## Numerical Simulation of Non-Spherical Particle Interactions in Quiescent and Turbulent Fluids



#### Jacob Paul Anderson

The University of Leeds EPSRC Centre for Doctoral Training in Fluid Dynamics School of Computer Science

Submitted in accordance with the requirements of

Doctor of Philosophy

December 2024

## Declaration

The candidate confirms that the work submitted is his own, except where work which has formed part of jointly authored publications has been included. The contribution of the candidate and the other authors to this work has been explicitly indicated below. The candidate confirms that appropriate credit has been given within the thesis where reference has been made to the work of others.

#### Methodology, Results Chapter 1 and Results Chapter 2:

- Anderson, J.P., Mortimer, L.F., Hunter, T.N., Peakall, J., Fairweather, M. : *Numerical Simulation of the Agglomeration Behaviour of Spheroidal Particle Pairs in Chaotic Flows.* Flow, Turbulence and Combustion, 2025. Vol. 114, pp. 941-965. *Journal Article.*
- Anderson, J.P., Mortimer, L.F., Hunter, T.N., Peakall, J., Fairweather, M.: *Agglomeration and Collision Behaviour of Non-Spherical Particles in Turbulence*. In: Tenth International Symposium on Turbulence, Heat and Mass Transfer, 2023, Rome. *Conference*.
- Purvis, J., Mortimer, L.F., Anderson, J.P., Fairweather, M.: *Fully resolved binary particle interactions in isotropic turbulence using hard and soft sphere collision models*. In: Tenth International Symposium on Turbulence, Heat and Mass Transfer, 2023, Rome. *Conference*.
- Anderson, J.P., Mortimer, L.F, Hunter, T.N. Peakall, J., Fairweather, M.: Non-Spherical Particle Interactions in Isotropic Turbulence Using an Immersed Boundary Approach. In: Proceedings of the 14th International ERCOFTAC Symposium on Engineering Turbulence Modelling and Measurements – ETMM14, 2023, Barcelona. Conference.

For publications that were co-authored, the other authors contributed to weekly supervisory meetings, providing guidance and sharing ideas for progression – except for where the present author's name appears third. In this case, the present author provided the simulation code, and the listed first author generated the simulation results under the guidance of all others including the present author. In the cases where the present author appears first listed, the role of all other authors was to provide guidance, and to proof-read and edit the papers. The simulation, data processing, analysis, and writing of the original drafts and revisions are attributed to the present author.

This copy has been supplied on the understanding that it is copyright material and that no quotation from the thesis may be published without proper acknowledgement. The right of Jacob Paul Anderson to be identified as Author of this work has been asserted by him in accordance with the Copyright, Designs and Patents Act 1988.

## Acknowledgements

I would like to take the opportunity to thank several people. Inevitably I will forget someone, but I am greatly appreciative of all of the support I have received over the last few years without which I would not have been able to do this.

Firstly, I would like to thank my supervision team, from whom I have learned a great deal. My first thanks go to Professor Michael Fairweather. As my main supervisor, he has offered a wealth of experience and knowledge during the project, which has been invaluable, with constant encouragement, guidance, ideas and pragmatism, as well as giving me the opportunity to develop my own ideas and skills. This has been supported in equal measure by Dr Lee Mortimer who offered many solutions and ideas, helping tackle difficult challenges faced throughout the project. Lee taught me an awful lot, especially at the start of the project, and showed great patience over this period for which I am very grateful. I greatly thank Professor Jeffrey Peakall, who provided a unique scientific perspective throughout, always providing useful and honest feedback and insight, as well as offering consistent encouragement. Lastly, Dr Timothy Hunter provided useful guidance and support, and I add a special word of thanks for his part in recruiting me into this programme for which I am also very grateful. I feel very lucky to have had such an engaged, helpful, and down-to-earth team of supervisors, who have been a pleasure to work with.

Secondly, I wish to give immense thanks to the CDT in Fluid Dynamics – the management team, teaching staff, and students – for the opportunity, and for the first-class training provided, as well as to the wider Fluids Group and to the Nuclear Group at Leeds. I particularly thank Professor Peter Jimack for his valuable and honest advice and guidance during the MSc portion of my time studying at Leeds. As part of the CDT, I was lucky to experience this MSc/PhD journey with a great cohort of fellow researchers, who are now friends. I would also like to thank Sellafield Ltd for their part in funding this opportunity, it was a privilege to collaborate with industry to understand the usefulness of the work. My thanks also to EPSRC for the funding provided.

I wanted to also give a special mention to Professor Anvar Shukurov whose immense passion for the field of applied mathematics and fluid dynamics first set me on this career trajectory during my undergraduate degree in Newcastle.

I wish to express sincere gratitude to those in my personal life – my parents, siblings, grandparents, family, and friends. At times it has been difficult to find time for everybody whilst trying to balance many competing challenges, particularly over the past year, and yet everyone's support and encouragement has been unwavering. I am eternally grateful for this.

Finally, to Georgia – I cannot put into words just how supportive you have been throughout this project, or begin to express my gratitude, and so as a small gesture I dedicate the thesis to you.

## Abstract

The present thesis brings new understanding to the field of turbulent multiphase flows by exploring agglomeration phenomena involving non-spherical colloidal particles in quiescent and turbulent fluid flows. The work addresses a significant gap in the literature by explicitly capturing the full interplay between the disparate scales of turbulence and van der Waals induced agglomeration. The novelty of the work comes from combining particle-resolved multiphase fluid dynamics, orientationally-dependent ellipsoidal collision and agglomeration treatment, and high-fidelity fluid simulation. This creates a unique state-of-the-art simulation framework that allows the study of turbulent agglomeration processes at a new level of detail. Through application of the developed methods, contributions to both fundamental scientific understanding and industrial applications are made, of notable benefit to areas such as nuclear waste management and sedimentology.

Whilst theoretical frameworks exist for microscopic forces such as DLVO interactions between non-spherical particle pairs, their explicit integration into macroscopic turbulence simulations remains largely unexplored. To address this, a state-of-the-art multiscale and multiphysics solver was developed and integrated with the open-source high-fidelity flow solver Nek5000. The fully resolved fluid field and structure-resolved particle-fluid coupling regime are thus obtained through direct numerical simulation and an immersed boundary method.

Findings are presented across three distinct results chapters. The first investigates binary interactions in quiescent conditions, establishing a baseline understanding of the interactions in terms of threshold velocities, whilst also demonstrating the types of configurations that lead to agglomeration. For non-spherical particles, the distribution of surface curvature over the particle is shown to critically influence the interaction tendencies. Disks exhibit strong face-to-face interactions but reduced overall agglomeration due to their favouring of edge collisions at higher curvature regions on the surface that subsequently induce the lowest attractive forces. Conversely, needles demonstrate higher agglomeration rates due to interactions being distributed more evenly across their surfaces, under uniform random sampling.

The second study builds upon the first by introducing turbulence, employing a homogeneous isotropic turbulence field which is validated against two reference studies. Turbulence is found to reduce overall agglomeration rates by increasing the relative velocities of particles, despite the tendency of the turbulence to improve the alignment characteristics of the particles. Disks increase their resistance to agglomeration due to their greater surface area for the fluid forces to act upon, inducing greater disruptive accelerations. As the turbulence level increases, spheres begin to become the favourable morphology for achieving agglomeration, as they have more optimal drag characteristics, despite lower overall attraction than needles.

The final study tackles behavioural modifications and multi-particle systems, introducing novel algorithms for scalable simulations of agglomeration under turbulence. Parallelisation of particle modules and computational efficiency improvements enable simulations over extended time and length scales, beginning to capture interesting emergent crystal-like structures. The thesis concludes by laying the foundation for future exploration of multiparticle structures and dynamics, bridging the gap between the micro and macroscales. The relative advantages of hard- and soft-sphere schemes are also discussed for the specific application of turbulent agglomeration.

## Contents

1	Intro	duction1	
	1.1	1.1 Motivation	
	1.2	Behavioural Modification 4	
	1.3	Sellafield and the Nuclear Industry5	
	1.4	Non-Spherical Particle Systems	
	1.5	Aims, Objectives and Outline	
	1.6	Thesis Organisation	
2	Litera	ature Review	
	2.1	Fluid Phase 17	
	2.1.1	Turbulence	
	2.1.2	Turbulent Energy Cascade	
	2.1.3	Energy Spectra	
	2.1.4	Measuring Turbulence	
	2.1.5	Computational Fluid Dynamics23	
	2.1.6	Reynolds-Averaged Navier-Stokes	
	2.1.7	Large Eddy Simulation25	
	2.1.8	Direct Numerical Simulation	
	2.1.9	Further Solution Methods26	
	2.1.1	0 Homogeneous Isotropic Turbulence and the Taylor-Reynolds Number	
	2.1.1	1 Forcing of Homogeneous and Isotropic Turbulence	
	2.2	Particle-Phase	
	2.2.1	Introduction and Relevance	
	2.2.2	Non-Spherical Challenges	
	2.2.3	Collisions: soft- and hard-sphere	
	2.2.4	Van der Waals and Interparticle Forces	
	2.2.5	Numerical Implementation of Van der Waals Forces	
	2.2.6	Non-Spherical Interparticle Forces	
	2.3	Multiphase Flow	
	2.3.1	Background and Theory	
	2.3.2	Experimental Study of Particles in Turbulence	
	2.3.3	Numerical Study of Particles in Turbulence	
	2.3.4	Motion of Non-Spherical Particles	
	2.3.5	Agglomeration of Particles in Turbulence57	

	2.4	Conclusions	59
3 Methodology		62	
3.1		Fluid-Phase	63
	3.1.1	Numerical Solution of the Navier-Stokes Equations	63
	3.1.2	2 Nek5000 Flow Solver	64
	3.1.3	3 Simulation Domain	65
	3.1.4	Turbulence Forcing Scheme	68
	3.2	Particle-Phase Methodology	73
	3.2.1	Overview	73
	3.2.2	2 Morphology and Meshes	73
	3.2.3	3 Orientation Tracking and Quaternion Formulation	76
	3.2.4	Immersed Boundary Method and Coupling	79
	3.2.5	5 Collision Detection	83
	3.2.6	Collision Mechanics	89
	3.2.7	<sup>7</sup> Spherical DLVO Interaction Force	91
	3.2.8	3 Non-Spherical DLVO Interaction Force	94
	3.2.9	Non-Spherical Soft-Sphere Overlap Calculation	100
	3.2.1	0 Non-Spherical Soft-Sphere Collision Physics	115
4	Non	-Spherical Particle Interactions in a Quiescent Liquid	117
	4.1	Introduction	117
	4.2	Agglomeration Transition of Spherical Particles	118
	4.3	Agglomeration Transition for Non-Spherical Particles of a Fixed Initial Orier	ntation. 127
	4.4	The Role of Morphology in Non-Spherical Particle Agglomeration	137
	4.4.1	Monte-Carlo Simulation of Randomly Orientated Interacting Disks	141
	4.4.2	2 Monte-Carlo Simulation of Randomly Orientated Interacting Needles	145
	4.4.3	3 Direct Comparison of Morphologies	147
	4.5	Orientational Order in Particle Agglomeration	149
	4.5.1	Alignment of Symmetry Axes	150
	4 5 6	2 Surface Collision Points	150
	4.5.2		152
	4.5.2	Effect of Surface Curvature Distribution on Agglomeration	152
	4.5.2 4.6 4.7	Effect of Surface Curvature Distribution on Agglomeration	152 159 162
5	4.5.2 4.6 4.7 Non	Effect of Surface Curvature Distribution on Agglomeration Conclusions -Spherical Particle Interactions in Forced Homogeneous and Isotropi	152 159 162 c
5 Tı	4.5.2 4.6 4.7 Non urbuler	Effect of Surface Curvature Distribution on Agglomeration Conclusions -Spherical Particle Interactions in Forced Homogeneous and Isotropiate	152 159 162 c 166
5 Tı	4.6 4.7 Non urbuler 5.1	Effect of Surface Curvature Distribution on Agglomeration Conclusions -Spherical Particle Interactions in Forced Homogeneous and Isotropice Introduction	
5 Tı	4.6 4.7 Non urbuler 5.1 5.2	Effect of Surface Curvature Distribution on Agglomeration Conclusions -Spherical Particle Interactions in Forced Homogeneous and Isotropic ice Introduction Validation of Single-Phase Turbulence	

	5.2.2	A Priori Estimation of Forcing Parameters	174		
	5.2.3	Validation Case I: $Re\lambda = 65$	175		
	5.2.4	Validation Case II: $Re\lambda = 143$	177		
	5.2.5	Validation Case Comparison	178		
	5.3	Forcing Dimensional Turbulence Boxes	180		
	5.4	Simulation Set-Up	186		
	5.5	Qualitative Analysis of Collision and Agglomeration Dynamics	188		
	5.6	Particle Interactions in HIT Approximating the Viscous Sublayer	191		
	5.7	Particle Interactions in HIT Approximating the Buffer Layer Region	197		
	5.8	Particle Interactions in HIT Approximating the Bulk Flow Region	199		
	5.9	Influence of Turbulence Characteristics on Particle Agglomeration	202		
	5.10 In	fluence of Morphology on Particle Agglomeration in Turbulence	209		
	5.11 C	onclusions	221		
6	Mult	i-Particle Soft-Sphere Agglomeration and Behavioural Modification Effects . :	225		
	6.1	Introduction	225		
	6.2	Behavioural Modifications for Binary Particle Interactions	226		
	6.2.1	Behavioural Modification I: Collisional Energy	228		
	6.2.2	Behavioural Modification II: Strength of Attraction	236		
	6.2.3	Behavioural Modification III: Strength of Turbulence	242		
	6.2.4	Behavioural Modification Summary	243		
	6.3	Development of a Multi-Particle Agglomeration Framework for Non-Spherical			
	Particle	Particles			
	6.3.1	Background and Motivation	244		
	6.3.2	Development of an Overlap Depth Calculation Scheme	250		
	6.3.3	Comparative Study of Soft-Sphere Collisions	252		
	6.3.4	Demonstration of Agglomeration Capability	256		
	6.4	Validation Against Hard-Sphere Agglomeration Results	258		
	6.5	Multi-Particle Collisions and Agglomeration	260		
	6.5.1	Simulation Initial Conditions and Challenges	260		
	6.5.2	Particle Injection Algorithm	262		
	6.5.3	Initial Results	263		
	6.5.4	Application to Behavioural Modifications	266		
	6.6	Reintroduction of Turbulence and DLVO Forces	267		
	6.7	Solver Improvements: Parallelisation, Time-Stepping and Speed-Up	270		
	6.8	Soft-Sphere Agglomeration in Turbulence	272		
	6.9	Agglomerate Structures Formed by Interacting Non-Spherical Particles	274		

6	6.10	Conclusions	275
7 Conclusions and Suggestions for Further Work 2			277
7	7.1	Summary of Main Developments and Findings	277
7	7.2	Suggestions for Further Work	285
References			

## List of Tables

Table 3.1: Important quantities relating to the morphology of five relevant shapes. The
coefficients are nondimensionalised by the volumetrically equivalent radius and hence
volume and surface area are similarly nondimensional
Table 4.1: Calcite particle parameters used in the simulations
Table 4.2: Results of the non-spherical cut-off velocity study compared by morphology.
Table 4.3: The results of the randomly distributed orientation study.       146
Table 5.1: Forcing parameters used in the two validation cases
Table 5.2: Time-averaged results comparison against both validation cases of Chouippe
and Uhlmann (2015), denoted CU 176
Table 5.3: Forcing parameters used to achieve the desired properties in the dimensional
cases
Table 5.4: Resultant values observed in the dimensional simulations and used for the
following study, all variables are as before and $u^\prime$ is given in $mms-1.$
Table 5.5: Initial conditions used in the turbulence simulations
Table 5.6: Agglomeration results for the lowest turbulence level box.       194
Table 5.7: Agglomeration results for the intermediate turbulence level box
Table 5.8: Agglomeration results for the highest turbulence level box.       198
Table 6.1: Simulation matrix for the behavioural modification studies. Each column
represents a unique set of $\sim\!50$ simulations
Table 6.2: Agglomeration data for the spherical particles, with varied coefficient of
restitution
Table 6.3: Agglomeration data for the needles, with varied coefficient of restitution 231
Table 6.4: Agglomeration data for the disks, with varied coefficient of restitution 232
Table 6.5: Agglomeration data for the spheres, with varied Hamaker constant
Table 6.6: Agglomeration data for the needles, with varied Hamaker constant
Table 6.7: Agglomeration data for the disks, with varied Hamaker constant
Table 6.8: Collated data showing the agglomeration rate for the three morphologies
against increasing turbulence level

# List of Figures

Figure 1.1: SEM images showing agglomerated colloidal particles of corroded Magnox
sludge taken from the pond facilities at Sellafield (Gregson et al., 2011)
Figure 1.2: SEM of micron-scale formation of 'desert-rose' type structures in LiCoO2
samples. White scale bar indicates a length of 10 $\mu m$ . Image taken from Chen and Grey
(2008)
Figure 1.3: SEM photomicrograph of shallow-water oolitic carbonates, in particular a
aragonite needle matrix (Bennett et al., 1989)11
Figure 2.1: Energy spectrum of a turbulent flow, with sub-ranges separated by vertical
dashed bars, and the universal form indicated by a dotted line. The blue swirls illustrate
the reducing size of the eddies, with three key length scales indicated
Figure 2.2: The classification map for particle-laden turbulent flows, originally from
Elghobashi (1991), this is an updated version (Elghobashi, 2006). Particle shear Stokes
number is shown on the vertical axis, and volume concentration on the horizontal axis.
Figure 3.1: Schematic of the computational domain, with a binary particle interaction
taking place between two needles66
Figure 3.2: Demonstrative example of the computational meshes used in the
simulations; here, possessing 203 elements67
Figure 3.3: Slice of the computational mesh (left), with a zoom (right) demonstrating the
Gauss-Lobatto-Legendre structure of the nodes within each element
Figure 3.4: Demonstrative example of the Python implementation of an Eswaran and
Pope (1998) non-dimensional forcing field $F_u$ for one slice in the $xy$ -plane (top) and one
slice in the $yz$ -plane (bottom) evolving over a short period of time, with four evenly spaced
realisations (blue arrow indicates time direction)71
Figure 3.5: Contour plot of non-dimensional velocity magnitude for $Re_{\lambda} = 65$
Figure 3.6: Isosurface plot of non-dimensional velocity magnitude for $Re_{\lambda} = 65$ around
the non-dimensionalised root-mean-square velocity $u'=0.7$
Figure 3.7: Computational meshes of the particles used in the immersed boundary
simulations. Pictured from left to right: a needle, a sphere, and a disk

Figure 3.8: The three coordinate systems used to define the state of the body. The vector Figure 3.9: Schematic demonstrating the implementation of the ghost-cell method on a computational grid, with points used in the calculation colour-coded: interior nodes (red), exterior nodes (green), ghost nodes (yellow). The thick black line represents a slice Figure 3.10: Schematic of the approach to calculating the initial surface points via Figure 3.11: Variables used in the calculation of  $R_i$ . Variables with superscript '0' indicate initial values and those without represent the final values converged upon by the Figure 3.12: Plot of DLVO interaction potential normalised by thermal energy, for two interacting spherical particles as a function of separation distance normalised by particle Figure 3.13: Schematic demonstrating how the surface interaction is captured about the Figure 3.15: Non-spherical DLVO interaction potentials for the three available modelling approaches, compared also to a sphere-sphere interaction. The insert window shows the pertaining orientation. On the legend, Schiller refers to Schiller et al. (2011), and Everaers refers to Everaers and Ejtehadi (2003). Separation distance is normalised by the Figure 3.17: Definition of the overlap variables computed in Approach 1. ..... 104 Figure 3.18: Demonstration of the two combined methods for computing overlap  $\delta$  and separation *d* facilitating a soft-sphere collision between aspect ratio 5:1 needles. Both quantities are normalised by the volume equivalent radius r. SS is soft-sphere; HS is hard-sphere......105 Figure 3.20: Illustration of varying the shape parameters to give a reasonable starting 

Figure 4.1: Time evolution of inter-particle separation for interacting particles, illustrating how this looks for typical agglomeration (solid line) and bouncing (dashed line) events.

Figure 4.5: The modal configurations of non-spherical particles used in the velocity study.

Figure 4.6: Simulation trajectory data for the lowest vdW force-inducing needle configuration. Plots show separation distance over time (top left), relative velocity over time (top right), DLVO force over time (bottom left), and mean distance covered in a Figure 4.7: Simulation trajectory data for the highest vdW force-inducing needle configuration. Plots show separation distance over time (top left), relative velocity over time (top right), DLVO force over time (bottom left), and mean distance covered in a Figure 4.8: Simulation trajectory data for the lowest vdW force-inducing disk configuration. Plots show separation distance over time (top left), relative velocity over time (top right), DLVO force over time (bottom left), and mean distance covered in a Figure 4.9: Simulation trajectory data for the highest vdW force-inducing disk configuration. Plots show separation distance over time (top left), relative velocity over time (top right), DLVO force over time (bottom left), and mean distance covered in a Figure 4.10: Mean separation distance plotted as a function of initial velocity for the two 

Figure 4.11: Mean separation distance plotted as a function of initial velocity for the two
modal disk cases, compared with spheres
Figure 4.12: The incorrectly sampled distribution (left) and the amended distribution
(right), for uniformly sampled orientations
Figure 4.13: Temporal evolution of inter-particle separation distance for agglomerating
disks (left) and rebounding disks (right)141
Figure 4.14: Temporal evolution of relative velocity for agglomerating disks (left) and
rebounding disks (right)
Figure 4.15: Temporal evolution of DLVO force magnitude between particles for
agglomerating disks (left) and rebounding disks (right)
Figure 4.16: PDFs of inter-particle separation (left) and relative particle velocity (right)
compared for agglomerating and rebounding disks144
Figure 4.17: Temporal evolution of inter-particle separation distance for agglomerating
needles (left) and rebounding needles (right)145
Figure 4.18: Temporal evolution of relative velocity between particles for agglomerating
needles (left) and rebounding needles (right)145
Figure 4.19: Temporal evolution of DLVO force magnitude between particles for
agglomerating needles (left) and rebounding needles (right)
Figure 4.20: PDFs of inter-particle separation (left) and relative particle velocity (right)
compared for agglomerating and rebounding needles147
Figure 4.21: PDFs of inter-particle separation (left) and relative particle velocity (right)
compared for disk interactions (solid line) and needle interactions (dashed line) 147
Figure 4.22: PDFs of the maximum achieved DLVO force magnitude in each simulation
for disks (solid line) and needles (dashed line)
Figure 4.23: PDFs of relative symmetry axis orientations between particle pairs for disks
(left) and needles (right), broken down into cases of agglomeration (solid line) and
rebound (dashed line). Vertical line indicates $\pi/4$ or half-way between perpendicular and
parallel alignments
Figure 4.24: Meridian of the surface, describing the angle $\phi$ , for disks (red) and needles
(blue) 153
Figure 4.25: PDFs of collision points in terms of the parametric angles $ heta$ (left) and $\phi$ (right)
for disks

Figure 4.26: Scatter plot of the $( heta,\phi)$ parameter space, for agglomerating disks (red) and
bouncing disks (blue)
Figure 4.27: PDFs of collision points in terms of the parametric angles $ heta$ (left) and $\phi$ (right)
for needles
Figure 4.28: Scatter plot of the $(\theta, \phi)$ parameter space, for agglomerating needles (red)
and bouncing needles (blue)
Figure 4.29: PDFs of the collision points $ heta$ (left) and $\phi$ (right), for all disks (solid line) and
all needles (dashed line)
Figure 4.30: Normalised Gaussian curvature shown as a colour map for needles (left)
and disks (right)
Figure 4.31: Normalised radius of Gaussian curvature shown as a colour map for needles
(left) and disks (right)
Figure 4.31: Diagram to illustrate perturbation of orientational configuration for aspect
ratio 5:1 disks
Figure 5.1: Pseudocolour velocity magnitude plots demonstrating the onset of
turbulence, separated by intervals of 0.5 normalised time units
Figure 5.2: Pseudocolour non-dimensional velocity magnitude contours on the exterior
surfaces shown for the turbulence box with $Re_{\lambda} = 65$
Figure 5.3: Temporal convergence of the spatially averaged turbulence quantities for
validation case I, Taylor-Reynolds number (top left), Taylor microscale (top right) and
large eddy length scale (bottom). The black horizontal lines show the time-averaged
values reported by Chouippe and Uhlmann (2015)
Figure 5.4: Temporal convergence of the spatially averaged turbulence quantities for
validation case II, Taylor-Reynolds number (top left), Taylor microscale (top right) and
large eddy length scale (bottom). The black horizontal lines show the time-averaged
values reported by Chouippe and Uhlmann (2015) 178
Figure 5.5: Non-dimensional Taylor microscale (left) and Taylor-Reynolds number (right)
measured across a non-dimensional channel flow for $Re_{\tau} = 180$
Figure 5.6: Pseudocolour velocity magnitude contours on the exterior surfaces shown for
the viscous sublayer approximation turbulence box with $Re_{\lambda} = 2.5$
Figure 5.7: Pseudocolour velocity magnitude contours on the exterior surfaces shown for
the buffer layer region approximation turbulence box with $Re_{\lambda} = 12.1$ 182

Figure 5.8: Pseudocolour velocity magnitude contours on the exterior surfaces shown for the bulk flow region approximation turbulence box with  $Re_{\lambda} = 22.2...$  183 Figure 5.9: Temporal evolution of the viscous sublayer approximation box for the two controlled quantities, Taylor-Reynolds number (left) and velocity fluctuation (right). . 184 Figure 5.10: Temporal evolution of the buffer layer region approximation box for the two controlled quantities, Taylor-Reynolds number (left) and velocity fluctuation (right). . 184 Figure 5.11: Temporal evolution of the bulk flow region approximation box for the two controlled quantities, Taylor-Reynolds number (left) and velocity fluctuation (right). . 184 Figure 5.12: Temporal evolution of separation distance (left) and relative velocity (right) of Figure 5.13: Velocity magnitude (m/s) contour plot for a lone needle intersecting a turbulence field. ..... 189 Figure 5.14: Velocity magnitude (m/s) contour plot for a pair of disks interacting in a turbulence field. ..... 189 Figure 5.15: Temporal evolution of separation distance (left) and relative velocity (right) between spheres interacting in the viscous sublayer turbulence box. Each line represents an individual interaction simulation......192 Figure 5.16: Temporal evolution of separation distance (left) and relative velocity (right) Figure 5.17: Temporal evolution of separation distance (left) and relative velocity (right) Figure 5.18: Temporal evolution of separation distance (left) and relative velocity (right) Figure 5.19: Temporal evolution of separation distance (left) and relative velocity (right) Figure 5.20: Temporal evolution of separation distance (left) and relative velocity (right) between disks interacting in the buffer layer turbulence box...... 198 Figure 5.21: Temporal evolution of separation distance (left) and relative velocity (right) Figure 5.22: Temporal evolution of separation distance (left) and relative velocity (right) 

Figure 5.23: Temporal evolution of separation distance (left) and relative velocity (right) Figure 5.24: PDFs demonstrating the overall tendencies of the spheres with respect to separation distance (top left), relative velocity (top right) and van der Waals force Figure 5.25: PDFs demonstrating the overall tendencies of the needles with respect to separation distance (top left), relative velocity (top right) and van der Waals force Figure 5.26: PDFs demonstrating the overall tendencies of the disks with respect to separation distance (top left), relative velocity (top right) and van der Waals force Figure 5.27: PDFs demonstrating the overall tendencies of the viscous sublayer box with respect to separation distance (top left), relative velocity (top right) and van der Waals (bottom) for the three morphologies. Vertical lines indicate the statistical mean...... 212 Figure 5.28: PDFs demonstrating the overall tendencies of the buffer layer box with respect to separation distance (top left), relative velocity (top right) and van der Waals (bottom) for the three morphologies. Vertical lines indicate the statistical mean...... 213 Figure 5.29: PDFs demonstrating the overall tendencies of the bulk flow box with respect to separation distance (top left), relative velocity (top right) and van der Waals (bottom) Figure 5.30: Symmetry axis alignment (left) and collision point location (right) for needle Figure 5.31: Symmetry axis alignment (left) and collision point location (right) for disk interactions in the viscous sublayer box...... 217 Figure 5.32: Symmetry axis alignment (left) and collision point location (right) for needle interactions in the buffer layer box...... 217 Figure 5.33: Symmetry axis alignment (left) and collision point location (right) for disk Figure 5.34: Symmetry axis alignment (left) and collision point location (right) for needle Figure 5.35: Symmetry axis alignment (left) and collision point location (right) for disk 

Figure 6.1: Probability distributions of separation distance (left) and relative velocity (right) between spheres interacting in the viscous sublayer turbulence box, contrasted for three values of the coefficient of restitution. 229 Figure 6.2 Probability distributions of van der Waals force measured between spheres interacting in the viscous sublayer turbulence box, contrasted for three values of the coefficient of restitution. 232 Figure 6.3: Probability distributions of separation distance (left), relative velocity (right) and van der Waals force (bottom), measured between needles interacting in the viscous sublayer turbulence box, contrasted for three values of the coefficient of restitution. 233 Figure 6.4: Probability distributions of separation distance (left), relative velocity (right) and van der Waals force (bottom), measured between needles interacting in the viscous sublayer turbulence box, contrasted for three values of the coefficient of restitution. 234 Figure 6.4: Probability distributions of separation distance (left), relative velocity (right) and van der Waals force (bottom), measured between disks interacting in the viscous sublayer turbulence box, contrasted for three values of the coefficient of restitution. 234 Figure 6.5: Probability distributions of relative orientation, measured between needles (left) and disks (right) interacting in the viscous sublayer turbulence box, contrasted for three values of the coefficient of restitution. 234

Figure 6.6: Probability distributions of separation distance (left), relative velocity (right) and van der Waals force (bottom), measured between spheres interacting in the viscous sublayer turbulence box, contrasted for three values of the Hamaker constant. ...... 238 Figure 6.7: Probability distributions of separation distance (left), relative velocity (right) and van der Waals force (bottom), measured between needles interacting in the viscous sublayer turbulence box, contrasted for three values of the Hamaker constant. ...... 239 Figure 6.8: Probability distributions of separation distance (left), relative velocity (right) and van der Waals force (bottom), measured between disks interacting in the viscous sublayer turbulence box, contrasted for three values of the Hamaker constant. ...... 240 Figure 6.9: Probability distributions of relative orientation, measured between needles (left) and disks (right) interacting in the viscous sublayer turbulence box, contrasted for Figure 6.10: The two orientational configurations used for the illustrative studies..... 253 Figure 6.11: Distance and overlap calculation between particles colliding in configuration 1. (Top left: Method 1; top right: Method 2; bottom left: Method 3.) The non-zero metrics 

Figure 6.14: Separation distance evolution between particles agglomerating in configuration 1. (Top left: Method 1; top right: Method 2; and bottom: Method 3.) ..... 257 Figure 6.15: Relative velocity evolution (left) and force interaction (right) between particles agglomerating in Configuration 1 under Method 3. ..... 258 Figure 6.16: Separation distance (left) and relative velocity (right) time evolution for colliding spheres under the soft-sphere model. Distinct lines are colour-coded by initial Figure 6.16: Separation distance evolutions for colliding needles under the soft-sphere model, demonstrating the cut-off initial velocity for agglomeration in the case of minimal Figure 6.17: Demonstration of random injection (left) and the main collisions at the centre Figure 6.18: Cumulative system separation between particles (left) and cumulative Figure 6.19: Kinetic energy of the particle phase (left) and cumulative soft-sphere force Figure 6.21: Cumulative system separation between particles (left) and cumulative relative velocity (right), measured over time for three values of the coefficient of Figure 6.22: Particle-phase kinetic energy (left) and coordination number (right), Figure 6.23: Cumulative system separation between particles (left) and cumulative Figure 6.24: Coordination number (left) and force magnitude (right), measured over time.

## Nomenclature

Symbol	Definition
Ш	$3 \times 3$ identity matrix
$\boldsymbol{T}_{p=i,j}$	Advected collision points (Method 1)
IHI	Algebra of quaternions
$k_B$	Boltzmann constant
$Re_B$	Bulk Reynolds number
K <sub>C</sub>	Central wavenumber (forcing method)
$\Delta \boldsymbol{u}$	Change in relative velocity due to a collision at surface
$ ilde{L}$	Characteristic length scale
r	Characteristic particle radius
$\widetilde{U}$	Characteristic velocity scale
d	Closest distance vector between two particle surfaces
е	Coefficient of restitution
$P_{col}$	Collision contact point
$\pmb{n}_{col}$	Collision normal
$\boldsymbol{e}_i$	Complex random number
$\boldsymbol{x}_{0}$	Converged contact point (Method 3)
$\alpha_L$	Convergence parameter (Method 3)
$C^j$ and $P^j$	Convergence parameters for common normal method
$\boldsymbol{x}' = (x', y', z')$	Coordinate vector – local/body-fitted frame
$\boldsymbol{x}=(x,y,z)$	Coordinate vector – global frame
$\boldsymbol{x}^{\prime\prime} = (x^{\prime\prime}, y^{\prime\prime}, z^{\prime\prime})$	Coordinate vector – rotating frame
$\kappa_{cut}$	Cut-off wavenumber for forcing scheme
$ ho_F$	Density of fluid phase
$ ho_p$	Density of particle phase
L	Domain length
$D_{\varphi_i}$	Dot product between parametric tangent vector and closest distance vector
$L_E$	Eddy length scale

$\kappa_E$	Eddy wavenumber
$m_e$	Effective mass
Z <sub>e</sub>	Electrolyte ionic valency
$q_e$	Elementary electric charge unit
$\boldsymbol{a}=(a,b,c)$	Ellipsoidal coefficient vector
$\eta_T$	Estimated Kolmogorov length scale (forcing method)
$Re_{\lambda}^{T}$	Estimated Taylor-Reynolds number (forcing method)
$\boldsymbol{e}_{\theta} = (\phi, \theta, \psi)$	Euler angles
$u_{ex}$	External forces acting during collision
$\epsilon_F$	Fluid dissipation
p	Fluid pressure field
S <sub>ij</sub>	Fluid strain-rate tensor
Т	Fluid temperature
$oldsymbol{u}_{IB}$	Fluid velocity at surface of immersed boundary
$oldsymbol{u}_E$	Fluid velocity of exterior point to immersed boundary
$oldsymbol{u}_G$	Fluid velocity of immersed boundary ghost node
u	Fluid velocity vector
$\mu_F$	Fluid viscosity – dynamic
$ u_F$	Fluid viscosity – kinematic
$F_{f}$	Force acting at a given immersed boundary face
$F_{SS}$	Force due to soft-sphere collision
<b>F</b> <sub>DLVO</sub>	Force due to total DLVO interaction
$\pmb{F}_{ ext{EDL}}$	Force due to electric double layer interaction
$\pmb{F}_{ m vdW}$	Force due to van der Waals interaction
β	Forcing constant
$\epsilon^{*}$	Forcing dissipation parameter
$F_u$	Forcing field
$\widehat{m{f}}_{EP}$	Forcing source term expressed in wavenumber space
$oldsymbol{f}_{EP}$	Forcing source term for homogeneous and isotropic turbulence
$T_L$	Forcing timescale

$T_L^*$	Forcing timescale (non-dimensionalised)
$\sigma^2$	Forcing variance
G	Gaussian curvature
$\boldsymbol{\omega} = (\omega_x, \omega_y, \omega_z)$	Global particle angular velocity
$\Delta x$	Grid step size
$u_{\scriptscriptstyle B}$	Half of the measured average collision velocity between particles in the bulk flow region
Α	Hamaker constant
$R_G$	Harmonic mean of Radii of Gaussian curvature
Не	Hessian matrix (Method 3)
$A_f$	Icosphere mesh face area
Ι	Imaginary unit
<b>f</b> <sub>c</sub>	Impulsive collision force (unused)
$\delta Z$	Increment of update for Lagrangian system variables
Ι	Inertia tensor in the local coordinate system
i	Inertia tensor in the global coordinate system
$\pmb{R}_1^0$ and $\pmb{R}_2^0$	Initial guesses for surface points satisfying common normal condition
$\boldsymbol{\omega}^p$	Initial particle angular velocity
$v_{rel}^p$	Initial particle relative velocity
$ \boldsymbol{d} _{init}$	Initial particle separation
ę	Integral length scale
$T_\ell$	Integral timescale
$d =  \boldsymbol{d} $	Intersurface separation distance
κ	Inverse Debye length
J	Jacobian matrix (Method 3)
η	Kolmogorov length scale
$T_{\eta}$	Kolmogorov time scale
Λ	Lagrange multiplier
$\mu_L$	Lagrange multiplier (modified)
$L(\boldsymbol{x},\Lambda)$	Lagrangian function for minimisation (Method 3)
l	Large eddy length scale

T <sub>e</sub>	Large eddy turnover timescale
$p_c$	Linear momentum at colliding surface point
$\boldsymbol{\omega}' = (\omega_x', \omega_y', \omega_z')$	Local particle angular velocity (body-fitted)
$u_i'$	Local velocity fluctuations
$\kappa_0$	Lowest forced wavenumber
L <sub>c</sub>	Macroscopic forcing length scale
Н	Mean curvature
$\epsilon^*_T$	Non-dimensional dissipation rate (forcing method)
$\eta_n$	Normal dashpot constant
$k_n$	Normal spring constant
п	Number density of electrolyte ions
$N_c$	Number of collisional timesteps
$N_f$	Number of forced wavenumbers
nelx	Number of spectral elements in <i>x</i> -direction
nely	Number of spectral elements in y-direction
nelz	Number of spectral elements in z-direction
A	Orientation/rotation matrix
δ	Overlap depth between interacting particle surfaces
R	Parametric location on particle surface
<b>R</b> *	Parametric location on particle surface scaled by coefficient
$\pmb{R}_1$ and $\pmb{R}_2$	Parametric surface points satisfying common normal condition for particles 1 and 2
$\boldsymbol{\vartheta} = (\vartheta, \varphi)$	Parametric surface variables of particle
М	Particle mass
dv	Particle relative velocity
$ au_p$	Particle response time
$u_p$	Particle velocity at centroid
$\mathcal{E}_{S}$	Permittivity of a solvent
$\varepsilon_0$	Permittivity of a vacuum
q	Quaternion
$r_G$	Radius of Gaussian curvature

$C_i$	Random initial condition for particle velocity
Θ	Reduced surface potential
α	Relative orientation of particle pair about the vector of closest approach
$oldsymbol{ u}_{col}$	Relative velocity at contact point (Method 3)
$u_r$	Relative velocity between particles at the instant of collision
$\Delta \boldsymbol{u}_c$	Relative velocity of nearest points between two particle surfaces
Re	Reynolds number
$Re_{ au}$	Reynolds number based on the shear/friction velocity (channel flow)
u'	Root-mean-squared velocity fluctuation
$\boldsymbol{a}^* = (a^*, b^*, c^*)$	Scaled coefficient vector (Method 2)
$\mu$	Scaling used for coefficient vector (Method 2)
$F_i(x)$	Shape function of particle $i$ – global
$f_i(x')$	Shape function of particle <i>i</i> – local
$\Delta t$	Simulation timestep
$[\boldsymbol{r}_c]_{ imes}$	Skew symmetric matrix of $m{r}_c$
$\boldsymbol{Z} = (x, y, z, \mu_L)^T$	Solution variables to Lagrangian problem (Method 3)
St	Stokes number
$St_\eta$	Stokes number based on Kolmogorov scale
$\boldsymbol{T}_{0}$	Stored collision points at moment of collision (Method 1)
$\chi_{ij}\eta_{ij}$	Surface interaction term (orientationally dependent)
$\zeta_{p=i,j}$ and $\zeta_{p=i,j}'$	Surface orthogonal principal curvatures at point of nearest approach
$t_{arphi}$	Surface tangent vector in $arphi$ -direction
$t_artheta$	Surface tangent vector in $artheta$ -direction
К	Symmetric system mass matrix
λ	Taylor microscale
$Re_{\lambda}$	Taylor-Reynolds number
Т	Torque of particle
F	Total force acting over the surface of immersed boundary
k	Turbulence kinetic energy

b	Uhlenbeck-Ornstein random processes
$X_{i,j}$	Uniformly distributed random variable $\sim U(0,1)$
$\widehat{q}$	Unit quaternion
$r_{f}$	Vector pointing from particle centre to face centre
$\boldsymbol{r}_{c}$	Vector pointing from particle centre to particle surface location
Φ	Vectorised formulation of Lagrangian problem (Method 3)
$\boldsymbol{u}_{c}$	Velocity of point on particle surface
τ	Viscous stress tensor
$\phi_{\scriptscriptstyle P}$	Volume fraction of particle phase
$\omega_{fluc}$	Vorticity fluctuation amplitude
к	Wavenumber vector
Z	Zeta potential
$\mathbb{R}^{n}$	n-dimensional real Euclidean space

### 1 Introduction

### 1.1 Motivation

Multiphase flows appear frequently in both nature and industry. These flows take many different forms and the category of 'multiphase flow' is far-reaching. The term encompasses flows containing dispersed phases, such as particles, droplets or bubbles (Loth, 2000), it includes multiple interacting continuous phases, and in the most complex cases there may be combinations of each (Thorn et al., 1997; Kim and Lowengrub, 2005). Familiar examples of multiphase flows include aerosol sprays, the formation and transportation of sediment in rivers, the removal and transport of industrial wastes, or the dynamic interplay between air and water travelling through a pipe. The generation of a deepened understanding of these flows naturally leads to the ability to improve current industrial methods, due to those working in industry having ever more avenues of exploration at their fingertips for the innovation of new processes. Improved environmental safety and conservation processes, and improved medical treatments, are further benefits that arise in the same way. As such, there is great value in scientific investigations that aim to develop our understanding surrounding the complex phenomena occurring in this rich field of study.

