(a) 3D domain and vortex isosurfaces

(b) Backward time FTLE

(c) Forward time FTLE

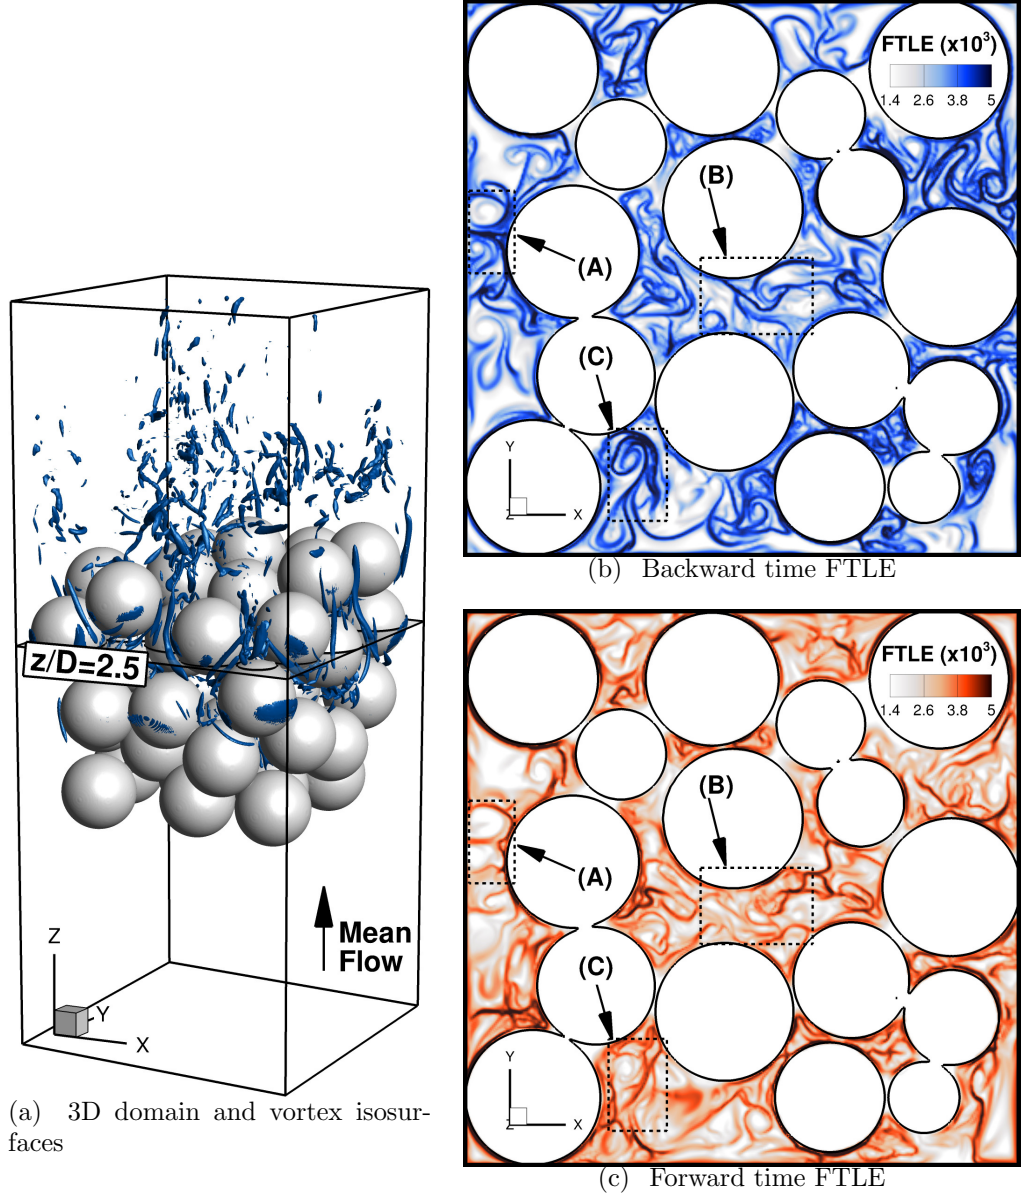Figure 4.10: Flow through the random sphere pack. (a) Shows the simulation domain and instantaneous isosurfaces of $\lambda_{ci}/\lambda_{ci}^{max} = 0.25$, and the $z/D = 2.5$ plane used for visualization of the FTLE fields. (b) and (c) show a snapshot of the FTLE fields for at the $z/D = 2.5$ cross stream slice. The LCS candidates in pores A, B, and C demonstrate distinct, time dependent behavior.

counter-rotating helical (vorticity aligned with the z axis, in and out of the page) vortices. The boundaries of the two vortex cores are well defined by the forward and backward FTLE fields. In time, these features wander slightly along the wall, but remain distinct throughout the simulation. (B) is a large open pore in the center of the packing, where several streams of fluid converge. The LCS candidates, which seem to distinguish the heads of jets as well as the boundaries of vortex cores, are highly transient and move throughout the pore as the simulation proceeds. Pore (C) is a large open region near the wall of the packing that is also highly transient. There is evidence of vortex shedding as the tail of the ridge feature inside the boxed region flaps periodically into the open area of the pore. These results give us confidence that the integrated approach can capture FTLE features in large scale complex configurations.

4.4.6   Computational Expense

By incorporating the FTLE calculation into our flow solver we have removed the need to post-process a sequence of velocity fields. However, in many ways, we have just shifted the computational overhead to the simulation, and it is important to quantify how much this will cost in terms of additional simulation time. If the cost of the FTLE computation were to become significantly large relative to the standard flow solver, the appeal of the integration we have described would diminish. In Figure 4.11 the expense of FTLE computations is plotted in terms of the percentage of the total simulation time that they consume for each case presented in this section. The expense for each case is further broken into contributions from (i) the Lagrangian update of the forward time flow map tracers, (ii) the semi-Lagrangian update of the backward time flow map scalar fields, (iii) the reconstruction of time T flow maps from the time h sub-steps including reading back the sub-steps from hard disk, and (iv) all other FTLE related computations including writing to the hard disk, composing the strain rate tensor, computing its max eigenvalue and enforcement of flow map boundary conditions. The double gyre case has been omitted, since the velocity field in that case is enforced artificially and the flow

solver itself does very little work. To make the comparison as equal as possible for all cases, the time step, Δt was chosen so that the maximum CFL number was 0.3, the integration time was set to $T = 2400\Delta t$, and the flow map sub-step was set to $h = T/12 = 200\Delta t$. All cases were run on the Lonestar super computer at the Texas Advanced Computing Center (TACC), which consists of 1888 nodes of 3.33GHz Intel 2 Hex-core (12 processing cores/node) Xeon 5680® processors connected with a 40Gbit/s Inifiniband Mellanox Switch. Each node of 12 cores shares 24GB of memory. The grids for each case were partitioned with between $38k$ and $56k$ control volumes per partition, resulting in parallel simulations using between 24 and 1200 processors.

It is encouraging to see that the computational time dedicated all FTLE related computations varies between a low of 14% (the vortex ring), and a high of 26.5% (the fixed cylinder) of the total simulation time. Variations between the cases have to do with differences in the load balancing of the parallel partition and the amount of iterations required by the flow solver for convergence, among other factors. It is evident that the implementation scales effectively to the larger, three dimensional simulations distributed in parallel on over 1,000 processors. It is important to note that the time required to construct the time T flow maps from the time h sub-steps scales almost linearly with the speedup factor, N. In this study we have chosen $N = 12$, as we have found it to be sufficient to resolve most FTLE features of interest in these flows.

Memory and hard disk usage is also of concern when performing either distributed memory or shared memory flow simulations. In our current implementation, the simulation memory increase is small because all flow map sub-steps are stored in temporary files on the hard disk. To do this requires

$$N_{cv} \times \frac{T}{h} \times \frac{6 \text{ doubles}}{cv} \times \frac{8 \text{ bytes/double}}{1024^3 \text{ bytes/GB}} \tag{4.23}$$

of temporary hard disk space. Thus, a 1 million cv simulation with $N = 12$ will require $0.54GB$. Our largest simulation, the 46 million cv simulation of flow
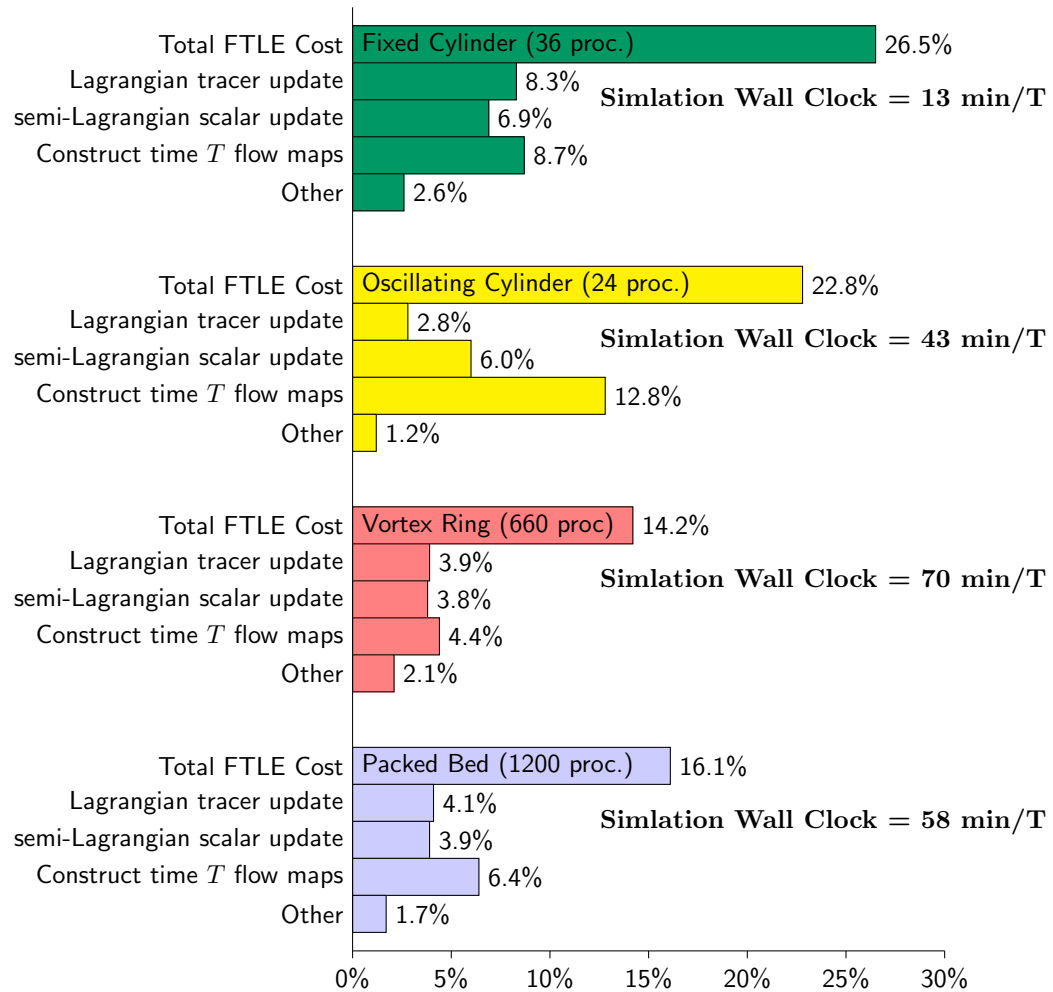
Figure 4.11: Computational cost of embedding the FTLE computations into the direct numeric simulation. Bars indicate percentage of the total simulation time dedicated to certain parts of the FTLE computations for each case studied.

through the packed bed with $T/h = 7$ required $22GB$ of temporary disk space, an amount readily available on most modern computing platforms.

To compare the computational overhead of our integrated approach with the post processing approach, 48 equally spaced velocity fields were output from the fixed cylinder simulation over an interval corresponding to the FTLE integration time, $T = 2400\Delta t$. The velocity fields were then interpolated to a uniform 1750×750 grid (1.3×10^6 total points) covering a subset of the total domain near the cylinder and its wake ($-10 \le x/D \le 25$, $-7.5 \le y/D \le 7.5$). Using a Lagrangian tracer grid of the same size, a single, backward time FTLE field was computed for integration time T using the freely available *LCS Matlab Toolkit*[84][2]. The computation took 17 minutes using a single core of a 2.83GHz Intel Core 2 Quad® processor, and consumed 0.47GB of memory. In comparison, a total of 24 FTLE fields (12 forward and 12 backward) were computed for the same number of cells during the simulation using the integrated approach. This required 13min of simulation time and 0.7GB of hard disk space to store the intermediate flow maps. The FTLE related computations accounted for 26.5% of the simulation time, meaning each FTLE field took about 9 seconds to compute. This represents a roughly 2 order of magnitude speedup over the post processing approach, and is due to both the parallelism of the integrated approach as well as the efficient flow map composition technique.

4.5   Conclusions

In this paper, we have demonstrated that both the forward and backward time FTLE fields may be computed, on-the-fly, during direct numerical simulation of the Navier-Stokes equations. Our implementation into a parallel, distributed memory, unstructured grid flow solver utilizes Lagrangian particle tracking to compose the forward time flow maps, and an Eulerian treatment of the backward time flow map along with a simple semi-Lagrangian advection scheme. To avoid redundant

---

[2]LCS MATLAB Kit Version 2.3 available from www.dabiri.caltech.edu/software.html

tracer integrations and scalar advections, time T flow maps are composed from a sequence of time h sub-steps, corresponding to the unidirectional single-tiered method of [9]. This enables us to visualize the evolution of the FTLE fields with high temporal resolution, without burdening the simulation excessively.

The implementation has been tested for several canonical and new flows where LCS theory has not yet been applied. We first showed that the Lagrangian and Eulerian treatments of the forward and backward time flow maps produce high quality and similar results using the analytic double gyre test case. We next showed the applicability to unstructured grids and bluff body flows (the fixed cylinder), rigid body motion and multiphase flow problems (the oscillating cylinder), three dimensional flows (the traveling vortex ring), and large scale simulations with complex boundaries (the packed bed). In all of these cases, the overhead related to the FTLE computations is small relative to the total simulation time. The implementation also scales well for simulations utilizing over 1000 processors.

The integration of these computations directly in the simulation eliminates the need for expensive post-processing of large data sets, but also has a number of additional benefits. Because the velocity field is available at the native space and time resolution of the simulation, integration errors in the flow maps are minimized. Having the FTLE fields available to the simulation makes exploring active control of the FTLE fields and their associated LCS much more accessible, and could prove important to a number of applications. Our future work will pursue the possibility of this. We also plan to integrate the extraction of the LCS based on new, more rigorous LCS theory [39, 27, 26, 40] as it continues to evolve. Additional exploration of LCS-like flow features such as inertial LCS, which are the attracting and repelling structures for inertial particles [84, 95], and burning invariant manifolds (BIMs), which are invariant barriers to reaction fronts [66] are also possible within a similar framework.

4.6   Acknowledgments

Chapter 5: Conclusion

This work investigated the porescale structure of inertial flows in porous media using direct numerical simulation data. The flow regime investigated was in the range of $10 \leq Re \leq 600$, covering the emergence of both steady and unsteady inertial flow features.

A major challenge associated with accurate porescale simulations of porous media and packed bed flows is how to accurately represent the complex solid-fluid interface while retaining computational efficiency. With this in mind, the relative performance of two simulation approaches were compared in Chapter 2. In the first and historical more utilized approach, an unstructured tetrahedral mesh is created that conforms to the solid boundaries of the porespace, and the no-slip boundary condition is enforced directly. In the second approach, a fictitious domain formulation is used wherein the entire domain, inclusive of the solid regions is treated as a single incompressible fluid. The solid fluid interface is represented on a regular, Cartesian grid by Lagrangian marker points with subgrid resolution. A synthetic force is then applied within the solid regions in order to satisfy the rigidity and no-slip constraints associated with the fixed, solid spheres. A number of systematic tests were performed to compare the two approaches performance for packed bed type applications. Both proved capable of capturing steady and unsteady porescale flow features. However, the fictitious domain approach demonstrated far superior performance in terms of the computational effort to obtain a fixed level of uncertainty. This is due to the regularity of the Cartesian grids, which also do not require expensive pre-processing (mesh generation). Moving forward, fictitious domain approaches should be the weapon of choice, when performing resolved simulation of porescale flow in packed beds.