Often, even with moderate flow rates, multiphase systems will be turbulent (Balachandar and Eaton, 2010). For flows with multiple continuous phases, there are many interesting phenomena at play such as heat and mass transfer effects due to turbulent mixing (Han et al., 2005; Corrsin, 1952; Morsi and Basha, 2015), or intricate non-linear dynamics occurring at the interface between phases (Roccon et al., 2023). In the case of particle-laden flows, similarly complex interactions occur between the turbulence field of the flow and the particles within (Squires and Eaton, 1990). The chaotic nature of the turbulence and the interdependent positions, orientations, and velocities of the particles creates a system that is challenging to study. With this challenge comes the opportunity for developments in important and seemingly unrelated fields such as medicine, sedimentology, and computing, or in mathematical disciplines like numerical analysis, chaos theory and statistics.

One powerful way to investigate such systems is through the use of numerical simulation. This approach is flexible and allows the investigator the ability to quickly cover a wide parameter space and probe systems in ways that are not possible through experiments. However, this relies on accurately capturing the physics at play. Since multiphase flows are so diverse, it can be challenging to work out exactly how to proceed in recreating the entire landscape of physical effects which can often operate on vastly different scales, with all such scales contributing to the overall behaviour of the system. In many particle-laden flows, the size of the particles is smaller than the smallest scales of turbulent motion, and yet their influence on the flow is significant (Mortimer et al., 2020; Wolde, 2023).

There exists no all-encompassing computational framework, and so there are many approaches available for any given problem, each with their own strengths and weaknesses. Even still, the developments that are made in seemingly narrowly scoped systems often turn out to be directly relevant elsewhere, which is an exciting aspect of the field. For example, the immersed boundary method of Peskin (2002), developed originally for the purpose of studying the human heart, has been successfully expanded to all kinds of multiphase flow problems, such as particles in a pipe, or bubbles in a channel (Uhlmann, 2005; Ardekani et al., 2016; Fröhlich et al., 2022). This has contributed to the overall capability of simulating many flows entirely unrelated to human biology and has been significant in providing a means to include aspects such as non-sphericity in fluid-particle simulations. It is clear that, from the considerable difficulty posed by these systems, there emerges scope for scientific contribution and development.

Wall-bounded multiphase flows are prevalent in many industries (Wiart et al., 2015), with problems usually involving solid boundaries in some important way. The three canonical wall-bounded flows – pipes, channels and spatially evolving boundary layers (El Khoury et al., 2013) – can all be observed regularly in industry: pipes and channels are encountered in transportation processes, for example those seen in the nuclear industry, whilst boundary layers are inherent to aerodynamic applications. However, the sheer scale of some industrial processes means that relevant investigation away from the site can be difficult. In that, both physical and numerical experiments will

struggle to recreate the full scale of the system – and, more importantly, the full range of scales involved.

In the case of an industrial-scale particle-laden pipe flow, a pipe diameter may be on the order of a metre, with dispersed particles on the order of microns or nanometres. Then, it likely contains billions of particles whose behaviour is dependent on a vast range of scales and many different physical phenomena. To recreate such systems for investigation, whilst capturing all of the important scales at once and keeping track of all of the particles' individual motions, poses an often-intractable challenge to both experimentalists and computational modellers alike. The path forward therefore is often to compartmentalise these problems into smaller subproblems that *can* be attacked. The acquired knowledge may then be used to inform and construct theories relating to the full-scale problem. Over time, as understanding deepens, methodologies develop, and resources increase, there can be continued steps towards fully understanding these systems. The alternative route is to model the full problem with extensive simplifications, and each approach has its place.

One thing that is not well understood pertains to the fundamental mechanical and chemical interactions taking place between particles in these turbulent flows. There are several interesting emergent phenomena, such as agglomeration and flocculation of particles, that are not captured by numerical studies unless the necessary extra physics are included, and for many systems these phenomena are extremely important to capturing the bulk-scale dynamics and overall multiphase behaviour (Bridgeman et al., 2009; Almohammed, 2018). It is therefore logical to interrogate these interactions at the particle-scale first, so that they can be understood in a fundamental way. This newfound knowledge can then be leveraged to inform and thus improve the modelling assumptions that are made to reduce computational complexity as a matter of necessity, so that full scale scenarios can be modelled effectively. Equally, the design of experimental rigs can be improved by better understanding the processes occurring within.

Driven by a desire to understand these systems at a fundamental level, this study aims to address the problem of particle interactions in turbulence from a first-principles basis, resolving as much of the physics as is computationally possible. As such, there is a need to compartmentalise and simplify the full systems being simulated. All of this points towards a certain flow system that allows detailed study of the interplay between particle-particle interactions and turbulence in a small-scale and isolated way: periodic boxes of homogeneous and isotropic turbulence. These idealised flows will allow the study of both the fluid and particle physics at the length scales required to generate fundamental understanding that can be used to improve simulations where the turbulence and particle-interaction dynamics are modelled, whilst ensuring that the simulations are computationally feasible. In this system, there will be ample opportunity to develop and test numerical methods for the modelling of particle interactions in turbulence.

### 1.2 Behavioural Modification

There are certain control measures that can be taken to improve industrial systems. Such procedures rely on scientific understanding to form their basis. An example is the agglomeration and break-up behaviour of particles in a particle-laden flow, which can be better understood to improve the effectiveness of various industrial processes. For instance, it may be desirable for particles to agglomerate so that they can be easily detected and separated from a carrier-phase fluid if the particulate matter is unwanted. Conversely, one might wish for the particles to remain separated so that the flow rate remains stable and the mass transfer through the pipe is uniform. Another such benefit of minimised agglomeration is the avoidance of deposition and settling of particles out of the flow leading to pipe blockage and failure, which carries significant economic and environmental implications in many industries. In either case, an improved understanding of how modifications to the material, chemical and mechanical properties of the different phases affect agglomeration can be leveraged to obtain these desirable outcomes in the overall system's behaviour – this leveraging is referred to as a *behavioural modification* (Mortimer and Fairweather, 2021).

For example, the Hamaker constant, which governs the strength of attractive particleparticle interactions, can be altered through the addition of salt to the fluid (Shahidzadeh et al., 1998) or through modification to its temperature (Bergström, 1997). These are two very accessible behavioural modifications that could influence the long-time behaviour of the industrial system in terms of its bulk agglomeration behaviour. Another approach to counteracting the attraction between particles and thus stabilising the suspension could be the matching of the refractive index between the particle and the fluid media; alternatively, an attempt can be made to introduce repulsive interactions that prevent aggregation, or polymers can be introduced into the situation to stabilise the suspension. By adsorption to the surface, polymers can create a volume exclusion effect. These ideas are discussed in detail from an experimental perspective by Mauleon-Amieva (2020).

It can be prohibitively difficult to probe these systems in a laboratory due to the challenge of performing intrusive particle-scale experiments within the context of a large-scale flow system. The natural step to take is to look to computational approaches to obtain the required insight. Mortimer and Fairweather (2021) laid out a simulation framework for this type of investigation. The approach taken is to utilise multiphase computational fluid dynamics over a wide parameter-space of different fluid (temperature, Reynolds number, pH) and particle (Hamaker constant, Debye length, coefficient of restitution, surface charge potential) properties to gain fundamental understanding of which properties the system and its agglomeration dynamics are most sensitive to. In the context of nuclear waste management, this approach was proven successful in outlining the best parameters for exploitation in the cited paper.

### 1.3 Sellafield and the Nuclear Industry

Nuclear power currently contributes 15% of the United Kingdom's electricity supply (World Nuclear Association, 2023), whilst supporting the country's medical and defence industries. This is achieved through 8 advanced gas-cooled reactors and one pressurised water reactor. Associated with this is the challenge of managing nuclear waste. There is legacy nuclear waste in abundance at Sellafield in Cumbria, UK where more radioactive material per square metre is processed than at any other site in Europe (Office for Nuclear Regulation, 2023). There are over 1000 nuclear facilities making it one of the most complex nuclear sites in the world. Sellafield currently functions as a decommissioning site having previously been an active nuclear power generation site between 1956 and 2003. Accounting for roughly two-thirds of the Nuclear Decommissioning Authority's annual spend, Sellafield is assigned a government funded budget exceeding £2 billion annually (Infrastructure and Projects Authority, 2023).

The decommissioning of legacy nuclear facilities is now a national priority – recognised as such by the UK Government (Nuclear Decommissioning Authority, 2021). In nuclear facilities, legacy waste has been sealed inside concrete containers and stored in ponds awaiting transfer to safer storage locations. The operation of transferring the waste from these legacy facilities to new storage locations is currently performed with extreme caution, which hinders the efficiency. Such caution is given due to the lack of understanding of how the flow will behave, hence investment is being made currently into scientific research to understand and predict the behaviours and mechanisms relevant to such flows.

In the context of the nuclear industry, streamlining and developing their methodologies leads to the improvement of economic efficiency, the reduction of environmental impacts, and the increase of safety – which are three motivating factors underpinning Sellafield Ltd.'s interest and investment in this research area. For Sellafield, it is important that the build-up of reactive or corrosive materials does not occur in their industrial processes. For example, if this leads to pipe failure, there is a significant cost incurred at the expense of the UK taxpayer. This material build-up is also a challenge faced by active nuclear reactor facilities, in which the accumulation of corrosion-related unidentified deposits (CRUD) in the coolant circuit can greatly impact the heat transfer performance of the reactor (Short et al., 2013).

The modification of certain physical and chemical parameters can help the nuclear industry to obtain desirable bulk flow properties and hence make operations safer and more controllable. Simple modifications to the bulk properties like fluid temperature and Reynolds number may be of great economic value, and alterations to the system can be further achieved through the injection of additives into the flow, like nanoparticles or polymers (Mortimer and Fairweather, 2022), and through the coating of particles (Gollwitzer et al., 2012).

### 1.4 Non-Spherical Particle Systems

There are endlessly many particle morphologies that can be classed as 'non-spherical' and might be of interest to study. The challenge of modelling them is to ensure the interesting effects of particle anisotropy are accurately captured without making the problem so complex that it is intractable. It is surprising just how complicated a system becomes under the simple addition of non-sphericity. In the context of numerical methods, considerable time and effort must be spent considering how to incorporate such effects in a reasonable way.

In industry, the accessibility of certain manufacturing techniques may influence the variety of shapes that are observed when it comes to studying the relevant multiphase systems, so too might an industry's desire for certain morphological properties, which naturally influences the choice of shapes in processes like packed-bed heat exchangers. Typical industrial shapes that are easy to manufacture, or that might be chosen for their morphological properties, are cubes, cuboids, ellipsoids or cylinders, and so these are the natural first choices for non-spherical modelling. However, each shape requires different treatment, with different approaches for the tracking of the surface required and different models for the relevant forces. This means a simple change to the shape requires an overhaul of the numerical approaches employed. One can track the surfaces of ellipsoids neatly with mathematical equations; however, this is not so simple with cubes, cylinders or cuboids, whose surfaces are comprised of several faces, with a non-smooth character – in such cases, one can turn to 'superellipsoid' and mesh-based approaches, but it becomes less clear how to precisely track and resolve forces between objects, for example during collisions.

As these simple shapes break, as they often will due to collisions and interactions within the system, new shapes may be formed. In a simple case, long cylinders that snap can form shorter cylinders with the same diameter, which can lead to interesting particle size distributions that are a function of a single axial length. Alternatively, a needle-type shape might break in a similar way, but this time the result is two completely new shapes, which are difficult to include into a particle model mathematically, or to keep track of statistically.

If the choice was made to include all effects – non-sphericity, polydisperse particle distributions, multiple different shapes, breakage processes, and so on – then the problem quickly runs away from the investigator and becomes impossible to study. So, investigations are typically limited to just one type of non-spherical shape, usually of one size, at a time. There is scope to go beyond this, but usually that would need to be the entire emphasis of the study. In the present work, there is a direct interest in non-sphericity, but also on the elucidation of the effects of turbulence and interaction

behaviours, which limits the scope of development in the direction of non-sphericity. Nevertheless, as will be seen in this thesis, there is a vast complexity that comes from including just one non-spherical shape at once, in terms of the entirely new phenomena that arise.

Spheroidal geometry is the simplest extension to a smooth and continuous nonspherical geometry, arrived at by 'stretching' or 'squashing' a sphere, giving 'needles' and 'disks', respectively. These non-spherical shapes can be thought of as representing a modal type of non-sphericity with either such operation leading to one mode of non-sphericity whose effect can thus be investigated. In Figures 1.1 - 1.3 it is demonstrated that these simple shapes occur regularly in both industry and nature. In particular, natural processes reliably lead to the promotion of shapes or structures of a certain kind. For example, in rivers, erosion can lead to long flat smooth pebbles, that are of a spheroidal or ellipsoidal shape (Koster et al., 1980). Or in very arid conditions, simple disk-shaped gypsum can aggregate to form beautiful desert rose structures (Al-Kofahi et al., 1993) - highlighting the complex emergent behaviour of spheroidal systems. In this case, the emergent complexity may be accessible from a modelling point of view by considering the disk-like constituent parts combined with a model for the interaction physics. This hypothesis underscores one of the hopes for the present work, where final investigations look to model the structures formed by the agglomeration process of non-spherical particles.

It is known from scanning electron microscope (SEM) images provided by Sellafield Ltd. that their Magnox sludge agglomerates resemble a kind of quasi desert rose structure, and so it is of interest to see if the numerical models developed throughout this thesis can capture this highly complex emergent behaviour. Figure 1.1 shows examples of SEM images where disk-like particles have agglomerated to form intricate structures possessing clear order: the disks agglomerate at an angle to form an open interlocking platelet structure, as described by Gregson et al. (2011). Developing numerical models capable of simulating these systems will allow assessment of what the necessary physics are for the resolution of the crystalline order observed in nuclear waste agglomerates, as well as generate fundamental understanding of the processes occurring within the system, which can aid control measures. In the case of gypsum desert roses, a precipitation process leads to their formation (Al-Kofahi et al., 1993). In the case of the aesthetically similar LiCoO2 samples shown in Figure 1.2, it is a much different process that leads to this structure forming, involving a chemical reaction (Chen and Grey, 2008). In the agglomeration process, it is likely that the observed crystalline structure arises due to charge properties of the surfaces. In the present work, surfaces are assumed to be of a homogeneous charge profile so it would be of interest to compare the structures that arise as a result of this assumption against experimental findings, in order to understand the importance and influence of anisotropic charge profiles.

There are important secondary effects relating to the mechanical forces which may bias the structures also. The geometric element of this interaction – i.e., how two non-spherical objects see each other and interact in 3-D space, given certain assumptions about their geometry – is scale independent, making the ideas developed in the thesis more widely applicable still, and thus the developed framework can potentially be applied to a wide range of physics and engineering problems.

Finally, the agglomerated structure of needle-like particles is also of interest. For example, the microstructure of clay sediments is often composed of needle-like elements and their arrangement has implications for the porosity of the sedimented structure (Bennett et al., 1989). The arrangement itself is formed as a result of the interparticle forces and understanding the relationship between such forces and the resultant structures can give insight to the field and develop fundamental understanding of the mechanisms underpinning important geological processes.


Figure 1.1: SEM images showing agglomerated colloidal particles of corroded Magnox sludge taken from the pond facilities at Sellafield (Gregson et al., 2011).



Figure 1.2: SEM of micron-scale formation of 'desert-rose' type structures in LiCoO2 samples. White scale bar indicates a length of 10µm. Image taken from Chen and Grey (2008).



Figure 1.3: SEM photomicrograph of shallow-water oolitic carbonates, in particular a aragonite needle matrix (Bennett et al., 1989).

# 1.5 Aims, Objectives and Outline

The high-level objective for the project is to contribute new understanding of multiphase flows, which will be achieved through investigations of particle agglomeration and non-sphericity. These two parameters are both challenging to implement individually, and in combination are very rarely studied. To the best knowledge of the author, there is no existing literature which incorporates both effects into high-fidelity simulations of turbulent flows, and so the novelty of the project is inherent to the methodology developed within. The aim is to devise a numerical framework that can form the basis for further expeditions into the field of turbulent non-spherical particle agglomeration and lay the groundwork for further computational developments. The methodology will then be used to investigate the motion and agglomeration of non-spherical particles at a range of Taylor-Reynolds numbers, which have been specifically chosen to represent regions observed in pertaining wall-bounded flows studied within the research group at the University of Leeds (Njobuenwu and Fairweather, 2015; Mortimer et al., 2019; Wolde, 2023).

Practicality and industrial relevance are achieved through regular consultation with the industrial partner, Sellafield Ltd., who relay the challenges faced currently within the nuclear industry and hence inform the simulation set-ups. Both the particulate-phase and fluid-phase are modified to understand the most important factors affecting agglomeration, break-up, and dispersion in a turbulent flow. This information will be highly relevant to industry where behavioural modifications can be made to improve flow control. However, the freedom very much exists for obtaining insight at a much more fundamental level as seen fit, and this will be necessary to the development of the described novel models.

In the homogeneous and isotropic turbulence box set-up used in this work, important features can be investigated in an isolated way which makes highlighting relevant mechanisms easier, as they are decoupled from other physical effects. Interactions can be studied in-depth to understand what is happening at the particle-scale – insight which can then be used to inform models that do not possess this level of fidelity. Specifically, boundary layers around the particle are resolved and their influence on particle motion studied. There will also be an implicit lubrication contribution due to the immersed boundary method. The particle interaction will be resolved in a time-

dependent manner by consideration of the interaction potential between particles, the collisions will induce torques, as will DLVO contributions, thus there are a multitude of competing effects occurring at the particle-scale.

DLVO theory, established by Derjaguin, Landau, Verwey, and Overbeek (Derjaguin and Landau, 1941; Verwey and Overbeek, 1948) describes inter-surface particle attraction and repulsion due to van der Waals and electrostatic forces. Such fidelity is not captured by conventional point-particle approaches and so the insight gained will be important in assessing the accuracy of approaches like Lagrangian particle tracking and providing recommendations on to how to improve those types of methods.

Each of the larger-scale simulations utilise an immersed boundary technique coupled to the highly accurate flow solver, Nek5000. With the use of the University of Leeds' high performance computing clusters ARC3 and ARC4, high-fidelity simulations can be performed at scale to cover a wide parameter space and obtain many realisations of these chaotic systems.

The current project used as its starting point an existing body of work (Mortimer, 2019). In existence already was an in-house immersed boundary code for spherical particles; this was coupled to forces facilitating (but having not yet resolved) permanent agglomeration in turbulence. There is significant scope to build upon this existing methodology. The code was written in such a way that it can be extended to non-spherical particles, but further developments are to be made to complete this. There are many further challenges associated with non-sphericity that require significant code development. These include collisions and particle-particle force interactions which require modelling approaches that are known to be difficult to implement, and must be compatible with one another and the fluid solver. Lastly for the developments, the fluid simulations previously undertaken utilised the 'linear forcing method', further develop the accuracy and breadth of capability of the simulation framework.

All of this will first be tested, visualised, interrogated, improved and validated locally in the Python programming language before being integrated with the direct numerical simulation flow solver Nek5000 (written in the FORTRAN 77 language). Upon implementation of the code in the full solver, quiescent fluid conditions can be first utilised to demonstrate the robustness of the method and highlight key elements of the model such as the influence of orientation and velocity on particle agglomeration and dynamics, where fluid effects are minimal. Then, homogeneous and isotropic turbulence will be introduced to the domain to study the interplay between non-spherical particle agglomeration events and turbulent advection of the particles.

The specific objectives of the project outlined at the beginning are as follows:

- Improve the fluid-forcing technique for the generation of homogeneous and isotropic turbulence by implementing a forcing scheme more suitable for particle-laden systems. This should be formulated as an acceleration term that can be added to the Nek5000 '.*usr*' file such that turbulence is forced at each time-step.
- Perform a validation of the improved forcing scheme at two separate Reynolds numbers ( $Re_{\lambda} = 65$  and 143), quantified and compared in terms of time-averaged turbulence quantities and the time evolution of the flow field.
- Implement a collision detection algorithm between non-spherical ellipsoidal particles and test its viability and robustness across the orientational-parameter space as a standalone piece of code in Python. Based upon this, develop the method with any adaptations necessary to improve convergence, robustness or applicability.
- Incorporate this new piece of code as a module to Nek5000 (FORTRAN 77) such that it can be run concurrently with the flow solver in multiphase simulations.
- Implement a hard-sphere collision algorithm that computes physically realistic collision outcomes between arbitrarily orientated ellipsoidal particles, built upon and coupled to the detection scheme. Test and adapt this in Python and then incorporate into the Nek5000 code.
- Generate non-spherical particle meshes for the desired aspect ratios. Provide the necessary extensions to the immersed boundary solver such that the anisotropic particle meshes are properly tracked and rotated. Generalise the function used in the code for the particle's inertia tensor such that arbitrary non-

spherical shapes can be included. Generalise the rigid-body framework, in particular the update of particle torques, for particles of a non-spherical nature. Generalise the particle-phase boundary conditions to non-spherical particles. Generalise the injection code to non-spherical particles to ensure non-overlap.

- Develop a methodology for including non-spherical DLVO forces into the particle advection module such that they experience realistic orientationallydependent forces. Verify the potentials formed between non-spherical particles across the orientational parameter space against existing literature.
- Improve the numerical experimental set-up such that simulations resolve particle agglomeration (and do not always break up due to turbulence strength). Investigate the process of agglomeration and how this interacts with the collision module.
- Investigate the role of orientations and velocities in particle agglomeration in periodic boxes of fluid with quiescent initial conditions for disks, needles, and spheres to provide insight into the role of morphology in agglomeration outcomes, as well as to provide a comparison case for turbulent simulations, and to better understand the proposed methodology in terms of the interaction between the modular components of the code.
- Investigate the role of the strength of turbulence on particle agglomeration for three particle morphologies (spheres, needles and disks) and for three different turbulent flow conditions with properties aiming to approximate a viscous sublayer, buffer layer and bulk region of an associated channel flow.
- Investigate the impact of behavioural modification techniques on the multiphase system in terms of the particle dynamics and interaction tendencies. These modifications should include alterations to the Hamaker constant, coefficient of restitution and the Taylor-Reynolds number of the flow. The same modifications will be made for both the spherical particle simulations and non-spherical simulations, which can then be compared in terms of their relative dependence on the varied parameters.
- Develop a soft-sphere method for particle collisions that can be generalised to non-spherical shapes, making a proper correction for orientational

dependencies. Implement the methodology to run concurrently with Nek5000 in place of the hard-sphere collision scheme, then validate the collision behaviours, and contrast with hard-sphere results.

• Develop a computational framework that allows for the simultaneous computation of many particle interactions. Investigate the types of structures formed when more than two particles agglomerate in turbulence.

# 1.6 Thesis Organisation

To give a brief overview of the content of the thesis, the next Chapter will provide a theoretical background to the computation of turbulent multiphase flows, before focusing on reviewing the relevant literature relating to particle-laden flow dynamics, agglomeration and their numerical implementation. Chapter 3 will cover the methodologies used for all studies presented in the results chapters. Significant code and methodology development has taken place herein and this has been outlined in terms of the theory used and the numerical implementations undertaken. Chapter 4 will be the first results chapter, and will demonstrate a validation of the particle-phase as well as present results on non-spherical particle interactions in a quiescent box. Chapter 5 will demonstrate a validation of the single-phase turbulence and present an investigation into the role of turbulence and morphology in non-spherical particle agglomeration. Chapter 6 will be the final results chapter, assessing the impact of behavioural modification techniques on non-spherical multiphase turbulent systems with recommendations for industry. This will be supplemented by a demonstration of significant development of a non-spherical multi-particle the soft-sphere implementation, with many mathematical and coding challenges overcome to reach this achievement described. Lastly, Chapter 7 will conclude the work and provide recommendations for future work in terms of extending the current approaches and developing new ones based on the insight obtained over the course of the project.

# 2 Literature Review

This chapter provides a review of the particle-laden multiphase flow literature, focusing on the interaction of particles in turbulence and the computational methods commonly used in their study. An effort has been made to first describe the key theory as a foundation for the scientific developments in this thesis. Therefore, the chapter begins with important fundamental fluid dynamics concepts, which can be thought of as a background, containing discussions on the numerical solution techniques of the Navier-Stokes equations and computational fluid dynamics (CFD), as well as the theory of turbulence and its relevant length and time scales. This sub-section is rounded-off with a review of the literature pertaining to the most relevant flow-case to the present thesis – that of homogeneous and isotropic turbulence (HIT).

Then, the focus turns to particles, addressing concepts such as motion, collisions, and interactions. Understanding these individual components in isolation is crucial to comprehending the complex interactions that occur when the two phases coexist. Finally, the review progresses to studies that focus on the interplay between particles and turbulence. This covers particle dispersion, clustering, and agglomeration findings, as well as the associated computational challenges and methodologies.

# 2.1 Fluid Phase

Fluid flows are encountered everywhere in nature and in industry: ranging from the gentle flow of water down a car windshield on a rainy day, to supersonic flows of burnt gases in a rocket nozzle. Fluids, which encompass liquids and gases, are composed of an unfathomable number of molecules at a microscopic level that behave like a continuous medium at the macroscopic level. The continuum approximation, which models fluid dynamic properties at a particular point in space and time as an ensemble average of molecular motion within a small local region, forms the basis of fluid dynamics. This ultimately allows for the formulation of the Navier-Stokes equations, which describe the motion of fluids to an exceptional degree of accuracy. The Navier-Stokes equations are a set of partial differential equations which arise from enforcing conservation of mass and momentum on a small fluid element. Their solution and

analysis underpin much of the research undertaken in the field of fluid dynamics, but they are not yet analytically solvable in the general case, and it remains an open question in mathematics as to whether smooth and globally defined solutions always exist. Fortunately, numerical solution of these equations *is* possible, often with exceptional accuracy, paving a route forward.

Naturally, numerical solutions present their own set of challenges because the equations are highly nonlinear and possess important information at a vast range of scales – often requiring significant computational resources or equation simplifications and models to achieve satisfactory results.

# 2.1.1 Turbulence

The nonlinearity of the Navier-Stokes equations gives rise to a fascinating phenomenon known as turbulence. Characterised by its chaotic and unpredictable motion at a range of scales, turbulence is highly irregular – and it remains poorly understood at a fundamental level, despite considerable attention. The mysterious nature continues to drive academic intrigue, and this is bolstered by its direct relevance to so many natural and industrial applications.

In laminar flows, fluid layers travel parallel to one another in a predictable way. However, as the velocity of the fluid increases (relative to other properties), or disturbances are encountered in the flow field, instabilities arise. It is from these instabilities that turbulence is born. Turbulent flows have greatly enhanced mixing due to their transportation of fluid perpendicular to the mean flow direction, in what are referred to as vortices or eddies, which are the swirling motions associated with turbulent flow. This mixing occurs much faster than achieved by regular diffusion. The increased shear stresses in turbulent flows perform deformation work and thus act to convert kinetic energy in the fluid to internal energy; hence, a turbulent flow is also much more dissipative (Tennekes and Lumley, 1972).

The relative importance of the individual terms in the Navier-Stokes equations (in particular, the inertial and viscous terms) helps researchers to understand conditions necessary for the onset of turbulence. From a physical viewpoint, there is a competition between the viscous forces which lead to the dissipation of energy and a dampening of the flow, and the inertial forces which can ultimately lead to the formation

of turbulent structures. The ratio of inertial to viscous forces is known as the Reynolds number and is defined as:

$$Re = \frac{\tilde{U}\tilde{L}}{\nu},\tag{1.1}$$

where  $\tilde{U}$  and  $\tilde{L}$  are characteristic velocity and length scales, respectively, and  $\nu$  is the kinematic viscosity of the fluid.

This dimensionless parameter helps to characterise a flow and predict its flow regime, i.e. whether it is going to be laminar, turbulent, or transitional. For low Reynolds numbers, a laminar flow regime is expected; whereas, for high Reynolds numbers, that exceed the critical Reynolds number, turbulence is expected to occur. This critical number is geometry specific but, for example, it tends to be around 2000 for pipe flows (Tennekes and Lumley, 1972).

## 2.1.2 Turbulent Energy Cascade

A series of very important developments in the theory of turbulence originate from the seminal work of Richardson (1922) and Kolmogorov (1941). Much of the modelling of turbulence is based on their ideas. It was noted by Richardson (1922) that turbulent flows consist of eddies of different sizes which interact with each other in a special way:

# 'Big whorls have little whorls that feed on their velocity; And little whorls have lesser and so on to viscosity.'

The different sizes of the eddies represent a range of length scales with corresponding time scales. Instabilities in the larger 'mother' eddies cause them to break down into smaller eddies – the 'children'. This breaking down of size due to instability allows the transfer of energy to gradually cascade from the larger scales to the smaller scales. This continues until the kinetic energy is contained within a small enough flow scale that molecular diffusion becomes important, converting it into thermal energy, which is termed 'viscous dissipation'.

Kolmogorov (1941) theorised that, for the smaller scales of turbulence, statistically the flow is locally homogeneous and isotropic, which is to say that the average motion

does not depend on position or direction. The structure of the small scales within all instances of turbulence are therefore similar, regardless of the various mechanisms that might have led to the inception of the turbulence. In this theory, the totality of the small-scale turbulence can be characterised by two parameters, the energy dissipation rate  $\epsilon_F$  and the kinematic viscosity  $v_F$ . Then, dimensional analysis leads to the following properties of a turbulent flow. The Kolmogorov length scale, the scale at which the fluid's kinetic energy is dissipated to heat, is  $\eta = \left(\frac{v_F^3}{\epsilon_F}\right)^{\frac{1}{4}}$  and the associated Kolmogorov time scale is  $T_{\eta} = \sqrt{\frac{v_F}{\epsilon_F}}$ . A real turbulent flow is fluctuating and so, locally, the fluid dynamics at these scales will be fluctuating too. Kolmogorov's theory

therefore only offers statistically representative scales typical of the given flow.

The assumption of local homogeneity and isotropy is justified by the vast difference in time and length scales of the largest motions in comparison to the small ones. In the timeframe of the large motions, the small ones are effectively in statistical equilibrium (George, 2013). To give an impression of scale, atmospheric flows with large-scale motions on the order of kilometres can have a Kolmogorov length scale on the order of millimetres. For an industrial pipe flow with a diameter of around 10cm the Kolmogorov scale could be expected to be on the order of 0.1mm, depending on Reynolds number. More precisely, Bailey et al. (2009) demonstrated experimentally that for a pipe with a diameter of 12cm and a Reynolds number based on this diameter of 70,000, the Kolmogorov length scale is measured at 0.22mm.

George (2013) outlines how Kolmogorov's picture of turbulence may be incomplete, referring to a more complex reality of turbulence that can be inferred from modern flow visualisations and direct numerical simulations. The author suggests that there is a more dynamic interplay at hand between small and large structures, which appear to be inherently intertwined, particularly when it comes to coherent flow structures. The smaller scales can, in reality, also transfer energy back up to the larger scales, but the net direction (of energy transfer) is in keeping with Kolmogorov's theory.

## 2.1.3 Energy Spectra

It has been established thus far that turbulent flows consist of a range of eddy sizes. Each of these eddy sizes contribute to the overall energy content of the flow and this energy is not distributed evenly across the various scales – rather, the energy content forms a distribution known as the energy spectrum. This spectrum provides insight into the relative importance of various length scales in the turbulent flow, indicating the energy content at each length scale and allowing identification of the energetically dominant scales of motion in the flow.

Rather than considering the length of the eddies directly, they can be characterised in terms of their spatial frequency, which is proportional to the inverse of their length. Eddies have a characteristic wavenumber  $\kappa$  based on their length scale  $L_E$ , defined by the simple relation  $\kappa_E = \frac{2\pi}{L_E}$ .





Figure 2.1 illustrates a typical energy spectrum for a turbulent flow which contains subdivisions into the energetic range, the inertial subrange and the dissipative subrange. The energetic range refers to the range of length scales where the bulk of the energy resides; this region is characterised by coherent structures, large-scale vortices, and includes the dominant modes of energy transfer within the flow. The

inertial subrange refers to the wavenumbers where the energy cascade occurs, transferring energy from the small wavenumbers to the large. At which point, the energy reaches the dissipative subrange where the viscous forces dominate, and the kinetic energy is converted to heat.

Kolomogorov's theory predicts a universal form for the energy spectrum in the inertial subrange. That is,  $E(\kappa_E) \sim \kappa_E^{-5/3}$ , which was noted independently by Obukhov (1941). This has been the subject of much validation which has increased the confidence in the overall theory (Dubrulle, 2019).

#### 2.1.4 Measuring Turbulence

Turbulence is nontrivial to quantify, owing to its wide range of scales, structures, and intricate dynamics. This vast spatial and temporal landscape poses an experimental challenge too, in that achieving a high level of resolution for the velocity field in laboratory experiments simultaneously for all scales is very difficult. One of the primary challenges for these measurements is the disruption of the turbulence which can be caused by inserting probes or sensors into the flow field. The presence of measurement devices thus alters the flow characteristics, leading to results that are not representative of the desired system.

To overcome this challenge, non-intrusive measurement techniques are sometimes utilised. Such techniques enable measurement without physical contact with the flow. Advanced imaging techniques, such as particle image velocimetry (PIV) or laser Doppler velocimetry (LDV), allow for the visualisation and analysis of turbulent flows without direct interference. For example, LDV was utilised by Den Toonder and Nieuwstadt (1997) to conduct an investigation of mean flow and turbulence statistics in a pipe, finding a Reynolds number dependence in their statistics. Mochizuki and Nieuwstadt (1996) present an exhaustive list of many experimental endeavours in the pursuit of analysing turbulent statistics and the Reynolds number independence of wall-bounded flows. The study of Den Toonder and Nieuwstadt (1997) is an example of high-quality experiments forming excellent benchmark cases for numerical simulations. It can be argued that progressing the field of fluid dynamics relies on an interconnection between physical and numerical experiments. In some cases, numerical investigation allows the researcher to go beyond what is possible with

current measurement techniques, but validation against experimental data is necessary to ensure the integrity of any such computations.

# 2.1.5 Computational Fluid Dynamics

Numerical modelling therefore offers an avenue to gain insight that is difficult to obtain through physical experiments: perhaps due to an experiment's prohibitive costs or the complexity of recreating the desired system in a laboratory. The arrival of computers powerful enough to compute numerical solutions of the Navier-Stokes equations has revolutionised the field of fluid dynamics. Since then, computing power has continued to increase, meaning an ever-greater applicability of this avenue of investigation (Agarwal and Lewis, 1992; Jamshed, 2015).

By discretizing the governing equations, the computational domain, and the evolution of time, computers can iteratively obtain solutions to the Navier-Stokes equations that become increasingly accurate as the resolution (the number of discrete points in space or time where the solution is obtained) is increased. This allows a spectrum of fidelities to be achieved at the discretion of the user. In the limit of increasing resolution, these solutions are bounded only by computational resources and the accuracy of the equations themselves.

However, it is important to balance this enticement with an appreciation for the limitations. Just like experiments, mistakes appear, either through programmer error or oversights in the formulation of the numerical method. Careful consideration must therefore be made such that obtained results can be taken forward with confidence. This is achieved by first assessing assumptions made in the modelling stage. These assumptions guide the formulation of the numerical model and help capture the essential physics of the problem at hand. Making any such assessment therefore requires understanding the underlying physics, including identifying physically realistic boundary conditions, and understanding the limitations or requirements of the numerical methods employed.

After this, it is important to validate against previous work, ideally experiments, wherever possible. Comparisons with experimental results allow for a critical evaluation of the model's ability to reproduce real-world phenomena. Sometimes a validation of the model against adjacent systems must suffice where direct validation

is not possible. In the case where such validation is not possible, it becomes doubly important to understand the system being modelled to the fullest extent possible, as well as to understand the common challenges associated with the implementation. This means that CFD research occurring at the forefront of current capabilities requires a real appreciation of the physics, since there often is no means of comparison or specific guidance.