A detailed characterization of vortical flow structures in porous media was made in Chapter 3. High resolution, time accurate simulations were used to capture both

steady and unsteady vortex dynamics at the pore scale. Using both qualitative visualization techniques along with and statistical comparisons of the porescale flow fields allowed for several key differences to be identified between vortical structures in randomly and arranged packings. There is a clear relationship between pore geometry and the types of pore-scale vortical structures observed. In arranged packings, the dominant flow features are multi-lobed vortex rings, with the vorticity aligned perpendicular to the mean flow. Conversely, the dominant flow feature for the Reynolds numbers studied in the random arrangement are helical vortices, elongated in the mean flow direction. The emergence of these different types of vortical features are confirmed through statistical analysis of velocity, vorticity and helicity in the porespace. Unsteady dynamics are also strongly linked to pore geometry. For Reynolds numbers near 600, spatially and temporally periodic Kármán like vortex shedding is observed at a single Strouhal number near 1 in a simple cubic arrangement. Flow through the random arrangement exhibits a wider range of time scales with Strouhal numbers between 0.29 and 1.57 but without any observed vortex shedding.

In order to study the effects of pore scale flow features on mixing and transport in complex porous media, a computational tool was developed in Chapter 4 to compute the finite-time Lyapunov exponent field. Rather than do this during a post processing procedure, the relevant computations were integrated into the DNS code, allowing for significantly reduced overhead and higher quality results. The approach was validated for a variety of flows including unsteady flow through a randomly packed bed of spheres at $Re = 600$. The results unambiguously define the boundaries of dynamically distinct porescale features including counter rotating helical vortices and jets, and capture time dependent phenomena such as vortex shedding at the pore level.

APPENDICES

Appendix A: Pointwise & Glyph Used to Automate Mesh Generation for Porous Media Flow

This brief article appeared in *The Connector*, a quarterly publication for Pointwise® users. March/April 2011. By Justin Finn.

At Oregon State University, researchers in the Computational Flow Physics Lab (CFPL) are using large scale, high resolution simulations to better understand flow through complex porous media. Both arranged and random sphere packs, such as the one shown in Figure A.1a, are being studied.
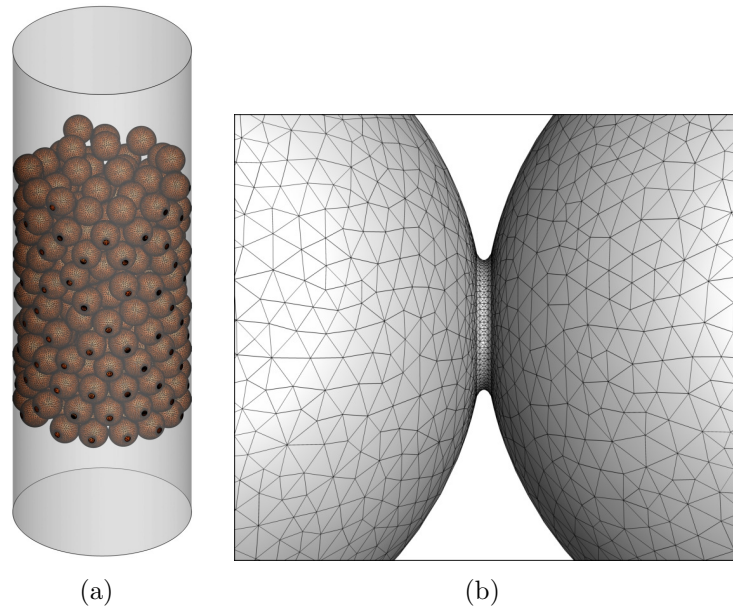

(a)   (b)

Figure A.1: A typical simulation domain consisting of a cylindrical tube packed randomly with contacting spheres. (a) The surface mesh on the spheres is shown. (b) Closeup of the contacting region. A smooth surface mesh transition provided by bridging two contacting spheres with a fillet.

Porous materials are building blocks for many of the natural and man-made systems that influence our daily lives, such as subsurface groundwater networks and packed bed reactors. Despite their importance, we know very little about flow dynamics in porous media, particularly at Reynolds numbers much greater than one. At these higher flow rates, the effects of fluid inertia may result in complex, unsteady, three dimensional flow features, including helical vortices, jets, and backflow regions. These pore scale features can dramatically effect properties of broader interest, such as pressure gradient, mixing, and dispersion in the medium.

Creating quality meshes for these types of geometries notoriously can be tricky and time consuming. The crux of the process is dealing with the sphere-to-sphere contact points, in which elements can become unmanageably small and have high aspect ratios. One way to mitigate this problem is to bridge the contact point with a small fillet or cylinder [56]. By removing the small regions of fluid near the contact points, a smooth, non-skewed surface mesh can be obtained, as shown in A.1b.

It quickly becomes challenging to generate this type of mesh for more than a few spheres by eye using a graphical user interface because of the large number of geometric entities and high surface area-to-volume ratio. The CFPL researchers realized that their problem was ripe for automation using the Pointwise/Glyph scripting capabilities. The geometry creation and mesh generation processes were parameterized easily, and a generalized algorithm was developed to handle arbitrary sphere positions and boundary types. Glyph, which provides the full functionality of Pointwise as an extension to the TCL scripting language, made the transition to automation surprisingly straightforward. In addition, a TK user interface, shown in Figure A.2, was designed to minimize the required user interaction. The CFPL group is now generating quality meshes for large scale sphere packs and has decreased their interaction time from days to minutes.

The flow simulations are performed using parallel research codes. Simulations utilizing up to 350 processors have been run on the NSF's TeraGrid supercomputers over a large range of Reynolds numbers. The high resolution meshes are able
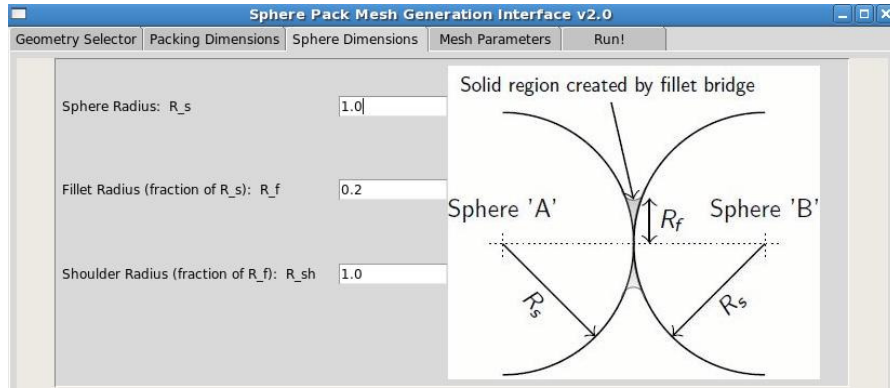
Figure A.2: The TK user interface for the Glyph scripts used for mesh automation.

to capture the important unsteady dynamics and vortex shedding at the pore scale. This is shown in Figure A, where the Lambda-2 vortex detection criteria is visualized as an isosurface for a Reynolds number of 529. Contour colors show regions of high and low pressure for a single instant in time.
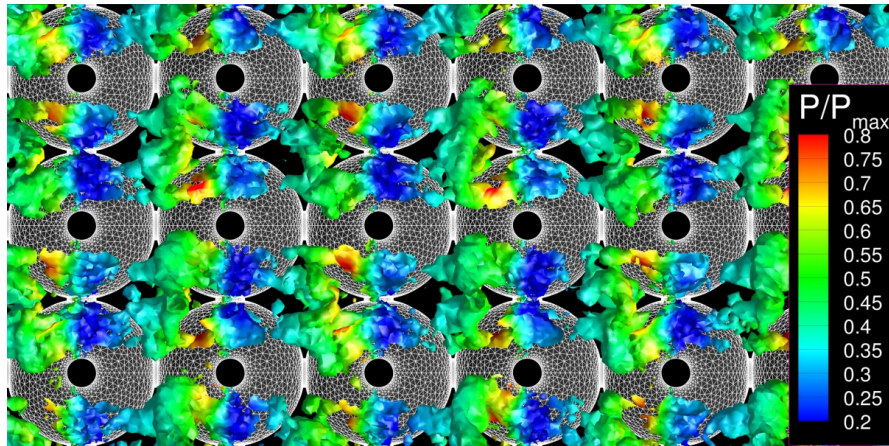
Figure A.3: Flow from left to right through a simple cubic packing of spheres at a Reynolds number of 529. A single isosurface of the Lambda-2 vortex detection criteria is shown. Color contours show the local pressure variation.