The vast range of scales inherent to turbulent systems means that significant resolution is required, but there are certain models and methodologies that look to circumvent this problem. Computational resources can be spared at the expense of accuracy, and it is up to the investigator to decide where to draw the line. Often this decision will be related to requirements of the academic or engineering challenge that is being investigated. For fast and instructive simulations aimed at determining key bulk phenomena, a lower degree of accuracy will suffice. To generate benchmark solutions or investigate fundamental phenomena, a greater level of accuracy is required. Sometimes, the very small scales of the flow have a large effect on the macroscopic behaviour of the system and so resolving these scales is necessary in obtaining sound results. This is the case with accurately resolving boundary layers on an aerofoil, or the dynamics of small particles immersed in a flow field.

One way to reduce computational load is through intelligent meshing – where the 'mesh' is the colloquial term for the numerical grid that overlays the domain to obtain a solution. Grid nodes, or mesh points, need not be evenly spaced across the domain and so it makes sense to have higher degrees of resolution at the points where it is needed, for example where the velocity gradients are greatest. Running test simulations and performing grid-independence studies can also help to instruct the design of the mesh, and so too can understanding the flow phenomena associated with certain geometries. Further still, one can use adaptive mesh refinement techniques which dynamically adjust local resolution to concentrate the computational efforts at the areas that require it (Hu et al., 2001).

Another simplification that is widely employed is the use of a turbulence model. Rather than requiring very high resolution to capture every velocity fluctuation, knowledge about turbulence can be used to construct models that impact the mean flow in the same way that turbulence would. Such a model will then operate on the same length and time scales as the bulk flow allowing for reduced computational requirements. These models essentially describe the averaged effects of turbulence.

## 2.1.6 Reynolds-Averaged Navier-Stokes

The Reynolds-Averaged Navier-Stokes (RANS) simulation technique is widely adopted in industrial settings and underpins many of the commercial CFD solvers that are used today, thanks to its fast compute time. The method relies on averaging the Navier-Stokes equations to obtain a set of equations that represent the mean flow, as alluded to in the previous subsection.

Specifically, a Reynolds decomposition is employed where the flow is split into a mean and a fluctuating component. Averaging gives the RANS equations, but the nonlinearity of the Navier-Stokes equations means that a term known as the Reynolds stress still contains a dependence on the velocity fluctuations. This is known as the *closure problem*. It is this term which requires modelling. Commonly used to close the equations are the  $k - \epsilon$  and  $k - \omega$  turbulence models, which have many subdenominations in wide use (McConkey et al., 2021).

With the small-scale behaviour averaged out, RANS allows the use of a relatively coarse numerical grid. This technique often drives towards a steady-state solution which means no requirement for time-stepping, and a further reduction in computational cost. Cases with time-dependence retained are known as unsteady RANS or URANS, utilised for applications with dynamic processes.

## 2.1.7 Large Eddy Simulation

As the name suggests, large eddy simulations (LES) remove the smallest eddies from the flow – through a filtering technique. These small scales require the most computational resources and so efficiency is increased. Much of the turbulence is still modelled, typically corresponding to around 80% of the turbulence energy content (Pope, 2000), and this is resolved all the way down to a certain cut-off length scale. Below this, a sub-grid scale model is employed which captures an averaged effect. This technique still obtains a reliable solution and lies between RANS and DNS in terms of accuracy.

# 2.1.8 Direct Numerical Simulation

Direct numerical simulation (DNS) involves no additional modelling of the flow. All time and length scales are resolved down to the Kolmogorov scales described earlier. This often requires a very dense mesh with many grid nodes and a very small timestep to recover the dynamics on the order of the Kolmogorov scales. If these requirements are met, then the generated results are closely aligned with what the Navier-Stokes equations predict. As such, they are a very strong investigative technique for studying fundamental fluid phenomena in detail and offer a good avenue for research into turbulence. Naturally, this is not a practical technique for industrial applications due to the very long compute times. Access to high performance computing systems can aid with the practicality but these are often not used in current industrial settings; however, their use is increasing (Government Office for Science, 2021).

# 2.1.9 Further Solution Methods

These three simulation techniques span the levels of resolution available for the computation of the Navier-Stokes equations. However, within these techniques there are further denominations and methodologies. The way the solutions are obtained in terms of the discretisations used, the mesh topologies, and the way information is passed around the mesh, are all areas of difference. For example, finite-volume, finite-difference, finite-element, and spectral methods (Boyd, 2001), are all numerical methods for the computation of partial differential equations and they all are used in CFD applications to generate solutions to the Navier-Stokes equations (Aref and Balachandar, 2018).

Further still, simulation techniques exist which do not solve the Navier-Stokes equations directly, like the Lattice-Boltzmann method which has wide applicability to complex geometries due to its loose boundary conditions (Chen and Doolen, 1998). Smoothed particle hydrodynamics offers another avenue and is a method that does not require a mesh, which comes with its own advantages and challenges (Monaghan, 1992).

# 2.1.10 Homogeneous Isotropic Turbulence and the Taylor-Reynolds Number

Homogeneous and isotropic turbulence is a type of idealised turbulence in which the time-averaged statistics are independent of position and direction (Taylor, 1935). It typically exists within a periodic box. These boxes of turbulence offer a unique way to study turbulent behaviour free from the influence of mean flow or geometry, which can be particularly relevant to multiphase flows, since the interaction between particles and turbulence can thus be studied in an isolated way for the purpose of generating fundamental insight.

In homogeneous and isotropic turbulence, it can be useful to introduce a parameter known as the Taylor microscale, used to characterise turbulence levels. The Taylor microscale represents the length scale at which the viscosity of a fluid significantly affects the dynamics of its eddies. This length falls within the inertial subrange and is larger than the Kolmogorov length scale (which falls in the dissipative subrange). The Taylor microscale (Tennekes and Lumley, 1972) is given by:

$$\lambda = \sqrt{\frac{15\nu_F}{\epsilon_F}} u', \tag{1.2}$$

where u' is the root mean squared (RMS) velocity fluctuation. This can then be used to define the Reynolds number based on the Taylor microscale, known as the Taylor-Reynolds number:

$$Re_{\lambda} = \frac{\lambda u'}{\nu_F}.$$
 (1.3)

#### 2.1.11 Forcing of Homogeneous and Isotropic Turbulence

As established in the background section, at the smallest scales of turbulence, energy is converted to heat and dissipated from the system. In the case of periodic boxes of homogeneous and isotropic turbulence, this mechanism eventually acts to dissipate all of the turbulence kinetic energy until the fluid is motionless at the continuum level. This is due to the absence of any inherent turbulence generating mechanisms like boundary-induced shear. Hence, for sustained turbulence, one must introduce a numerical scheme that replenishes the dissipated energy. Such efforts are known as *forced* approaches – in contrast to *decaying* approaches, which involves an initial condition for the velocity field being chosen possessing energy and, as time evolves, the field is allowed to decay.

The obvious disadvantage of a decaying turbulence field is the limited time window in which investigation can take place. In the context of particulate flows, this may be prohibitive to resolving desired phenomena such as multi-particle agglomeration structures, which can take a significant period of time to form (Spoelstra, 1989). Forced turbulence circumvents this problem and allows the turbulence field to become statistically stationary, often allowing the opportunity to select specific desired turbulence properties. This makes it a useful investigative tool. Outside of the context of periodic boxes, these same turbulence forcing techniques can be used to promote turbulence or to introduce turbulence-like randomness into RANS solutions, for example in aeroacoustic noise generation (Béchara et al., 1994; Bailly and Juve, 1999).

There are several successful approaches to the problem of forcing, most of which involve the addition of a forcing term to the Navier-Stokes momentum equation. One widely used method is that of Lundgren (2003) who introduced a forcing term proportional to the local velocity field. Lundgren considered the turbulence kinetic energy budget equation and noticed that the production term proportional to the velocity is absent in the case of homogeneous and isotropic turbulence, hence the author derived a forcing scheme that enhances the production of energy. The class of methods that have followed this approach are known as linear forcing schemes, for example that of Rosales and Meneveau (2005).

Lundgren's approach is simple to implement and effective at producing turbulence. However, it is well-established that the integral length scale of the turbulence is limited to 0.13 times the size of the box (Janin et al., 2021) which limits the range of turbulent systems available to the investigator. Beyond this, the turbulence is slow to become statistically stationary and has been observed to diverge over time particularly with the inclusion of particles (Chouippe and Uhlmann, 2015). Moreover, the transfer of kinetic energy from particles to the fluid is interfered with by linear forcing, since energy is introduced at all wavenumbers, including those at which the particles are having an influence (Mallouppas et al., 2013). This makes such a scheme problematic for particle-laden flows and inappropriate for two-way coupling regimes and beyond (Lucci et al., 2010).

An alternative approach to the linear forcing method, and to generating the forcing in physical space more broadly, is the construction of the forcing scheme in Fourier space. Kraichnan (1970) laid the mathematical work for many forcing methods that would follow, e.g. Mallouppas et al. (2013) and Janin et al. (2021). The idea of this method is to construct a synthetic velocity field in Fourier-space which can be used to generate a forcing field for the fluid. The random Fourier modes are summed over a spherical shell in such a way that statistical homogeneity and isotropy are recovered. The method takes as an input an energy spectrum which defines the distribution of energy added to the simulation through the forcing. One key advantage of this method is the ability to choose the wavenumbers of the fluid to which this energy is supplied. The common pertaining exploitation is to only supply energy to the smallest wavenumbers of the fluid, which is equivalent to driving the largest scale motions. The smallest scales of the turbulence are then developed by the energy cascade. In the case of multiphase flows, this removes the interference with length scales relevant to the particle-phase, demonstrated by Mallouppas et al. (2013).

Some forcing methods depend on the resolved velocity field at each timestep, in which case the method is said to be *deterministic*. One advantage of this approach is to have more control over fluctuations in the global statistics of the flow during the simulation. For example, if the turbulence kinetic energy (TKE) drops due to the natural cascade of energy, more energy can be supplied to rectify this immediately. This controlled modulation of the turbulence field allows a reduction in the oscillations of such properties as the kinetic energy. Janin et al. (2021) introduced a synthetic source term into the Navier-Stokes equations based on the method proposed by Schmidt and Breuer (2017). In their work, the amount of energy to be injected at each timestep is drawn from an energy spectrum, and so they introduced a selective forcing scheme that only draws from the pre-defined spectrum up to a cut-off wavenumber, such that the correct amount of energy is supplied. This cut-off wavenumber can be dynamically controlled by considering the balance of required energy with the dissipation in the system. The result demonstrated by the authors was vast reductions in the oscillations of TKE in the system over time.

Another advantage of deterministic schemes is faster convergence towards the desired TKE levels, which is particularly pronounced in forcing schemes that use a modulation coefficient based on this desired energy, e.g. in the methods proposed by Carroll and Blanquart (2013), Mallouppas et al. (2013) and Bassenne et al. (2016).

Eswaran and Pope (1988) and Alvelius (1999) are two examples of approaches that generate the forcing field independently of the resolved fluid velocity. Such schemes are said to be *stochastic*. Eswaran and Pope (1988) rely on the use of Uhlenbeck-Ornstein random processes to drive the evolution of the field. The advantage of this approach is most obvious in application to multiphase flows, wherein a source of instability is removed that arises from the forcing being dependent upon the velocity (Chouippe and Uhlmann, 2015). Secondly, it is immediately clear what alterations to the (sufficiently time-averaged) fluid quantities were a result of the introduction of the particle-phase, since there is no interaction with the forcing scheme.

# 2.2 Particle-Phase

## 2.2.1 Introduction and Relevance

There are various numerical simulation techniques for capturing particle behaviour on many scales, such as molecular dynamics, the discrete element method (DEM), dissipative particle dynamics, and stochastic methods (Español and Warren, 2017). The work herein involves macroscopic particles treated as discrete and distinct entities and so the discrete element method (Cundall and Strack, 1979) and the pertaining literature is of relevance. There, particles are tracked individually with motions updated according to Newton's second law, with contact mechanics resolved. Whilst the present work considers the interplay between the motion of particles and a fluid, the modelling techniques used to study particles in isolation can be instructive and, in some cases, directly applicable. Two key cases where this holds true are for the modelling of particle-particle collisions and particle-particle attractive interactions. Moreover, non-sphericity remains challenging to numerical studies, and many of the first steps to include the physics of non-sphericity have taken place in a simpler context that does not include a fluid.

# 2.2.2 Non-Spherical Challenges

The modelling and study of non-spherical particles is much more challenging than the spherical counterpart because of the orientational degrees of freedom and the additional complexity this introduces; for example, how the forces and torques experienced by the particles depend upon specific orientations. Further still, these new degrees of freedom can interact with the translational degrees of freedom in nonlinear and complex ways.

The first challenge lies with tracking the orientation of the particles over the course of a simulation. For rigid-body particles, this is typically done through the use of Euler angles (Lu et al., 2015) which require an associated rotation convention (Goldstein, 1950). There are multiple such conventions of choice, but the result is typically the construction of a rotation matrix which can be used to update the orientations of rigid particles. Evans and Murad (1977) developed an approach with a quaternion formulation of the orientation matrix, using Euler parameters, that allows for the robust and efficient tracking of particles. The use of quaternions in favour of Euler angles is

preferable since there are singularities inherent to the Euler angles which make them inefficient to compute, whilst there is no such problem with quaternions (Fan and Ahmadi, 1995).

The quaternion approach was further developed by Zhao and van Wachem (2013) such that there is no requirement for an orientation matrix and rotations can be computed directly from the quaternions. These methods have been successfully applied to research problems such as the discrete element modelling of non-spherical particles (Langston et al., 2004) and two-phase channel flows (Mortensen et al., 2008; Njobuenwu and Fairweather, 2013a). The rotations and angular velocities of the rigid particles will then typically be calculated according to Euler's rotation equations (Lu et al., 2015).

## 2.2.3 Collisions: soft- and hard-sphere

When two bodies come into contact there is an exchange of momentum and a loss of kinetic energy associated with the collision, this needs to be modelled in a simulation with particles to ensure the kinetic energy of the particle-phase remains physically realistic (Machado et al., 2012). There are many options when it comes to computing this process. At the most basic level of modelling, instead of computing collision mechanics, particle overlaps can be prevented without a physical model at all. Particles found to be overlapping could simply be moved such that they are no longer overlapping, or a non-physical repulsive force can be implemented normal to the particle surfaces preventing particles getting close enough to touch (e.g. Shardt and Derksen, 2012). At the other end of the spectrum, full tracking of the particle surface and its deformation under load can be computed and even coupled to a fluid-phase in the case of fluid-structure interaction problems. The most common approaches lie between these extremes. They are referred to as *soft-sphere* (Tsuji et al., 1993) and hard-sphere (Hoomans et al., 1996) collision models. Both look to model the kinetic energy loss in a collision event. These approaches are effective and computationally efficient hence why they predominate. In the literature, hard-sphere collisions may also be referred to as 'event-driven' and a soft-sphere collision as 'time-driven', in keeping with their respective implementations being instantaneous and temporal, respectively.

Generally, hard-sphere collisions do not permit the overlap of particle surfaces for more than a single time-step and some models do not allow an overlap at all. In the latter case, a collision is deemed to have occurred when the distance between particle surfaces is less than a small prechosen value. Momentum conservation is used to determine the mechanics of the collision, and this is computed in a classical way, reminiscent of pool-ball calculations undertaken in a high-school physics class. Except now, a parameter known as the 'coefficient of restitution' is included to model energy loss. The calculation takes place over a single time-step and can therefore be thought of as an impulse force model. Despite being a simple model, it is effective for materials which do not deform much under contact and have a known coefficient of restitution obtained through experiment. It is worth noting that a complexity arises for anisotropic morphologies (for both hard and soft models), where properties such as the coefficient of restitution are dependent upon the exact position on the surface that an impact occurred (Wynn, 2009).

The soft-sphere model offers slightly more resolution and does permit overlap. Collisions under this model will take place over a number of timesteps with the strength of the collision force depending upon the depth of the overlap. A Hertzian contact model is employed under the assumption that deformations to the surface can be approximated to first order. The implementation is typically handled through a linear spring-dashpot system in which the coefficients used in the model are parameterised according to the material properties of the particles (Costa et al., 2015).

It is not clear how to effectively handle multiple collisions in a single time-step between many particles in the hard-sphere model and so the soft-sphere approach holds an advantage for denser systems. It is for this reason that soft-sphere modelling is more widely used in DEM simulations (Ma et al., 2022) and commercial packages where many particles are present, for example in granular flows or the formation of sedimentary beds. Since the compute time that is required scales with the collision frequency, hard-sphere methods are preferable in the context of dilute systems. Buist et al. (2016) were able to combine the advantages of both in their creation of a hybrid scheme that uses an adapted hard-sphere methodology for binary collisions and a soft-sphere methodology for collisions involving multiple particles.

Non-spherical collisions are significantly more complex to resolve due to their orientational dependency and the breaking down of many of the assumptions underpinning spherical collision modelling. The computation of the resultant physics

depends on the direction of the vector normal to the tangent surface at the collision point, and this is nontrivial for anisotropic particles. For a typical collision between nonspherical particles, a torque will be induced by the fact that the normal vector along which the collision force acts very rarely points towards the centre of mass. Further, this creates confusion relating to the coefficient of restitution which is typically defined as the ratio of the post-collision normal velocity to the pre-collision normal velocity and as such, energy conservation in terms of the rotational degrees of freedom is often not properly considered for non-spherical particles, unless direct consideration is made of the contact point location which can be difficult to define. It also creates a challenge when it comes to choosing a coefficient of restitution to match a specific material, when it is very likely that this value has been obtained for normal spherical collisions.

Regarding the numerical implementation, the first key element of a collision algorithm is the contact detection scheme. For spheres, this is very straightforward. One can simply subtract the sum of the radii from the distance between centres to find the distance between surfaces. It is much more complicated in the case of non-spherical particles, where determination of the contact can take around 80% of the particle-phase computation (Williams and O'Connor, 1999). If the position of the particle centres is fixed in space, the distance between surfaces remains dependent on the relative orientation of the bodies. Hence, it is necessary to include an algorithm that can determine the contact, or closest distance vector, between particles. This vastly increases the computational requirement of the collision module in comparison to that of spheres. However, this may not greatly increase the overall compute time in multiphase systems if the solution of the fluid-phase is orders of magnitude more intensive; for an extreme example, imagine the case where just two particles are coupled to an entire channel flow DNS.

For cases where there are many particles interacting, and potentially colliding, it is known that the number of particle interactions scales with the number of particles squared. There are, however, efficient algorithms that reduce the computational complexity and ought to be utilised for calculations with many particles. For example, Breuer et al. (2012) presented an efficient contact detection search algorithm that uses two staggered grids, reducing the computational complexity by an order of magnitude such that it is on the order of the number of particles.

In a review of recent advances in non-spherical DEM modelling, Ma et al. (2022) outline two key categories for collision detection algorithms between super-ellipsoidal (an extension of ellipsoidal) particles. These are the common normal and geometric potential approaches. In the former approach, the geometrical fact that a common normal vector, tangent to both surfaces, must be shared at the closest point between convex bodies, can be leveraged to devise iterative schemes that drive towards the points satisfying this condition (Jain et al., 2019) using a variation of gradient descent. Alternatively, the problem can be written as an optimisation problem to be solved by Newton's method (Wellman et al., 2008). The geometric potential methods do not always satisfy the aforementioned condition which can lead to different results between approaches, but for small overlaps they tend to be closely aligned.

Throughout the literature, the geometric potential method is preferred for its quick convergence (Dziugys and Peters, 2001; Gan et al., 2015). The idea of the geometric potential method is to devise a function – the geometric potential – that is minimised by the points of deepest penetration in opposing particles. This function is less than one for overlapping particles, equal to one for contacting particles, and greater than one for particles that are not touching. The minimum of the function in these cases then defines the deepest overlap point, the contact point, and the closest point of approach, respectively.

Donev et al. (2007) used an overlap potential method for contact detection and resultant forces in the context of studying jammed packings of hard ellipsoids. Gan et al. (2015) used the geometric potential method for determining overlaps, contacts, and distances between attractive fine ellipsoidal particles in CFD-DEM simulations of fluidisation and again in simulations of packing (Gan et al., 2016); however, no attempt is made to account for particle shape in the resulting contact mechanics.

There are also approaches to the problem of contact detection that use particle meshes which already exist for use elsewhere in the simulation. Nagata et al. (2019) used image points reflected over the particle boundary, which were already there as part of their ghost-cell immersed boundary method. These image points were reused as probes to check for collisions between non-spherical particles. Shardt and Derksen (2012) simulated dense suspensions (~45% volume fraction) of red blood cells with resolution of their bi-concave shape, and with collisions accounted for by distributing

normal vectors across the surfaces of the particles. These normal vectors were iterated over and checked for proximity to other normal vectors on neighbouring particles. Any distances below a small threshold were deemed to represent a contact and a repulsive force was applied along these normal vectors proportional to a computational parameter chosen to represent the strength of the collision force. The same normal vectors were also used in their immersed boundary scheme, saving computational resources. A full order of magnitude difference was observed in the sedimentation rates measured in their calculations when compared to a reference medical experiment. This was attributed to an absence of agglomeration-inducing forces and serves to highlight the importance of including such forces in the modelling of systems where they are acting.

The robust and effective collision models developed for spheres are often extended and applied to non-spherical particles. There is not one agreed upon way to do this, and so different methods can generate different results (Ma et al., 2022). When it comes to non-spherical soft-sphere collisions and their requirement for an overlap, it is not agreed upon how the contact point should be defined during overlap (Lu et al., 2015). The two foremost approaches are the use of the geometric centre of the overlapping volume, or the midpoint of the line that connects the two deepest points of overlap on the respective bodies.

Upon identification of a contact, there remains the challenge of computing the resultant mechanics. For non-spherical particles, there lacks an all-encompassing framework for this computation and there is scope for the development of such. Jain et al. (2019) published an approach to modelling non-spherical hard-sphere collisions in a viscous fluid, with analysis undertaken to determine the requisite forces and torques, this included a novel contact detection scheme using a common normal method. This was complemented by a model for lubrication forces acting on a scale not resolved by the computational mesh wherein non-sphericity was accounted for using the local curvature of the body. Their initial findings demonstrated an orientational dependency in the resolved coefficient of restitution and an interplay between this parameter and the lubrication forces. Jain et al. (2020) successfully applied this method in a further study for the investigation of bedload sediment transport. In a comparison of collision detection schemes, Girault et al. (2022) reported issues with the convergence of this scheme for small separation distances, for which the original authors published an

addendum (Jain et al., 2022). In this addendum, the use of the scheme was not recommended due to the computational resources required at small separations.

In the algorithm comparison paper of Girault et al. (2022), further methods for computing the distance between ellipsoidal particles were compared and assessed in terms of their applicability to particle-resolved direct numerical simulations. As well as the above common normal method, they contrast a Newton-Coulomb method (Abbasov, 2015), the Gilbert-Johnson-Keerthi (1988) algorithm, the penalty function method (Tasmasyan, 2014), an algebraic method (Uteshev, 2008), and the moving balls method (Lin et al., 2001). The comparison resulted in the Newton-Coulomb approach being ruled out for its divergent behaviour. The penalty function method and the common normal approach of Jain et al. (2019) were ruled out for computational efficiency reasons. The moving balls and Gilbert-Johnson-Keerthi methods were found to be comparable in terms of their robustness and speed, and well-suited to particle-laden DNS.

In a study of non-spherical particles in a channel flow, van Wachem et al. (2014) constructed their non-spherical particles entirely from spheres in the collision detection step and checked for collisions between the spheres that represented their abstract shape – as did Langston et al. (2004). The former was coupled to a soft-sphere approach, citing the spherical approach of Mindlin and Deresiewicz (1953), under an axisymmetric assumption about the contact points. It is not stated if there was any effort to include non-sphericity into the collision mechanics.

The literature on soft-sphere methodologies for non-spherical particles is quite sparse for multiphase flows where a proper correction has been made to the soft-sphere force to account for shape and orientation. In the context of an immersed boundary method, Ardekani et al. (2016) implemented a non-spherical soft-sphere model based on one derived for spheres by Costa et al. (2015), with a correction made for the respective curvatures at the contact point on the interacting particles. This curvature was used to generate dissimilarly sized spheres to be used in the linear spring-dashpot model which in turn approximate the correct interaction magnitude.

Wynn (2009) presented an approach for elastic ellipsoids colliding with a wall, using a Hertzian formulation. Unlike typical hard-sphere and spherical approaches, the normal stiffness of the body changes with local curvature in this model. Using this framework,

it was shown that implementing a constant coefficient of restitution over the entirety of a non-spherical body is a significant simplification to the collision mechanics – highlighting a limitation of hard-sphere approaches, such as that of Jain et al. (2019), and soft-sphere approaches that rely on a constant coefficient of restitution.

#### 2.2.4 Van der Waals and Interparticle Forces

The van der Waals (vdW) force is a distance-dependent attractive interaction between particles. It can hold matter together at a range of scales, from the atomic to the colloidal. The force is fairly weak in comparison to those involved in chemical bonds and, as such, particles combined under this force are more susceptible to being separated. However, there are certain systems where this force will dominate over other typically stronger interaction forces.

To understand the origin of van der Waals attraction, quantum mechanics must first be considered. The attraction arises due to the probabilistic nature of the electron cloud of a particle. Randomly, a quantum fluctuation can occur whereby one side of a particle becomes more positively or negatively charged due to an instantaneous local increase in electron density, thus creating a dipole. A nearby particle will have a dipole induced by the already slightly polarised particle and hence the particle pair will experience a positive attraction. Under these circumstances, two particles can combine without covalent or ionic bonding (Winterton, 1970).

Throughout the literature, the term 'van der Waals force(s)' is used loosely. Sometimes it refers solely to the London dispersion force, which is described in the previous paragraph, and sometimes to that force plus the Keesom forces and Debye forces. In the literature, 'van der Waals force' also indiscriminately refers to solitary interactions between pairs of molecules, which can be referred to as *microscopic vdW forces*, as well as to effects at a larger scale. In the latter case, many such microscopic interactions are taking place between collections of molecules forming colloidal bodies; if the net effect of the sum of these interactions is considered, then this can be thought of as the *macroscopic vdW force* which operates on a larger scale. The present work considers interactions between particles at the colloidal scale and so hereafter 'van der Waals force' is used in the macroscopic sense. Specifically, to refer to the net attraction experienced between colloids as a result of averaging the molecular-scale van der Waals forces across all molecular interaction sites.

In a seminal paper, London (1930) showed that the interaction energy of two similar molecules is proportional to the inverse of the distance to the sixth power, through a quantum mechanical approach. This gives an approximate solution to the force experienced between two non-polar molecules at a microscopic level. Subsequently, Hamaker (1937) was able to derive a macroscopic expression for the interaction energy between two colloids suspended in a liquid, wherein the interaction energy is found to be proportional to the inverse of the surface separation distance. This derivation relies on the assumption of pairwise additivity, wherein the interaction between two atoms is treated independently of the influence from neighbouring atoms. For a pair of macroscopic bodies, these interaction pairs are summed to obtain an approximate total interaction – this is only 'approximate' because each of these interactions are interconnected with one another and not truly separate. Despite this simplification, this theory has been found to agree well with experiment (Magan and Sureshkumar, 2005).

Homogeneously charged particles dispersed in a fluid can also form an electrical double-layer of ions. The first layer is adsorbed onto the surface from the surrounding fluid creating a charged surface. This charge then attracts counterions via the Coulomb force and hence a double layer is formed. Israelachvili (1992) gives a model to account for this effect based on the chemical properties of the system, with an exponentially decaying separation distance dependence. This model was successfully used by Fujita and Yamaguchi (2007) to study the behaviour of suspensions of nanoparticles.

When the effects of the van der Waals attractive forces and electric double layer repulsive forces are considered under a single framework, with both effects combined to form a single distance-dependent interaction potential, this is known as DLVO theory, named after the important contributions of Derjaguin and Landau (1941) and Verwey and Overbeek (1955).

#### 2.2.5 Numerical Implementation of Van der Waals Forces

Including vdW forces to multi-particle simulations can lead to the resolution of emergent phenomena such as agglomeration (Abbasfard et al., 2016) and fluidisation (Ye et al., 2004; Gu et al., 2016). The numerical implementation of these forces between colloidal particles under the DLVO framework provides a means to

approximate the physics. However, there are some limitations with the current modelling approaches. The most recognised issue is the requirement for an intersurface cut-off distance in the vdW potential. As two particles approach zero distance, the potential approaches infinity, since the expression is unbounded. To alleviate this, a distance is selected below which the expression will no longer increase in magnitude for decreasing separation distance. This cut-off distance is somewhat arbitrarily chosen throughout the literature, typically taking on values between 0.1-1.0 nm (Abbasfard et al., 2016). The lack of an established universal methodology for selecting this parameter means that this is a big source of discrepancy between studies and places a limitation on the repeatability of numerically obtained results.

Moreover, Abbasfard et al. (2016) demonstrated that the critical velocity, beyond which two particles will bounce rather than agglomerate, is strongly dependent upon the value of the vdW cut-off distance. Mihajlovic et al. (2020) highlighted that it is the full interplay between this parameter, the Hamaker constant, and the numerical resolution of the simulation in terms of how well sampled the various forces are during the particle interaction, that determines whether two particles will rebound or agglomerate. If the necessary parameters are available, perhaps through accompanying experiments, then the cited study proposes a methodology for defining the cut-off distance. Based on the maximal values of the van der Waals force observed in pertaining experiments, and with a given Hamaker constant, the proposal is that the cut-off distance can be tuned such that the maximal van der Waals forces are matched between numerical and physical experiments. Sadly, without this full set of parameters from experiments, there is still the same uncertainty around how to select the cut-off distance and it becomes a computationally tuned parameter. Even theoretical approaches to rectifying this issue are destined to fall short, since the surface roughness proves to be very important in the determination of maximal van der Waals forces, and analytical derivations typically assume smooth surfaces which can increase the strength of the maximal forces by orders of magnitude.

A further challenge is introduced when considering soft-sphere collisions: the particle surfaces can overlap and the DLVO potential does not account for this. Mihajlovic et al. (2020) investigated this problem in the context of fluidisation. They developed an approach in which the vdW force is excluded during particle overlap, guided by the theoretical work of Rietema et al. (1993). It was shown that including the vdW force

during overlap made the determination of simulation parameters unclear and therefore the best approach was said to be turning attractive forces off at the point of collision. It was further recommended that particle overlaps should be no more than 1% of the radius of the particle and the distance a particle travels in a single timestep should ideally be 10 times smaller than the cut-off distance for the vdW force, such that the maximal forces can be properly sampled. Similarly, Kobayashi et al. (2013) recommended a dynamic setting of the adhesion force in related studies, since smaller values of spring stiffness in the soft-sphere model can cause overestimation of adhesion effects.

Another challenge presented when modelling these systems is the disparity between time and length scales associated with the different forces. If particles are mostly advected by a macroscopic behaviour, like through hydrodynamic forces, then for a general simulation the period that DLVO forces are important will be incredibly small in comparison to the macroscopic time scales. The vdW forces will only begin to spike when particles are at a very short separation distance, much smaller than the length of the radii. This poses a challenge with respect to computational efficiency if one is to fully resolve the physics and avoid resorting to a model for the short-range behaviour. It is unnecessary to have a very high level of resolution to calculate the interparticle forces when this is only required for a few timesteps, i.e. when the DLVO forces are significant. Equally, it is important not to choose a timestep that adequately resolves the particle advection due to the fluid behaviour but misses the level of fidelity required to resolve these short-range interaction phenomena. One solution is to use a variable timestep based on interparticle distance; another is to use a separate timestep for the fluid- and particle-phases. The first solution is only applicable to systems with a limited number of particles. For dense systems with many particles, the timestep is almost always very small and thus the computational load is barely reduced in comparison to just setting a very small timestep to begin with (Mihajlovic et al., 2020).

#### 2.2.6 Non-Spherical Interparticle Forces

As evidenced, the modelling of interparticle forces poses a significant challenge. This is increased further by the orientational degrees of freedom associated with non-sphericity. The interparticle forces under investigation herein are reliant upon the surface characteristics of the interacting particles and so models can be developed

based on these properties; in particular, on the surface curvature properties at the nearest points of approach, where non-spherical corrections can be applied to spherical models (Zhou et al., 2011).

Hamaker's (1937) derivation for the interaction between macroscopic bodies is valid only for spheres and involves carrying out a double integral, or sum, over the molecular interaction sites between two particles. It is not immediately clear how to incorporate arbitrary orientations into the evaluation of this integral. Everaers and Ejtehadi (2003) conducted a systematic approach to this problem for spheroids where they first evaluated the integrals using a Gay-Berne potential for pole-pole interactions where the axes of the spheroids were aligned, which simplified the problem and allowed for analytical solutions. They used these modal cases to inform their full model over the full orientational parameter space. This approach uses the analysis of White (1983) in deriving a full expression for arbitrary orientational configurations. The result is a continuous model that accounts directly for spheroidal shape.

Schiller et al. (2011) also used the analysis of White (1983) and concepts from differential geometry to derive pair potentials for interactions between spheroidal bodies. This is generalised and thus not only applicable to van der Waals forces but also to capillary, depletion, and electrostatic forces. The paper reports that attractive forces between spheroidal bodies leads to an alignment between particle symmetry axes, whilst repulsive forces discourage this. Preliminary experiments performed in the present work demonstrate that the approaches of Everaers and Ejtehadi (2003) and of Schiller et al. (2011) are equivalent for vdW forces between arbitrarily orientated spheroids.

# 2.3 Multiphase Flow

## 2.3.1 Background and Theory

The complexity of turbulence and the associated challenges have been introduced. Adding a dispersed-phase only increases this complexity, with many interdependent non-linear emergent behaviours. How the particles influence the fluid must be considered, as must how the fluid impacts individual particles and the particle-phase as a collective, as well as how the particles interact with one another. The bulk behaviour of particulate flows is often quantified in terms of the particle dispersion, deposition, resuspension, and interactions, as well as the turbulence augmentation or attenuation as a result of the particle-phase. There are many interesting emergent phenomena that have been observed through consideration of these parameters. Each of which is impacted by the flow regime.

The bulk parameters of a multiphase flow allow a classification of the flow regime; these parameters include volume fraction, particle-to-fluid density ratio, particle distribution, and particle size. The volume fraction,  $\phi_P$ , is simply the ratio of the volume occupied by the particles to the volume occupied by the fluid and particles, or the total volume. Despite its simplicity, it is a powerful parameter giving direct insight to how the particles are coupled to the flow. Using this tool paired with the Stokes number (a measure of particle response time, precisely defined below) one can predict well the flow regime of a given particulate flow (Elghobashi, 1994), and hence make *a priori* predictions as to how the multiphase system will behave.

Generally, increasing volume fraction leads to an increased effect upon fluid turbulence. For the case  $\phi_P \leq 10^{-6}$ , the flow is said to be *one-way coupled* (Elghobashi, 1991). This is where the impact of particles on the fluid turbulence (and the flow in general) is so minor that it can safely be ignored. The dynamics of the particles are determined by the fluid, but the particle-phase's momentum is too insignificant to impact back upon the flow field. Beyond this lower bound, particles do have a momentum transfer effect on the fluid. Namely, in the interval  $10^{-6} \leq \phi_P \leq 10^{-3}$ , the flow is said to be *two-way coupled*. In this region, sufficient momentum transfer is occurring to alter the flow field, but the occurrence of particle-particle interaction is infrequent enough that it can be neglected.

For  $\phi_P \ge 10^{-3}$  particle-particle collisions are significantly impacting both particle dynamics and fluid turbulence. A system including these interactions is said to be *fourway coupled*. In this regime, particle agglomeration and break-up can also be considered.

The particle response time  $\tau_p$  plays an important role in classifying the flow and predicting the mechanisms within it. For the Stokesian flow regime (very low Reynolds number), the parameter is defined by  $\tau_p = \rho_p d^2 / 18 \rho_f v_F$  where *d* is the particle

diameter,  $v_F$  is the kinematic viscosity of the fluid, and  $\rho_p$  and  $\rho_f$  are the densities of the particle and fluid phases, respectively. This is a reduced version of the full equation; however, for particle Reynolds numbers up to unity it remains a good approximation (Fessler and Eaton, 1994). The ratio of this response time to the fluid timescale can be used to define an important parameter  $St_\eta$  which is the Stokes number based on the Kolmogorov timescale,  $St_\eta = \frac{\tau_p}{\tau_\eta}$ . Whilst the definition of particle response time is strictly valid for lower particle Reynolds number, it is widely used to get an impression of the multiphase behaviour. Once particle Reynolds numbers increase significantly, the drag force deviates from Stokes' law and thus adaptations to the described relation are made which incorporate a more accurate drag coefficient (Fessler and Eaton, 1994) in order to retain its validity in different regimes, for example the drag relation of Oseen (1927).

The Stokes number governs the behaviour of particles with respect to the flow field, where a larger Stokes number indicates a larger heavier particle. At low Stokes numbers  $St_{\eta} \ll 1$ , particles follow the streamlines very closely. Conversely, for large Stokes numbers  $St_{\eta} \gg 1$ , the response time of the particle far outweighs the timescales of vortical motions in the flow and hence the turbulence field has a limited effect on the particles. For  $St_{\eta} \approx 1$  the impact of vorticies on the particles is maximal: since the scales are well-matched. Significant coherent motions can be induced in the particulate-phase by similarly sized vortical structures.

Interesting behaviour related to this parameter has been observed in multiphase turbulent flows – for example, the preferential accumulation of particles near to the wall in canonical wall-bounded flows. Originally it was believed that, without an external force, the stochastic forcing on the particles due to the turbulence field would ultimately lead to a uniformity in the particle distribution over time. This has been shown not to be the case with the observation of preferential accumulation of heavier-than-fluid particles in regions where there is high strain-rate and low vorticity (Fessler and Eaton, 1994). The dispersion of the particles is believed to be the result of the centrifuging of particles away from the centre of eddies. This effect is strongest for Stokes numbers close to unity, due to the comparable timescales of the fluid and particle response, as described above. Conversely to larger particles ejecting away from eddies, lighter-than-fluid particles such as bubbles will accumulate in these

regions of intense vorticity because they do not possess the momentum to overcome the advection of the fluid.

The motion of particles immersed in a turbulent flow field is largely thought to be dominated by the turbulent fluctuations and flow structures occurring at the length-scale of the particle (Jiang et al., 2022). As such, the dynamics of particles and the associated scaling laws can be thought of in the context of Kolmogorov's theory outlined in Section 2.1.2, particularly for tracer-like particles, e.g. Voth et al. (2002) took this approach in an experimental context to understand particle acceleration.

Turbophoresis is the name given to the underlying process where particles migrate preferentially according to decreasing turbulence kinetic energy. This was predicted mathematically by Reeks (1983) via derivation from the particle kinetic equation. The result of this effect in wall-bounded flows is particles accumulating close to the wall (Kuerten and Vreman, 2005; Marchioli et al., 2008). Contrastingly to the 'lighter-than-fluid' particle example, turbophoresis requires that the inertia of the particle is sufficiently large that its motion can decouple from the fluid streamlines, hence crossing them and migrating towards the near-wall region where the turbulence level is comparatively low. A resultant effect is that wall collisions are seen to increase with increasing Stokes number. For Stokes numbers around 0.1, based on viscous scales, the probability of a particle-wall interaction taking place is negligible, whilst for larger values of the order 100 this is a common occurrence (Mortimer et al., 2019).

The result of heavy particles being ejected from the vortices is that the TKE and its dissipation are reduced. For microparticles, where the volume fraction is high, the opposite effect is seen wherein TKE will decay at a lesser rate and energy dissipation can be increased. A balance between these two extremes exists where 'ghost particles' found in the range  $0.1 < St_K \le 1$  have little effect at all on the TKE levels. These regions are shown diagrammatically in Figure 2.2.


Figure 2.2: The classification map for particle-laden turbulent flows, originally from Elghobashi (1991), this is an updated version (Elghobashi, 2006). Particle shear Stokes number is shown on the vertical axis, and volume concentration on the horizontal axis.

The aforementioned effects fall under the category of turbulence modulation. For dilute systems there are several mechanisms whereby particles affect turbulence. One such effect is the enhanced turbulence energy dissipation due to a significant presence of particles and their drag. Another effect is how the addition of particles to a fluid increases the effective viscosity and can therefore reduce turbulence; similarly, the increased inertia can also reduce turbulence.

There are also effects that increase the turbulence levels of a fluid. For example, when a particle's kinetic energy is transferred to the fluid, increasing the fluid's kinetic energy. The occurrence of vortices in the particle's wake region increases the velocity fluctuation in that area, also increasing turbulence through vortex shedding. On top of this, Mittal (1999) presented another effect of significance when there is a high level of free stream fluctuations. That is, that the sphere acts like an oscillator and a resonance effect returns high levels of kinetic energy to the surrounding fluid. An excellent review of turbulence modulation is given by Loth (2023), as well as by Eaton (2006, 2009).

# 2.3.2 Experimental Study of Particles in Turbulence

Experimental research for wall-bounded multiphase flows generally looks to quantify the mean velocity, Reynolds stresses, and velocity fluctuations of the fluid and particles, as well as particle concentration. Another point of focus is the deposition and resuspension properties of the particulate-phase, which become of increasing importance for higher volume fractions where significant deposition can affect throughput in a pipe and potentially lead to blockages, as motivated in the introduction.

There are a number of challenges that prevent progress in experimental particle-laden studies. For example, measurements of the fluid-phase are impeded by the presence of these particles in the flow. Similarly, many commonly used fluid measurement techniques require the inclusion of additional particles that follow the fluid streamlines, in order to visualise the flow field through the post-processing of camera-taken images.

Laser-Doppler anemometry (LDA) measures the fluid-phase by inserting tracer-like particles that are assumed to follow the motion of the fluid. Then, measurements are taken using the Doppler shift of a laser beam reflected from a particle and fringe distance to estimate velocity. The tracer particles are assumed to have a negligible effect on the fluid-phase behaviour. For a multiphase experiment, the particulate-phase particles are generally much larger than the tracer particles. Hence, they scatter more light, allowing for the particulate-phase to be separated by signal-processing.

An early use of this method by Lee and Durst (1982) demonstrated interesting behaviours of glass particles in a vertical duct flow. At the time of publishing, there were discrepancies in the literature between theoretical predictions and experimental measurements. This work was able to propose a new consideration to improve the theoretical approach: beyond a certain cut-off frequency a particle becomes unresponsive to the turbulent fluctuations of the carrier-phase. Kulick et al. (1994) investigated the interactions between small dense particles and fluid turbulence in a channel flow of  $Re_B = 13800$  using LDA, where  $Re_B$  is the bulk Reynolds number. The particles were smaller than the Kolmogorov length scale and followed some but not all of the turbulent motions. The velocity fluctuations were higher in the streamwise direction for the particles compared to the single-phase fluid, but the reverse was true in the transverse direction which is contrary to the general findings in the experimental literature (Shokri et al., 2016). The addition of particles results in attenuation of the turbulence field. The attenuation increased with increasing particle Stokes number, particle mass loading and distance from the wall. Turbulence attenuation was found to be dependent upon the frequencies of the turbulence: there were preferred frequencies for which attenuation was seen with greatest effect.

Caraman et al. (2003) investigated particles of the same size in a downward gas flow using phase-Doppler anemometry. They demonstrated that the particulate-phase had higher streamwise velocity fluctuations than the single-phase fluid and comparable velocity fluctuations in the radial direction. The same set-up was utilised in Borèe and Caraman (2005) for bidisperse particles of a higher concentration; here, particle velocity fluctuations in the radial direction were found to be higher than those of the fluid. In the literature, there is agreement that the streamwise fluctuations are higher for the particulate-phase compared to the fluid-phase, but there is no such consensus for the radial fluctuations. Where radial fluctuations were found to be lower for the particulate-phase (Kulick et al., 1994; Varaksin et al., 2000), the validity of the results have been questioned in relation to insufficient pipe-length, and insufficient consideration of electrostatics and wall-roughness.

Particle image velocimetry (PIV) resolves the velocity instantaneously in a plane; again, using flow tracer particles. It holds the advantage over LDA of producing a vector field representation of the flow, as opposed to measuring velocity at a point. When the concentration of particles is low enough such that a single particle can be tracked from image-to-image, then it is referred to as particle tracking velocimetry (PTV). A useful overview of the dilute particle-laden pipe flow experimental literature is given by Shokri et al. (2016), with particular reference to PIV and earlier by Westerweel (1997) for digital PIV.

Noticing a limitation in the range of Reynolds numbers being investigated in the experimental literature ( $Re_B < 30,000$ ), Shokri et al. (2016) investigated a pipe flow at a high Reynolds number ( $Re_B = 320,000$ ) using a combination of PIV and PTV. This was motivated by the fact that current experimental Reynolds numbers were well below what is directly relevant to industrial applications. A vertical pipe orientation is utilised to avoid the axisymmetric particle concentrations due to gravity. In the core region ( $r_{loc}/R_{pipe} < 0.85$ , with  $r_{loc}$  the location in the radial direction, and  $R_{pipe}$  the pipe radius), the particles are observed to have a lower mean velocity than the fluid. Whereas, in the near-wall region ( $r_{loc}/R_{pipe} > 0.85$ ), the particulate-phase mean velocity is higher than that of the fluid. This is attributed to the absence of a no-slip condition at the wall for particles. Also noted in the paper is the influence of the lift force on particle concentration is given by Balachander and Eaton (2010) and the limitations of PIV are discussed by Sciacchitano and Wieneke (2016).

Bellani et al. (2012) were able to experimentally study turbulence modulation by large neutrally buoyant non-spherical particles in the context of homogeneous and isotropic turbulence, produced by directing flow through a grid. The chosen particles were of the length of the Taylor microscale. They found that spheres promoted TKE dissipation from the system as compared to the single-phase flow by 15%, whilst prolate spheroids of aspect ratio 2 only increased the TKE dissipation by 3%. The volume fraction was matched for this comparison; however, surface area was not matched.

#### 2.3.3 Numerical Study of Particles in Turbulence

Particle-laden flows can be studied numerically by treating the particles in an Eulerian or Lagrangian way. The Lagrangian approach tracks the position of a particle in the flow and calculates the unique dynamics on each local particle by solving the Newtonian equation of motion (Maxey and Riley, 1983; Maxey, 1987). The Eulerian approach is more holistic wherein a function is devised describing the bulk behaviour of the particles as if this phase were a continuum (Durst et al., 1984). From this, volume-averaged values for relevant parameters can be calculated. The use of an Eulerian approach can be less computationally expensive but the intricacies of the dynamics are lost (Gouesbet and Berlemont, 1999). This method is more suited to

dense particulate flows of a high volume fraction, whereas the Lagrangian approach is more suited to dilute flows (Shirolkar et al., 1996).

Within the Lagrangian framework, there are two main denominations of methods. The first is the point-particle approach, wherein particles are tracked as a single point within the fluid and advected based upon interpolations of the local fluid velocities to a point representing the particle's location, (with necessary modifications made to accurately reproduce the coupling effects). The other route, offering more fidelity and resolution, is the class of particle-resolving methods. As the name suggests, the full surface of the particle is resolved in this case and the interaction with the fluid is calculated across this boundary, in terms of the stresses acting on the surface due to the fluid. The former approach is much less computationally expensive, whilst the latter offers an improved approximation of the finer-scale physics. The computational requirements of the fluid solution, in the case of particle-resolved methods such as the immersed boundary method, increases with decreasing particle diameter, as there is a greater requirement to resolve smaller and smaller flow scales. As such, the majority of particle-resolved studies involve particle diameters much greater than the Kolmogorov length scale; whereas, the majority of point-particle studies are restricted to particle diameters below the Kolmogorov length scale (Schneiders et al., 2019). This restriction is due to the fact that there is no longer a scale separation between the particle and fluid phases (Balachandar and Eaton, 2010) and thus the typically used equations of Maxey and Riley (1983) are no longer valid. Even still, some studies have attempted to extend the point particle approach beyond the Kolmogorov length scale (Marchioli and Soldati, 2013; Marchioli et al., 2016) and work by Fröhlich et al. (2018) suggests that point particle methods have the potential to work for particle diameters on the order of the Kolmogorov length scale. Due to the computational requirements, the immersed boundary method typically is limited in the number of particles that can be dealt with (Van Wachem et al., 2015). More computationally expensive still is the arbitrary Lagrangian-Eulerian method (Hu, 1996; Hu et al., 2001) wherein particles are defined by unstructured meshes that evolve with time. This adaptive element requires remeshing which introduces much of the computational overhead.

Each of the single-phase approaches (e.g. RANS, LES and DNS) can be used to simulate the carrier-phase in multiphase systems. They each have advantages and disadvantages beyond simply the single-phase resolution achieved. For example,

since LES does not capture all scales, dense particulate flows are not well modelled (Kuerten, 2016). This is because it is unknown how the two-way coupling influences TKE on the unresolved scales and one loses the benefits of resolving the large-scale turbulent structures through the loss of information most pertinent to the particle-phase (Elghobashi, 1994).

In the same way, for a RANS approach, a correction would need to be included that improves the chosen turbulence model to account for the presence of particles and how particles affect the turbulence length scales and strength. The advantage of DNS is that with sufficient resolution one can resolve the boundary layers around resolved particles and more accurately describe the interplay between turbulence and particles at a local level, without complicated models. Unfortunately, this comes with high computational cost and such simulations are mostly restricted to small numbers of particles or reduced solution domains – for example, isotropic turbulence boxes, rather than full channels or pipes. Even if the surface of the particle is not resolved, for example in the case of Lagrangian particle tracking (LPT) simulations, the small-scale fluid fluctuations are still resolved and this will greatly increase the accuracy of the particle-phase behaviour, which is strongly influenced by these scales.

The literature for particle-laden wall-bounded flows spans a limited range of Reynolds numbers, with much of the focus being centred at and around  $Re_{\tau} = 180$ . Beyond  $Re_{\tau} = 360$  the presence of particle data (in terms of shear stresses and velocity fluctuations) is more limited. To achieve Reynolds numbers of orders of magnitudes higher than this is difficult due to the exceptional level of mesh refinement that is required. With increasing computational resources, this will ultimately become possible, but for now it is a limiting factor on single-phase wall-bounded flows and thus also on the pertaining multiphase flows.

One application of significant interest is how the particles and turbulence fields interact, and how this is influenced by the coupling scheme. Rani and Vanka (2000) conducted a study of two-way coupling effects in particle-laden pipe flow at a shear Reynolds number of 360 using DNS. The particle equation of motion considered only drag, and the volume fraction of the system was  $6.84 \times 10^{-5}$ . The study concluded that mean velocity profiles and RMS statistics were unaffected by two-way coupling. However, energy spectra showed that turbulence augmentation was occurring at both

the pipe centre and at the wall, with augmentation falling rapidly away from these two extremes. The simulations showed particle accumulation along low-speed streaks and showed that two-way coupling reduced the strength of hairpin vortices in the flow, which are responsible for those streaks. So significant turbulence augmentation was said to be occurring at the smaller dissipative scales.

Rani et al. (2004) conducted a parametric study investigating the influence on the turbulence field due to parameters such as the volume fraction, Stokes number and settling velocity. This time, interparticle collisions were also considered (four-way coupling). The result of this inclusion was that the particle concentration in the near-wall region was reduced by a factor of 25%. So, particle-particle collisions were found to reduce the levels of turbophoresis but it remained a dominant mechanism in the flow, nevertheless. This effect is more prominent for particles with higher response times which has been observed by Yamamoto et al. (2001). As a result of this effect, more particles are found in the central region of the pipe. Hence, greater streamwise kinetic energy is imparted by the particles to the fluid at these locations, which is demonstrated by flattened streamwise and radial RMS plots. For locations at around 20% of the pipe diameter, there is a clear increase in the velocity fluctuations (RMS). This is further exemplified by a reduction in the turbulent energy dissipation term here.

DNS was coupled to LPT for a similar investigation by Marchioli et al. (2003) wherein a shear Reynolds of 337 was studied in an upward pipe flow. In this case, drag, lift and gravity were considered but the system was only one-way coupled. The simulations showed the expected accumulation of particles in the near-wall region, demonstrating that particle transfer towards the wall is more efficient than transfer away. The paper gives a useful description of the physics of particle migration, stating that particle transfer is achieved through preferential pathways corresponding to advective motions known as sweeps and ejections (Best, 1992). These motions are representative of the instantaneous Reynolds stresses, and they contribute to positive turbulence production. A sweep corresponds to a local increase of shear stress at the wall, whilst ejection corresponds to an analogous decrease.

Vreman (2007) investigated via DNS the characteristics of vertical air-solid pipe flow for small heavy particles at a shear Reynolds number of 277. A Lagrangian approach was again utilised and the effects of particle collisions were included. The investigated mass loadings covered a wide range from 0.1 to 30. A model for wall-roughness was incorporated and this was found to have a greater effect on the quality of the predictions in relation to experimental results than the inclusion of collision parameters or the lift force.

At a fundamental level, the particle-phase's interaction with turbulence can be studied in boxes of homogeneous and isotropic turbulence (HIT). Many studies have utilised this academic approach to elucidating some of the key mechanisms at play, as well as for the purpose of developing scaling laws. Squires and Eaton (1991) studied particle dispersion in both decaying and forced boxes of homogeneous and isotropic turbulence. They were able to demonstrate that particles accumulate in regions of high strain rate and low vorticity, due to particle inertia biasing the trajectories, which was predicted by Maxey (1987). Elghobashi and Truesdell (1989a, 1989b) studied also the effects of particle dispersion in decaying homogeneous and isotropic turbulence which evolved over the course of the simulation from  $Re_{\lambda} = 25$  down to  $Re_{\lambda} = 16$ . This was compared to measurements by Snyder and Lumley (1971) with good agreement.

Elghobashi and Truesdell (1993) further used their decaying homogeneous and isotropic turbulence boxes to investigate the effects of the particle-phase on the turbulence field due to two-way coupling, at a volumetric loading ratio of  $5 \times 10^{-4}$ . They found that, at high wavenumbers, their particles increased the fluid turbulence kinetic energy – as well as increased the viscous dissipation rate and the transfer of energy from the larger scales. They concluded that in the absence of turbulence forcing or sustaining mechanisms, the inclusion of particles increases the speed at which the turbulence decays, even though the initial injection may decrease it momentarily.

Most studies of particle-laden homogeneous and isotropic turbulence focus on neutrally buoyant particles, but interesting anisotropic effects can be observed through the influence of gravity forces, such as a reverse cascade (Elghobashi, 1994). More recently, Chouippe and Uhlmann (2015) studied the effects of gravity in HIT in the context of particle settling, making use of elongated boxes for their calculations and adapting their forcing scheme accordingly. This study used the immersed boundary method, one of the most prominent particle-resolving methods.

The immersed boundary method (IBM), developed by Peskin (1977, 2002) was originally devised for application to the human heart. This first formulation allowed for deformation of the immersed boundary; whereas, it is now often used alongside rigid-body theory for non-deformable objects, particularly in the context of particulate flows (Uhlmann, 2005; Jain et al., 2020; Mortimer and Fairweather, 2021). The original framework involves determining a force distribution at arbitrary Lagrangian points across the boundary and applying this force back to the Eulerian phase through a regularised Dirac delta function, with the force relying upon the elastic properties of the material and a generalised form of Hooke's law. This approach is named the distributive immersed boundary method and is generally regarded as first-order accurate (Mortimer and Fairweather, 2021).

The non-distributive class of immersed boundaries involve enforcing the no-slip condition directly at the particle boundary or through direct solution of the Navier-Stokes equations (Mark and van Wachem, 2008). To this end, Mohd-Yusof (1997) proposed a 'momentum-forcing method' that enforces the no-slip condition at the Lagrangian points of the immersed boundary by mirroring the velocity field over the surface of the immersed boundary such that a smooth gradient is achieved, precisely recovering the no-slip condition at the boundary. The direct advantage being the second-order accuracy achieved in comparison to the first-order distributive forcing methods. This method was subsequently taken forward by Tseng and Ferziger (2003) who devised the ghost-cell mirroring technique that ensures the no-slip condition is met by careful selection of interior and exterior points across a particle mesh. Whilst obtaining second-order accuracy and thus improving upon the original method, issues with mass conservation were reported under this approach, to which Kim and Choi (2006) provided a rectification using a mass source term. Finally, the ghost-cell method was further adapted by Mark and Van Wachem (2008) wherein an efficient implementation was demonstrated that does not allow for mass flux across the boundary.

Schneiders et al. (2019) presented the first particle-resolved study on turbulence modulation for ellipsoidal particles on the order of the Kolmogorov scale through an immersed boundary technique, focusing on the energy balance between phases. An analysis was undertaken that considered the relative effect of the particles rotation rates on the kinetic energy of the fluid. The necessity was highlighted for resolving orientationally dependent drag and lift forces in non-spherical point-particle models to accurately reproduce the multiphase global energy balance. Further recommendations and analyses pertinent to point-particle models were presented, for which the inclusion of non-sphericity remains a predictive challenge. The study thus also highlighted the strengths of IBM approaches versus point-particle approaches.

## 2.3.4 Motion of Non-Spherical Particles

Analytic expressions for the motion of, and torques acting on, ellipsoidal tracer particles were derived in the seminal work of Jeffery (1922) which is often used as a reference case for validation of studies involving ellipsoidal particles. As yet, there are no universal analytic expressions for spheroidal motion accounting for fluid inertia (Voth and Soldati, 2017) which emphasises the importance of numerical simulations as an investigative tool.

One of the interesting behaviours of spheroidal particles is the preferential orientations and alignments in turbulent flows. In turbulent shear flows, there is an observed preferential alignment near to the wall, observed by Mortensen et al. (2008), Marchioli and Soldati (2013), and Zhao et al. (2015). There is a clear divergence between the behaviour of prolate and oblate spheroids: the prolate spheroids align with their symmetry axis in the streamwise direction and the disks align in the wall-normal direction, which can be understood through Jeffery orbits (Jeffery, 1922; Voth and Soldati, 2017), which are the dynamics arising from the solution of the Stokes equations for spheroids and arbitrary bodies of revolution (Bretherton, 1962; Ishimoto, 2023). The particles essentially enter a repeated dynamic motion that is unique to such particle morphologies.

Counterintuitively, particle alignment is still observed even when there is no mean flow. In the case of homogeneous and isotropic turbulence, it has been observed that particles that are nearby will locally align with one another through the action of the local velocity gradient tensor. Pumir and Wilkinson (2011) observed this in the case of slender rods. Using an immersed boundary method, Jiang et al. (2022) studied the dynamics of spheroids in HIT, they also observed a preferential alignment with local flow structures for particles that belonged to the inertial subrange, extending previous observations which had been limited to particles of the size of the dissipative range. The authors highlighted the importance of resolving the interface in the case of nonspherical particles, since much of the information relating to particle rotation comes from computation of the particle's boundary layer.

There have been experimental studies of the motion of non-spherical particles. For example, fibres have been studied extensively by the following authors using a range of techniques: Carlsson et al. (2006) studied the orientations of fibres in a shear flow; Dearing et al. (2012) studied the same in a turbulent pipe flow using PIV; and Hakansson et al. (2013) conducted similar experiments in a half-channel using cameras. Each study corroborated the fact that fibres accumulate near the wall in low speed streaks, aligning with the mean flow.

Njobuenwu and Fairweather (2013a, 2013b, 2014a) studied one-way coupled ellipsoidal particles using Lagrangian particle tracking and a quaternion method, assessing the particle's orientational and translational behaviour, and later assessing the role of particle shape on this behaviour (Njobuenwu and Fairweather, 2014b), finding that the forces and torques experienced by a particle are strongly dependent upon shape.

Non-spherical particle motion in viscous fluid is often studied in terms of how the particles rotate with respect to their major axes. These rotation rates are important to a particle's associated drag and lift coefficients and how the motion of a non-spherical particle can be expected to evolve over time. Njobuenwu and Fairweather (2015) extended their methodology to study the dynamics of ellipsoidal particles in a turbulent channel flow through the use of LES, leading to the observation of five modes of ellipsoidal particle motion: periodic, steady rotation, tumbling, precessing, and nutating. Disk-like particles were found to change between these states most frequently, as compared with needles. At a higher shear Reynolds number (467 as opposed to 300 in the previously cited study) Qi and Luo (2002, 2003) observed the following modes for needles: tumbling, precessing, nutating, log-rolling and inclined rolling. Huang et al. (2012) studied shear Reynolds numbers up to 700 and observed seven rotational modes and eight periodic modes.

A review paper by Mandø et al. (2007) gives an appreciation for the breadth of nonspherical particle motion in two-phase flow, with an emphasis on characterising and describing particle shape and motion. Ten years on, Voth and Soldati (2017) presented a comprehensive review paper on non-spherical particles in turbulence, with particular emphasis on particle dynamics and modelling approaches.

# 2.3.5 Agglomeration of Particles in Turbulence

The introduction of turbulence to particle-fluid systems promotes collisions. This affects agglomeration in two ways. Firstly, the probability of an agglomerate structure forming is increased in a given time period purely through an increased number of interactions and collisions – often understood in the agglomeration literature through collision kernels (Saffman and Turner, 1956). Conversely, the shear experienced across agglomerates due to locally fluctuating velocity gradients leads to the capacity for breakage of said structures. On top of this, turbulence can lead to breakage of agglomerates through more frequent and stronger impacts with other freely moving particles and aggregates. The relative importance of these three mechanisms is impacted by the properties of the turbulence field such as the size of the flow structures and the strength of the velocity gradients (Breuer and Khalifa, 2019). Chen and Li (2020) showed, using DEM coupled to a DNS flow solver, that violent collisions and breakages occur in straining sheets where particles are rapidly ejected from vortices.

As an agglomerate forms in the turbulent flow, its size increases and so the associated Stokes number changes. Hence, the effect of the turbulence on the particles differs over time as the size distribution evolves. This further increases the complexity of the turbulence-agglomeration interplay and highlights its dynamic nature. Ho and Sommerfield (2002) demonstrated the importance of modelling agglomeration effects when possible since there is an effect on the carrier-phase as well as the significant impact on the particle-phase.

Turbulent agglomeration is an important area for study due to its wide applicability to industry and industrial processes. For example, agglomeration promoted through turbulence was found to be an effective pretreatment to fine particle extraction in an electrostatic precipitator (ESP) by Bin et al. (2018). A combination of turbulence-enhanced and chemically-promoted agglomeration led to the ESP's extraction efficiency increasing from 75.5% to 83.1%, which has relevance to the present thesis as both techniques fall into the category of potential 'behavioural modifications'. The study found inclusion of turbulence to impact the size distribution by decreasing the concentration of small particles and increasing that of large particles. The

agglomeration efficiency was most pronounced for the smaller particles since their Stokes number is lower and hence the effect of the turbulence field is relatively stronger. Lastly, the presence of larger particles within the base particle distribution was seen to be a very important agglomeration mechanism, as these larger particles acted as a nucleus for the smaller particles to adhere to, highlighting that polydispersity has an important role in real systems.

There have been efforts to incorporate agglomeration physics into LPT simulations, this is typically done as a 'switch', where a binary determination is made at the point of particle collision as to whether two point particles will agglomerate based on their velocities and other relevant properties. Breuer and Almohammed (2015) presented an effective methodology that considers an energy balance of the particles, integrated with a hard-sphere collision model. This was successfully applied to turbulent channel flows by Mortimer and Fairweather (2020) to study agglomeration phenomena. Almohammed (2018) separates LPT agglomeration models into two main categories: energy-based (Löffler and Muhr, 1972; Hiller, 1981) and momentum-based (Weber et al., 2004). The thesis of Almohammed (2018) offers a thorough evaluation of these approaches.

Using a DNS-DEM coupled approach, Chen et al. (2019) investigated the agglomeration of particles in HIT using a simple adhesion parameter that can be tuned to control the overall 'sticking' efficiency of the particles. This was applied to study the collision-induced breakage of agglomerates in HIT (Chen and Li, 2020) and was further extended to include electrostatic effects, specifically the long-range Coulomb repulsive force (Ruan et al., 2021), which was found to reduce agglomeration without altering the form of agglomerate structures.

A few studies attempt to fully resolve the physics acting at the particle scale. This approach is of value for particles whose size is comparable to the fluctuating length scales in the turbulent flow, since velocity gradients across the surface of the particle will be significant. This is especially true for non-spherical particles where increased aspect ratio can lead to greater rotations, further complicating the problem. Fully-resolved particle aggregation in HIT was studied by Derksen (2011) for mono-sized spherical particles using an IBM. The author studied particles in a periodic turbulence box for which particle diameters were between two and 10 times the Kolmogorov

length scale. The turbulence was sustained using the linear forcing method implemented in a lattice Boltzmann solver. The particle interactions were calculated using a combination of the hard-sphere approach of Yamamoto et al. (2001) and a simple square-well potential for the attractive particle interaction. A square-well is a very crude approximation to the actual potential experienced between adhesive particles; nonetheless, this offered an opportunity to study the basic interplay between turbulence and the aggregation process. The study found stronger turbulence leads to smaller aggregates with less constituents, implying a reduction in agglomeration and/or a promotion of breakage.

The inclusion of a more physically realistic interaction potential was achieved by Mortimer and Fairweather (2021) through use of DLVO theory; in particular, distancedependent force equations were incorporated into a turbulent system with finite-sized boundary-resolved particles. The equations used were those previously employed by Fujita and Yamaguchi (2007) to study the organisation of nanoparticles in a liquid solvent without turbulence. Similarly to Derksen (2012), an immersed boundary technique was used by Mortimer and Fairweather (2021) with turbulence in the periodic box sustained through the use of linear forcing, which was outlined as an area for improvement due to its unstable characteristics and slow convergence. The level of fidelity of the flow-field was comparatively greater, since a high-order spectral element method was used for direct numerical simulation of the fluid-phase. This allowed for the resolution of important fluid phenomena at the scale of the particle's surface. A wide parameter study of binary particle interactions highlighted the most effective areas for exploiting agglomeration to optimise industrial processes; namely, the coefficient of restitution, the Debye length, and the turbulence strength, thus demonstrating further the applicability of the present investigations to industry. Whilst the DLVO forces had noticeable differences on the interaction behaviour, agglomerates did not form for extended periods, possibly due to the chosen turbulence levels being too great.

# 2.4 Conclusions

As evidenced, significant progress has been made towards the simulation of physically relevant turbulent agglomerating systems, but the state-of-play is simplistic and

idealised in some cases. Further steps that could be taken beyond the existing literature are as follows:

- Firstly, through the inclusion of non-spherical particles, rough surface particles, or polydisperse particles, which are more representative of the true characteristics of physical systems. Simulation of these systems will increase the understanding of the accuracy and limitations of the spherical, smooth and monodisperse assumptions. Attempts have been made to include these elements at basic levels of fidelity, but there is clear scope for further development, as well as the fact the entire parameter space of these variables (non-sphericity, roughness, polydispersity) has not been investigated.
- Secondly, there lacks entirely a computational framework that includes all of the elements of an agglomeration-resolved non-spherical and turbulent system. There are very few studies that attempt to include agglomeration at a particleresolved level even for spheres, whilst also resolving fluid turbulence. This is due to the disparate time and length scales involved. With access to high performance computing, steps can be made in this direction.
- Thirdly, improvements could be achieved through increased realism in the fluidphase: whilst HIT boxes offer the chance to study the interaction between turbulence and agglomeration, true agglomeration will include effects owing to mean-flow and solid boundaries. The challenge there is that the full wallbounded flows of industrial significance are dimensionally huge in comparison to the scale of the particles and their interactions. The number of particles and thus computational requirements are thus also huge if particle-scale physics is to be resolved. A smaller step forwards would be the inclusion of an improved forcing method for the turbulence which thus far rely on the effective but limited linear forcing method which is not suited to particle-laden studies. The review of forcing schemes in Section 2.1.11 indicates the best route forward is to use the scheme of Eswaran and Pope (1988).
- Collision modelling between non-spherical particles lacks a clear framework. There is a need for a soft-sphere model between spheroidal particles that does not rely on assumptions of sphericity. A first step towards this would be in generating a simple and effective way to track and compute particle overlap.

Lastly, there is scope for investigation of many particles interacting in turbulence with DLVO forces included, rather than a square well, and the effect on the turbulence field. This last point of investigation is only possible, in fact, with an improved forcing method, since the linear forcing is known to impact upon the particle-fluid coupling such that it is difficult to determine modulations on the turbulence field that occurred as a result of the particles as opposed to the forcing itself.

In the present thesis, attempts have been made to provide solutions to the mentioned gaps in the literature. Non-spherical particle interactions have been tackled in detail with novel modelling contributions, coupling the methodologies previously used (without consideration of fluid) to a turbulent flow. In doing so, the existing methodology of Mortimer and Fairweather (2021) has been expanded, wherein developments have taken place to improve the turbulence forcing method, and to incorporate non-sphericity into all elements of the modelling. Lastly, developments have taken place facilitating non-spherical particle collisions and agglomeration under a soft-sphere framework, which is not seen in the turbulent particle-laden flow literature.

# 3 Methodology

This chapter discusses the numerical methods used for the simulation of the multiphase flows presented in the results chapters. It encapsulates both the existing methodologies that have been implemented, as well as the range of new developments made over the course of the project. Much of the work involved in this project has been focused on the development of novel codes for the purpose of simulating non-spherical particle agglomeration in homogeneous and isotropic turbulence. This has been approached using first principles modelling, with consideration of other relevant studies and how best to combine and improve upon them. This effort has mostly been in the direction of the particle-phase modelling, since non-sphericity introduces significant additional complexity; however, the fluid-phase has also received attention in terms of an implementation for turbulence forcing, as well as coupling with the particle-phase. All of this is outlined in detail throughout this chapter.

Particle non-sphericity has been included in a way that accurately resolves the fluid dynamics very close to the surface of the particle – through an immersed boundary method (IBM) and an exceptionally high-fidelity turbulence-resolving direct numerical simulation (DNS). Combining this powerful simulation approach with a model for interparticle attractive and repulsive interaction forces is relatively unchartered territory in the existing literature, where a handful of studies have used this approach for spheres (Derksen, 2011; Mortimer et al., 2021) and, to the author's knowledge, no such studies exist for spheroids. There exists a range of numerical studies that model each of the important components in isolation or with low fidelity techniques. For example, non-spherical particles interacting without resolution of the fluid-phase (Schiller et al., 2011), non-spherical immersed boundaries in turbulence without agglomeration forces (Jain et al., 2021), or immersed boundary spherical particles interacting in turbulence with attractive forces in a very limited number of cases (Mortimer et al., 2021). By combining these elements, more accurate simulation approaches are developed, and hence a better understanding of the fundamental behaviours of interacting non-spherical particles in real physical turbulent systems is provided. This forms the main methodology used throughout the thesis.

Finally, significant developmental work went into producing a soft-sphere capability that can handle small-scale DLVO agglomeration events in turbulence over much wider simulation times with many simultaneously interacting particles. Two new ideas were proposed here for the computation of overlap between interacting particles, and a third method was also implemented from an existing methodology, adapted for spheroids.

# 3.1 Fluid-Phase

#### 3.1.1 Numerical Solution of the Navier-Stokes Equations

The technique employed for the solution of the fluid-phase is direct numerical simulation. Time-dependent solutions to the incompressible Navier-Stokes equations of a Newtonian fluid are computed. Thus, there are two key assumptions: that of incompressibility, which is justified by the low Mach number of the present flows; and that of a Newtonian fluid, which is justified since the simulations involve water as the continuous medium. The DNS computations are facilitated by an existing open-source spectral-element method based CFD code, Nek5000 (Fischer et al., 2008), the numerical approach of which is described in detail here. The dimensional Navier-Stokes equations solved in this process are:

$$\frac{\partial \boldsymbol{u}}{\partial t} + \boldsymbol{u} \cdot \nabla \boldsymbol{u} = -\frac{1}{\rho_F} (\nabla p - \nabla \cdot \boldsymbol{\tau}) + \boldsymbol{f}_{EP}, \qquad (3.1a)$$

$$\nabla \cdot \boldsymbol{u} = 0, \tag{3.1b}$$

where, u(x,t), p(x,t), and  $\rho_F$  are the fluid velocity, pressure, and density, respectively, and  $\tau(x,t)$  is the viscous stress tensor. A forcing source term  $f_{EP}$  is included in the equations to introduce and maintain turbulence, which is described in detail in Section 3.1.4.

The viscous stress tensor is defined by:

$$\boldsymbol{\tau} = \mu_F [\nabla \boldsymbol{u} + \nabla \boldsymbol{u}^{\mathrm{T}}], \qquad (3.2)$$

with  $\mu_F = \nu_F \rho_F$  being the dynamic viscosity. The divergence is implemented for a Newtonian fluid (due to constant viscosity) as:

$$\nabla \cdot \boldsymbol{\tau} = \mu_F \nabla^2 \boldsymbol{u}. \tag{3.3}$$

DNS necessitates the resolution of all temporal and spatial scales involved in the flow, requiring a sufficiently fine grid. During the generation of the numerical mesh, it is ensured that both the Kolmogorov length scale  $\eta$  and the integral length scale  $\ell$  are captured; whilst the corresponding Kolmogorov and integral time scales  $T_{\eta}$  and  $T_{\ell}$  are also captured through consideration of the time-step and simulation time, respectively. For the small-scale particle interactions seen in the present thesis, the full effects of the turbulence field are sampled across the many simulations making up the Monte-Carlo study.

Such a strong level of fidelity (i.e. DNS) is chosen since particles of a size that are known to be strongly influenced by the small turbulent motions of the fluid are investigated (Elghobashi, 1991; Mortimer and Fairweather, 2017). Moreover, the aim is to resolve phenomena such as boundary layers in the vicinity of the particle, which requires a very high level of resolution. This approach enables investigation of the fundamental and small-scale interaction behaviours.

#### 3.1.2 Nek5000 Flow Solver

The spectral-element solver Nek5000 (Fischer et al., 2008) was employed for the solution of the fluid-phase. It is chosen in part for its excellent scalability, allowing for parallel computations to take place on the University of Leeds' high-performance computing clusters, ARC3 and ARC4, which have been utilised for all the present simulations. These computations were typically distributed over 32 cores for forced turbulence simulations and 8 cores for initially quiescent liquid simulations, with runtimes ranging from 0.5 hours to 45 hours, depending on the specific problem. Nek5000 is written in FORTRAN 77 which gives fast execution times and efficient

memory usage; hence, this is the language used for each of the additional fluid and particle modules that were added to the code. The interfacing with the highperformance computing clusters took place in a Linux environment.

Nek5000 relies on the spectral element method (Boyd, 2001) for its computations, where the solution to the Navier-Stokes equations is approximated as a sum of weighted basis functions. In typical finite element methods (FEMs), the chosen basis functions are often piecewise continuous hat functions, whereas spectral methods employ polynomials as basis functions. Consequently, FEMs are local in nature, while spectral methods are global. The spectral approach enables higher-order approximation of the solution and exhibits the desirable property of "spectral accuracy," where the computational error decays exponentially with increasing polynomial order (Aref and Balachandar, 2018). The spectral element method combines the advantages of both approaches by partitioning the domain into structured elements and constructing the polynomial basis functions within each element.

Within these elements, the spectral element method operates on a non-uniform grid, with carefully chosen node distributions such that the quadrature rules ensure accurate integration of polynomials up to a specified degree. These choices facilitate mathematical conveniences that lead to faster convergence with fewer degrees of freedom compared to FEMs and other numerical approaches. For these reasons, the spectral element method is well-suited for high-fidelity fluid simulations, particularly for periodic domains with a simple global structure (Trefethen, 2000), as in the cases considered here.

In Nek5000, these polynomials are constructed using Lagrange interpolants defined on nodes corresponding to the Gauss-Lobatto-Legendre quadrature points, ensuring integration and differentiation are highly accurate and numerically stable.

#### 3.1.3 Simulation Domain

The general simulation domain is a cubic periodic cell with side-length *L* (Figure 3.1). The value of *L* is different for the fluid-phase validation and for the simulations involving particles. For the validation,  $L = 2\pi$ , was chosen to match the reference study of Chouippe and Uhlmann (2015). When particles are introduced, the simulations take place in a dimensional domain with L = 1 mm, to recreate the relative dimensions of the nuclear waste system being modelled and so that the particle-phase does not require non-dimensionalisation. This is best avoided since it is unclear how to non-dimensionalise the DLVO equations presented in Sections 3.2.7 and 3.2.8 such that the relative strengths of the fluid and DLVO forces remain consistent. The domain size is chosen to recreate a small region within a turbulent channel or pipe flow that might be found in industry (Mortimer et al., 2019). Within these small regions, the direct interplay between small numbers of particles and the turbulence field can be studied to better understand the fundamental interaction behaviours and dynamics. By reducing the simulation domain to this extent, exceptional resolution at the scale of the particles is achieved.



Figure 3.1: Schematic of the computational domain, with a binary particle interaction taking place between two needles.



Figure 3.2: Demonstrative example of the computational meshes used in the simulations; here, possessing 20<sup>3</sup> elements



Figure 3.3: Slice of the computational mesh (left), with a zoom (right) demonstrating the Gauss-Lobatto-Legendre structure of the nodes within each element.

The mesh is generated using the *genbox* tool in the Nek5000 software, with the results of this operation displayed in Figures 3.2 and 3.3. This allows the user to specify the number of elements in each orthogonal direction, defined by:  $nelx \times nely \times nelz$ . In the present work, these elements are distributed uniformly and equally in the three

directions. The simulation domain encloses  $24 \times 24 \times 24$  or  $30 \times 30 \times 30 \times 30$  elements depending on the study, which are further discretised into the internal Gauss-Lobatto-Legendre (GLL) nodes, the non-uniform spacing of which can be seen in Figure 3.3. The elements are evenly spaced, the internal nodes are not. This leads to a mesh totalling 2.7 million or 4.7 million equivalent grid points, respectively. Thus, an exceptional level of fidelity is achieved in the fluid-phase for which all length-scales of the turbulence field are resolved. The boundary conditions are set to be periodic in all three directions for both the fluid- and particle-phases.