Appendix B: Parallelization of the flow map construction scheme

The "single tiered unidirectional" reconstruction procedure proposed by [9] is used to compose the time T flow maps from the sequence N time h flow map substeps. Here, we describe the parallelization of this step in our distributed memory, unstructured grid flow solver. Let t be the current simulation time which is a multiple of the flow map sub-step, h ($\mod (t, h) = 0$), and assume that we are interested in constructing the time T forward flow map, $\boldsymbol{\Phi}_{t-T}^{t}$ from the substeps which have already been written to the hard disk. Let t_C be the current *construction time*. We initialize $t_C = t - T$, so that our construction begins at the beginning of our forward time interval, $[t - T, \ t]$, and initialize a temporary construction of the flow map, $\boldsymbol{\Phi}_{t-T}^{t_C} = \mathbf{x}_{cv}$ everywhere. We then perform the following steps in a loop from $n = 1$ to $n = N$:

- Each processor reads its own time h flow map, $\boldsymbol{\Phi}_{t_C}^{t_C+h} = \boldsymbol{\Phi}_{t-T+(n-1)h}^{t-T+nh}$ into memory from the binary file located on the hard disk.

- The goal is to update the flow map construction, $\boldsymbol{\Phi}_{t-T}^{t_C}$, to account for the current sub-step. This involves interpolating the values of $\boldsymbol{\Phi}_{t_C}^{t_C+h}$ which is known at the grid points, to the non-mesh points of the current flow map construction, $\boldsymbol{\Phi}_{t-T}^{t_C}$

$$\boldsymbol{\Phi}_{t-T}^{t_C+h} = \mathcal{I}\boldsymbol{\Phi}_{t_C}^{t_C+h} \circ \boldsymbol{\Phi}_{t-T}^{t_C} \tag{B.1}$$

It is likely that the non-mesh points of the current construction will lie across processor boundaries, meaning a coordinated search and retrieve operation is required by all processors to perform this update. To do this, each processor **sends** a query to all processors which could potentially contain each non-mesh point, $\boldsymbol{\Phi}_{t-T}^{t_C}$. This is determined simply by comparing each non-mesh point to the minimum bounding box of the grid partition belonging to all other processors.

- Each processor **<u>receives</u>** a list of non-mesh points, which may lie within its grid partition, to search for. These are the points where the value of the current flow map sub-step is needed by another processor. A bounding box Oct-Tree search algorithm is used to efficiently determine the cell (if any) which contains each of these points. If found, the value of $\mathbf{\Phi}_{t_C}^{t_C+h}$ is interpolated to the non-mesh point using the interpolation kernel described in section 4.3.1. These values of are sorted and **<u>sent</u>** back to the processor which requested each search.

- Each processor **<u>receives</u>** values of the updated flow map $\mathbf{\Phi}_{t-T}^{t_C+h}$ from potentially all other processors.

- The construction time is advanced to $t_C = t_C + h$.

At the end of this loop, $\mathbf{\Phi}_{t-T}^{t_C}$ is an approximation of time T flow map $\mathbf{\Phi}_{t-T}^{t}$. The procedure can be repeated to compute the backward time flow map, $\mathbf{\Phi}_{t}^{t-T}$, by starting at $t_C = t$ and advancing the construction time in the negative time direction.

As the number of sub-steps, N, increases this portion of the algorithm can become computationally expensive due to the relatively naive search procedure for non-mesh points in our unstructured grids. In the case of a shared memory structured grid flow solver, these steps could be simplified significantly. Nonetheless, we find that the total expense of this step remains small (see section 4.4.6) relative to the simulation time for moderate values of N.

Appendix C: Turbulent Flow in a FCC Unit Cell

In this section, preliminary results are presented for turbulent flow through a face centered cubic (FCC) unit cell. This configuration is shown in figure C.1 and is one of two (HCP being the other) possible homogeneous sphere packings with the lowest possible porosity, $\epsilon = 0.26$. It consists of 6 half spheres located on each face, and 8 quarter spheres located at the corners. As the flow travels through the unit cell in the X direction, it experiences dramatic accelerations and decelerations due to expansions and contractions in the pore cross sections. A decade ago, Hill and Koch [48] studied this configuration in detail for moderate Reynolds numbers[1] in the unsteady and weakly turbulent regimes ($81 < Re < 286$) using a Lattice Boltzmann approach. They observed a very early transition to unsteady flow at $Re = 81$, and a transition to weakly turbulent flow at $Re \approx 220$. At the time, their study pushed the envelope of available computational power, and they were unable to perform systematic grid convergence studies, or study flows in the fully turbulent regime. This study aims to pick up where they left off, by demonstrating grid convergence requirements, and explore the characteristics of three dimensional turbulent flow in the FCC configuration.

The fictitious domain approach outlined in Chapter 2 is used to compute the flow field. Accordingly, a uniform Cartesian grid is generated, as well as marker points that are used to locate the solid surface. The flow is then started from rest, and driven in the positive X direction by a constant pressure gradient. This pressure gradient is defined for a specific Reynolds number from the observations of Hill & Koch. They developed the following correlation for drag force on the

---

[1]Note the Reynolds number definition used by Hill & Koch in their series of papers on porous media is different from the one used in this dissertation. Specifically, they define $Re = UD/\nu$, where U is the superficial velocity. For the sake of consistency, their values have been converted for this discussion to pore Reynolds number definition of Dybbs & Edwards.
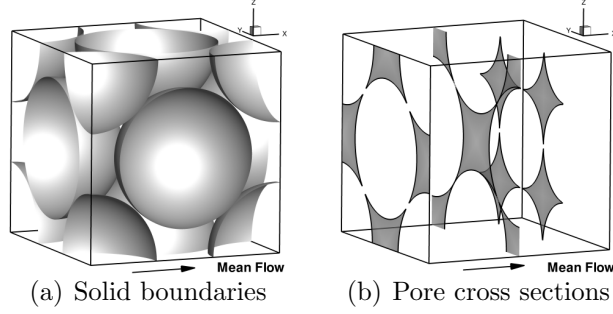
(a) Solid boundaries   (b) Pore cross sections

Figure C.1: The face centered cubic (FCC) unit cell configuration.

spheres:

$$F = 462 + 3.65Re \tag{C.1}$$

This can be related to the pressure gradient as:

$$\overline{\nabla P} = 18\mu(1 - \epsilon)F\overline{U}/D^2 \tag{C.2}$$

Where $\overline{U}$ is the time averaged *superficial* velocity, q/A. Three Reynolds numbers, $Re = 288, 530, 950$ have been chosen based on prior work and available computational resources: The highest Reynolds number Hill and Koch were able to simulate was $Re = 288$, and present computational resources will likely restrict our study to $Re \lesssim 1000$.