#### 3.1.4 Turbulence Forcing Scheme

A forcing source term  $f_{EP}$  is introduced to the equations to maintain homogeneous and isotropic turbulence according to the methodology of Eswaran and Pope (1988). The force is generated independently of the flow field and is parameterised *a priori* such that the long-time behaviour of the forced fluid converges upon the desired statistically stationary turbulence level and Taylor-Reynolds number.

The generation of the field is achieved through the use of stochastic processes which drive the evolution of the forcing field  $f_{EP}$ . Namely, six independent Uhlenbeck-Ornstein (U-O) random processes are used for each of the forced wavenumbers in the fluid, corresponding to a real and imaginary part for each component of the wavenumber vector  $\boldsymbol{\kappa} = (\kappa_1, \kappa_2, \kappa_3)$ . The forcing field is constructed in Fourier-space and is denoted by  $\hat{f}_{EP}(\boldsymbol{\kappa}, t)$ . Only small wavenumbers are excited, those below the cut-off wavenumber  $|\boldsymbol{\kappa}| \leq \kappa_{cut}$ , which corresponds to only exciting large-scale motions. Then, the natural turbulence energy cascade develops the smaller scales of the turbulence field until the full turbulent energy spectrum is resolved and sustained.

The U-O random processes  $\hat{b}(\mathbf{\kappa}, t)$  are numerically updated using a finite-difference equation:

$$\hat{b}_{i}^{j+1} = \hat{b}_{i}^{j} \left( 1 - \frac{\Delta t}{T_{L}} \right) + e_{i}^{j} \left( 2\sigma^{2} \frac{\Delta t}{T_{L}} \right)^{\frac{1}{2}},$$
(3.4)

68

where the subscript *i* refers to the wavenumber direction, superscript *j* refers to the time level,  $\Delta t$  denotes the fluid timestep,  $T_L$  is the forcing timescale and  $\sigma^2$  is its variance. The forcing parameters  $\kappa_{cut}$ ,  $\sigma$  and  $T_L$  are chosen to match those used by Chouippe and Uhlmann (2015), in order to perform a validation. This validation, found in Chapter 5, demonstrates the validity of these choices with respect to the resolution achieved in the simulations.

The term  $e_i(\kappa, t)$  represents a complex random number that follows a standard normal distribution, generated using the Box-Muller transform approach during the simulation. The continuity condition in Fourier-space is  $\hat{f}_{EP}(\kappa, t) \cdot \kappa = 0$  which is satisfied by:

$$\hat{f}_{EP}(\boldsymbol{\kappa},t) = \hat{\boldsymbol{b}}(\boldsymbol{\kappa},t) - \boldsymbol{\kappa}(\boldsymbol{\kappa}\cdot\hat{\boldsymbol{b}}(\boldsymbol{\kappa},t))/(\boldsymbol{\kappa}\cdot\boldsymbol{\kappa}), \qquad (3.5)$$

representing a projection of  $\hat{b}$  onto the plane normal to  $\kappa$ . Then, an inverse Fourier transform is performed to obtain the entire forcing field  $f_{EP}$  used in the solution of the Navier-Stokes equations:

$$\boldsymbol{f}_{EP}(\mathbf{x}_{ijk}) = \sum_{l=-N_f}^{N_f} \sum_{m=-N_f}^{N_f} \widehat{\boldsymbol{f}} \exp\left(l\kappa_{1,l} \boldsymbol{x}_i\right) \exp\left(l\kappa_{2,m} \boldsymbol{y}_j\right) \exp\left(l\kappa_{3,n} \boldsymbol{z}_k\right),$$
  
$$\forall i, j, k.$$
(3.6)

This transform is achieved in practice by three successive one-dimensional transforms to reduce the computational complexity (Chouippe and Uhlmann, 2015); however, a spectral grid is forced, so a transform is first performed onto a uniform grid, before interpolating the correct forcing to the computational mesh:

$$\boldsymbol{A}_{mni} = \sum_{l=-N_f}^{N_f} \hat{\boldsymbol{f}}_{lmn} \exp(l\kappa_{1,l} \boldsymbol{x}_i), \quad \forall \, m, n, i$$
(3.7)

$$\boldsymbol{B}_{nij} = \sum_{m=-N_f}^{N_f} \boldsymbol{A}_{mni} \exp(I\kappa_{2,m} y_j), \quad \forall n, i, j$$
(3.8)

$$\boldsymbol{f}_{EP}(\mathbf{x}_{ijk}) = \sum_{n=-N_f}^{N_f} \boldsymbol{B}_{nij} \exp(I\kappa_{3,n} \boldsymbol{z}_k). \ \forall i, j, k$$
(3.9)

Equations (3.7) to (3.9) represent the successive one-dimensional transformations, at grid points  $x = (x_i, y_j, z_k)$ , with  $I = \sqrt{-1}$  denoting the imaginary unit and  $\kappa_{\alpha,\gamma}$  being the wavenumber vector, given by:

$$\kappa_{\alpha,\gamma} = 2\pi\gamma/L_{\alpha},\tag{3.10}$$

for  $\gamma \in \{-N_f, ..., N_f\}$  with zero forcing contribution at  $\gamma = 0$  and  $\alpha \in \{1, 2, 3\}$  being the wavenumber direction.

As is demonstrated in Figure 3.4, where a time evolution of two plane slices of the forcing field is presented, the stochastic processes are correlated in time: each iteration of the forcing field is dependent upon the last. It is also observed that the forcing structures generated are significant with respect to the size of the domain, these large structures begin to form the correct energy cascade and spectrum once coupled to a Navier-Stokes equation solver. This is the field used in the generation of the Taylor-Reynolds number  $Re_{\lambda} = 65$  ( $T_L = 1.67$ ,  $\epsilon^* = 0.00298$ ) validation study described in Chapter 5: Section 2. Each realisation is separated in time by 0.3 non-dimensional time units.



Figure 3.4: Demonstrative example of the Python implementation of an Eswaran and Pope (1998) non-dimensional forcing field  $\mathbf{F}_u$  for one slice in the *xy*-plane (top) and one slice in the *yz*-plane (bottom) evolving over a short period of time, with four evenly spaced realisations (blue arrow indicates time direction).



Figure 3.5: Contour plot of non-dimensional velocity magnitude for  $Re_{\lambda} = 65$ .



Figure 3.6: Isosurface plot of non-dimensional velocity magnitude for  $Re_{\lambda} = 65$ around the non-dimensionalised root-mean-square velocity u' = 0.7.

In Figures 3.5 and 3.6, it is shown how this forcing scheme manifests itself once coupled to the Navier-Stokes equations using Nek5000. The white contour lines in the former figure demonstrate the range of scales in the flow as would be expected given the descriptions of Chapter 2: Section 1, whilst the shading of the surface in Figure 3.5 represents the local velocity magnitude. The strongest velocity magnitudes often correspond to regions near to the centre of the turbulent eddies, where strong swirling motions are occurring. This range of scales can only be realised through the DNS solution of the Navier-Stokes equations, as clearly the smallest scales are not present in the raw forcing function (Figure 3.4). Lastly, the isosurface plot demonstrates locations at which the RMS velocity centres at u' = 0.7, helping to visualise the coherent structures present in the flow about the mean fluctuation magnitude.

# 3.2 Particle-Phase Methodology

## 3.2.1 Overview

The modelling of the particle-phase is achieved by the addition of extra modules to the Nek5000 code. These modules are called and executed at every timestep and run concurrently with Nek5000; to achieve this, they are added to the base fluid solver code upon compilation. The particle meshes are read in as *.raw* files at the beginning of the simulation, storing all of the necessary information about faces and vertices, which are then accessible through matrices defined in the solver at initialisation.

## 3.2.2 Morphology and Meshes

By using a mesh to represent the particle, one has access to any shape that can be made in the meshing software, but incorporating interesting physics typically relies on models that have a precise definition of the surface – ideally a continuous one – and this is where the mathematical complexity increases with non-spherical shape. It is possible to model cubes, cuboids and cylinders continuously using 'superellipsoids', but for the present work, ellipsoidal particles were chosen. Specifically, spheroids, which is the special case of an ellipsoid where two of the principal axes are of the same length. This allows for the study of disks and needles, which are two shapes found regularly as products in nuclear waste transport systems, and in nature (Li et al., 2020; Slootman et al., 2023). This finds a good balance between mathematical tractability and the relevant morphological anisotropy.

The quadric equation for the surface of a general unrotated ellipsoid is:

$$\frac{(x')^2}{a^2} + \frac{(y')^2}{b^2} + \frac{(z')^2}{c^2} = 1,$$
(3.11)

where a = (a, b, c) contains the axial lengths of the particle in the respective orthogonal directions of the particle's body-fixed frame x' = (x', y', z'). With b = c, the equation reduces to that of a spheroid:

$$\frac{(x')^2}{a^2} + \frac{(y')^2 + (z')^2}{b^2} = 1.$$
(3.12)

In this formulation, a > b defines a prolate spheroid, referred to from hereon as a 'needle'. The case a < b defines an oblate spheroid, referred to as a 'disk'. Note that a = b reduces Equation (3.12) to the equation for a sphere.

Equation (3.12) is exploited throughout the current work for its mathematical simplicity. In practice, the coordinate system can always be transformed to the body-fixed frame of a given particle, where the exact location of every point on the surface of the particle is accessible through the given equation. This is convenient for operations such as calculating normal vectors, distances between surfaces, or the local curvature which proves to be very useful in non-spherical particle modelling.

The term 'symmetry axis' (of the spheroid) will refer always henceforth to the principal axis of the spheroid, aligned with the x'-direction. Then, the asymmetry axes are those aligned with the y'- and z'-directions. The prime notation is referring to any variables, or later operators, defined with respect to the body-fitted coordinate system – the local coordinate system of a given particle. The unique length of the particle in the direction of the symmetry axis, as compared to the asymmetric axes, makes for geometric anisotropy. There is a circular symmetry about the x'-axis, and elliptic shape about the y'- and z'-axes. This anisotropy creates orientational dependency in all elements of the modelling, further complicating: particle orientation tracking, the inertia tensor and resulting torques, collision detection and resolution, calculation of inter-particle separation, the dynamics resulting from the flow field, the assessment of results, the inter-particle force interaction model, and more. In each case, steps have been taken to deal with these additional complexities which will be explained throughout.

For use in the simulations, icosphere meshes comprised of 320 triangular faces were generated in the graphics software tool, *Blender* (Hess, 2010). The number of faces was chosen to be in line with Mortimer (2019), which forms the basis of the present work. The cited work demonstrated that this level of resolution for the particle mesh was required to accurately reproduce the correct drag and lift properties in a validation study of the IBM. In Blender, the relative lengths of the icosphere's principal axes are

altered to create spheroids of two new morphologies. The morphologies, demonstrated in Figure 3.7, have a disk and needle shape. This is achieved by stretching or compressing the icosphere in the direction of the particle's symmetry axis and adjusting the scale accordingly such that volume is kept constant. The original volume is that of a sphere with  $r = 50 \mu m$ , which is a characteristic length for the particles and is sometimes referred to as the 'volumetrically equivalent particle radius'.



Figure 3.7: Computational meshes of the particles used in the immersed boundary simulations. Pictured from left to right: a needle, a sphere, and a disk.

Table 3.1 demonstrates the surface area increase for increasing aspect ratio spheroids. It is worth noting that, as aspect ratio increases, the surface area of the disk diverges away from that of the volumetrically equivalent needle of the same aspect ratio, which could have implications when comparing results.

hence volume and surface area are similarly nondimensional.				
Shape	Aspect Ratio	<b>Coefficients</b> ( <i>a</i> , <i>b</i> , <i>c</i> )	Volume	Surface Area
Sphere	1:1	(1.00, 1.00, 1.00)	4.19	12.6
Needle	2.5:1	(1.84, 0.737, 0.737)	4.19	14.2
Disk	2.5:1	(0.543, 1.36, 1.36)	4.19	14.8
Needle	5:1	(2.92, 0.585, 0.585)	4.19	17.1

(0.342, 1.71, 1.71)

4.19

20.2

Table 3.1: Important quantities relating to the morphology of five relevant shapes. The coefficients are nondimensionalised by the volumetrically equivalent radius and hence volume and surface area are similarly nondimensional.

## 3.2.3 Orientation Tracking and Quaternion Formulation

5:1

Disk

The chosen tool to compute and update the orientation of the rigid bodies, as well as rotate any vector quantities between world space and their body space, is a quaternion formulation; specifically, the Euler parameters (Goldstein, 1950; Mortensen et al., 2008). Quaternions are an extension of the complex number system which allow for convenient computation of three-dimensional rotations. They are represented by a four-dimensional system comprised of a scalar and vector part,  $\boldsymbol{q} = (q_0, \mathbf{v})$ , with  $\mathbf{v} = (q_1, q_2, q_3)$ ,  $\boldsymbol{q} \in \mathbb{H}$ ,  $q_0 \in \mathbb{R}$ ,  $\mathbf{v} \in \mathbb{R}^3$ .

Specification of the orientation state of the rigid body firstly relies on the Euler angles  $(\phi, \theta, \psi)$  which represent an ordered set of three rotation operations. The order is important since each individual rotation rotates the coordinate system about which the next operation will take place. For a visual description of this process the reader is referred to Goldstein (1950). It is possible to use these Euler angles in isolation to define the orientation of the body; however, a quaternion system is preferable since firstly it is not susceptible to gimbal lock (where two of the three rotational axes become parallel and thus the system loses a degree of freedom and its ability to represent all orientations), and secondly it is not singular (Fan and Ahmadi, 1995).

From the Euler angles, the Euler quaternion parameters q can be constructed:

$$q_0 = \cos\left(\frac{1}{2} \left[\psi + \phi\right]\right) \cos\frac{\theta}{2}, \qquad (3.13a)$$

$$q_1 = \cos\left(\frac{1}{2} \left[\psi - \phi\right]\right) \sin\frac{\theta}{2},\tag{3.13b}$$

$$q_2 = \sin\left(\frac{1}{2} \left[\psi - \phi\right]\right) \sin\frac{\theta}{2}, \qquad (3.13c)$$

$$q_3 = \sin\left(\frac{1}{2} \left[\psi + \phi\right]\right) \cos\frac{\theta}{2}.$$
(3.13d)

In turn, a rotation matrix A can be constructed which transforms between the orientational reference frame of the particle and the global frame. Specifically, the rotation matrix can be thought of as an operator which acts on the unprimed system to transform it into the primed system (Goldstein, 1950):

$$\mathbf{A} = \begin{pmatrix} 1 - 2(q_2^2 + q_3^2) & 2(q_1q_2 + q_0q_3) & 2(q_1q_3 - q_0q_2) \\ 2(q_1q_2 - q_0q_3) & 1 - 2(q_1^2 + q_3^2) & 2(q_2q_3 - q_0q_1) \\ 2(q_1q_3 + q_0q_2) & 2(q_2q_3 + q_0q_1) & 1 - 2(q_1^2 + q_2^2) \end{pmatrix}$$
(3.14)

This gives two equivalent ways to transform a vector p between coordinate systems: either by using the quaternion method  $p' = qpq^*$ , where  $q^*$  is the conjugate of q; or by using the transformation matrix p' = Ap. Then, the inverse rotation matrix, which can be used to transform in the opposite direction, is given by both the transpose and the inverse of A, since the matrix is orthogonal. This process reads as p = $A^T p' = A^{-1} p'$ ; naturally, it is simpler to compute the transpose. In the case of the quaternions, this operation can be performed through  $p = q^* p' q$ .



Figure 3.8: The three coordinate systems used to define the state of the body. The vector C joins the origin to the particle centre in the global frame.

In Figure 3.8, it is demonstrated graphically how the orientational and positional state of the body is defined and tracked. In the *global coordinate frame* x - y - z there lies a fixed coordinate system against which it is possible to define the others, this is sometimes referred to as the *laboratory frame*. In the *body-fixed particle frame* x' - y' - z' the system moves with the particle through space, allowing tracking of the particle's intrinsic properties as it rotates and moves, and allowing relation of these properties back to the global frame using the methods described above. Finally, there is an intermediate frame that rotates with the particle x'' - y'' - z'', the *Eulerian* or *Orientational frame*, which is not translated into its position in space *C* but rather just keeps track of the orientational state.

It is important to transform between these coordinate frames in the correct order. For convenience, most of the operations conducted on individual particles are implemented in the particle's body-fixed frame. When moving the particle back into its correct configuration in the global frame, it must first be rotated and then translated. To achieve the opposite transformation, an inverse translation must first be performed, followed by the inverse rotation.

The particle orientation can hence be fully described by the quaternion vector q and an associated matrix A which describes the orientation of the rigid body. The next consideration is how to evolve these. The particles possess an angular velocity  $\omega_p$ which is used to determine the rotation of the body over time. First-order differential equations are solved to time-evolve the Euler parameters (Mortensen et al., 2008); in this work, this is performed using a standard Euler integration method. A Runge-Kutta scheme is also implemented and available; however, the Euler scheme was sufficient for the present simulations, given the extremely small timestep, as will be discussed later. The following system of differential equations:

$$\frac{\mathrm{d}\boldsymbol{q}}{\mathrm{d}t} = \begin{pmatrix} \frac{\mathrm{d}\boldsymbol{q}_{0}}{\mathrm{d}t} \\ \frac{\mathrm{d}\boldsymbol{q}_{1}}{\mathrm{d}t} \\ \frac{\mathrm{d}\boldsymbol{q}_{2}}{\mathrm{d}t} \\ \frac{\mathrm{d}\boldsymbol{q}_{2}}{\mathrm{d}t} \\ \frac{\mathrm{d}\boldsymbol{q}_{3}}{\mathrm{d}t} \end{pmatrix} = \frac{1}{2} \begin{pmatrix} \boldsymbol{q}_{0} & -\boldsymbol{q}_{1} & -\boldsymbol{q}_{2} & -\boldsymbol{q}_{3} \\ \boldsymbol{q}_{1} & \boldsymbol{q}_{0} & -\boldsymbol{q}_{3} & \boldsymbol{q}_{2} \\ \boldsymbol{q}_{2} & \boldsymbol{q}_{3} & \boldsymbol{q}_{0} & -\boldsymbol{q}_{1} \\ \boldsymbol{q}_{3} & -\boldsymbol{q}_{2} & \boldsymbol{q}_{1} & \boldsymbol{q}_{0} \end{pmatrix} \begin{pmatrix} \boldsymbol{0} \\ \boldsymbol{\omega}_{x}' \\ \boldsymbol{\omega}_{y}' \\ \boldsymbol{\omega}_{z}' \end{pmatrix},$$
(3.15)

describes this quaternion evolution over time. The method of integrating the quaternions is effective but incurs error due to numerical drift, which ultimately can lead to the quaternions losing their unity during a simulation. It is imperative that the quaternions retain their unitary nature for accurate rotations. As such, a renormalisation is carried out every timestep to ensure unit quaternions are always dealt with:

$$\widehat{q} = \frac{q}{\sqrt{q_0^2 + q_1^2 + q_2^2 + q_3^2}}.$$
(3.16)

#### 3.2.4 Immersed Boundary Method and Coupling

An immersed boundary method is utilised to couple the particle-phase to the fluid using the ghost-cell method described in Mark and van Wachem (2008). This allows the inclusion of finite-sized particles of a complex geometry in the flow. The IBM code used in the present work was an existing one – originally implemented for spherical particles

by Mortimer (2019) but written in such a way that it could be extended to non-spherical geometries. It has received the necessary modifications herein to make it properly compatible, for example in generalising the inertia tensor for use with ellipsoidal bodies, reformulating the boundary conditions, and extending the torque calculations' compatibility to ellipsoids. The method runs concurrently with Nek5000 to couple the immersed particle meshes with the fluid-phase solver. Initially, a structured grid of surface points is read in which defines the locations of all the faces (made up of three vertices each) of the particle in the form of a .raw file. These data are then scaled during the particle injection module such that it is of the correct size for the simulation: of characteristic length r.

The no-slip condition is enforced at the surface of the particle mesh discretely using a direct boundary condition imposition method. This is achieved by a ghost-cell mirroring technique (Tseng and Ferziger, 2003), in which fluid 'ghost nodes' inside the particle boundary are manually set to values which ensure the requisite boundary condition. Namely, it is ensured that the fluid velocity at the surface of the particle matches exactly with the velocity at the face of the particle by forcing the fluid very near to the surface. This requires multiple fluid points for each triangular face of the particle mesh. Inside the boundary are the ghost nodes; these are points in the fluid mesh that lie directly adjacent to the boundary, i.e. there are no fluid points that lie between a ghost node and the nearest particle boundary. The location of an 'exterior point' is then determined from a ghost node. This is done by reflecting the ghost node across the particle face to the outside of the particle. In doing so, the midpoint of the line formed between these two nodes represents a point directly on the particle surface. The schematic in Figure 3.9 demonstrates this for a smooth particle surface shown in 2-D on a simple uniform grid for explanatory purposes. The yellow ghost nodes are predetermined from the position and orientation of the particle with respect to the underlying fluid mesh.



Figure 3.9: Schematic demonstrating the implementation of the ghost-cell method on a computational grid, with points used in the calculation colour-coded: interior nodes (red), exterior nodes (green), ghost nodes (yellow). The thick black line represents a slice of the particle surface.

Typically, this process will not align the exterior point with a point also in the fluid mesh and so the exact velocity value at this point must be determined through spectral interpolation of the nearby fluid points. Then, the no-slip condition is enforced by:

$$\boldsymbol{u}_G = 2\boldsymbol{u}_{IB} - \boldsymbol{u}_E, \qquad (3.17)$$

with  $u_G$  being the fluid velocity of the ghost node,  $u_{IB}$  the fluid velocity at the surface of the immersed boundary and  $u_E$  the fluid velocity of the exterior point. Lastly, there are 'interior nodes' that lie inside the particle boundary but do not lie adjacent to the surface – these are set to match the velocity at the nearest boundary.

Advection of the particle is computed by integrating the total hydrodynamic forces and torques acting on the surface. The surface of the icosphere mesh is naturally divided into its 320 constituent triangular faces. The area of each can be calculated using the relation:  $Area = \frac{1}{2}(\overrightarrow{PQ} \times \overrightarrow{PR})$ , for points P, Q, R that make up the vertices of the triangular face.
The total surface forces acting on the immersed boundary are given by the surface integral of the stress tensor on the particle:

$$\boldsymbol{F} = \int_{IB} (-p\boldsymbol{n} + \boldsymbol{\tau} \cdot \boldsymbol{n}) dS, \qquad (3.18)$$

with the first term on the RHS being the pressure forces and the second the viscous forces, evaluated numerically across the 320 faces through a summation:

$$\boldsymbol{F} = \sum_{N} (-p_f \boldsymbol{n}_f \boldsymbol{A}_f + \boldsymbol{\tau} \cdot \boldsymbol{n}_f \boldsymbol{A}_f), \qquad (3.19)$$

where the subscript *f* represents a quantity evaluated at a given face, *N* is the total number of faces, *n* is the face unit normal, and  $A_f$  is the area of the face. The pressure and viscous stresses are interpolated from their respective fields defined on the fluid mesh such that they are evaluated at the centre of the face they are acting on. The torque, *T*, can also be calculated by:

$$T = \sum_{N} r_f \times F_f, \qquad (3.20)$$

where  $r_f$  is the vector pointing from the particle centre to the face centre. The total torque of the body can then be used to update the angular velocity of the particle according to Euler's rotation equations in the particle reference frame, as a rigid-body is assumed.

This was extended from the simplified spherical equation  $(I \frac{d\omega}{dt} = T)$  to give the spheroidal case (Goldstein, 1950):

$$I_{x} \frac{d\omega_{x}}{dt} = T_{x} + \omega_{y}\omega_{z}(I_{y} - I_{z}),$$

$$I_{y} \frac{d\omega_{y}}{dt} = T_{y} + \omega_{x}\omega_{z}(I_{z} - I_{x}),$$

$$I_{z} \frac{d\omega_{z}}{dt} = T_{z} + \omega_{x}\omega_{y}(I_{x} - I_{y}),$$
(3.21)

where  $\boldsymbol{\omega}$  is the particle's angular velocity and  $\boldsymbol{I}$  its inertia tensor, given for a general ellipsoid by:

$$I = \frac{M}{5} \begin{pmatrix} b^2 + c^2 & 0 & 0\\ 0 & a^2 + c^2 & 0\\ 0 & 0 & a^2 + b^2 \end{pmatrix},$$
(3.22)

where *M* is the mass of the particle and *a*, *b*, *c* are the radii of the unrotated spheroid in the x', y', z' directions, respectively, with b = c.

The particles receive a force from the fluid in their advection, they exert an implicit force on the fluid through their boundary, and they interact with one another through collisions and inter-particle forces. In this way, the system is *four-way coupled*.

#### 3.2.5 Collision Detection

For spheres, the inter-surface distance is found trivially by subtracting twice the particle's radius from the distance between the two particle centres. For ellipsoids, with that same distance between their centres as with the spheres, the inter-surface distance now depends on the relative orientation of the particles. Therefore, a more complex algorithm is required to determine distance or contact.

The full collision implementation, which can be subdivided into the detection module and the mechanics module, is based on the work of Jain et al. (2019), chosen due to the fact that a closest distance vector between a pair of surfaces, as well as the nearest surface points, are explicitly solved for. These quantities are useful in the calculation of the inter-particle forces herein since they arise due to surface-surface interactions.

To give a brief overview, the method involves a common normal approach for the detection algorithm and a hard-sphere collision model for the resolution of the mechanics. A thorough description of the method is given but the reader is advised to refer to the original text for further details. The implemented modifications are outlined as follows.

The collision search begins when the particle centres are within a distance of twice the longest particle radii, since this is the first instance a collision could occur between approaching particles. This is sometimes referred to as 'bounding spheres' and it prevents redundant calculations between particles that have no possibility of collision. When this condition is satisfied, the closest distance between the two surfaces is then searched for. For this, the parametric equations for an ellipsoid are used to determine a point R on the surface in the body frame:

$$\boldsymbol{R} = \begin{pmatrix} a \cos\varphi \sin\vartheta \\ b \sin\varphi \sin\vartheta \\ c \cos\vartheta \end{pmatrix}, \tag{3.23}$$

with polar angle  $\varphi \in [0, \pi]$ , azimuthal angle  $\vartheta \in (-\pi, \pi]$  and coefficient vector  $\mathbf{a} = (a, b, c)$ . Equation (3.23) is valid only for an unrotated body and so transformations between coordinate frames are conducted frequently in the implementation of the following algorithm. Differentiating  $\mathbf{R}$  w.r.t.  $\varphi$  and  $\vartheta$  gives the tangent vectors  $\mathbf{t}_{\varphi}$  and  $\mathbf{t}_{\vartheta}$ . The closest distance between two surfaces is found when the vector  $\mathbf{d}$  joining the two surfaces is perpendicular to both tangent vectors, for both particles. The surfaces share a common tangent plane at these locations, both normal to  $\mathbf{d}$ . The implementation involves monitoring the four dot products:

$$D_{\varphi_i} = \mathbf{t}_{\varphi_i} \cdot \boldsymbol{d}, \qquad i \in \{1, 2\}, \tag{3.24a}$$

$$D_{\vartheta_i} = \mathbf{t}_{\vartheta_i} \cdot \boldsymbol{d}, \qquad i \in \{1, 2\}, \tag{3.24b}$$

until  $D_{\varphi_i}, D_{\vartheta_i} \to 0$  or, in practice, are smaller than some small value  $\varepsilon$ , taken to be  $10^{-6}$  in the present calculations.

Before starting the iterative procedure, initial guesses for the nearest surface points  $R_1^0$  and  $R_2^0$  (for particles 1 and 2, respectively) are made as follows. A vector  $d_0$  is created between the two particle centres. The points on this line which intersect the particle surfaces are chosen as an initial guess. For a sphere, this would immediately give the closest points. The geometry is shown in Figures 3.10 and 3.11.

To find the first initial point  $\mathbf{R}_1^0$ , the system is transformed into the reference frame of particle 1, i.e. its body-fixed frame centred at the origin. Then, the location of the centre of particle 2 after this transformation is none other than the vector  $\mathbf{d}_0$  in the body frame of particle 1 which can loosely be denoted as  $\mathbf{d}'_0$ . The quadric equation for an unrotated ellipsoid can then be exploited. Namely, parametrically iterate along the vector  $\mathbf{d}'_0$ , from the origin (corresponding to the centre of the particle) with each iteration giving new points (x', y', z') that eventually satisfy Equation (3.11) once on the surface. This gives an initial  $\mathbf{R}_1^0$  and the same process with 1 and 2 swapped gives  $\mathbf{R}_2^0$ .



Figure 3.10: Schematic of the approach to calculating the initial surface points via coordinate transform.

The following iterative process drives  $\varphi_i$  and  $\vartheta_i$  from initial guesses towards the points  $R_1(a, \varphi_1, \vartheta_1)$  and  $R_2(a, \varphi_2, \vartheta_2)$  where the tangent vectors are normal to *d*:

$$\varphi_{i}^{j+1} = \varphi_{i}^{j} + C \frac{\mathbf{t}_{\varphi_{i}}^{j} \cdot \mathbf{d}^{j}}{|\mathbf{t}_{\varphi_{i}}^{j}||\mathbf{d}^{j}|}, \qquad i \in \{1, 2\},$$
(3.25a)

$$\vartheta_{i}^{j+1} = \vartheta_{i}^{j} + C \frac{\mathbf{t}_{\vartheta_{i}}^{j} \cdot \boldsymbol{d}^{j}}{\left| \boldsymbol{t}_{\vartheta_{j}}^{j} \right| \left| \boldsymbol{d}^{j} \right|}, \qquad i \in \{1, 2\},$$
(3.25b)

with *C* being the volume-equivalent diameter of a sphere divided by the principal axis' coefficient in Jain et al. (2019), C = 2r/a. However, this term was modified and seen to improve convergence in higher aspect ratio morphologies, described below. In testing of the algorithm after implementation, its performance significantly reduced with increasing aspect ratios in terms of the algorithm's ability to converge on the correct surface points at short distances. The solution to this problem was firstly to reduce the initial increments by which the angles changed, and then further reduce this increment when convergence was not achieved within 30 iterations of the algorithm. After each set of 30 iterations, the value of *C* halves. The new coefficient can therefore be written as:

$$C^{j} = \frac{\pi}{180} P^{j}, \tag{3.26}$$

with  $P^{j} \in \{1, \frac{1}{2}, \frac{1}{4}, \frac{1}{8}, \frac{1}{16}, \frac{1}{32}, \frac{1}{64}, \frac{1}{128}\}$  being the parameter which slows the increment by which the angles are updated.

The authors released a corrigendum (Jain et al., 2019) to their work *after* the presently described modification had been implemented, which addresses the same issue. Their solution is to introduce an *ad hoc* appearance of *a* in the denominator of the dot product, i.e., the denominator becomes  $a + |t_{\varphi_i}^j| |d^j|$ . It was found that this also fixes the convergence issues described above. This issue may arise from the fact that the

increment that the angles are increased by does not consider the location on the surface. So, when moving parametrically through  $[0, \pi]$ , equal increments of an angle do not correspond to equal distances traversed on the surface of the particle. Perhaps a more robust algorithm would account for this. It was attempted to directly account for this in the present work, such that increments were properly scaled, to test this hypothesis and potentially significantly improve the method. To do this properly, the need to compute elliptic integrals was found, which are notoriously difficult. In the interest of reducing computational complexity, this approach was abandoned in favour of the described approach wherein the angular update is halved.

When the dot products in Equation (3.24) are approximately zero on both particles, the search is terminated, and the closest distance vector d between the two particle surfaces is obtained. At each timestep where a collision is possible, d is calculated. That is, until the closest distance between the two particle surfaces is less than a small value  $\varepsilon$  – at which point, a collision is said to have occurred.



Figure 3.11: Variables used in the calculation of  $\mathbf{R}_i$ . Variables with superscript '0' indicate initial values and those without represent the final values converged upon by the algorithm.

The collision detection algorithm typically does not converge for overlapping surfaces and so it is imperative that the surfaces are not allowed to overlap in the simulations. This requires setting a sensible value of  $\varepsilon$  that collides particles at the right moment. It is important not to set  $\varepsilon$  to be too small to run the risk of no collision being detected, which can occur if the relative velocity of two particles is great enough to bypass this small window in a single timestep, allowing particles to move inside of one another, causing the algorithm to break down. Conversely,  $\varepsilon$  must not be so large that the particles are rebounding and moving apart too early in their interaction, such that they do not get close enough to experience the maximum values of the van der Waals (vdW) force.

This touches on an important oversight that must be avoided where the timestep is so great that the maximal vdW forces are experienced only for a small number of timesteps, if at all. Then, given the large variation in the vdW force across small distances, the absolute value of this force experienced is effectively random. 'Random' in the sense that the separation distance falls somewhere in the interval  $(0, \varepsilon)$  which is as dependent upon where the particles were located in the previous timestep, as much as it represents any true physics. This interval is a small increment in terms of distance but it corresponds to a wide range of vastly different DLVO magnitudes (covered in Sections 3.2.7 and 3.2.8) – and the exact value of the maximal vdW force experienced has a very highly weighted impact on the resolved outcome of agglomeration. So, in work that considers agglomeration, it is very important that these maximal values are properly sampled to ensure that the outcome of the agglomeration procedure is deterministic. To properly resolve this, it must be ensured that the timestep is low enough for proper sampling. When coupling the simulations to a fluidphase, it is advisable to have a variable timestep or at least a separate timestep for the particle- and fluid-phase.

In typical hard-sphere approaches,  $\varepsilon$  is constant; however, in the present work, a dynamic choice of this variable was implemented based upon the velocity of the particles at the points of closest approach. Note that this velocity is not necessarily the same as the translational velocity of the particle, since the non-spherical particle has a rotational component that contributes to the local velocity of a surface point. Denoting the velocity at the surface point as  $u_c$ , the minimum allowed separation distance is given by:

$$\varepsilon(t) = |\Delta \boldsymbol{u}_c(t)| \Delta t, \qquad (3.27)$$

where  $|\Delta u_c(t)| \Delta t$  constitutes the maximum possible distance the two bodies can travel relative to one another, thus preventing overlap. Once a collision has been determined, the collision physics are modelled using the hard-sphere approach, which gives an impulsive collision force  $\mathbf{f}_c$  that acts over the interval  $\Delta t$ . The change in particle velocity is assumed to be instantaneous. In practice, this dynamic setting means that as particles lose speed (e.g. due to a collision) they are able to get successively closer and are thus more likely to sample larger values of the van der Waals force, allowing for repeatable results.

A coefficient of restitution e is included to model the kinetic energy loss in the collision as a result of surface deformation and other dissipative forces. In the present work, a value of 0.4 is typically chosen for e, which aligns with the material properties of calcite (Mortimer et al., 2019). This means that only 40% of the relative velocity is retained upon collision; although, the resolved value of e is mildly orientationally dependent under this model. This orientational dependency is enhanced by the attractive interparticle forces which act to accelerate the approach and attenuate the speed of rebound – and are themselves orientationally dependent.

#### 3.2.6 Collision Mechanics

Once contact is determined, the collision forces  $\mathbf{f}_c$  can be computed according to a hard-sphere model and applied to the particles. The impulsive collision force is applied only at one timestep and acts over the time interval  $\Delta t$ . Calculations to determine the force  $\mathbf{f}_c$  are made at the point of contact and variables computed at this contact point are denoted by the subscript *c*, whereas global variables that hold information of the full body are denoted by a subscript *p*.

Just before collision, the velocity at the eventual contact point is:

$$\boldsymbol{u}_c = \boldsymbol{u}_p + \boldsymbol{\omega} \times \boldsymbol{r}_c, \tag{3.28}$$

where  $r_c$  is the vector joining the contact point and the particle centre.

At time step n, the quantity  $u_c$  is related to the previous velocity by:

$$\boldsymbol{u}_c^n = \boldsymbol{u}_c^{n-1} + \boldsymbol{u}_{ex} + \boldsymbol{K} \cdot \boldsymbol{p}_c, \qquad (3.29)$$

where  $p_c$  is the linear momentum and  $u_{ex}$  is the velocity due to external forces besides the collision force, acting during the interval  $\Delta t$ . The derivation of this can be found in Jain et al. (2019). This can then be used to derive the force via:

$$\mathbf{f}_c = M \frac{\boldsymbol{u}_c^n - \boldsymbol{u}_c^{n-1}}{\Delta t}, \qquad (3.30)$$

where M is the mass of the particle. This can be applied to the particle mesh and hence directly included in the force balance. However, a different approach is adopted here. Since the particles are rigid bodies, the angular velocity is defined globally, allowing derivation of the particle angular velocity from:

$$\boldsymbol{\omega}_p = \boldsymbol{\omega}_c = \frac{r_c \times u_c}{|r_c|^2},\tag{3.31}$$

which can be used to update the particle velocities by rearranging Equation (3.28):

$$\boldsymbol{u}_p = \boldsymbol{u}_c - \boldsymbol{\omega}_p \times \boldsymbol{r}_c. \tag{3.32}$$

In Equation (3.29), the symmetric system matrix is given by:

$$\mathbf{K} = m^{-1} \mathbb{I} + [\mathbf{r}_c]_{\times}^{\mathrm{T}} \cdot \mathbf{i}^{-1} \cdot [\mathbf{r}_c]_{\times}, \qquad (3.33)$$

Where  $[r_c]_{\times}$  is the skew symmetric matrix of  $r_c$  and **i** is the inertia matrix in the global coordinate system arising from the transformation  $\mathbf{i} = A \cdot I \cdot A^{\mathrm{T}}$ , with I being the body frame inertia tensor defined in Equation (3.22). Neglecting  $u_{\mathrm{ex}}$  and using Newton's third law to determine how the linear momentum should be shared between particles (equal and oppositely), the following equations are derived:

$$\boldsymbol{u}_{c,1}^{n} = \boldsymbol{u}_{c,1}^{n-1} - \mathbf{K}_{1} \cdot \mathbf{p}_{c}, \qquad (3.34a)$$

$$\boldsymbol{u}_{c,2}^{n} = \boldsymbol{u}_{c,2}^{n-1} + \mathbf{K}_{2} \cdot \mathbf{p}_{c}.$$
(3.34b)

Taking the difference of these two equations gives the relation:

$$\mathbf{p}_c = -\mathbf{K}_{12}^{-1} \cdot \Delta \boldsymbol{u},\tag{3.35}$$

with  $\mathbf{K}_{12} = \mathbf{K}_1 + \mathbf{K}_2$ . By defining the relative velocity between the two particles at the point of collision as  $\boldsymbol{u}_r$ , the change in relative velocity due to a collision can be defined,  $\Delta \boldsymbol{u} = \boldsymbol{u}_r^{n-1} - \boldsymbol{u}_r^n$ , in a way that can be implemented using known quantities:

$$\Delta \boldsymbol{u} = \boldsymbol{u}_r^{n-1} + e(\boldsymbol{u}_r^{n-1} \cdot \mathbf{n})\mathbf{n}, \qquad (3.36)$$

where the second term on the RHS comes from using the *Poisson hypothesis* and *e* is the normal coefficient of restitution.

#### 3.2.7 Spherical DLVO Interaction Force

As introduced in Section 2.2.4, van der Waals forces are important in the interactions between neutrally-charged particles at the nano- and micron-scale. The forces have been included, along with electric double layer (EDL) repulsive forces, through the DLVO framework (Derjaguin and Landau, 1941; Verwey and Overbeek, 1948). The following equation for the attractive van der Waals forces between two homogeneously charged spherical particles of the same radius r is well-established, derived by Hamaker (1937):

$$\boldsymbol{F}_{\rm vdW} = -\frac{Ar}{12d^2},\tag{3.37}$$

where A is the Hamaker constant and d is the inter-surface separation. Less wellestablished, EDL repulsive forces are treated in several different ways throughout the literature. For the present work the following expression is used, from Israelachvili (1992), used also by Fujita and Yamaguchi (2007):

$$F_{\rm EDL} = \frac{64\pi r n k_B T \Theta^2 e^{-\kappa d}}{\kappa},\tag{3.38}$$

where *n* is the number density of electrolyte ions,  $\Theta$  is the reduced surface potential,  $k_B$  is the Boltzmann constant, *T* is the fluid temperature and  $\kappa$  is the inverse Debye length. These forces are then linearly combined by:

$$F_{\rm DLVO} = F_{\rm vdW} + F_{\rm EDL} \,. \tag{3.39}$$

The reduced surface potential is calculated using:

$$\Theta = \tanh\left(\frac{z_e q_e Z}{4k_B T}\right),\tag{3.40}$$

where  $z_e$  is the electrolyte ionic valency,  $q_e$  is the elementary electric charge, and Z is the zeta potential. Each of these quantities has been chosen to match calcite particles in water at 20°C, typical conditions used for this nuclear waste simulant. The full list of parameters is given in Chapter 4: Table 1. Introducing  $\varepsilon_0$  as the permittivity of a vacuum and  $\varepsilon_s$  as the relative permittivity of solvent, the inverse Debye length can be defined as:

$$\kappa = \sqrt{\frac{2nz_e^2 q_e^2}{\varepsilon_0 \varepsilon_s k_B \mathrm{T}}} \,. \tag{3.41}$$

The force is evaluated at every timestep between particle pairs and is included in both interacting particles' force balances with equal magnitude and opposite direction along the vector of closest approach. The charge distribution is assumed to be uniform across the particle surface, for both spherical and non-spherical particles, in line with the assumptions made in Hamaker (1937) theory.

Finally, due to the divergent nature of Equation (3.37), a cut-off distance must be chosen for the van der Waals force beyond which the force does not increase (Abbasfard et al., 2016). In this work, the parameter is set to be 5 *nm*. This is a source of discrepancy between studies and agglomeration proves to be strongly dependent upon this parameter in the following chapters. Experimental validation would be very useful to optimise this parameter correctly. Currently, it somewhat overrides the importance of the Hamaker constant (which ideally should govern the strength of the vdW force alone) and is thus a limitation on the numerical modelling of these forces.

The resultant DLVO potential for two interacting spheres is shown below in Figure 3.12, where the vdW cut-off is seen acting to prevent divergence. Clearly, changing the location of this cut-off distance has significant effects on the overall potential.



Figure 3.12: Plot of DLVO interaction potential normalised by thermal energy, for two interacting spherical particles as a function of separation distance normalised by particle radius.

#### 3.2.8 Non-Spherical DLVO Interaction Force

The DLVO forces are important at a length-scale much shorter than the radii of the particles, which is one way of stating the Derjaguin approximation. This is evidenced by Figure 3.12, where the *x*-axis has been appropriately scaled. In fact, due to their very short-range nature, most of the contribution to the force occurs around the points of nearest approach. Different relative orientations of the particles can lead to greatly different geometric properties at these points.

One way of accounting for this orientational dependency and its effect on the surface interaction is to consider the local curvature of the surface at the points of closest approach on the respective particles. The *Gaussian curvature* is a measure of local surface curvature and it can be defined at each point on a surface. The calculation of this quantity for an ellipsoid in its body-fixed frame is given by:

$$G = \left\{ abc \left[ \frac{(x')^2}{a^4} + \frac{(y')^2}{b^4} + \frac{(z')^2}{c^4} \right] \right\}^{-2},$$
(3.42)

which allows for a definition of the radius of Gaussian curvature through:

$$r_G = \frac{1}{\sqrt{G}}.\tag{3.43}$$

This approach is used by Jain et al. (2019) for the calculation of lubrication forces between ellipsoidal particles and appears also in Ardekani et al. (2016). Since DLVO forces are also a product of a surface-surface interaction, a similar method can be constructed. Namely, the spherical radius r can be replaced with a representative radius  $r_G$ , approximating the amount of surface interaction taking place for a given orientational configuration. These radii are calculated at the points of closest approach on the surfaces of the interacting particles, obtained through the collision detection scheme, with those points previously referred to as  $R_1$  and  $R_2$ . Considering two interacting bodies which are orientated differently to one another in general, each particle will have its own radius of Gaussian curvature at its own point of closest approach  $r_{G,1}$ ,  $r_{G,2}$ , as in Figure 3.13. Since forces are being discussed, it is required that their magnitudes are shared equally and oppositely, and so these radii must be combined to form a characteristic length which represents the correct interaction magnitude. This is done by taking the harmonic mean of the radii of Gaussian curvature:

$$R_G = \frac{2 r_{G,1} r_{G,2}}{r_{G,1} + r_{G,2}}.$$
(3.44)

This characteristic length  $R_G$  then replaces the volumetrically equivalent spherical radius *r* in Equations (3.37) and (3.38).



Figure 3.13: Schematic demonstrating how the surface interaction is captured about the point of closest approach in Equation (3.43) – not drawn to scale.

It was found that this approach performed well in introducing orientational dependencies into the DLVO force modelling for interactions between normal configurations; however, there is a further degree of freedom not accounted for in the above model whose effect becomes pronounced when particles interact in off-normal

configurations. After derivation and implementation of the above model was completed, it was then discovered by the author that Everaers and Ejtehadi (2003) had proposed a term  $\chi_{ij}\eta_{ij}$  that accounts for the orientational dependency based upon the local surface curvature, *with* consideration for the omitted degree of freedom. This degree of freedom is the rotation of the bodies relative to one another about their shared normal vector, depicted diagrammatically in Figure 3.14.

The  $\chi_{ij}\eta_{ij}$  term was included in the equations for two interacting spherical particles to give the final equation for the force between two interacting non-spherical particles and it was this new equation that was ultimately taken forward into the generation of the results presented herein. The first derived model was found to generate the same results as the second in the quiescent fluid velocity study, as demonstrated in Chapter 4; however, across the full spectrum of orientations, discrepancies do arise – which is a result of the differing interaction potentials demonstrated at the end of this subsection in Figure 3.15(d).

Including the term given by Everaers and Ejtehadi (2003), the final equations are written as follows, where the first bracketed term on the RHS represents the vdW attraction and the second term describes the EDL repulsion:

$$\boldsymbol{F}_{\text{DLVO}} = -\chi_{ij}\eta_{ij} \left(\frac{A\sigma}{12|\boldsymbol{d}|^2} - \frac{64\pi\sigma nk_B T\Theta^2 e^{-\kappa|\boldsymbol{d}|}}{\kappa}\right) \boldsymbol{\hat{d}}, \qquad (3.45)$$

$$\chi_{ij}\eta_{ij} = \frac{2\sigma^{-1}}{\sqrt{\left(\frac{1}{\zeta_i} - \frac{1}{\zeta_i'}\right)\left(\frac{1}{\zeta_j} - \frac{1}{\zeta_j'}\right)\sin^2(\alpha) + \left(\frac{1}{\zeta_i} + \frac{1}{\zeta_j}\right)\left(\frac{1}{\zeta_i'} + \frac{1}{\zeta_j'}\right)}}$$
(3.46)

Here, *d* is the vector of closest approach,  $\zeta_{p=i,j}$  and  $\zeta'_{p=i,j}$  are the surface's principal curvatures at the point of closest approach, and  $\alpha$  is the relative orientation of the particles about the shared closest approach vector. Finally,  $\sigma$  is the characteristic length-scale of the particles, taken simply to be *r* and not  $R_G$  (to avoid accounting for the surface curvature twice). A complete description of  $\chi_{ij}\eta_{ij}$  is given by Schiller et al.

(2011), where it is derived geometrically. As expected, the effect of  $\chi_{ij}\eta_{ij}$  on the potential is that the well deepens for particles experiencing greater surface interaction as a result of their configuration, and shallows for the reverse.

The surface's principal curvatures can be calculated directly, from the *first fundamental form* (Yan Bin-Jia, 2020), or indirectly using the Gaussian and mean curvatures. Since the Gaussian curvature is already calculated in the initial implementation, an indirect approach was chosen, which is conceptually simpler. Hence, the mean curvature can be calculated by (Bektas, 2017):

$$H = \frac{|(x')^2 + (y')^2 + (z')^2 - a^2 - b^2 - c^2|}{2 a^2 b^2 c^2 \left[\frac{(x')^2}{a^4} + \frac{(y')^2}{b^4} + \frac{(z')^2}{c^4}\right]^{3/2}}.$$
(3.47)

From this, the principal curvatures  $\zeta_{p=i,j}$  and  $\zeta'_{p=i,j}$  can be determined by exploiting the fact that both the Gaussian and mean curvatures can be expressed in terms of the principal curvatures:  $G = \zeta \zeta'$  and  $H = \frac{(\zeta + \zeta')}{2}$ , by definition. These relations are then used to form a quadratic equation for  $\zeta$ , from which the desired variables can be found:

$$\zeta^2 - 2H\zeta + G = 0, \qquad (3.48a)$$

$$\zeta' = 2H - \zeta \,. \tag{3.48b}$$



Figure 3.14: Diagram depicting the definition of  $\alpha$  for two interacting spheroids.

In Equation (3.46),  $\alpha$  is described as the relative orientation of the particles about their closest approach vector, d. This relative orientation is defined in terms of the symmetry axes of the particles. To understand this, the closest approach vector connecting the particles is considered in Figure 3.14 and it may be noted that the rotation of the particles relative to one another about this vector preserves the surface properties at the points of closest approach – if the angle made between  $\hat{x}'$  and d is preserved. However, the surface properties local to (but not directly at) this point will change as a result of orientation. In order to calculate  $\alpha$ , the angle that the  $\hat{x}'$  vector of the particle makes in the plane normal to  $\hat{d}$  is first calculated for both particles, then the difference in these two angles is computed.



Figure 3.15: Non-spherical DLVO interaction potentials for the three available modelling approaches, compared also to a sphere-sphere interaction. The insert window shows the pertaining orientation. On the legend, Schiller refers to Schiller et al. (2011), and Everaers refers to Everaers and Ejtehadi (2003). Separation distance is normalised by the characteristic particle radius.

In Figure 3.15, the effect of incorporating the orientational dependencies into the model for aspect ratio 5:1 needles is illustrated, where the various approaches mentioned are contrasted with a sphere-sphere interaction. The first three subfigures, (a) to (c), examine the pole interactions, where the curvature radii are maximal and minimal. It is observed that end-point-to-end-point interactions are energetically unfavourable, with a well depth roughly an order of magnitude lower than spheres of the same volume. This is in contrast to the configuration in Figure 3.15(c), where the well depth more than doubles for this alignment, as compared to spheres.

It is noted that for these normal interactions, the adopted model (Everaers and Ejtehadi, 2003) produces identical results to those used by the authors noted above; however, taking an arbitrary configuration of the particles in Figure 3.14(d)

demonstrates how the adopted model begins to differ in the predicted well depth. Lastly, it is noted that the orientation terms that appear in the approaches of Everaers and Ejtehadi (2003) and Schiller et al. (2011) are identical, with their lines falling precisely on top of one another in the plots, despite their formulation being expressed differently in the cited papers. They are both guided by the mathematical approach of White (1983). These plots were generated by implementing all three proposed models in Python, and that of Everaers and Ejtehadi (2003) was therefore used in the full solver.

# 3.2.9 Non-Spherical Soft-Sphere Overlap Calculation

The full description including the reasoning and motivation behind developing softsphere approaches, as well as the challenges associated, are discussed in detail in Chapter 6. For brief context, the key challenge with soft-sphere models for nonspherical particles is that the method necessitates an overlap of the particle surfaces, which is nontrivial to keep track of for arbitrarily orientated bodies.

When two ellipsoidal particles are overlapping, the problem of finding the two nearest surface points becomes a problem of finding the deepest points of overlap, from which one can work out the overlap distance and thus the requisite collision force. The optimisation method that found the closest points of approach is no longer useful for the overlapping problem, since the problem itself is no longer *convex* – and so the iterative procedures no longer converge.

The overlap problem itself is different in nature to the computation of the closest distance and is more difficult to solve. The previously described *common normal* condition uniquely defines the precise closest distance vector between non-overlapping surfaces. There is no simple analogue for overlapping bodies.

For a soft-sphere scheme to be implemented between two colliding bodies, the chosen method must return,

- 1. A shared contact point, to which the force interaction is assumed localised,
- 2. A shared collision normal, along which the force acts, and
- 3. A penetration or overlap depth.

The following subsection presents three different approaches for the calculation of the listed quantities. Additionally, since the chosen soft-sphere scheme also dissipates energy, it is necessary to compute the relative velocity at the collision point, which is also described.

# Approach 1: Collision Point Tracking

The first method relies on a simple assumption: that the deepest collision points remain the same for the duration of the collision. This is a reasonable assumption for small overlaps in normal collisions, which take place over a relatively short interval with respect to the other physics of the problem. This assumption has been used in the literature for other collision-based approaches (Wynn, 2008). However, to the best of the author's knowledge, the presently described approach is new.

If the stated assumption holds, then there is a simple way forward, because the location of the collision points an instant before collision are already defined within the present framework. Those are the points previously used to compute the hard-sphere collision. The implicit second assumption here is that the deepest overlap points, which this method aims to track, are also the correct points to be using to define the location of the collision points.

It is possible to perform no hard-sphere collision and instead track the location of these fixed points for the duration of a soft-sphere collision to easily and efficiently calculate the elusive parameter. Dependent on the overlap depth, a penalty force is incurred by the overlapping particles, which acts to separate them – so the overlap depth majorly influences the motion of the bodies. Since the particles are rigid bodies, the location of points at the surface will obey the rigid body equation and follow the same motion that is applied to the full body due to the soft-sphere force. Thus, although the common normal algorithm for computing these points no longer works after overlap, the location of the assumed points is still known to the same accuracy as that of the particle's motion.

The developed algorithm is described below.

- 1. Calculate non-overlapping variables:
  - Calculate closest surface points  $R_1$ ,  $R_2$  between the ellipsoid surfaces using Jain et al.'s (2022) common normal detection scheme.
  - Calculate the vector of closest approach  $d = R_2 R_1$ .
  - Calculate the surface separation distance d = |d|.
  - Calculate the velocity at points of closest approach on the surface using the following rigid body equation,  $v_c = v_p + \omega_p \times r_c$
- 2. Calculate the relative velocity at the surface points:  $v_{rel} = |v_{c,2} v_{c,1}|$ . Determine whether overlap will occur in the next time-step by checking whether the distance travelled in the next update exceeds the separation distance, i.e.,  $v_{rel} \cdot dt > d$ .
  - If No, return to Step 1.
  - If Yes, store collision points  $R_1$ ,  $R_2$  by setting  $T_1$ ,  $T_2 = R_1$ ,  $R_2$ .
  - Advect  $T_1, T_2$  separately at the end of the timestep, once all forces are calculated using  $T_i^{j+1} = T_i^j + v_{c,i}^j \cdot dt$ .
  - Proceed to Step 3.
- 3. An overlap was predicted in the previous timestep for the new time level. Advance time level:  $t^{j+1} = t^j + dt$ . Confirm the overlap. The following approach was used to confirm the overlap in the initial implementation used for the results in Chapter 6. However, it was later realised that it is much simpler to just evaluate the candidate overlap point inside the shape function of the other body – this subsequently is implemented and described in *Approach 3*.
  - Calculate the distance to centre for overlapping point. That is: the distance between deepest overlap point T<sub>2</sub> and particle centre C<sub>1</sub>. Compare this with the distance to surface, the distance between C<sub>1</sub> and S<sub>1</sub> where S<sub>1</sub> must lie on the line C<sub>1</sub>T<sub>2</sub>. Shown diagrammatically in Figure 3.16.

•  $S_1$  is an unknown, to find this point, rotate back to particle 1 reference frame. Then, use a line search along the line  $(\overline{C_1T_2})' = T'_2$  updating location x' = (x', y', z') until

$$\left(\frac{x'}{a}\right)^2 + \left(\frac{y'}{b}\right)^2 + \left(\frac{z'}{c}\right)^2 - 1 = 0$$
 is satisfied.

• If the distance from  $C_1$  to  $S_1$  is greater than  $C_1$  to  $T_2$ , then overlap is confirmed, proceed to step 4. Else, return to step 1 and discard  $T_1, T_2$ .



Figure 3.16: Diagram showing the approach taken to confirm overlap.

- Overlap is now confirmed, return zero for the interparticle separation calculations. Turn off the DLVO forces as recommended by Mihajlovic et al. (2020). Calculate the required soft-sphere collision variables (with reference to Figure 3.17):
  - Collision point  $P_{col}$  is the midpoint of  $\overline{T_1T_2}$ .
  - Collision normal  $n_{col} = \frac{\overline{T_1 T_2}}{|\overline{T_1 T_2}|}$ .
  - Overlap distance  $\delta = |\overline{T_1 T_2}|$ .
  - Relative velocity at collision point  $v_{c,rel} = v_{c,2} v_{c,1}$ .
  - Velocity acting along collision normal is  $v_{col} = (v_{c,rel} \cdot n_{col}) n_{col}$



Figure 3.17: Definition of the overlap variables computed in Approach 1.

- 5. At the end of the timestep, update  $T_1, T_2$  using the new velocities. Do this until the overlap condition is no longer met.
  - Important subtlety: in the code, the tracking of the collision continues after the points are no longer inside the opposite shape for a couple of timesteps because other nearby points may still be overlapping after the completion of the theoretical overlap. Here,  $T_1, T_2$  will not necessarily be the deepest points at the final collision step due to torques induced. If the common normal algorithm will not converge, it is implied there must still be a small overlap and as such a small force is added to move the particles apart. That can be seen in Fig 3.18, where a small overlap contribution needs to be added at t = 0.029s.
  - Once the separation distance moves beyond the distance that the particles can travel in a single timestep based on their current velocities, discard  $T_1$ ,  $T_2$  and the soft-sphere collision is complete. Return to Step 1.



Figure 3.18: Demonstration of the two combined methods for computing overlap  $\delta$ and separation *d* facilitating a soft-sphere collision between aspect ratio 5:1 needles. Both quantities are normalised by the volume equivalent radius *r*. SS is soft-sphere; HS is hard-sphere.

Figure 3.18 shows how the soft-sphere scheme is operating for a case with an arbitrarily chosen orientation between two needles. When particles are non-overlapping,  $\delta$  is zero and when particles are overlapping *d* is zero. There is a smooth transition between the two schemes for calculating distance, as the overlap begins. The transition from overlap back to non-overlap is not as smooth. This holds the potential to destabilise simulations, although this did not happen in testing between two particles.

The method works well for binary particle collisions where the initial velocities, and hence collisions, are normal. Where there is more rotation at the point of impact, or large torques are induced by the collision, the assumption that  $T_i$  are the deepest points will become worse as the collision progresses, and the transition back to non-overlapping conditions will be less robust. That means for many particles interacting, with multiple torques coming from the mechanics of multiple collisions and the background turbulence, this may not be the suitable choice. A full evaluation is made in Chapter 6.

## Approach 2: Surface Scaling

The second novel approach that was developed uses an idea from the geometric potential framework. There, a family of similar surfaces are constructed with each level surface representing a variation of the value of a 'potential'. Roughly, the amount the potential must be varied to find a contact point (black marker in Figure 3.19) can be used to define the overlap depth of the bodies. Some examples of this method being used were highlighted in Chapter 2 Section 2.3.



Figure 3.19: Illustration of varying through a family of similar surfaces to simplify the overlap problem. The outer curves represent the real definition of the ellipsoidal surfaces. The dotted lines are the similar surfaces, with the innermost curves therefore representing the solution to the relaxed problem.

In *Approach 2*, this scaling idea is applied to the existing tools available in the code. Thus, the surfaces were varied by a scaling factor  $\mu$ , which was applied to the parametric form of the surface equation (Equation (3.23)) to give the following,

$$\boldsymbol{R}^{*} = \begin{pmatrix} a^{*} \cos \varphi \sin \vartheta \\ b^{*} \sin \varphi \sin \vartheta \\ c^{*} \cos \vartheta \end{pmatrix},$$
(3.49)

where  $a^* = (a^*, b^*, c^*) = \mu \cdot (a, b, c)$  is the scaling applied to the surfaces through the principal axes vector a. The effect of this is illustrated in Figure 3.19 by the dotted lines. Then, at each level, the common normal algorithm was conducted between the scaled down surfaces until the scheme converged.

Once the solution converged, points  $R_1^*$  and  $R_2^*$  were identified as the points that satisfied the common normal constraint for a similar surface. Therefore, this gives a solution to a relaxed version of the common normal problem but does not directly satisfy it. This approximation improves for smaller overlaps and varies as a function of position on the respective surfaces.

The useful output from the above procedure is the parametric locations of the surface collision points,  $\varphi_i$  and  $\vartheta_i$ , which can be substituted back into Equation (3.23) and transformed to the global frame to give the location of the collision points in that frame. In the same way as *Approach 1*, the shared contact point is deemed to be the midpoint between these two surface points, and the collision normal is the unit normal connecting these two surface points.

In terms of how to vary  $\mu$  through the interval (0,1), or equivalently how to scale the surfaces from one iteration to the next, two approaches are proposed. The first trialled approach was to simply scale  $\mu$  by a factor close to one, e.g.,  $\mu^{p+1} = 0.99\mu^p$  or  $\mu^{p+1} = 0.999\mu^p$ , with a trade off between speed and precision. This biases the sampling of  $\mu$  towards values closer to one, which is reasonable given the small overlap requirement, but is a somewhat arbitrary method and can lead to long converge times and lower precision in cases where the penetration depth is greater. This was the method chosen initially and is thus the method that was used in the results shown in Chapter 6, where 0.995 was the chosen scaling factor. This typically required hundreds of total iterations to converge.

A second, more intelligent, approach would be to use a bisection algorithm, taking advantage of the fact that the algorithm does not converge when the surfaces overlap. Thus, there is a condition (i.e., whether the algorithm converged or not) that allows one to halve the search interval repeatedly, leading to comparatively fast convergence. The problem that was ultimately seen with the attempted implementations of this idea, irrespective of the search method used, was that the formulation is too heavily reliant on the common normal algorithm either converging or not to determine the overlap state of the bodies. Whilst this is an extremely strong predictor, it is not entirely robust, as sometimes the algorithm will converge on erroneous points when in a state of overlap. This was no great problem for the two particle case, but when moving to multiple particles this edge case was able to cause problems as there are many simultaneous overlaps to evaluate.

The solution to this issue is to pair the algorithm with a precise mathematical definition of the overlap state of the bodies, which is typically done through constructing a fourth-order polynomial characteristic equation (that describes the orientational state of the bodies) and then interrogating the discriminant to understand whether they are in a state of overlap (Pazouki et al., 2012). Alternatively, one can construct a matrix system that describes the problem and this is either degenerate or not depending on the overlap state.

Whilst feasible, at this point the methodology was becoming complex without guarantee of the desired level of robustness, as well as getting very close to a standard geometric potential implementation. So, it was reasoned that a more pragmatic approach was to find a demonstrably robust implementation for the multi-particle case and proceed with that. It will be demonstrated in Chapter 6 that both concepts work for the two-particle case, including for particle agglomeration. As future work, it would be of interest to see whether the two proposed ideas can be made reliable for multi-particle simulations with minor modifications.

#### Approach 3: Optimisation Technique

The following method adapts the work of Podlozhnyuk et al. (2017), a non-spherical soft-sphere scheme which is used in the open-source discrete element method (DEM) solver LIGGGHTS (Kloss et al., 2012), where its robustness has been repeatedly demonstrated. Their formulation is for the general case of superquadric particles, whereas here the model is implemented into a FORTRAN 77 code that is capable of handling the specific case of ellipsoids (and in the present thesis only spheroidal particles are investigated.) This new capability is then included as a module to the

existing code which sits within a Nek5000 coupling framework as before. However, with this robust capability developed, it is possible to extract the full particle module outlined in the methodology for the purpose of stand-alone DEM simulations of non-spherical particles should future work choose to go in this direction.

## 3A. Problem Formulation

The general idea of the method is for an optimisation technique to converge on a contact point that fulfils two competing constraints. This is achieved by a Lagrange multiplier approach.

For the twice continuously differentiable interacting spheroidal particle surfaces, defined by their global shape functions,  $F_1(x)$  and  $F_2(x)$ , it is possible to formulate the contact detection problem as seeking a point that is midway between particles 1 and 2, and closest to both (Houlsby, 2009):

Minimise 
$$F_1(x) + F_2(x)$$
, subject to  $F_1(x) = F_2(x)$ . (3.50)

The formulation of the Lagrange multiplier approach will be presented momentarily, but it is important to define some of the key building blocks of the method. Firstly, the shape function of a spheroid with its symmetry axis aligned with x' in the body frame can be taken from Equation (3.12), rewritten here as

$$f(\mathbf{x}') = \frac{(x')^2}{a^2} + \frac{(y')^2 + (z')^2}{b^2} - 1.$$
 (3.51)

This gives a convenient way of confirming overlap, as well as attacking certain subproblems within the method:

- $f(x'_0) < 0$  defines a point  $x'_0$  inside the spheroid,
- Whilst  $f(x'_0) = 0$  implies a point on the surface,
- And  $f(x'_0) > 0$  implies the point lies outside of the surface.

Thus, a transformation back to the body frame always allows evaluation of the penetration state of a given coordinate with respect to the spheroid.

Such a transform can make use of the raw quaternion method described earlier, but to align closely with the cited paper, this was performed via

$$\boldsymbol{x}' = \boldsymbol{A}^T \cdot (\boldsymbol{x} - \boldsymbol{x}_c). \tag{3.52}$$

using the rotation matrix A defined in Equation (3.14).

Thus,

$$F(\boldsymbol{x}) = f(A^T \cdot (\boldsymbol{x} - \boldsymbol{x}_c)), \qquad (3.53)$$

represents an evaluation of the global shape function, using the local analogue with a suitable transformation.

The gradients and second derivatives are required in the algorithm to form the Hessian matrix. Again, these are much easier to compute in the local body frame and so the gradient in the body frame is computed as,

$$\nabla' f = \begin{pmatrix} 2x'/a^2 \\ 2y'/b^2 \\ 2z'/b^2 \end{pmatrix},$$
(3.54)

and the second derivatives form a constant local Hessian matrix as follows:

$$He(f) = \begin{pmatrix} f_{xx} & f_{xy} & f_{xz} \\ f_{yx} & f_{yy} & f_{yz} \\ f_{zx} & f_{zy} & f_{zz} \end{pmatrix} = \begin{pmatrix} 2/a^2 & 0 & 0 \\ 0 & 2/b^2 & 0 \\ 0 & 0 & 2/b^2 \end{pmatrix},$$
(3.55)

where the subscripts are used to denote a partial derivative of the function with respect to the written variable. Using methods described previously, this allows the transformation back to global coordinates via,

$$\nabla F(\mathbf{x}) = A \cdot \nabla' f, \qquad (3.56)$$

$$He(F) = A \cdot He(f) \cdot A^{T}.$$
(3.57)

110

This offers a useful simplification to the general framework described by Podlozhnyuk et al. (2017) which has cross terms in the Hessian and further complexities relating to later curvature computations. The present code has been structured (in terms of the subroutine variables and data) such that this can be extended to superquadrics quite straightforwardly if required, which would further allow the simulation of cubes and cuboids.

Returning to the optimisation problem and introducing a Lagrange multiplier  $\Lambda$ , the Lagrangian for this problem is

$$L(\mathbf{x},\Lambda) = F_1(\mathbf{x}) + F_2(\mathbf{x}) + \Lambda \left(F_1(\mathbf{x}) - F_2(\mathbf{x})\right).$$
(3.58)

The general theory gives  $\nabla_{x,\Lambda} L(x,\Lambda) = 0$  as the location of the stationary point of the system which can be computed and simplified to give the following system of equations,

$$\nabla F_1(x) + \mu_L^2 \nabla F_2(x) = 0, \qquad (3.59a)$$

$$F_1(x) - F_2(x) = 0, (3.59b)$$

with  $\mu_L^2 = (1 - \Lambda)/(1 + \Lambda)$ . This gives four equations and four unknowns, and the  $x_0$  that satisfies this system is said to be the 'contact point'. In the present framework, this is used to identify the collision variables in the state of overlap only. The methodology for determining the inter-surface distance described earlier is retained for any pair of particles that are not overlapping, as with the previous two approaches, even though the present method also converges for spheroids not overlapping and could thus be used to define a distance function.

The previous method of Jain et al. (2019) gives the point at which the surface normal vector is shared, which is the geometrically precise way to define the shortest distance between convex shapes. The presently described formulation is not as precise, but the problem of overlap is more complex, and this method has the advantage of being proven to robustly and reliably define contacts. Paired together, there is a strong framework developed for the complete definition of a dynamic system of non-spherical particles.

Even though the model assumptions change between the three approaches, all three are successful in solving the problem, with different advantages and disadvantages addressed in Chapter 6.

#### 3B. Numerical Solution

To solve the system formulated in Equation (3.59) an iterative scheme based on Newton's method is employed. Following Podlozhnyuk et al. (2017), the problem can be written as the linear matrix system

$$J \cdot \delta \mathbf{Z} = -\boldsymbol{\Phi} \tag{3.60}$$

where the solution variables  $\mathbf{Z} = (x, y, z, \mu_L)^T$  are incrementally updated through solving the above system for  $\delta \mathbf{Z}$  and computing

$$Z^{n+1} = Z^n + \alpha_L \cdot \delta \mathbf{Z}. \tag{3.61}$$

Here,  $\alpha_L$  is a scaling parameter that ensures convergence. Specifically, it relaxes the increment of update at each iterative step in a subloop until  $|\Phi^n| > |\Phi^{n+1}|$ , which was implemented in the same way as the heuristic approach developed in Section 3.26 to fix the convergence issues of Jain et al.'s (2019) common normal approach (i.e., repeatedly halving the value of  $\alpha_L$  when the described condition is not met and recomputing).

The vector  $\boldsymbol{\Phi}$  stores the formulation of the problem from Equation (3.59), thus

$$\boldsymbol{\Phi} = \begin{pmatrix} \Phi^{(1)} \\ \Phi^{(2)} \\ \Phi^{(3)} \\ \Phi^{(4)} \end{pmatrix} = \begin{pmatrix} \partial_x F_1 + \mu_L^2 \partial_x F_2 \\ \partial_y F_1 + \mu_L^2 \partial_y F_2 \\ \partial_z F_1 + \mu_L^2 \partial_z F_2 \\ F_1 - F_2 \end{pmatrix},$$
(3.62)

with  $\partial_{x_i}$  used as shorthand for the partial derivative operator with respect to  $x_i \in \{x, y, z, \mu\}$  and a bracketed superscript is used to represent the entries of the vector  $\boldsymbol{\Phi}$ . In Equation (3.60), *J* is the associated Jacobian of  $\boldsymbol{\Phi}(x, y, z, \mu)$ ,

$$J = \begin{pmatrix} \partial_x \Phi^{(1)} & \cdots & \partial_\mu \Phi^{(1)} \\ \vdots & \ddots & \vdots \\ \partial_x \Phi^{(4)} & \cdots & \partial_\mu \Phi^{(4)} \end{pmatrix},$$
(3.63)

$$\therefore J = \begin{pmatrix} \partial_x^2 F_1 + \mu_L^2 \partial_x^2 F_2 & 0 & 0 & 2\mu_L \partial_x F_2 \\ 0 & \partial_y^2 F_1 + \mu_L^2 \partial_y^2 F_2 & 0 & 2\mu_L \partial_y F_2 \\ 0 & 0 & \partial_z^2 F_1 + \mu_L^2 \partial_z^2 F_2 & 2\mu_L \partial_z F_2 \\ \partial_x [F_1 - F_2] & \partial_y [F_1 - F_2] & \partial_z [F_1 - F_2] & 0 \end{pmatrix}$$

This gives a  $4 \times 4$  matrix that needs to be constructed and then inverted to give the key ingredient of the solution at each iterative step,  $\delta Z$ .

Inverting a  $4 \times 4$  matrix is the bottleneck. In the special case of spheroids, the fixed sparsity of the matrices can be exploited, seen above in Equation (3.63), for a small advantage. Thus, a tailored implementation of *LU*-decomposition was employed, ensuring no redundant calculations.

This approach splits the matrix *J* into lower *L* and upper *U* diagonal matrices such that

$$J = LU, (3.64)$$

through Gaussian elimination principles (Doolittle's method), allowing

$$J^{-1} = (LU)^{-1} = U^{-1}L^{-1}$$
(3.65)

which is comparatively simple to compute for triangular matrices L and U via forward and back substitution, respectively.

It is noted here that since FORTRAN 77 is the native language of Nek5000, no external libraries were used for these calculations and every algorithm in the eventual code base was written from the ground up utilising only fundamental functions (e.g., trigonometric, exponential, absolute value, etc.) and standard program procedures, which greatly increased the amount of development work required, but served as useful training.

In contrast to the other two approaches, this method solves directly for the contact point, rather than for the two surface points on the respective colliding bodies, which is not sufficient to determine the overlap depth and collision normal. Therefore, it is necessary to construct an overlap normal via  $\mathbf{n} = \nabla F_1 / |\nabla F_1|$  or  $\mathbf{n} = -\nabla F_2 / |\nabla F_2|$ . One can then iterate from the collision point along this normal direction until the evaluation

of the shape function at the iterative point is approximately zero, to identify the colliding surface points. This was again done using a bisection method, exploiting the fact that evaluation of the shape function changes sign whenever the surface is crossed. This provides a rule through which to halve the search domain, i.e., evaluate the shape function at the midpoint and see if the sign changed: if it did, then the root is contained within this half of the search domain; if not, then it lies in the other half of the domain. There are faster algorithms for finding such roots, but this worked quickly in the present context.

Finally, Podlozhnyuk et al. (2017) described a procedure for guaranteed convergence of the method, wherein for the first contact between a particle pair, the problem was solved from the equivalent spherical case to form an initial guess of the contact point and then iteratively the particles were 'grown' into the correct shape by gradually increasing the shape and blockiness parameters of the superquadric surfaces towards their true values. This was implemented in the present code, with the iterative procedure first being conducted on spheres with radius equal to the smallest spheroid radius and then incrementally grown to the full spheroidal shape through 10 equal intervals. Ultimately, the growth procedure was seen to be unnecessary and was turned off for the results generated in Chapter 6, saving computational resources as the method converged anyway. Full details of the more complex approach are given in the referenced paper. In the present work, using the analogous spherical case with radius equal to the smallest spheroidal radius equal to the smallest spheroidal radius equal to the smallest spheroidal radius min(a) as the initial guess for the Newton iteration led to convergence in all cases.



Figure 3.20: Illustration of varying the shape parameters to give a reasonable starting value for the iterative method, ensuring convergence in the full case.

# 3.2.10 Non-Spherical Soft-Sphere Collision Physics

For simplicity, and to be consistent with the hard-sphere implementation, only normal collisions were considered, i.e., the tangential force component was neglected, but this could be included straightforwardly as an extension to the method without the need to compute any further variables.

The spring-dashpot framework is used for the soft-sphere collision physics, where a spring force and a damping force are used to separate the particles and reduce the velocities post-collision, respectively. The normal collision soft-sphere force is thus,

$$\boldsymbol{F}_{SS} = -k_n \cdot \delta \cdot \boldsymbol{n}_{col} - \eta_n \cdot \boldsymbol{\nu}_{col}$$
(3.66)

with  $n_{col}$  again the collision normal,  $\delta$  the overlap depth and  $v_{col} = (v_{c,rel} \cdot n_{col}) n_{col}$  the relative velocity at the contact point in the direction of the collision normal. The model coefficients are  $k_n$  the spring constant and  $\eta_n$  the dashpot coefficient. These are derived by Van Der Hoef et al. (2004) and represent the solution of a linear harmonic oscillator which enforces the following relationship: that the magnitude of the relative velocity post-collision is equal to the coefficient of restitution times the pre-collision relative velocity, which introduces *e* back into the model.

$$k_n = \frac{m_e(\pi^2 + \ln^2 e)}{(N_c \Delta t)^2} \cdot \sqrt{R_G}$$
(3.67a)

$$\eta_n = -\frac{2m_e \ln e}{N_c \Delta t} \tag{3.67b}$$

where  $m_e = \left(\frac{1}{M_1} + \frac{1}{M_2}\right)^{-1}$  is the effective mass, and  $N_c$  is the number of timesteps the collision is computed over, effectively determining the collision duration. Set too high, the collisions take unrealistically long to complete, and the overlap depth becomes large. Set too low, there is a loss of accuracy as the force varies largely between timesteps. In the present work, a value of  $N_c = 40$  was used, which is much higher than typical approaches, e.g.,  $N_c = 8$  in Ardekani et al. (2016), but the timestep of the simulation is extremely low to accommodate the van der Waals forces, so the collision duration should not be made vanishingly small in real terms.

Finally, the spring constant is scaled by the radius of Gaussian curvature to reflect the difference in stiffness along an ellipsoidal body. The square root of this quantity was chosen in line with Zheng et al. (2013). This quantity is implemented using the same relations described in the development of the non-spherical DLVO force model, which can be seen in Equations (3.42) – (3.44) and thus the existing code was able to be leveraged again here.

# 4 Non-Spherical Particle Interactions in a Quiescent Liquid

# 4.1 Introduction

The thesis to this point has described the development of a unique methodology for the numerical simulation of particle agglomeration within initially quiescent and turbulent fluid flows, and there was a significant development period to reach this point. Now, beyond testing the developed techniques, it is important to investigate the sensitivity of the agglomeration model to the relevant degrees of freedom: velocity and orientation. These two properties affect the total kinetic energy and potential energy, respectively, and thus govern the agglomeration outcomes of interaction events between colliding particle pairs.