C.1 Grid Requirements

In order to estimate grid resolution requirements for this case, an analogy is made to turbulent channel flow. Standard practice for DNS studies of channel flow is to place the first grid cell within $\delta^+ = 1$ of the solid boundary, so that the viscous sublayer is resolved by the simulation. In channel flow there is no form drag, and δ^+ can be estimated directly if the applied pressure gradient is known. In random sphere packings, there is significant form drag, flow separation occurs, and the local boundary layer thickness is difficult (if not impossible) to predict a-priori.

The estimates made here can be viewed as a starting point to understand grid requirements for DNS in packed beds.

The pressure drop per unit length, $\frac{\Delta P}{L}$, through an arrangement of spheres with porosity ϵ and specific surface $\mathcal{S} = \frac{6}{D}$ is balanced by the total stress on the sphere surface, $\langle \tau \rangle$.

$$\langle \tau \rangle = \frac{\Delta P}{L} \frac{\epsilon}{\mathcal{S}(1-\epsilon)} \tag{C.3}$$

Dividing both sides of equation C.3 by ρU_p^2, and introducing the modified friction factor, $f = \frac{\Delta P}{\rho U_p^2} \frac{D}{L} \frac{\epsilon}{1-\epsilon}$ [21], we obtain

$$\langle \tau \rangle = f \frac{\rho U_p^2}{6} \tag{C.4}$$

For DNS, we want the wall coordinate, $\delta^+ = \delta u_\tau / \nu = 1$, where $u_\tau = \sqrt{\langle \tau_{ws} \rangle / \rho}$. In this relationship $\langle \tau_{ws} \rangle = \alpha \langle \tau \rangle$ is the area averaged wall shear stress, and δ is dimensional distance from boundary. The parameter α relates the area average wall shear stress to the area averaged surface stress, $\langle \tau \rangle$, and may be conservatively estimated as $\alpha = 0.15$ [113, 35] for random and arranged packed beds. The normalized grid spacing for $\delta^+ = 1$ can then be expressed in terms of Reynolds number, porosity, and the friction factor.

$$\frac{D}{\delta} = Re \cdot \sqrt{\frac{\alpha}{6} f} \cdot \frac{1-\epsilon}{\epsilon} \tag{C.5}$$

Where, Dybbs & Edward's'[21] definition of the pore Reynolds number has been used.

$$Re = \left(\frac{\rho_\beta U_p D_{sp}}{\mu_\beta} \right) \left(\frac{\epsilon}{1-\epsilon} \right) \tag{C.6}$$

Ergun's correlation [24], $f = \frac{150}{Re} + 1.75$, can be used to determine an estimate for the friction factor in porous media, allowing an estimate of the required grid spacing to be made with equation C.5. In Figure C.1, the required resolution for the FCC unit cell has been plotted as a function of Reynolds number for a grid resolution corresponding to $\delta^+ = 1$ and $\delta^+ = 3$. This estimate shows that a grid
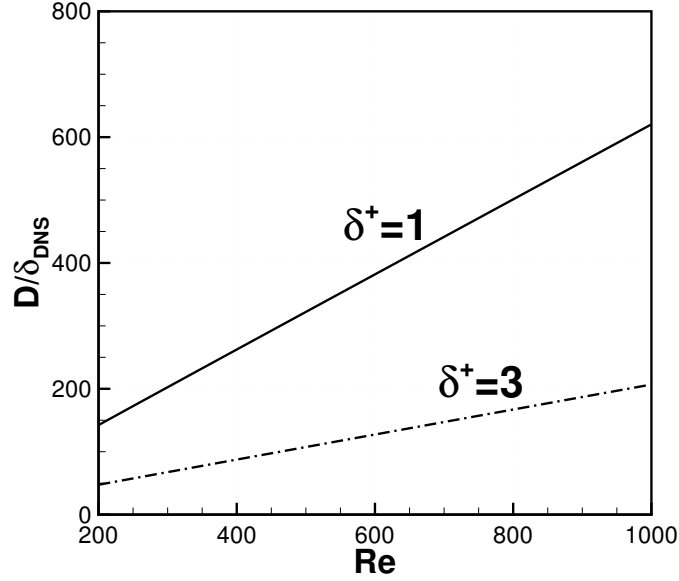
spacing of $D/\delta = 200$ will provide "near DNS" resolution for Reynolds numbers between 200 and 1000.

C.2 Preliminary Results

As a first step, grid convergence studies were performed at each Reynolds number to determine the grid resolution required to capture all space and time scales of the turbulence. Both integrated quantities such as the bulk flow rate, and local quantities such as time average velocity and turbulent kinetic energy (TKE) distributions in the porespace are being monitored to determine the grid requirements for DNS at each Reynolds number. A sample measurement of the time averaged TKE distribution made at different grid resolutions is shown in figure C.2 for $Re = 950$. As can be seen here, the major features of the TKE distribution are captured at all grid refinement levels but the magnitude of TKE is not consistent between any grids, suggesting that all turbulent scales are not yet resolved. Based on the present available results, *rough* estimates for grid requirements are summarized in table C.1 for each Reynolds number to be considered. These will be
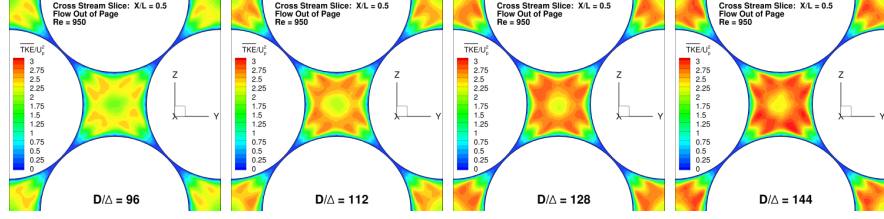
Figure C.2: Convergence of Mean streamwise velocity and turbulent kinetic energy with increased grid resolution

Table C.1: Grid size estimates for DNS of turbulent flow through the FCC unit cell

Re	Maximum Current D/Δ	Estimated Required D/Δ	Estimated Grid Size
288	176	176	15M
530	144	250	44M
950	176	250+	44M+

adjusted as needed based on future quantitative comparison of the results on grids with resolution up to $D/\Delta = 250$.

The preliminary results available do allow for a qualitative comparison of the porescale structure of turbulence at the three different Reynolds numbers considered. Figure C.3 shows instantaneous isosurfaces of the swirling strength vortex detection criteria, λ_{ci} [115, 14] for each of the three Reynolds numbers. In each case, the isovalue is chosen to be roughly 25% of the maximum values in the domain. Qualitatively, one can see a larger scale separation at higher Reynolds number. Also, the vortical structures are less isotropic at the lower Reynolds numbers. At these flow rates, there seems to be a preference for the alignment of the vortex tubes in the mean flow direction, meaning that there is significant helicity in the system (vorticity alignment with the mean flow). This is not as evident at the highest Reynolds number (950). Further quantitative analysis of the relationship between helicity and Reynolds number is warranted. The scale separation of the porescale flow features is also shown in figure C.4, where an instantaneous vector field in shown in a cross-stream slice on top of the λ_{ci} contours for each Reynolds number.

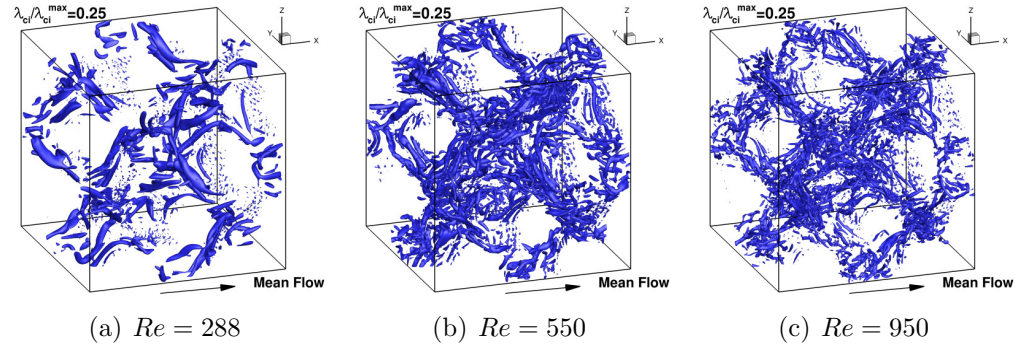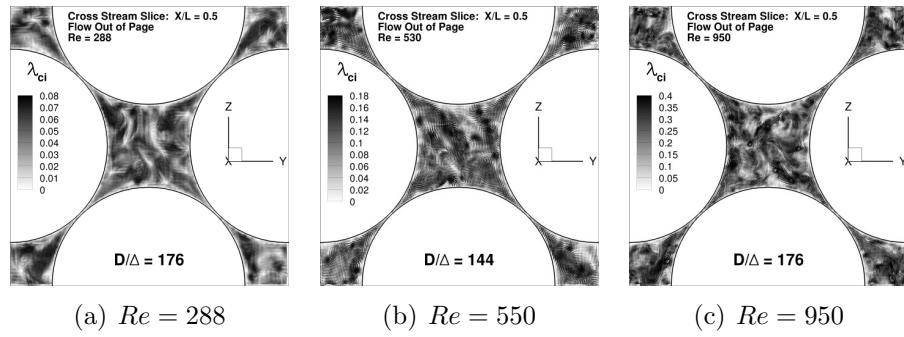(a) $Re = 288$ (b) $Re = 550$ (c) $Re = 950$

Figure C.3: Isosurfaces of λ_{ci}

(a) $Re = 288$ (b) $Re = 550$ (c) $Re = 950$

Figure C.4: Midplane slice showing instantaneous vector field. Greyscale contours show swirling strength.

Bibliography

[1] P. Andrigo, R. Bagatin, and G. Pagani. Fixed bed reactors. *Catalysis today*, 52(2-3):197–221, 1999.

[2] S.V. Apte, M. Martin, and N.A. Patankar. A numerical method for fully resolved simulation (FRS) of rigid particle-flow interactions in complex flows. *Journal of Computational Physics*, 228(8):2712–2738, 2009.

[3] J.K. Arthur, D.W. Ruth, and M.F. Tachie. Piv measurements of flow through a model porous medium with varying boundary conditions. *Journal of Fluid Mechanics*, 629(1):343–374, 2009.

[4] T. Atmakidis and E.Y. Kenig. CFD-based analysis of the wall effect on the pressure drop in packed beds with moderate tube/particle diameter ratios in the laminar flow regime. *Chemical Engineering Journal*, 155(1):404–410, 2009.

[5] Kenneth E Bencala. A perspective on stream-catchment connections. *Journal of the North American Benthological Society*, pages 44–47, 1993.

[6] G. Berkooz, P. Holmes, and J.L. Lumley. The proper orthogonal decomposition in the analysis of turbulent flows. *Annual Review of Fluid Mechanics*, 25(1):539–575, 1993.

[7] FJ Beron-Vera, MJ Olascoaga, and GJ Goni. Oceanic mesoscale eddies as revealed by Lagrangian coherent structures. *Geophysical Research Letters*, 35(12):L12603, 2008.

[8] M.J. Blunt. Flow in porous media–pore-network models and multiphase flow. *Current opinion in colloid & interface science*, 6(3):197–207, 2001.

[9] S.L. Brunton and C.W. Rowley. Fast computation of finite-time Lyapunov exponent fields for unsteady flows. *Chaos*, 20(1):017503, 2010.

[10] J. Cadafalch, CD Pérez-Segarra, R. Consul, and A. Oliva. Verification of finite volume computations on steady-state fluid flow and heat transfer. *Journal of Fluids Engineering*, 124:11–21, 2002.

[11] HPA Calis, J. Nijenhuis, BC Paikert, FM Dautzenberg, and CM Van Den Bleek. CFD modelling and experimental validation of pressure drop and flow profile in a novel structured catalytic reactor packing. *Chemical Engineering Science*, 56(4):1713–1720, 2001.

[12] M.B. Cardenas and J.L. Wilson. Dunes, turbulent eddies, and interfacial exchange with permeable sediments. *Water Resources Research*, 43(8):W08412, 2007.

[13] I.B. Celik, U. Ghia, P.J. Roache, et al. Procedure for estimation and reporting of uncertainty due to discretization in CFD applications. *Journal of Fluids Engineering-Transactions of the ASME*, 130(7):078001, 2008.

[14] P. Chakraborty, S. Balachandar, and R.J. Adrian. On the relationships between local vortex identification schemes. *Journal of Fluid Mechanics*, 535(1):189–214, 2005.

[15] I.C. Christov, J.M. Ottino, and R.M. Lueptow. From streamline jumping to strange eigenmodes: Bridging the Lagrangian and Eulerian pictures of the kinematics of mixing in granular flows. *Physics of Fluids*, 23:103302, 2011.

[16] C. Conti, D. Rossinelli, and P. Koumoutsakos. GPU and APU computations of finite time Lyapunov exponent fields. *Journal of Computational Physics*, 231(5):2229–2244, 2012.

[17] A.G. Dixon and M. Nijemeisland. CFD as a design tool for fixed-bed reactors. *Ind. eng. chem. res*, 40(23):5246–5254, 2001.

[18] A.G. Dixon, M. Nijemeisland, and E.H. Stitt. Packed tubular reactor modeling and catalyst design using computational fluid dynamics. *Advances in Chemical Engineering*, 31:307–389, 2006.

[19] H. Dong, M. Bozkurttas, R. Mittal, P. Madden, and GV Lauder. Computational modelling and analysis of the hydrodynamics of a highly deformable fish pectoral fin. *Journal of Fluid Mechanics*, 645:345–373, 2010.

[20] H. Dütsch, F. Durst, S. Becker, and H. Lienhart. Low-Reynolds-number flow around an oscillating circular cylinder at low Keulegan–Carpenter numbers. *Journal of Fluid Mechanics*, 360:249–271, 1998.

[21] A. Dybbs and RV Edwards. A new look at porous media fluid mechanics–Darcy to turbulent. *Fundamentals of Transport Phenomena in Porous Media*, pages 199–254, 1984.

[22] L. Eça, M. Hoekstra, and P.J. Roache. Verification of calculations: An overview of the Lisbon workshop. In *23rd AIAA Applied Aerodynamics Conference*, pages 1–12, 2005.

[23] B. Eisfeld and K. Schnitzlein. The influence of confining walls on the pressure drop in packed beds. *Chemical Engineering Science*, 56(14):4321–4329, 2001.

[24] S. Ergun. Fluid flow through packed columns. *Chem. Eng. Prog*, 48(2):89–94, 1952.

[25] R. Falgout and U. Yang. hypre: A library of high performance preconditioners. *Computational Science ICCS 2002*, pages 632–641, 2002.

[26] M. Farazmand and G. Haller. Computing Lagrangian coherent structures from their variational theory. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 22(1):013128–013128, 2012.

[27] M. Farazmand and G. Haller. Erratum and addendum to "A variational theory of hyperbolic Lagrangian coherent structures" [Physica D 240 (2011) 574–598]. *Physica D: Nonlinear Phenomena*, 241:439–431, 2012.

[28] J. Finn, E. Shams, and S.V. Apte. Modeling and simulation of multiple bubble entrainment and interactions with two dimensional vortical flows. *Physics of Fluids*, 23(2):3301, 2011.

[29] Justin Finn and S.V. Apte. Characteristics of vortical structures in random and arranged packed beds of spheres. *Proceedings of the ASME Fluids Engineering Summer Meeting*, (FEDSM2012-72354), 2012.

[30] Justin Finn and S.V. Apte. Relative performance of body-fitted and fictitious-domain simulations of flow through porous media. *ASME Fluids Engineering Summer Meeting*, FEDSM2012-72354(FEDSM2012-72355), 2012.

[31] C. Garth, F. Gerhardt, X. Tricoche, and H. Hagen. Efficient computation and visualization of coherent structures in fluid flow applications. *Visualization and Computer Graphics, IEEE Transactions on*, 13(6):1464–1471, 2007.

[32] R. Glowinski, TW Pan, TI Hesla, DD Joseph, and J. Periaux. A fictitious domain approach to the direct numerical simulation of incompressible viscous flow past moving rigid bodies: application to particulate flow. *Journal of Computational Physics*, 169(2):363–426, 2001.