The relative orientation of the particles affects the amount of surface interaction taking place at the closest point of approach and therefore alters the depth of the DLVO potential well. The relative velocity of the particles affects their kinetic energy and thus determines whether particles can escape the potential well. The simulations performed and discussed in this chapter are designed to investigate the agglomeration and collision behaviours with respect to these variables in a quiescent liquid so as to decouple the turbulence effects that will become important later.

To this end, a parameter study is undertaken to investigate the role of initial velocity in particle interactions between spheres and between four representative orientation cases of spheroidal particles. Further, a Monte Carlo study is used to investigate the full orientational parameter space, generating probability density functions (PDFs) that describe the overall system behaviour. Analysis of the PDFs takes place with respect to the separation distance and relative velocity tendencies, as well as to particle alignment. Further breaking down of these results into the distinct categories of agglomerating and bouncing particles helps to understand which orientations, velocities, and system properties lead to an agglomeration event, rather than a rebound.
## 4.2 Agglomeration Transition of Spherical Particles

The potential well governing a particle interaction remains the same for all orientations of interacting spheres. Hence, the first relevant dynamic variable in the interaction is the approach velocity, or the initial kinetic energy, of the particles. At some fixed velocity, a transition should occur where particles no longer agglomerate upon collision. This is referred to herein as the *cut-off velocity*, below which particles agglomerate and above which they rebound. The precise value of this velocity is dependent upon the properties of the system, e.g. the Hamaker constant, the Debye length, the zeta potential at the surface, the fluid temperature, and so on.

For all of the present simulations, the particles under investigation are composed of the material calcite, and they are suspended in water at 20 °C, since this is an informative analogue to nuclear waste systems where calcite can be used as a simulant due to its similar physical and chemical properties (Hussain et al., 2021). The full list of properties for the quiescent system is presented in Table 4.1.

Parameter	Value	Units
Particle density, ${m  ho}_p$	2710	$kg m^{-3}$
Volume-equivalent spherical radius, $r$	50	μ <b>m</b>
Restitution coefficient, <i>e</i>	0.4	-
Hamaker constant, A	22.3	zJ
Inverse Debye length, $\kappa$	328947368	<i>m</i> -1
Surface charge density, $\Theta$	0.00037	Cm <sup>-2</sup>
Van der Waals cut-off distance	5	nm

Table 4.1: Calcite particle parameters used in the simulations.

To investigate the value of the cut-off velocity for spheres, a parameter study over a suitable velocity range was undertaken. Two particles were injected into the quiescent box of liquid at a fixed separation of  $5\mu m$ , which is one tenth of a particle radius. This small separation was chosen since the velocity field is quiescent, and so the particle motion does not change significantly on approach. The particles were given equal and opposite velocities such that a normal collision occurs at the centre of the box. The

range of velocities were chosen to represent all collision velocities observed in a channel flow at  $Re_{\tau} = 180$  (Mortimer et al., 2020); here, the system properties were also chosen to be consistent with the cited study. An additional benefit of this present study was thus to obtain a validation of the cut-off velocity chosen in the cited paper's Lagrangian particle tracking (LPT) agglomeration model, which itself is an adaptation of Breuer and Almohammed's (2015) method. In their deterministic energy-based model, and in LPT agglomeration models in general, the physics are being resolved at a lower fidelity on account of the agglomeration determination procedure acting over a single timestep as a binary calculation. The deterministic instantaneous model is employed upon detection of a collision. There, a determination is made, based on the colliding particle's energies, as to whether an agglomerate is formed or not. A formed agglomerate of two spherical particles then becomes a larger spherical particle that preserves the volume of the particle-phase. The present methodology offers a significant increase on this level of detail and thus the results should be instructive to LPT models, perhaps offering a means of validation to the cut-off velocities chosen therein.

For the present study, 14 separate simulations were conducted with initial relative velocities assigned over the range  $[0.175, 2.45] mm s^{-1}$ . The total number of timesteps for each simulation was chosen such that a collision occurred at half of the total number of timesteps. In the absence of interaction forces, this would mean the particles always finish at the same separation at the end of the simulation, irrespective of initial velocities; however, with interaction forces present, one can look at the final separation distance or mean distance covered to ascertain the influence of the interaction. Particles finishing with approximately zero separation have agglomerated and those with appreciable separation have rebounded. The mean separation (averaged over the respective simulation) can then be plotted as a function of the initial relative velocity to determine the so-called 'cut-off velocity', and this will be the means used to determine cut-off velocity (e.g., in Figure 4.3).



Figure 4.1: Time evolution of inter-particle separation for interacting particles, illustrating how this looks for typical agglomeration (solid line) and bouncing (dashed line) events.

In Figure 4.1, the time evolution of the separation distance between particle surfaces is shown for two example cases of particle agglomeration and rebound. For this illustration, the same initial velocity is chosen in both cases, but for the *Bounced* case, the attractive forces are not resolved, thus making clear precisely what effect the attractive forces are having.

Particle pairs finishing with roughly zero separation after the full simulation time form what is constituted in this work as an agglomerate, where the interacting pair of particles have succumbed to attractive forces. In this state, particles lie nearby at approximately zero but finite separation with approximately zero relative velocity. They are bound in this state by the attractive part of the DLVO forces, or the van der Waals (vdW) force, requiring thereafter a sufficient injection of kinetic energy to break their bond. In a real physical system, this may come about from the surrounding fluid or a tertiary particle collision. In Figure 4.2, this separation evolution is plotted for all velocity cases (for which there are 14 equally spaced initial velocities investigated), as well as graphs of: particle relative velocities, the magnitude of the DLVO forces experienced between particles, and the relative velocity normalised by the initial velocity.



Figure 4.2: Trajectory data for the 14 initial velocity cases of spherical particle collisions. Labelled left to right, top to bottom: (a) inter-particle surface separation distance, (b) particle relative velocity, (c) total DLVO interaction force, and (d) normalised relative velocity. The legend remains consistent across subfigures.

Of the 14 initial velocities studied, 10 pairs of particles rebounded and four agglomerated – the lowest two velocity cases, as might be expected. The very lowest initial velocity case resulted in an agglomerate forming almost immediately upon collision, whilst the second lowest case displayed more interesting behaviour. There, the particles are seen to initially collide and begin to move apart, but the kinetic energy retained after the collision is not sufficient for the particles to escape the potential well and thus the van der Waals attractive force successfully acts to accelerate the particles back towards one another, reversing their velocities.

Non-instantaneous agglomeration events suggests that there is interesting interparticle behaviour occurring in a 'transition region' that lies between the cases that directly rebound and the cases that quickly agglomerate. These intermediate velocity cases display a more dynamic interplay over the interaction period – an effect which will be missed by simpler models. The ramifications of this interplay are best described in terms of the velocity evolution, discussed shortly.

Before discussing this, it is important to make explicit a key mechanism shown by this study, which is the interplay between DLVO forces and the hard-sphere collision scheme. With the given physical parameters, the contribution of the electric double layer (EDL) to the total DLVO forces is low enough that the total force remains positive, or attractive, for all separation distances (which is not the case for all systems), this is evidenced by Figure 3.12. The result of *necessarily attractive* DLVO forces is guaranteed collisions between particle surfaces that approach one another unimpeded and, as such, the post-collision velocity rather than the pre-collision velocity determines whether particles will agglomerate or rebound, as this is the lower energy state.

The observed mechanism of agglomeration is that the particles collide, losing energy, but they retain some post-collision velocity allowing them to begin to separate. At which point, the attractive forces act to bring the particles back together until another collision occurs, losing further kinetic energy, and this continues until the pair reach a form of equilibrium. In other words, post-collision, a competition ensues between the retained kinetic energy of the particle pair and the DLVO forces which act to bring the particles back together. There is a constant exchange of kinetic and potential energy observed across the present simulations in this way.

Even in this quasi-equilibrium state of agglomeration, with particles very close and barely moving, the DLVO force contribution remains non-zero and thus the particles are always being accelerated toward one another, which will be demonstrated in the subsequent discussion. From a computational perspective, there is then a continual need to compute collisions to counteract this and prevent overlap – exacerbated if there is some background disturbance to the particle's positions within the potential well, such as from a turbulence field. For some non-spherical orientations, the amount that the DLVO forces can accelerate the particles leads to velocities that surpass the

collision detection distance and an overlap can occur between agglomerated particles, which is a source of numerical instability (given that the detection scheme necessarily requires non-overlap), albeit this is rare. Special treatment needs to be taken for this case, which can either be to artificially separate the particles by moving them a very small amount when overlap is detected, or by increasing the collision detection distance such that it accounts for the maximal single timestep acceleration due to interparticle forces. The former option being chosen here, separating the particles by 1.0 nm.

In Figure 4.2(a), the particles approach along constant lines, where the gradient corresponds to the velocity. Upon reaching zero separation and undergoing a collision, particles retain 40% of their velocity (equivalently, 16% of their kinetic energy) as per the user-defined coefficient of restitution, which is chosen based on available data for calcite particles (Njobuenwu and Fairweather, 2017). The velocity change can be observed clearly in Figure 4.2(b). A loss of energy because of the impact occurs instantaneously in the collision model. So, the velocity graph should be a step function in the absence of DLVO interaction forces. However, it is observed that the lower the initial relative velocity, the more the graph deviates from this expected shape.

On the one hand, as the particles approach one another just before impact, and the DLVO force shown in Figure 4.2(c) increases, the particles are slightly accelerated. This effect is more pronounced for lower initial velocities, because the lower velocity cases allow the DLVO forces to act for a longer period. Importantly, the acceleration is only occurring over a very short distance, just before collision, where this extra kinetic energy is gained from the DLVO potential energy.

Secondly, and most importantly for the simulations in this thesis, the *observed coefficient of restitution* varies as a function of initial velocity, as a result of the DLVO forces. Where a more prolonged interaction takes place, more kinetic energy is lost overall and so the observed coefficient of restitution is lower once the particles separate, despite the collision model specifying the coefficient of restitution (COR) as e = 0.4 indiscriminately. (The observed COR can be defined here as the ratio of the observed post-interaction velocity to the observed pre-interaction velocity.) This finding indicates that agglomeration processes have a strong influence on the global energy balance of the particle-phase for systems where such interactions are prominent. For

macroscale models this implies that the enforced coefficient of restitution in a collision should vary as a function of the experienced vdW force.

Figure 4.2(d) emphasises this point by normalising the relative velocity evolutions by their respective initial relative velocity. For simulations absent of DLVO interactions, these collision events would then collapse onto one curve post-collision on this graph, but something very different is observed. That is, as the initial relative velocity is reduced, the observed COR is also reduced and the acceleration of the particles precollision is increased. The horizontal dashed line indicates the value of the COR set in the simulation, and it is therefore clear that only the fastest moving particle cases approach the set value. The slower moving cases reduce below this line monotonically until reaching the cases in which particles agglomerate. Such cases could be thought of as zero observed COR. The acceleration experienced pre-collision is so pronounced for the initially slowest moving particles that the velocity before collision is doubled compared to the prescribed initial velocity. Even still, these particles agglomerate. This shows that for the lower velocities the DLVO forces have much more control over the overall energy evolutions of the particles. In other words, the prominence, or relative importance, of these forces in a given system is a function of the initial kinetic energy of the particle phase. Said another way, fast moving particles will effectively behave independently of these forces such that their inclusion in a model is not required to the same degree.

Finally, returning to Figure 4.2(c), the temporal evolution of the DLVO forces can be seen. From this plot, it is made abundantly clear just how short the range of these forces is. As alluded to earlier, the contribution of the EDL is negligible for the present system and so the total DLVO contribution is dominated by the vdW attraction. This attraction only begins to spike in magnitude for particles extremely close to one another. The graphs themselves are only displaying data starting from  $1\mu m$  of separation, which is 1/50 of the particle radius. This poses a significant computational challenge, since the scales of the DLVO forces are miniscule in comparison to the overall particle motion (as well as the fluid timescales, demonstrated in Chapter 5). The increase in DLVO force is resolved smoothly and the simulations remain stable; however, a timestep of  $dt = 1.0 \times 10^{-6}$  was required to achieve this. For comparison, this is two orders of magnitude lower than the timestep used for the single-phase turbulence. It is later described that this becomes even starker in contrast when

replacing spherical particles with non-spherical particles. It was found that reducing the vdW cut-off parameter further increased the requirement for timewise resolution. It is therefore clear why models are typically used for multiphase fluid simulations that include agglomeration, rather than explicitly resolving these forces each timestep, but this is at the expense of resolving the previously described phenomena.



Figure 4.3: Mean particle surface separation distance over the interaction simulation for the 14 initial velocity cases.

To summarise the preceding analysis, the overall relationship between initial velocity and agglomeration is shown in Figure 4.3. It is seen that initial relative velocities in the range [0.175, 0.7]  $mm s^{-1}$  define the agglomeration range, with 0.7  $mm s^{-1}$  therefore defining the cut-off velocity for this system. Beyond this, there exists a transition region observed in the range [0.7, 1.575]  $mm s^{-1}$  where the vdW forces are still reducing the distance travelled in a simulation to varying degrees, through their action upon the post-collision velocities. Beyond 1.575  $mm s^{-1}$  the distance travelled by the particles tends towards a stable value indicating little influence from the DLVO forces.



Figure 4.4: Short time-scale view of agglomerating particles reaching a quasiequilibrium state.

Figure 4.4 demonstrates, over a very short time window, the effect that including a dynamic collision detection separation has, which was introduced in Section 3.2.5. After each collision, in which particles lose 60% of their speed and 84% of their kinetic energy, the separation distance at which a collision is detected becomes smaller, proportional to the loss of velocity. Ultimately, this allows the particles to become much closer to one another than if a fixed collision detection separation was set.

The advantage is twofold. Firstly, over time there is a much greater likelihood of the highest vdW forces being sampled by the interacting particles. Importantly, whether an agglomeration event occurs or not is strongly dependent upon the highest vdW forces sampled in an interaction. Thus, this approach increases the reliability in predicting deterministic agglomeration outcomes. In particular, when the final separation before collision can fall in a wide band between 0 and  $\varepsilon$  (the dynamic collision detection distance) as a result of low temporal resolution, there enters a degree of randomness as to whether an agglomerate forms based on exactly what this separation is. Allowing progressively smaller  $\varepsilon$  values alleiviates this without having to dramatically reduce the timestep – which would be the required approach if  $\varepsilon$  was constant, in order to recover the necessary resolution. Secondly, the

agglomerate that is formed is more stable to perturbations from the flow field because there is a greater distance for the particles to travel for them to leave the potential well.

# 4.3 Agglomeration Transition for Non-Spherical Particles of a Fixed Initial Orientation

To begin to understand the role of non-sphericity and orientation in the agglomeration outcomes, the same study was conducted again for non-spherical particles with 'normal' or 'modal' orientational configurations. Specifically, two relative orientations were chosen: those that induce the maximal vdW interaction forces and those that induce the minimal. For needles, the former occurs when the  $\hat{x}'$  axes are aligned but their centres are separated perpendicularly to these vectors, and the latter occurs when the  $\hat{x}'$  axes are aligned and their centres both lie on this same line. Simply put, either there is a long-edge-to-long-edge interaction, or there is an end-to-end interaction. For disks, the reverse is true: the maximal forces are induced by a faceto-face interaction, whilst the minimal are induced by edge-to-edge interactions. This is because for the disks, the faces represent the regions of least curvature, or greatest surface interaction; whilst for the needles, this same region is found along the long edge. For now, the degree of freedom wherein particles can rotate about the closest approach vector is excluded, although this would induce even lower forces for the disks, but would not change the needles, since the minimal configuration is symmetric about this vector. The full parameter space, including this variable, is addressed later in the orientational study, as well as in the boxes of homogeneous and isotropic turbulence in Chapters 5 and 6. These configurations can be seen in Figure 4.5, whilst the potential wells formed by the configurations were shown in Section 3.2.8.



Figure 4.5: The modal configurations of non-spherical particles used in the velocity study.

Figures 4.6 and 4.7 demonstrate the evolution of dynamic properties during the needle interactions. The behaviour is qualitatively similar to the spherical case where a transition occurs from agglomerating to bouncing events; however, the points at which this occurs are different. In other words, the cut-off velocity is a function of orientation, with disparate behaviour observed when comparing the two configurations. The minimal force inducing configuration has a cut-off velocity of les than 0.350 mm s<sup>-1</sup>; whereas the maximal configuration experiences agglomeration until an initial velocity of 1.225 mm s<sup>-1</sup> is reached.

In the time evolution plots, a relative increase in the effect of the DLVO interaction is seen for the maximal vdW orientation case as compared to the minimal. This is evidenced by the sharper increase in relative velocities which in turn affects the evolution of the separation distances. Specifically, the change in velocity before collision becomes appreciable enough so as to be seen in the separation graphs: the lines begin to deviate from a straight line, becoming curved before collision. The same effect is seen post-collision where rebounding particles have deflected separation lines owing to the strong attractive vdW component of the DLVO force. The increased strength is evidenced by the DLVO plots, wherein the graphs peak at different DLVO force values. In the maximal case, a peak of 10.5 nN is observed and for the minimal case this value lies at 0.79 nN, which is greater than a full order of magnitude difference; it can hence be determined that relative orientation is very important to the interaction strength and the subsequent interaction dynamics under the given model. This is particularly the case here because the chosen aspect ratio is 5: 1, which allows for significant deviation in the local surface properties across the particle.

The values at which the interaction forces peak in each set of plots remains the same for all velocity cases, since the orientation remains the same. (For exactly normal collisions, which these are, no torque is induced and so there is not a dynamic component of orientation.) The exact value at which the DLVO plots peak is thus a direct consequence of the chosen vdW cut-off distance, as well as the physical system parameters.

Finally, a very slight dip can be observed in the DLVO strength, best seen in the DLVO plot of Figure 4.7 for the slowest velocity case. This is owing to the EDL contribution. The vdW force reaches an artificial maximum at  $5 \times 10^{-9} m$  separation but the EDL force is not capped until a  $1 \times 10^{-11} m$  separation is achieved. For agglomerating particles, the separation distance continues to evolve within this range as the previously described energy exchange takes place; as such, there is a small amount of room for manoeuvre where the EDL force exercises minor influence on the system. This takes a small period to reach equilibrium, at which point the DLVO plots become stationary, indicating system equilibrium. The EDL contribution acts to separate the surfaces and this reduces the overall attractive magnitude of the DLVO interaction.



Figure 4.6: Simulation trajectory data for the lowest vdW force-inducing needle configuration. Plots show separation distance over time (top left), relative velocity over time (top right), DLVO force over time (bottom left), and mean distance covered in a simulation as a function of initial velocity (bottom right).



Figure 4.7: Simulation trajectory data for the highest vdW force-inducing needle configuration. Plots show separation distance over time (top left), relative velocity over time (top right), DLVO force over time (bottom left), and mean distance covered in a simulation as a function of initial velocity (bottom right).



Figure 4.8: Simulation trajectory data for the lowest vdW force-inducing disk configuration. Plots show separation distance over time (top left), relative velocity over time (top right), DLVO force over time (bottom left), and mean distance covered in a simulation as a function of initial velocity (bottom right).



Figure 4.9: Simulation trajectory data for the highest vdW force-inducing disk configuration. Plots show separation distance over time (top left), relative velocity over time (top right), DLVO force over time (bottom left), and mean distance covered in a simulation as a function of initial velocity (bottom right).

Very similar trends are seen for the interacting disks in Figure 4.8 and Figure 4.9, as those described for the needles. The maximally attractive configuration once again leads to much greater spikes in the relative velocities and forces. Deflection of the separation lines occurs again for the maximal case. All of this leads to the same divergence in the behaviour for which a greater cut-off velocity is obtained by the maximal configuration.

For the disk morphology, the cut-off velocity for the minimally orientated particles lies at  $0.525 mm s^{-1}$ , and at  $2.10 mm s^{-1}$  for the maximally orientated particles. So, for both needles and disks there is an approximate quadrupling of the velocity required to facilitate post-collision rebounds, between the minimal and maximal cases. A key result here is that the disks possess the capacity for agglomeration at the highest

velocities, as long as the particles are maximally orientated, whilst needles can avoid agglomeration at the lowest velocities so long as they are minimally orientated.

Investigating this further, the DLVO force of the disks reaches peak values of 31 nN and 1.2 nN for the maximal and minimal configurations, respectively, meaning the maximal force is triple that of the needles, whilst the minimal peak DLVO force for the disks is around double that of the needles. This immediately indicates that disks will form the strongest agglomerates, requiring the most energy to break apart, whilst the needles will create the weakest. This is owing to their geometric properties, as will be analysed in Section 4.6. Whilst this gives a strong indication of what is happening at the extremes, this gives no insight as to how the behaviour interpolates between these extremes, and so further investigation must take place to understand the full orientational parameter space.

In Figures 4.10 and 4.11, the overall results are shown for the three morphologies, where the spherical results are also presented again in contrast with the two extreme non-spherical configurations. Both the disks' and the needles' extreme cases are bisected by the spherical results. The curves of these two extreme cases and the spheres do not intersect and thus the spherical case remains bounded between them. This is to be expected, since the local surface curvature at all points on a sphere is greater than the maximal case and lower than the minimal case of a spheroid with equal volume, and the non-spherical DLVO model is primarily a function of local surface curvature when velocity is held constant.

As the mean separation distance over the simulation converges to an approximately stable value, with increasing initial velocity, the minimal curve for both needles and disks appears to roughly converge on that of the spheres. Whereas the maximal curves for both needles and disks appear to converge at a significantly lower value, and, even with increasing initial velocity, the separation appears to remain reduced. This indicates that even when the effect upon the relative velocity is not great, the distance travelled in the simulation remains lower. Contrary to these findings, it may be expected that the curves would converge upon one another as velocity increases because the impact of DLVO interactions becomes negligible. Perhaps this is the ultimate trend that would be seen if a greater range of velocities is sampled.

It is possible that this finding is a result of a prolonged interaction taking place, due to the DLVO forces but it is not clear what the mechanism is causing this.

It can also be observed from the preceding figures that the transition region – i.e. the range of velocities for which DLVO forces are having a significant impact, but not significant enough for agglomeration – spans a wider range of velocities for the maximal configurations in both morphological cases. This is particularly pronounced for the maximally orientated disks that retain significant DLVO effects even beyond the initially chosen velocity range. This was partly also the reason for expanding the velocity range in this case. For comparative purposes and completeness, a final run was undertaken in the minimal and spherical cases at a higher velocity, where convergence of their respective mean distances is further confirmed. The implication therefore is that a binary on/off agglomeration model at a macroscopic level is a better approximation for minimal vdW force configurations than for the maximal ones.

Finally, the maximally orientated disks formed a good calibration case to understand what the maximum timestep was that could be used in simulations that would span the full orientational parameter space. The very first simulations undertaken herein matched the timestep of the fluid-phase,  $dt = 1 \times 10^{-4} s$ , which was immediately shown to be a poor choice by the lack of determinism in the mean distance plots. It was found that a timestep of at most  $1 \times 10^{-6} s$  was required to ensure that the spherical results were fully deterministic, otherwise the DLVO forces were not properly sampled. The issue was that agglomeration events were seemingly as reliant on the quasi-random discrete separations they found themselves at just before collision, as much as the actual energy they possessed. The collision detection distance is a function of how far the particles can travel relative to one another in a single time step, which means that for a low temporal resolution simulation the particles can undergo a collision well before they enter the effective DLVO range, thus missing what should have been an agglomeration event. Conversely, for a small fixed separation distance, the particles would indeed experience DLVO forces more often, but there would be a large risk of overlap occurring which causes simulations to diverge.

In the case of the maximally orientated disks, this timestep needed to be lowered even further still to  $dt = 2 \times 10^{-7}$  since instabilities arose for particle pairs sitting in a potential well for extended periods. It was consistently observed that for this configuration, particles 'vibrate' against each other much more after agglomerating. The DLVO forces always remain active, and so they continue to supply significant energy to the system, facilitating motion at a very close range. For the maximal interaction cases, this supply of energy is significant enough to cause problems with the collision algorithm, because particles can accelerate enough to bypass the collision distance and cause overlap. When controlled, particles are still accelerated over a small number of timesteps, but they then lose this energy again through a hard-sphere collision and the simulation remains stable. The relative velocity plot in Figure 4.9 illustrates this where the line on the graph appears thicker after agglomeration as a result of this 'vibration' or constant exchange of energy at high frequency.



Figure 4.10: Mean separation distance plotted as a function of initial velocity for the two modal needle cases, compared with spheres.



Figure 4.11: Mean separation distance plotted as a function of initial velocity for the two modal disk cases, compared with spheres.

The key findings of this velocity study are recorded in Table 4.2, demonstrating the reliance on orientation in determining the cut-off value of velocity, and underscoring a need to resolve non-sphericity to capture the peak forces in the system.

Result	Sphere	Needle (5:1)	Disk (5:1)
Cut-off velocity, maximal config. $(mm \ s^{-1})$	0.7	1.225	2.10
Cut-off velocity, minimal config. $(mm \ s^{-1})$	0.7	0.35	0.525
Peak vdW force, maximal config. ( <i>nN</i> )	3.69	10.5	31.1
Peak vdW force, minimal config. ( <i>nN</i> )	3.69	0.79	1.20

Table 4.2: Results of the non-spherical cut-off velocity study compared by morphology.

# 4.4 The Role of Morphology in Non-Spherical Particle Agglomeration

Having varied the relative velocity of the particles, the disparate cut-off velocities between morphologies and between modal configurations indicates a dependence on the shape of the particle, as well as the orientation. To isolate these contributions, it makes sense to keep the initial relative velocity fixed and randomise the initial orientations such that the full parameter space of configurations is sampled. Then, the overall results can be compared between disks and needles to isolate the effect of morphology. Naturally, spheres cannot be included in this analysis, since their DLVO contribution is not a function of orientation. Since these are deterministic simulations, the simulation result is the same every time for spheres of a fixed initial velocity, unlike the results of Chapter 5 where the chaos of the turbulence field is able to introduce significant variability to the spherical case.

To obtain randomly orientated spheroidal particles is not straightforward, in the sense that one cannot just uniformly randomly distribute the Euler angles  $e_{\theta}$ : this generates oversampling towards the poles and thus generates bias in the distribution of configurations. To alleviate this, a special distribution must be given to the angles such that sampling of orientations is uniformly random over the surface of a sphere. If this is achieved, then the orientational parameter space being sampled for each particle is uniformly random, as required by the Monte Carlo study.

The required distribution for  $e_{\theta}$  is as follows (Simon, 2015):

$$\phi_i = 2\pi \cdot X_{i,1},\tag{4.1a}$$

$$\theta_i = \arccos(1 - 2 \cdot X_{i,2}), \tag{4.1b}$$

$$\psi_i = 2\pi \cdot X_{i,3},\tag{4.1c}$$

where  $X_{i,j} \sim U(0, 1) - i.e.$ , the three unique realisations of  $X_{i,j}$ , for particles  $i \in \{1,2\}$ , are uniformly randomly distributed in the interval [0,1], thus giving six uniform random variables per simulation. Obtaining these random variables was achieved through the *drand()* command built into FORTRAN 77 where the seed was set uniquely for each simulation through a pertaining bash script that generates the file structure required to run the many realisations on the University of Leeds' HPC system, ARC4. Figure 4.12 demonstrates the application of this procedure to a unit normal vector  $\hat{x} = (1, 0, 0)$  for 50,000 realisations obtained in Python, where a clear oversampling is pronounced on a band around the sphere's equator for the incorrect case, which is not present for the case that uses Equations (4.1).



Figure 4.12: The incorrectly sampled distribution (left) and the amended distribution (right), for uniformly sampled orientations.

Within the present computational framework, the injection location and orientation of the particle is prescribed by the user in terms of the location of the centre of the particle and the three Euler angles defining the initial orientation. However, this does not allow one to specify an initial closest surface point separation, since this is dependent upon relative orientation, for non-spherical particles. Therefore, to ensure that particle surfaces begin at a fixed separation, for all orientations, a further algorithm was required.

This injection algorithm places the particles equidistantly about the xz-plane with a separation equal to 2.1 times the length of the longest particle principal radii, to ensure no overlap is possible. Here, the separation distance is measured using the existing collision detection algorithm, and the particles are moved along the closest approach vector d by the correct amount such that the required initial separation is met.

Like the methodology adopted in the previous section, two non-spherical particle meshes of the same morphology and size are injected. Initially, particles were once again moved to 5  $\mu m$  apart but the preliminary simulation results indicated that even this small separation is an unnecessarily large distance, given how short-range the interaction forces are. Hence, a separation of 2  $\mu m$  was chosen, which is justified in the following figures, wherein the final state of the simulation is reached in advance of the total simulation time for all realisations. The goal with the analysis is to generate

PDFs that describe the average interaction behaviour of all realisations of the system and so the key is to resolve the entirety of the interaction phenomena and nothing thereafter. Any surplus data accumulated in the results simply serves to obscure the interaction behaviour in the final PDFs. The PDFs were generated by applying a Monte Carlo approach, wherein many cases were simulated to capture statistical variations. To ensure robustness, histograms of the resulting data were first plotted in Python, and kernel density estimation (KDE) was applied to obtain smooth probability density functions. The KDE plots were used exclusively in the thesis, as they provided a clear and continuous representation of the distribution. The number of instances was deemed sufficient once the PDFs remained stable upon the addition of new data, implying convergence.

As described above, upon injection, the closest approach vector between the two particles was calculated and it was along this direction that the particles were set to approach, with an equally shared relative velocity of  $0.875 \, mm \, s^{-1}$ . The injection velocity was chosen as a representative value in accordance with observations of relative collision velocities made in a  $Re_{\tau} = 180$  channel flow (Mortimer et al., 2020), which corresponds to the mean particle-phase collision velocity in the bulk flow region. Although the particles are initially separated about the *xz*-plane, directing their initial velocities normal to this plane would not ensure that the particles collide after an equal time for all simulations, hence the two particles are directed to approach along the closest approach vector *d*. As before, the simulations were allowed to run for an amount of time pre- and post-collision such that the full interaction behaviour could be simulated.

Presented below are the same time evolution plots as introduced in the previous subsection. This time, the results have been subdivided into two groups: agglomeration events and bouncing events. The determination of these categories was made by the separation distance at the end of the respective simulation: if a particle pair were at a separation distance that lies within the pertaining potential well at the final timestep, particles were said to have agglomerated, with this determination taking place in the post-processing step. Both disks and needles received 120 realisations, which was the required amount to obtain smooth PDFs for the overall behaviour of the respective morphologies once simulations had been subdivided into

agglomeration and bouncing events. The PDFs and histograms of the Euler angles were also monitored to ensure no oversampling was present for any range of angles.



4.4.1 Monte-Carlo Simulation of Randomly Orientated Interacting Disks

Figure 4.13: Temporal evolution of inter-particle separation distance for agglomerating disks (left) and rebounding disks (right).



Figure 4.14: Temporal evolution of relative velocity for agglomerating disks (left) and rebounding disks (right).



Figure 4.15: Temporal evolution of DLVO force magnitude between particles for agglomerating disks (left) and rebounding disks (right).

From the results of Figures 4.13 to 4.15, the contrast in the behaviour between agglomerating particles and rebounding particles is immediately clear for interacting disks. For the agglomeration cases, the separation plots are comparatively homogeneous. Particles pairs have all agglomerated well before the full simulation time, with some attempting, unsuccessfully, to escape the potential well, giving rise to the small bump after the collision point – the particles beginning to rebound are recaptured by the attractive vdW component of the DLVO force.

In contrast, the separation plots of the bouncing particles serve to highlight quite clearly how significant the particle orientation is on the observed coefficient of restitution. Recalling that each particle pair receives the same relative initial velocity, it is clear that there is a wide spectrum of CORs resolved by varying the orientations, evidenced by the wide range of gradients of the separation distance time evolution plots. This is further evidenced in the relative velocity plots which also show a spectrum of post-collision velocities.

As with the separation plots, the relative velocity plots display a stark contrast in the behaviour of agglomerating and rebounding particles. The 'agglomerated' relative velocity plots are noisy, with a significant number of collisions taking place for any given agglomeration event. It is interesting how there is significant noise and continual behaviour in the relative velocity plots that does not translate to significant behaviour observed in the separation distance plots. As alluded to earlier, this is because whilst the DLVO forces can significantly increase the relative velocity quickly, the particles

also lose this velocity quickly through the hard-sphere collision step and so particles remain trapped close to one another.

It is also clear that there is a much wider variety of interaction strengths taking place when varying the orientational configurations of the particles, as compared with the velocity study. Observing the DLVO plots, there is a wide range of values resolved, but the way these values are distributed is not uniform, despite the variation of the orientations being uniform. This indicates that there is an importance in the way that the DLVO force magnitude is distributed according to the shape of the surface of the particle, which will be analysed in Section 4.6. It also could indicate that despite the orientations being uniformly distributed for both particles, there could still be bias for which points on the surface of the non-spherical particles experience collisions most frequently, due to the anisotropic nature of the particle geometries and how this influences the collision point.

Interestingly, the very maximum values of the DLVO forces are rarely observed for the disks, indicating that the probability of a face-to-face collision is low. So even though disks were shown to facilitate the strongest interaction strengths, this precise configuration is rarely observed in the collisional parameter space. Further still, interaction magnitudes within its vicinity are also rarely observed. Rather, there is significant clustering of observed interaction strengths in the range [0.5, 5] nN, which is a significantly reduced range compared to the full range, demonstrated in the previous section to be [1.2, 31.1] nN. The fact that peak values are resolved even lower than the so-called minimal range, as mentioned earlier, occurs since introducing rotations about the vector of closest approach allows for even smaller forces to be experienced.

Finally, it is possible to analyse the overall system behaviour for the disk morphology by combining the trajectory data into PDFs, as seen in Figure 4.16. As would be expected by the preceding analysis, a clear divergence between agglomerating and bouncing particles is observed. Naturally, the separation distances spike towards zero more for the agglomerating particles than with rebounding ones. There is also much greater variability in the plot of the bouncing particles, as is to be expected.

The relative velocity PDFs also demonstrate behaviour that is in keeping with expectations. There is a natural spike in the PDFs at the initially prescribed value of

0.875  $mm s^{-1}$ , and then the agglomerating particles retain less velocity post-collision and ultimately approach an equilibrium state, biasing the PDFs towards low relative velocities. In contrast, the rebounding particles retain greater proportions of their initial kinetic energy – enough to escape the DLVO potential well. The variance in the agglomerating particles may have been expected to be lower, since in real physical systems, a pair of agglomerated particles would not retain any relative velocity, but the quirk of using a hard-sphere method with discrete energy exchanges, is that the velocity is able to increase for a short while, which allows a range of perhaps unphysical velocities to be sampled, thus showing up in the PDFs. A soft-sphere methodology (as presented in Chapter 6) circumvents this problem by allowing collisions, or energy exchanges, to take place more smoothly between particles, and therefore a more stable equilibrium state can be reached upon agglomeration.



Figure 4.16: PDFs of inter-particle separation (left) and relative particle velocity (right) compared for agglomerating and rebounding disks.



4.4.2 Monte-Carlo Simulation of Randomly Orientated Interacting Needles

Figure 4.17: Temporal evolution of inter-particle separation distance for agglomerating needles (left) and rebounding needles (right).



Figure 4.18: Temporal evolution of relative velocity between particles for agglomerating needles (left) and rebounding needles (right).



Figure 4.19: Temporal evolution of DLVO force magnitude between particles for agglomerating needles (left) and rebounding needles (right).

The results shown in Figures 4.17 to 4.19 demonstrate that the general behaviour of the particle interactions for the case of the needles is qualitatively similar to the disks. In much the same way, there is a clear divergence in the behaviour between agglomerating and bouncing particles.

However, there are a few key differences in the trajectory plots, compared with the disks. The first of which is that there are far fewer bouncing events for the needles. This appears to arise from the different ways in which the DLVO interaction magnitudes are distributed, which is the second key difference. Specifically, the bulk of the peak interaction magnitudes appear to fall in the range [2, 12] nN, which is twice as wide as the disks' range. Unlike the disks, the needles are in the vicinity of the maximal possible interaction strengths much more frequently, indicating that the probability of a long-edge-to-long-edge interaction is much more likely than the analogous face-to-face interaction in the disks. Furthermore, there are clearly not many interactions sampled where the minimal interaction forces are induced for the needles. This again indicates the importance of investigating the role of the surface geometry, which will be considered further in Sections 4.5 and 4.6.

One final important observation relates to the computational method rather than the physics. That is the fact that particles do not collide at exactly the halfway point with the same reliability seen with the disks. The slopes on the separation distance graph for around one third of the results reduce in gradient as these particle pairs approach collision, but no such reduction in the corresponding relative velocity is observed. So, this is not a physical effect but it is the accuracy of the collision detection algorithm (used to compute the closest approach vector) breaking down at very short separations, which introduces a degree of uncertainty to this separation – as reported by Girault et al. (2022). The computation remains stable and smooth and is only impactful at extremely small distances. However, this cannot be ignored since it allows the DLVO forces to act for slightly longer, which likely further promotes agglomeration for the needles.



Figure 4.20: PDFs of inter-particle separation (left) and relative particle velocity (right) compared for agglomerating and rebounding needles.

The overall system behaviour in terms of PDFs of the inter-particle separation distance and the relative particle velocity for the needles is shown in Figure 4.20, displaying similar trends to that of the disks. Simply, agglomerating particles spend longer close together with lower separation dx and with lower relative velocity dv. Whilst this seems self-evident for collisions in a stagnant flow, this approach to analysing the system will prove to be more useful when averaging out the turbulence effects in Chapter 5. There is not quite the same difference between agglomerating and bouncing particles in the relative velocity plots, compared with the disks, but the sample size being very low for the bouncing particles makes it difficult to draw any specific conclusions on this.





Figure 4.21: PDFs of inter-particle separation (left) and relative particle velocity (right) compared for disk interactions (solid line) and needle interactions (dashed line).

Contrasting the system behaviour for the two morphologies is possible by comparing the overall PDFs from the total collected data. In Figure 4.21, it can be observed that the needles have a greater propensity to agglomerate across the orientational parameter space for the present initial velocity. This is evidenced in the PDFs by the needles remaining closer together, with a lower relative velocity on average.

In the case of relative velocities, the rightmost PDF shows a sharp peak at the prescribed initial velocity but the interesting behaviour occurs on the left of the plot where the post-collision behaviour is accumulating. From the variance of this plot, it is possible to get a feel for the vdW force acting to alter the observed coefficient of restitution and how this is different across morphologies. There is a wider variation in the disks which can be partly attributed to fewer agglomeration events, allowing for a greater weighting of non-zero post-collision velocities in the sampling, since for agglomeration events there is an oversampling of very low relative velocities.



Figure 4.22: PDFs of the maximum achieved DLVO force magnitude in each simulation for disks (solid line) and needles (dashed line).

An interrogation of the vdW forces in Section 4.3 indicated that disks have the potential to resolve the highest attractive force at the point of collision, but the results of the preceding section indicate that the needles overall have much higher average values of the vdW force across orientations. To confirm this, only the maximum obtained vdW attractive force registered in each simulation was taken as an individual data point and PDFs were generated for each morphology based on this data. Shown in Figure 4.22, a very stark discrepancy is observed between the two shapes.