[33] M.A. Green, C.W. Rowley, and A.J. Smits. The unsteady three-dimensional wake produced by a trapezoidal pitching panel. *Journal of Fluid Mechanics*, 685(1):117–145, 2011.

[34] A. Guardo, M. Coussirat, F. Recasens, MA Larrayoz, and X. Escaler. CFD study on particle-to-fluid heat transfer in fixed bed reactors: Convective heat transfer at low and high pressure. *Chemical Engineering Science*, 61(13):4341–4353, 2006.

[35] P.R. Gunjal, V.V. Ranade, and R.V. Chaudhari. Computational study of a single-phase flow in packed beds of spheres. *AIChE Journal*, 51(2):365–378, 2005.

[36] AM Guzman and CH Amon. Dynamical flow characterization of transitional and chaotic regimes in converging–diverging channels. *Journal of Fluid Mechanics*, 321(1):25–57, 1996.

[37] S. Haeri and JS Shrimpton. On the application of immersed boundary, fictitious domain and body-conformal mesh methods to many particle multiphase flows. *International Journal of Multiphase Flow*, 40:38–55, 2012.

[38] G. Haller. Distinguished material surfaces and coherent structures in three-dimensional fluid flows. *Physica D: Nonlinear Phenomena*, 149(4):248–277, 2001.

[39] G. Haller. A variational theory of hyperbolic Lagrangian coherent structures. *Physica. D*, 240(7):574–598, 2011.

[40] G. Haller and F.J. Beron-Vera. Geodesic theory of transport barriers in two-dimensional flows. *Physica D Nonlinear Phenomena*, 241:1680–1702, 2012.

[41] G. Haller and T. Sapsis. Lagrangian coherent structures and the smallest finite-time Lyapunov exponent. *Chaos*, 21(2):023115, 2011.

[42] G. Haller and G. Yuan. Lagrangian coherent structures and mixing in two-dimensional turbulence. *Physica D: Nonlinear Phenomena*, 147(3-4):352–370, 2000.

[43] F. Ham and G. Iaccarino. Energy conservation in collocated discretization schemes on unstructured meshes. *Center for Turbulence Research Annual Research Briefs*, pages 3–14, 2004.

[44] Frank Ham, Sourabh Apte, George Constantinescuf Herrmann, Krishnan Mahesh, and Parviz Moin. Unstructured LES of reacting multiphase flows in realistic gas turbine combustors. *Annual Research Briefs-2003*, page 139, 2003.

[45] Y.A. Hassan. Large eddy simulation in pebble bed gas cooled core reactors. *Nuclear Engineering and Design*, 238(3):530–537, 2008.

[46] S.E. Hieber and P. Koumoutsakos. A Lagrangian particle level set method. *Journal of Computational Physics*, 210(1):342–367, 2005.

[47] R.J. Hill and D.L. Koch. Moderate-Reynolds-number flow in a wall-bounded porous medium. *Journal of Fluid Mechanics*, 453:315–344, 2002.

[48] R.J. Hill and D.L. Koch. The transition from steady to weakly turbulent flow in a close-packed ordered array of spheres. *Journal of Fluid Mechanics*, 465:59–97, 2002.

[49] R.J. Hill, D.L. Koch, and A.J.C. Ladd. The first effects of fluid inertia on flows in ordered and random arrays of spheres. *Journal of Fluid Mechanics*, 448:213–241, 2001.

[50] R.J. Hill, D.L. Koch, and A.J.C. Ladd. Moderate-reynolds-number flows in ordered and random arrays of spheres. *Journal of Fluid Mechanics*, 448:243–278, 2001.

[51] NA Horton and D. Pokrajac. Onset of turbulence in a regular porous medium: An experimental study. *Physics of fluids*, 21:045104, 2009.

[52] Alice Y. L. Huang, Michelle Y. F. Huang, Herve Capert, and Rong-Her Chen. Optical measurements of pore geometry and fluid velocity in a bed of irregularly packed spheres. *Experiments in Fluids*, 45:309–321, 2008.

[53] J. Jeong and F. Hussain. On the identification of a vortex. *Journal of Fluid Mechanics*, 285(1):69–94, 1995.

[54] D. Kim and H. Choi. A second-order time-accurate finite volume method for unsteady incompressible flow on hybrid unstructured grids. *Journal of Computational Physics*, 162(2):411–428, 2000.

[55] Donald L Koch and John F Brady. Dispersion in fixed beds. *Journal of Fluid Mechanics*, 154(1):399–427, 1985.

[56] M. Kuroki, S. Ookawara, and K. Ogawa. A high-fidelity CFD model of methane steam reforming in a packed bed reactor. *Journal of Chemical Engineering of Japan*, 42(Supplement.):73–78, 2009.

[57] YE Kutsovsky, LE Scriven, HT Davis, and BE Hammer. Nmr imaging of velocity profiles and velocity distributions in bead packs. *Physics of Fluids*, 8:863, 1996.

[58] F. Lekien and J. Marsden. Tricubic interpolation in three dimensions. *International Journal for Numerical Methods in Engineering*, 63(3):455–471, 2005.

[59] F. Lekien and S.D. Ross. The computation of finite-time Lyapunov exponents on unstructured meshes and for non-Euclidean manifolds. *Chaos*, 20(1):017505–1, 2010.

[60] F. Lekien, S.C. Shadden, and J.E. Marsden. Lagrangian coherent structures in n-dimensional systems. *Journal of Mathematical Physics*, 48:065404, 2007.

[61] S. Leung. An Eulerian approach for computing the finite time Lyapunov exponent. *Journal of Computational Physics*, 230(9):3500–3524, 2011.

[62] D. Lipinski and K. Mohseni. A ridge tracking algorithm and error estimate for efficient computation of Lagrangian coherent structures. *Chaos*, 20(1):017504–1, 2010.

[63] P. Magnico. Hydrodynamic and transport properties of packed beds in small tube-to-sphere diameter ratio: pore scale simulation using an eulerian and a lagrangian approach. *Chemical Engineering Science*, 58(22):5005–5024, 2003.

[64] P. Magnico. Pore-scale simulations of unsteady flow and heat transfer in tubular fixed beds. *AIChE Journal*, 55(4):849–867, 2009.

[65] K. Mahesh, G. Constantinescu, and P. Moin. A numerical method for large-eddy simulation in complex geometries. *Journal of Computational Physics*, 197(1):215–240, 2004.

[66] J. Mahoney, D. Bargteil, M. Kingsbury, K. Mitchell, and T. Solomon. Invariant barriers to reactive front propagation in fluid flows. *EPL (Europhysics Letters)*, 98(4):44005, 2012.

[67] I. Malico and P.J.S.A. Ferreira de Sousa. Modeling the pore level fluid flow in porous media using the immersed boundary method. *Numerical Analysis of Heat and Mass Transfer in Porous Media*, pages 229–251, 2012.

[68] M. Mathur, G. Haller, T. Peacock, J.E. Ruppert-Felsot, and H.L. Swinney. Uncovering the Lagrangian skeleton of turbulence. *Physical Review Letters*, 98(14):144502, 2007.

[69] R. McDermott and SB Pope. The parabolic edge reconstruction method (PERM) for Lagrangian particle advection. *Journal of Computational Physics*, 227(11):5447–5491, 2008.

[70] P. Miron, J. Vétel, A. Garon, M. Delfour, and M.E. Hassan. Anisotropic mesh adaptation on Lagrangian coherent structures. *Journal of Computational Physics*, 2012.

[71] P. Moin and SV Apte. Large-eddy simulation of realistic gas turbine-combustors. *AIAA Journal*, 44(4):698–708, 2006.

[72] G.E. Mueller. Radial void fraction distributions in randomly packed fixed beds of uniformly sized spheres in cylindrical containers. *Powder technology*, 72(3):269–275, 1992.

[73] G.E. Mueller. Numerically packing spheres in cylinders. *Powder technology*, 159(2):105–110, 2005.

[74] B.L. Nelson. *Scaling analysis for the pebble bed of the very high temperature gas-cooled reactor thermal hydraulic test facility*. PhD thesis, Oregon State University, 2009.

[75] M. Nijemeisland and A.G. Dixon. CFD study of fluid flow and wall heat transfer in a fixed bed of spheres. *AIChE journal*, 50(5):906–921, 2004.

[76] C. O'Farrell and J.O. Dabiri. A Lagrangian approach to identifying vortex pinch-off. *Chaos*, 20(1):017513–017513, 2010.

[77] J.M. Ottino. *The kinematics of mixing: stretching, chaos, and transport.* Cambridge University Press, 1989.

[78] S. Ovaysi and M. Piri. Direct pore-level modeling of incompressible fluid flow in porous media. *Journal of Computational Physics*, 229(19):7456–7476, 2010.

[79] N. Park and K. Mahesh. Numerical and modeling issues in LES of compressible turbulence on unstructured grids. *AIAA Paper*, 722(2007):85, 2007.

[80] NA Patankar and DD Joseph. Modeling and numerical simulation of particulate flows by the Eulerian-Lagrangian approach. *International Journal of Multiphase Flow*, 27(10):1659–1684, 2001.

[81] V.A. Patil and J.A. Liburdy. Optical measurement uncertainties due to refractive index mismatch for flow in porous media. *Experiments in Fluids*, 53(5):1453–1468, 2012.

[82] Vishal Patil. *Study of creeping, inertial and turbulent flow regimes in porous media using particle image velocimetry.* PhD thesis, Oregon State University, 2012.

[83] T. Peacock and J. Dabiri. Introduction to focus issue: Lagrangian coherent structures. *Chaos*, 20(1):017501, 2010.

[84] J. Peng and JO Dabiri. Transport of inertial particles by Lagrangian coherent structures: application to predator-prey interaction in jellyfish feeding. *Journal of Fluid Mechanics*, 623:75, 2009.

[85] C.S. Peskin. *Flow patterns around heart valves: A digital computer method for solving the equations of motion.* PhD thesis, Sue Golding Graduate Division of Medical Sciences, Albert Einstein College of Medicine, Yeshiva University, 1972.

[86] C.S. Peskin. Numerical analysis of blood flow in the heart. *Journal of Computational Physics*, 25(3):220–252, 1977.

[87] M. Pilotti. Viscous flow in three-dimensional reconstructed porous media. *International Journal for Numerical and Analytical Methods in Geomechanics*, 27(8):633–649, 2003.

[88] M. Rashidi, L. Peurrung, AFB Tompson, and TJ Kulp. Experimental analysis of pore-scale flow and transport in porous media. *Advances in Water Resources*, 19(3):163–180, 1996.

[89] PJ Roache. Perspective: a method for uniform reporting of grid refinement studies. *Journal of Fluids Engineering-Transactions of the ASME*, 116(3):405–413, 1994.

[90] P.J. Roache. Quantification of uncertainty in computational fluid dynamics. *Annual Review of Fluid Mechanics*, 29(1):123–160, 1997.

[91] P.J. Roache. Conservatism of the grid convergence index in finite volume computations on steady-state fluid flow and heat transfer. *Journal of Fluids Engineering-Transactions of the ASME*, 125(4):731–732, 2003.

[92] D.J. Robbins, M.S. El-Bachir, L.F. Gladden, R.S. Cant, and E. von Harbou. CFD modeling of single-phase flow in a packed bed with MRI validation. *AIChE Journal*, 2012.

[93] Alexandre M Roma, Charles S Peskin, and Marsha J Berger. An adaptive version of the immersed boundary method. *Journal of computational physics*, 153(2):509–534, 1999.

[94] R. Samelson. Lagrangian motion, coherent structures, and lines of persistent material strain. *Annual Review of Marine Science*, 5(1), 2012.

[95] T. Sapsis, G. Haller, et al. Clustering criterion for inertial particles in two-dimensional time-periodic and three-dimensional steady flows. *Chaos*, 20(1):7515, 2010.

[96] D Seguin, A Montillet, and J Comiti. Experimental characterisation of flow regimes in various porous media-i: Limit of laminar flow regime. *Chemical engineering science*, 53(21):3751–3761, 1998.

[97] S.C. Shadden, J.O. Dabiri, and J.E. Marsden. Lagrangian analysis of fluid transport in empirical vortex ring flows. *Physics of Fluids*, 18:047105, 2006.

[98] S.C. Shadden, F. Lekien, and J.E. Marsden. Definition and properties of Lagrangian coherent structures from finite-time Lyapunov exponents in two-dimensional aperiodic flows. *Physica D: Nonlinear Phenomena*, 212(3-4):271–304, 2005.

[99] E. Shams, J. Finn, and SV Apte. A numerical scheme for Euler-Lagrange simulation of bubbly flows in complex systems. *International Journal for Numerical Methods in Fluids*, 67:1865–1898, 2010.

[100] N. Sharma and N.A. Patankar. A fast computation technique for the direct numerical simulation of rigid particulate flows. *Journal of Computational Physics*, 205(2):439–457, 2005.

[101] J.A. Simeonov and J. Calantoni. A pressure boundary integral method for direct fluid-particle simulations on Cartesian grids. *Journal of Computational Physics*, 230(5):1749–1765, 2011.

[102] P.K. Smolarkiewicz and C. Larrabee Winter. Pores resolving simulation of Darcy flows. *Journal of Computational Physics*, 229(9):3121–3133, 2010.

[103] TH Solomon and JP Gollub. Chaotic particle transport in time-dependent Rayleigh-Bénard convection. *Physical Review A*, 38(12):6280, 1988.

[104] G. Sridhar and J. Katz. Effect of entrained bubbles on the structure of vortex rings. *Journal of Fluid Mechanics*, 397(1):171–202, 1999.

[105] A. Staniforth and J. Côté. Semi-Lagrangian integration schemes for atmospheric models- a review. *Monthly Weather Review*, 119(9):2206–2223, 1991.

[106] M. Stohr, K. Roth, and B. Jahne. Measurement of 3D pore-scale flow in index-matched porous media. *Experiments in Fluids*, 35(2):159–166, 2003.

[107] J. Strain. Semi-Lagrangian methods for level set equations. *Journal of Computational Physics*, 151(2):498–533, 1999.

[108] T. Suekane, Y. Yokouchi, and S. Hirai. Inertial flow structures in a simple-packed bed of spheres. *AIChE Journal*, 49(1):10–17, 2003.

[109] GI Taylor. On the decay of vortices in a viscous fluid. *Phil. Mag*, 46(6):671–674, 1923.

[110] Hendrik Tennekes and John L Lumley. *A first course in turbulence.* MIT press, 1972.

[111] M. Uhlmann. An immersed boundary method with direct forcing for the simulation of particulate flows. *Journal of Computational Physics*, 209(2):448–476, 2005.

[112] A. Vargas, R. Mittal, and H. Dong. A computational study of the aerodynamic performance of a dragonfly wing section in gliding flight. *Bioinspiration & Biomimetics*, 3:026004, 2008.

[113] C.A. Wentz Jr and G. Thodos. Total and form drag friction factors for the turbulent flow of air through packed and distended beds of spheres. *AIChE Journal*, 9(3):358–361, 1963.

[114] B.D. Wood. Inertial effects in dispersion in porous media. *Water Resources Research*, 43(12):12, 2007.

[115] J. Zhou, RJ Adrian, S. Balachandar, and TM Kendall. Mechanisms for generating coherent packets of hairpin vortices in channel flow. *Journal of Fluid Mechanics*, 387(1):353–396, 1999.

[116] L. Zhu and C.S. Peskin. Simulation of a flapping flexible filament in a flowing soap film by the immersed boundary method. *Journal of Computational Physics*, 179(2):452–468, 2002.

[117] A.A. Zick and G.M. Homsy. Stokes flow through periodic arrays of spheres. *Journal of Fluid Mechanics*, 115:13–26, 1982.