As described, the disks indeed facilitate the greatest possible vdW forces and, in support of this, a single datapoint is registered at around 20 nN, which is much greater than the rightmost registered value of the needles. However, the overall distribution for the disks has a much narrower character, indicating a lower variance despite a wider range. The bulk of the disk data is accumulating at the low end of the available vdW force range. Both of the distributions are right-skewed, but the needles are much less concentrated around lower vdW values, and vdW forces beyond 5 nN are resolved with much greater frequency.

Figure 4.22 is particularly insightful since these PDFs should be almost independent of velocity, short of considering lubrication and fluid effects. That means that the figure gives a good approximation to the distribution of forces experienced across the collisional parameter space for all velocities (within reason) and thus indicates the forces that should be expected between aspect ratio 5:1 spheroids irrespective of velocity.

To conclude, this analysis helps to explain why 73 disks agglomerated and 47 rebounded, at an agglomeration rate of 61%, whilst 114 needles agglomerated and just 6 rebounded, at a rate of 95%. The results are tabulated below.

Result	Needle (5:1)	Disk (5:1)
Agglomerated	114	73
Bounced	6	47
Agglomeration Rate	95%	61%

Table 4.3: The results of the randomly distributed orientation study.

### 4.5 Orientational Order in Particle Agglomeration

To this point, a picture is emerging that there is difference in the behaviour of disks and needles that arises from some interplay between the respective geometries and the orientations at the point of collision. To further understand this, PDFs have been generated based on the alignment of the particle's symmetry axes. A perfect alignment between symmetry axes facilitates a face-to-face interaction for the disks, which is the maximal energy configuration; however, this alignment also allows for a minimal configuration, which occurs for an edge-to-edge interaction. The same is true for the needles, where a collision at the end points occurs only if the particles are aligned, and a collision with the strongest force also occurs when the particles are aligned when the long edges collide. So, a priori it is not immediately obvious if there ought to be any order to this parameter, but Figure 4.23 demonstrates that there is.

#### 1.8 1.8 Probability Density Function Probability Density Function Agglomerated Agglomerated 1.6 1.6 Bounced Bounced 1.4 1.4 1.2 1.2 1.0 1.0 0.8 0.8 0.6 0.6 0.4 0.4 0.2 0.2 0.0 L 0.0 0.0 L 0.0 0.6 0.8 1.0 1.4 1.0 0.2 0.4 1.2 0.2 0.4 0.6 0.8 1.2 Relative Symmetry Axis Orientation (rads) Relative Symmetry Axis Orientation (rads)

#### 4.5.1 Alignment of Symmetry Axes

Figure 4.23: PDFs of relative symmetry axis orientations between particle pairs for disks (left) and needles (right), broken down into cases of agglomeration (solid line) and rebound (dashed line). Vertical line indicates  $\frac{\pi}{4}$  or half-way between perpendicular and parallel alignments.

It can be observed in Figure 4.23 that for both disks (left) and needles (right), there is a clear dependency on the relative orientation of the particle symmetry axes when it comes to determining whether a particle agglomerates. To highlight this, the data has been broken down into categories of agglomerating and rebounding particle pairs and plotted over  $\left[0, \frac{\pi}{2}\right]$ , which is the periodic region given the symmetry of the shapes. A further line has been drawn at  $\frac{\pi}{4}$  which highlights the half-way point between parallel and perpendicular configurations. The left half of the plot can be thought of as the more parallel alignments, with the right half the more perpendicular. For both graphs this roughly defines where there is a crossover from agglomeration being dominant, in the left half of the parameter space, corresponding to the more aligned configurations, to the right half of the domain where bouncing events dominate.

For disks, this is extremely pronounced: there is practically no representation of bouncing particles in the more parallel region of the alignment domain and a significant

1.4

overrepresentation, with a clear peak, in the perpendicular region. This indicates that the outcome of the collision, in terms the particle pair's agglomeration state, can be predicted reliably from a knowledge of the particle symmetry axis alignment, so long as the disks' alignment falls in the interval  $\sim [0.25, 0.65]$  radians.

The disk-like particles that agglomerated tended to be aligned; however, perpendicular alignments still have a significant representation for agglomerating particles. In perpendicular arrangements, the edge of one disk will collide close to the face of another. So, the representation of perpendicular arrangements in the agglomeration distribution is perhaps due to the significant strength of the contribution coming from a point that is located near to the centre of a disk's face, which appears to outweigh the reduced contribution from the corresponding edge location.

This touches on a subtle detail that comes from the design of the numerical experiment, which can be considered a limitation: by starting with particles at fixed locations in space, then randomly orientating them, and then directing them normally to one another, the particles will never experience an edge-edge collision in a perpendicular arrangement because the vector along which they must approach in this case prevents this. This is due to the particle centres being fixed in space at the point of injection. Hence, not all possible collisions are sampled under this methodology, but the results remain instructive nonetheless since such a wide area of the parameter space is considered. To describe this precisely, the full orientational parameter space is sampled, but the full *collisional* parameter space is not, which is all possible pairs of collision points.

In the next chapter, with particles placed in a turbulence field, the added degree of chaos will naturally impact the normal vectors along which the particles approach, and widen the collisional sample space studied, hence helping to move towards the goal of understanding binary particle collisions as fully as possible.

Returning to Figure 4.23, for needles, the picture is less divergent than for the disks, since both PDFs retain significant representation over the length of full parameter space. There is, again, a favouring of parallel configurations for agglomeration events. The PDFs intersect one another close to the half-way point, and then bouncing events are prevalent for the more perpendicular alignments. Practically all values of particle alignment are traversed by both the agglomeration and rebound distributions at a non-

zero value. This lack of disjointedness in the distributions comes from the fact that there are more degrees of freedom at play than just this one, indicating further analysis is required to fully understand non-spherical particle agglomeration.

Overall, parallel configurations appear to be favoured for agglomerating particles for both morphologies. This is in line with the results of Schiller et al. (2011) who studied attractive spheroidal particles without resolution of a fluid, concluding that attractive forces favour parallel alignment whilst repulsive forces favour perpendicular alignment. As alluded to, this phenomenon occurs partly because alignment of the symmetry axes facilitates the greatest surface interaction cases and hence greater attractive vdW forces in the present case.

Finally, an interesting effect is seen in the case of the disks at exactly parallel alignment, for which the rebound line slightly overtakes the agglomeration line. In the limit of a sufficiently large sample size, these two lines would perhaps meet here, since there is an equal number of ways to sample a perfect edge-edge collision in disk configurations as there is a perfect face-face collision. The reason this region of the graph is noteworthy, is that the peak of the agglomeration distribution therefore does not occur at the configuration that facilitates maximal forces but rather just tilted away from this.

In conclusion, the distributions reflect both the influence of attractive forces on the system, as well as subtle geometrical features of the unique morphologies, which influences the probabilities of certain arrangements.

### 4.5.2 Surface Collision Points

The initial orientations of the particles, and the orientations of the particles at the point of collision, only tell a part of the story. The points on the surfaces at which the collision takes place can also be monitored and analysed to further understand non-spherical particle agglomeration.

The following results use the parametric form of the ellipsoidal equation, introduced in Equation 3.2.3, wherein the unrotated ellipsoid is parameterised in terms of two angles,  $0 < \theta < \pi$  and  $-\pi < \phi < \pi$ , which define a unique location on the surface. For a spheroid whose symmetry axis is in the *x'* direction, varying  $\theta$  for constant  $\phi$  defines the *parallels* of the surface, whilst varying  $\phi$  and keeping  $\theta$  constant defines the

*meridians* of the surface; these are the *lines of curvature*. Traversing a parallel of the surface retains the same curvature, since these lines represent the circular symmetry about the so-called symmetry axis. Whereas traversing the meridians will lead to a variation in the local surface curvature, since they are of elliptical shape, with such a cross-section drawn in Figure 4.24.

In theory, agglomeration outcomes and interaction strengths ought to be a function of  $\phi$  alone when considering the collision points on the surface. The strength of this interaction is further mediated by the alignment of the symmetry axes, as demonstrated in Section 4.5.1, as well as the relative orientations about the vector of closest approach. Finally, as established throughout the preceding analysis, the velocity of the particles ultimately decides whether there is sufficient kinetic energy to escape the potential well, and therefore avoid agglomeration, given the configurational and collisional conditions.



Figure 4.24: Meridian of the surface, describing the angle  $\phi$ , for disks (red) and needles (blue).


Figure 4.25: PDFs of collision points in terms of the parametric angles  $\theta$  (left) and  $\phi$  (right) for disks.



Figure 4.26: Scatter plot of the  $(\theta, \phi)$  parameter space, for agglomerating disks (red) and bouncing disks (blue).

First analysing the disks, presented in Figure 4.25, the PDFs of collision points on the surfaces of the colliding particles are shown. For each interaction there are two pairs of datapoints collected, which is the  $\phi$  and  $\theta$  values on the respective surfaces, defining the points at which the collision took place. This means that each datapoint relies on another specific datapoint to determine its agglomeration state. A datapoint appearing in a region where attractive forces are largest can still not be an agglomerate, since its partnering datapoint might be in the lowest attractive force region, and vice versa.

The results displayed above agree with the prediction made in the previous paragraph; that is, there is no clear dependence on  $\theta$ , with the small deviation between lines probably being attributable to sample size. Whereas a very significant trend is

observed for  $\phi$ . In particular, the total collision points on the surface of the disk are collecting at around  $\phi = \pm \frac{\pi}{2}$ , which is the location of the edge points of the disks. At these locations, the DLVO forces induced are minimal since the surface interaction is minimal. As expected, then, at these locations there is a clear predominance of bouncing events, rather than agglomeration events. Conversely, near to the face of the disk, at around  $\phi = 0$ , where DLVO forces being induced are maximal, there is a predominance of agglomeration events.

Figure 4.26 shows the full parameter space on a single plot and further visualises that which is described in the previous paragraph. Across  $\theta$ , the collision points appear to be roughly uniform with no predominance of red or blue datapoints. Whereas there is a clear stratification seen in  $\phi$ , in that all points accumulate around the edge locations, whilst at the very edge of the shape, bouncing events accumulate.



Figure 4.27: PDFs of collision points in terms of the parametric angles  $\theta$  (left) and  $\phi$  (right) for needles.



Figure 4.28: Scatter plot of the  $(\theta, \phi)$  parameter space, for agglomerating needles (red) and bouncing needles (blue).

The breakdown for needles, seen in Figures 4.27 and 4.28, does not carry the same insight for agglomeration as it did with the disks, on account of the fact there is not a great enough sample size of bouncing events to make any meaningful conclusions. The agglomeration distribution very closely follows the total distribution, as would be expected given the weighting of agglomeration events in the overall collision sample; however, a pattern is perhaps emerging where there is an overrepresentation of agglomeration at  $\phi = \pm \frac{\pi}{2}$ , and a very slight underrepresentation at  $\phi = 0$ , which is the opposite of the disks. Further investigation is required to confirm this.

Plotting the 6 bouncing events was avoided given the limited sample size and the lack of statistical confidence that can therefore be placed in the trends shown.

Turning the focus to the total distributions of disk collisions and needle collisions, there is a clear difference in the locations on the surfaces that ultimately experience a collision, irrespective of DLVO forces. The probability of colliding at the various locations on the particle surfaces is different based on morphology, even when uniformly randomly sampling the orientations, as is the case here. This raises the question of *why*.

One hypothesis can be constructed by considering particle shape. Disks are shown to be much more likely to collide in the vicinity of the circular band that defines the 'edge' of the shape, and needles are much more likely to collide along their long edges. Under a stretching transformation, the spheroids remain topologically equivalent, and the collision points accumulate in similar regions of the  $(\theta, \phi)$  parameter space. However, the physical mechanisms driving these accumulations differ due to the anisotropic geometry and curvature distributions of the spheroid. Specifically, collisions peak about the circle, or the central parallel line of curvature, that lies normal to the x' axis at x' = 0 for both spheroid types. In a disk morphology, this circle manifests itself as the so-called 'circular band' defining its outermost meridian – and in needles it is simply the central circle perpendicular to the so-called 'long-edge', where DLVO forces happen to be maximal.

Relating this back to Section 4.4.2, there were not many events where minimal forces were induced for the needles, nor maximal forces induced for the disks, which led to results contrary to initial expectation. A good explanation for this finding, therefore, is linked to this effect. Since there are only two points on a needle where the forces are minimised, the end-points, or the poles, it stands to reason that these points are not often frequented in collisions over the entire orientational parameter space. This runs in contrast to disks, where the minimal forces are induced on the circular band. Analogously, the maximal forces on the disks are only induced at two points: the direct centres of the faces, and the least frequented locations. This is in contrast to needles, where the maximal forces exist on a band around the centre of the shape, the most frequented locations.

This is an adequate first explanation, however it does not fully capture why disks do not still frequently induce large vdW forces in the regions close to those central points on the face: nearby to those points, the surface curvature is still low and thus the induced forces are high.

An explanation is hypothesised here, and the results that follow give credence to the idea. Since two particles are involved in the interaction, both contribute surface collision points to the interaction. The Gaussian curvature model involved combining the radii observed at said collision points to give an equal and opposite strength of force experienced by both particles, based on a weighted contribution of either surface curvature. The way the mathematics is set up there biases the scaling towards the smaller radius involved, which disproportionately reduces the interaction strength in general. If the statistical sampling of surface collisions points is such that minimal

locations are often frequented and maximal locations are rarely frequented, then the probability of achieving two near-maximal surface locations is very low, and in fact the total orientational parameter space that facilitates such maximal interactions is clearly low for the disks: small perturbations away from a maximal configuration quickly presents the edge of one disk into the interaction, which thusly scales the interaction magnitude down significantly. There is actually a very mathematically rich problem here that looks to investigate the emergent probability distribution of colliding geometrical objects, irrespective of the present application.



Figure 4.29: PDFs of the collision points  $\theta$  (left) and  $\phi$  (right), for all disks (solid line) and all needles (dashed line).

The described effects can be considered in greater depth by directly comparing the overall distributions of  $\theta$  and  $\phi$  for the two morphologies on a single plot, as is shown in Figure 4.29. Once again, there is no order to  $\theta$  and the information about spheroidal collisions in terms of the surface points can be described fully by  $\phi$ . The overall qualitative nature of the graphs is very similar: both distributions have two local maxima located at  $\pm \frac{\pi}{2}$  and a local minimum at zero, with the distributions falling away at  $\pm \pi$ . As has been illustrated by the agglomeration findings, it is where these distributions differ which greatly impacts agglomeration. Specifically, the needles vastly outweigh the disks at  $\phi = \pm \frac{\pi}{2}$  and the reverse is true at  $\phi = 0$  and  $\phi = \pm \pi$ . This confirms that needles are sampling their maximally attractive configurations much more frequently than the disks are equivalently sampling theirs.

Lastly, the distribution of the disks is more uniform than the needles, arising from the local minima being less extreme. This indicates that the frequency of experiencing the

respective lowest probability configurations is greater for the disks than for the needles. In other words, it is found to be more likely to collide with the centre of a disk, than with the end point of a needle.

# 4.6 Effect of Surface Curvature Distribution on Agglomeration

The surface area of a disk is greater than a needle for a fixed volume, and so the results to this point run contrary to what might be an initial expectation. One might expect forces that depend on surface interaction to be, in general, more prominent across shapes of greater surface area.

Beyond this, the *total curvature* of the surface – defined as the integral of the Gaussian curvature over the entire surface – is constant for spheroids. The geometric properties important in governing the interaction behaviour of spheroids is therefore encoded in the way that this curvature is distributed across the surface. Recall, crucially, that the strength of the interaction is scaled by the local curvature at the points of closest approach in the present model.

Observing Figures 4.30 and 4.31, it is clear that the curvature accumulates at the extremities of the shape. For disks, this means that most of the curvature of the shape is located on the circular band defining its boundary; whilst the curvature accumulates at the end points of the needles, and thus so do the minimal values of the radius of Gaussian curvature. (The radius of Gaussian curvature crudely approximates the magnitude of the surface interaction coming from a given particle in the present agglomeration model.) The curvature is seen to be comparatively homogeneous across the rest of surface, compared to the extremes. This means that the penalty incurred, in terms of the resolved force, for colliding at a location near to these extremities is stark.

Moreover, in the limit of increasing aspect ratio, the extent to which the curvature gets concentrated at the extremes increases, and the curvature across the rest of the body decreases, to conserve total curvature. This has interesting implications for very high aspect ratios. In the current literature, aspect ratios of 5:1 are high for spheroidal investigations, particularly for immersed boundary studies that must rely on a certain number of fluid points to be resolved in every direction within the particle mesh, and

thus scale badly with increasing aspect ratio from a computational perspective. Related to this topic, for discrete shapes like non-spherical particle meshes, the curvature is concentrated at the vertices, since the actual triangular faces of the mesh possess zero curvature. This means that mesh-based approaches for resolving forces that depend on local curvature, have a limitation placed on them by the resolution of the mesh, so too do they face a challenge in how to calculate and implement this local curvature. The continuous model used herein is therefore well suited to agglomeration modelling.



Figure 4.30: Normalised Gaussian curvature shown as a colour map for needles (left) and disks (right).



Figure 4.31: Normalised radius of Gaussian curvature shown as a colour map for needles (left) and disks (right).

The visualisations of the surface curvature bring more cognisance to the results obtained since disks were found to collide frequently near to the visualised penalty region, whilst needles avoided theirs generally. The upshot of this is for needles to agglomerate at a higher rate than disks: 95% compared to 61%. As mentioned, there is an obvious difference in the probability of colliding with a band rather than a point, on account of the former occupying much more of the domain than the latter. But there is another geometric effect which compounds this. This is the matter of which points on the surface protrude the most, which relates to both the morphology and the respective curvature of the surface.

To illustrate this, consider just one rotational degree of freedom,  $0 < \eta_1 < \pi/2$ , shown in Figure 4.32, which will be the amount of relative tilt towards one another for two spheroids that are initially perfectly aligned in their maximally attractive configuration, i.e., face-to-face for the disks or long-edge-to-long-edge for the needles. As before, the spheroids approach one another along their vector of closest approach. For no 'tilt', or  $\eta_1 = 0$ , the spheroids will collide at the centre of the faces (disks) and at the centre of the long edges (needles). Tilting the particles towards one another, there will come a point where the edge overtakes the front of the shape and becomes the colliding point, or the point closest to the other particle. The amount of so-called 'tilt' required for the edge, or minimal force locations, to become the point at which a collision occurs, is significantly less for the disks than for the needles, owing to the fact that the stomach of the needle protrudes further than that of the disk. In this case a greater proportion of the  $\eta_1$  parameter space is thus occupied by minimal attractive interactions for disks, than analogously for needles. This is a simplification in one rotational degree of freedom aiming to hypothesise as to why the observed distributions differ as a result of geometry.

This is further compounded. For both shapes that have undergone some rotation  $\eta_1$ , there is a second rotational degree of freedom which is symmetric and will not affect the forces (rotation about the x' axis) but there is another rotation  $\eta_2$  that still carries significance. Performing this  $\eta_2$  rotation for the needles will always immediately move the end-point, or minimal force location, away from the other particle, allowing the end-points to avoid another during approach, resulting in a collision location that thus increases the attractive force; whereas, for disks there is no such guarantee. This further favours the agglomeration propensity of needles. Overall, the way the same

volume is distributed between the two shapes, clearly leads to overall differences in the points of collision under uniform sampling.



Figure 4.31: Diagram to illustrate perturbation of orientational configuration for aspect ratio 5:1 disks.

Increasing the aspect ratio extremises these effects by concentrating the curvature at the extremities of the shape and reducing the protrusion of the shapes in the direction of the smallest radii. However, there is no reason to assume that each of these effects behaves linearly with respect to the others, when changing the aspect ratio. It would seem there is rich potential for mathematical (both geometric and statistical) investigation in this direction and this is perhaps even necessary to fully understand the collisional and attractive/repulsive interaction behaviours of spheroidal particle pairs – as the results are clearly strongly sensitive to the particles' changing morphologies.

### 4.7 Conclusions

The spherical and non-spherical elements of the particle model have been investigated to demonstrate robustness. All elements of the model have been shown to be numerically stable, with the only fault being in the determination of extremely small separations between needles, in a limited number of cases. A velocity study was first undertaken for spherical particles, to understand DLVO force interaction in detail under the present framework, giving insight into the precise value at which particles no longer agglomerate, labelled the 'cut-off velocity'. The velocity parameter space investigated was chosen to match collision velocities in a related channel flow. The cut-off velocity was shown to be located at around 0.7  $mm s^{-1}$ , which is in keeping with Mortimer et al. (2020), providing extra validity to the choice of value in the LPT simulations undertaken in the nuclear research group at the University of Leeds.

After this, the focus turned to investigating how the cut-off velocity changes for nonspherical particles of two distinct morphologies. Namely, disks and needles of a 5:1 aspect ratio. This same study was undertaken for two modal configurations of the respective morphologies, with interesting variation in the results. Face-to-face disks were found to continue to agglomerate at the highest velocities, with this configuration facilitating maximal vdW force magnitudes of triple those observed in the maximal needle configuration. Collisions at the end points of needles were found to induce the lowest attractive forces in the configurations considered, and thus particle rebound occurred at the lowest velocity for this case.

Following this, orientations were varied uniformly such that random collisions could take place, generating insight into the role of morphology and orientation on agglomeration outcomes under the present model, at a fixed velocity. The study indicated that needles have a significantly higher propensity to agglomerate, and both morphologies tend to show an alignment of their symmetry axes when undergoing agglomeration. A further important finding relates to the observed coefficient of restitution, where it was seen that this outcome is strongly dependent upon the orientation of the shape at the point of collision.

The findings were further understood in terms of the geometry of the shapes. It was demonstrated that the curvature of the particles accumulates at the extremities of the surface, which creates regions where the resolved attractive forces are very different to the overall shape average. Pairing this understanding with an investigation into the probability with which certain locations on the surface experience a collision, these results generate new understanding relating to binary spheroidal particle collisions under uniformly randomised initial conditions. Namely, disks are found to be much

more likely to collide at locations close to their edge, which coincides with the greatest curvature of the shape and thus the lowest attractive forces. Conversely, needles are found to collide at locations away from the accumulation of surface curvature, thus facilitating stronger interactions on average. This ultimately leads to needles favouring agglomeration more than disks, despite the greater surface area of disks and the greater maximal van der Waals forces obtainable in theory. In fact, the van der Waals force distribution highlighted this point most clearly, where needles were shown, contrary to expectation, to resolve much higher average forces.

The preceding analysis is limited in two ways. Firstly, the analysis only considers spheroids of a 5: 1 aspect ratio. The protrusion of the shape was theorised to place a bound on the extent to which it must be tilted towards its counterpart to result in an edge-based collision. In the limit of increasing aspect ratio, this so-called protrusion reduces, but the distribution of surface curvature is also changed as a function of aspect ratio. Therefore, there are uncharacterised relations there rich for mathematical investigation, and, by extension, the results herein may not translate to significantly greater aspect ratios.

Secondly, the restriction for particles to approach along their closest approach vectors allowed for a simple means to ensure the initial distance, and simulation times, remained consistent. However, this aspect of the numerical experiment places a slight restriction on which locations on the surface can experience a collision on a second particle, given some fixed orientation of the first, which will be alleviated by introducing turbulence on approach.

Overall, agglomeration has been shown to be dependent upon: particle velocity, relative orientation of the particles' symmetry axes, relative orientation of the particles about the closest approach vector, the location of surface collision points, particle shape (or, equivalently, the distribution of surface curvature), and it is predicted that aspect ratio is also crucial. A further analysis could look at this from a mathematical perspective to understand the relationship between curvature, aspect ratio, and the relation to collision points.

Translating these findings to applicable physical scenarios allows the formulation of a few potential ideas. For disks and needles with the present geometries, the results suggest that systems of needles should aggregate faster than disks, and that the

resulting aggregates would be more stable with respect to potential shear-induced disaggregation or breakage since their components are associated with stronger interactions on average. Agglomeration, and the resultant sedimentation or deposition, assuming the fall velocity of the formed aggregates is sufficient for net sedimentation, will be faster for needles. This suggests that shape in aggregating systems has a first order effect on sedimentation rates. However, there are further effects to investigate before having conviction in this statement, since multiple interacting particles can complicate things considerably. At the very least there is the strong indication that the formation of doublet agglomerates is considerably more likely for a system of 5:1 aspect ratio needles than the analogous disks.

Disks were seen to agglomerate over less of the total orientational parameter space, which may imply a restriction of the types of structures disks can realistically form, if they are reliant on a certain configuration to agglomerate. Relating this back to the observation about likelihoods – there are only two locations on a disk that facilitate the maximal configuration. Once these are occupied by other particles already agglomerated, the regions of the disks that are still exposed and thus candidate locations for another particle to join, are also the regions of the shape that do not facilitate agglomerates to form only in long chains with the faces adhered. Such a chain would become increasingly unstable in a turbulent flow as it lengthened.

## 5 Non-Spherical Particle Interactions in Forced Homogeneous and Isotropic Turbulence

## 5.1 Introduction

In the previous chapter, steps were taken to understand the way that the fundamental properties of the particles and the fluid affect collisions, interaction behaviours and agglomeration propensities. From this base, there is now a general understanding of the particle-phase behaviour in relative isolation, from which it is possible to approach the more dynamic problem of *turbulent* particle agglomeration.

Industrial processes that are impacted by agglomeration will very often be turbulent, for example in pipe-based transportation of particulate matter (Ho and Sommerfeld, 2002; Reeks, 2014; Bin et al., 2018). As motivated earlier, the particles observed in industry will almost always be non-spherical in nature. Current methodologies and studies in the literature rarely address these two important points, and there are no existing studies that simultaneously address both non-spherical particle agglomeration and turbulent particle advection, as demonstrated in the literature review in Chapter 2. Further still, any such endeavours into particle agglomeration do not utilise a fully resolved DNS-IBM approach, as is used here. These studies may therefore miss the particle-scale dynamics and interactions that can influence the bulk behaviour of the system. Hence, this chapter presents a state-of-the-art simulation approach aiming to obtain insights at a deeper, more fundamental, level than previous work.

A theme throughout the thesis relates to DLVO forces operating on very short length and time scales – smaller than those of the fluid – and the fact that this creates a challenge with respect to their modelling. To resolve both the fine length scales required for a DNS of turbulence and the short time scales required to deterministically resolve DLVO-induced agglomeration events, (shown in the previous chapter to be  $dt \sim O(10^{-7}) s$  for non-spherical shapes), there is a heavy computational cost incurred both spatially and temporally for the respective elements of the simulation. This partly motivates the following studies being restricted to periodic boxes of homogeneous and isotropic turbulence (HIT), in which the particle meshes can be made sufficiently large compared to the computational fluid domain, such that the restrictions on the immersed boundary technique can be met without requiring unfeasible numbers of grid points, or adaptive meshing, as would be the case in an industrial scale pipe or channel geometry. In the work by Chouippe and Uhlmann (2015) for example, 16 nodes were required across a particle diameter to accurately simulate immersed spherical boundaries. In the present work, 15 nodes are used.

A second simplification is employed: *pairs* of particles will be studied at first to reduce computational overhead, as opposed to the millions of particles present in realistic industrial systems. As alluded to, for a realistic size pipe or channel geometry, the particles are so small in relation to the full domain that the number of mesh points would become so large as to make the present study unfeasible, whilst retaining the DLVO and IBM requirements on the constituent particles. However, as computational resources continue to improve, for example with the accelerating field of GPU processing (which is particularly well suited to particle methods), these requirements will be more readily met. So, there is great value in developing these approaches now, such that progress further down the line is faster. Naturally, the large particle numbers of real systems give rise to many important effects due to multiple simultaneous forces acting at once from the nearby particles. Hence, it is of value to develop this method further to accommodate many interacting particles, which will be demonstrated in Chapter 6.

The key benefit of these initial simplifications is the opportunity to study in detail the process of collision and agglomeration from a fundamental perspective. An isolated interaction between two particles is of great interest since the models most frequently employed in the literature treat particle interactions as distinct, often relying on algorithms that visit each particle pair consecutively rather than simultaneously. Therefore, a deeper understanding of individual pair interactions can improve those models. Lastly, this problem set up allows an examination of the influence that turbulence has over the collision and agglomeration process. By resolving all flow scales, it is possible to examine how the finest flow structures can affect the close-range interaction as it happens, whilst also capturing the particle-scale turbulent

structures' effect on the larger scale particle motions. The boxes of HIT allow a range of flow scales to be resolved, and so a wide and chaotic parameter space is explored.

Employed throughout this chapter and the next, there is an attempt made to match the system properties to channel flow LPT-DNS calculations performed at the University of Leeds (Mortimer et al., 2020). The idea, rather than picking the turbulence and particle-phase properties arbitrarily, was to choose those most consistent with the different flow regions as seen when traversing the width of a channel, to get a complete picture of the background fields two agglomerating particles may be subject to when immersed in a canonical wall-bounded flow. The flows chosen in the cited work were representative of those encountered in nuclear waste processing flows, and so is the present work.

This matching introduces some cognisance to the studies by imparting representative fluid forces and turbulent flow structures on the particles, but it is important to note what will be colloquially referred to as the 'viscous sublayer', 'buffer' and 'bulk flow' regions are just approximations based upon the mean statistics calculated in those regions and cannot capture some of the more intricate fluid physics that take place, due to the homogeneity and isotropy restrictions. For example, larger scale coherent motions cannot form in boxes of HIT, nor is there a mean flow direction or geometric wall effects, and these two things are very important to the overall particle dynamics in industrial scale flows.

The approximation aims to simulate particle collision processes in the reference frame of a collision taking place in a typical nuclear waste flow. Some of the effects of the mean flow on the particles are therefore included in the initial relative velocities as they collide, as measured in supplementary channel flow calculations. This reduces the need to consider the mean flow, leaving the fluctuating component of the velocity field to be approximated by the HIT set-up. In this frame of reference, the turbulence is matched in terms of the Taylor-Reynolds number and the root mean squared velocity fluctuations to create the types of conditions experienced by particles colliding and agglomerating in industrially relevant scenarios.

## 5.2 Validation of Single-Phase Turbulence

As described in Section 3.1.4, the forcing scheme of Eswaran and Pope (1988) was chosen to force and sustain the turbulence in the fluid. The choice of forcing function was made due to two main factors. Firstly, since the turbulence will not be ultimately used for solely a single-phase domain, but rather with particles for a multiphase study, it was important to choose a method that did not interfere with the particle-scale forcing, and vice versa. This is achieved through forcing the larger scale modes in the flow, i.e., the smallest wavenumbers, as introduced earlier. Described by Mallouppas et al. (2013), such an approach ensures the nonlinear transfer function, as well as the redistribution of energy, remain unaffected by the interaction between the forcing of the fluid turbulence and of the particle-phase. Then, for numerical studies it becomes easier to isolate the effect of the particles on the fluid-phase.

Secondly, the method of Eswaran and Pope (1988) was selected in part due to the arguments given in Chouippe and Uhlmann (2015). The latter work showed the method is applicable to immersed boundary particle-laden HIT. The specific arguments informing their choice of method related to the following four criteria: small-scale statistics should be repeatable and agree with previous studies; long-time integration should remain stable in the presence of two-way coupled particles; the forcing should be efficient to compute, even within a non-pseudospectral framework; the scheme should allow for *a priori* estimation of the resolved turbulence based on the chosen forcing parameters. Each of these criteria were demonstrated within the reference material, giving further confidence to the selection of methodology.



Figure 5.1: Pseudocolour velocity magnitude plots demonstrating the onset of turbulence, separated by intervals of 0.5 normalised time units.

Figure 5.1 demonstrates the successful implementation of the forcing scheme achieved in the present work. The figure helps to show how the development of turbulence within the domain is achieved over time through the action of the forcing function. This figure and subsequent ones make use of the large eddy turnover timescale  $T_e = {u'}^2/\epsilon_F$  for normalisation of the time units. Values for the forcing parameters used here are  $T_L = 1.67389$  and  $\epsilon^* = 0.00298$  in line with case I of the validation presented shortly hereafter. This corresponds to a resolved Taylor-Reynolds number of  $Re_{\lambda} = 65$ .

From a quiescent initial condition shown in the first plot, a forcing function is applied to the incompressible Navier-Stokes momentum equations. The result is that after 0.5 time units, as shown in the second plot, a non-zero velocity field emerges.

The largest scales, corresponding to the smallest wavenumbers, are exclusively targeted by the forcing function – which is visible in the second subplot, where the scales are similar in size to that of the domain. It can be seen in the subsequent three subplots that these larger scale structures break down into smaller eddies, in line with

Kolmogorov's theory of turbulence (Chapter 2.1.2), at which point the full turbulence energy spectrum has emerged, and this remains stable throughout the simulation. With a periodic box of homogeneous and isotropic turbulence successfully sustained, particles can be added to the domain with the knowledge that the fluid effects experienced over a simulation will be on average the same due to the ergodicity of the stochastic forcing field. However, within each individual simulation considered below, the intricacies of motion resulting from the chaotic turbulence field are explored.

#### 5.2.1 Calculating Statistical Quantities

Two validation cases were conducted to demonstrate the correct implementation of the forcing function in the code. The natural choice for validation was the work of Chouippe and Uhlmann (2015), rather than Eswaran and Pope (1988), because of the more comparable computational power employed presently as well as the availability of validation data for two different Taylor-Reynolds number cases.

As described earlier, the Taylor-Reynolds number characterises the turbulence levels experienced within a region of turbulent flow. It is important to demonstrate that the correct Taylor-Reynolds number is met, as this is the chief number characterising the behaviour of the system (Equation (2.2)). As with a typical Reynolds number, it combines the characteristic velocity, length and kinematic viscosity of the fluid system. In this case, the characteristic velocity is taken to be the root mean square of the velocity fluctuation u' and the length scale is the Taylor microscale  $\lambda$  (Section 2.1.10).

The Taylor microscale is estimated from the local fluctuating strain rate field, which can then be spatially averaged over the simulation domain to give the global value at each timestep. The local velocity fluctuations are averaged in the same way, whilst the viscosity is constant, giving that which is required for the calculation. This is tracked temporally over the simulation and time-averaged to demonstrate convergence.

The relation between the fluctuating strain-rate field and the Taylor microscale  $\lambda$  is given by Equation (5.1) (Hinze, 1975), expressed in one dimension without loss of generality since the flow is isotropic. In practice, all three spatial directions are used with equal weighting in the averaging to improve the sample size, after ensuring isotropy. It is possible to use the same field to estimate the dissipation in the system and this relation is given in Equation (5.2):

$$\overline{\left(\frac{\partial u'}{\partial x}\right)^2} = \frac{\overline{u'^2}}{\lambda^2},$$
(5.1)

$$\epsilon_F = 2\nu_F \overline{S_{\iota J} S_{\iota J}},\tag{5.2}$$

where  $\epsilon_F$  is the turbulent fluid kinetic energy dissipation rate,  $\nu_F$  is the fluid kinematic viscosity,  $S_{ij}$  is the strain-rate tensor, and an overbar is used to represent a spatially averaged quantity. Similarly, the turbulence kinetic energy k of the fluid is calculated via:

$$k = \frac{1}{2}\overline{u_i'u_i'} = \frac{3}{2}\overline{u'^2},$$
(5.3)

where  $u'_i$  are the local velocity fluctuations in the three Cartesian directions, and u' is the root mean squared velocity fluctuation defined earlier, with the simplification again owing to isotropy.

These variables are explicitly computed at each timestep in the simulation and outputted as data, whereas the rest of the turbulence quantities computed in the validation are derived from this live data set, using the well-known relations of Kolmogorov (Chapter 2.1.2). Eswaran and Pope (1998) describe that such estimates are most reliable when the Reynolds number is high enough for the small scales to decouple from the larger scales.

The forcing function was described in Section 3.1.4 in terms of the general ideas of the method as well as the implementation. It is possible to describe also some of the key properties of the method, starting with the complex vector-valued random processes  $\boldsymbol{b}$  that drive the time evolution of the forcing field.

Recall the method takes place in wavenumber space with the vector  $\boldsymbol{\kappa} = (\kappa_1, \kappa_2, \kappa_3)$  describing the pertaining wavenumbers. These wavenumbers only receive contributions of non-zero forcing in a limited band. In particular, the lowest wavenumber band corresponding to the largest scale motions in physical space, where the Navier-Stokes equations are solved. In the present simulations, the condition  $0 < |\boldsymbol{\kappa}| \le \kappa_{cut}$  is enforced which ensures that only wavenumbers within a

spherical shell are excited, excluding the origin. So,  $\kappa_{cut}$  is the maximum wavenumber of the forced modes.

As for the properties of the process: firstly, each random process exhibits a mean of zero in the limit  $\Delta t \rightarrow 0$ . The implication for the forced velocity field is that it also exhibits zero mean, as is required by HIT. Secondly, each realisation of the forcing field is correlated in time since the stochastic process itself is correlated. These properties are laid out in Equations (5.4) and (5.5), respectively:

$$\langle \boldsymbol{b}(\boldsymbol{\kappa},t)\rangle = 0, \tag{5.4}$$

$$\langle b_i(\boldsymbol{\kappa}, t) b_j(\boldsymbol{\kappa}, t+s) \rangle = 2\sigma^2 \delta_{ij} \exp(-s/T_L), \tag{5.5}$$

where  $\sigma^2$  is the variance of the random process,  $\delta_{ij}$  is the Kronecker delta,  $T_L$  is the forcing timescale, *s* represents a dummy variable arbitrarily shifting the temporal variable, and angular brackets indicate temporal averaging.

In practice, the variance is not prescribed directly and is instead calculated from the relation  $\epsilon^* = \sigma^2 T_L$ , where the non-dimensional quantity  $\epsilon^*$  is introduced and referred to in the original text as the non-dimensional dissipation rate. In the limit that  $T_L \rightarrow 0$ , the simulation becomes white noise (Eswaran and Pope, 1988) and the mean energy input trivially tends to zero (Chouippe and Uhlmann, 2015). As such,  $T_L$  and  $\epsilon^*$  must be chosen independently.

In the work of Chouippe and Uhlmann (2015), it is reported that since the number of forced wavenumber modes is small (less than 100) the cost associated with a computation of the turbulence forcing is small in comparison to the Navier-Stokes solver step. The same choice was made in the present work in terms of the number of forced modes, and so a similarly small relative cost was observed. In the present work, the random processes are updated once globally per time step and passed to the separate processors to be evaluated locally on the physical grid.

#### 5.2.2 A Priori Estimation of Forcing Parameters

To pre-calculate the quantities that should be expected out of the simulation, given a set of input forcing parameters, there were two relations available: one from each of the two main sources followed.

Eswaran and Pope (1988) define new non-dimensional parameters which are constructed from their forcing input parameters  $(T_L, \epsilon^*, \kappa_{cut})$ . These are then used to generate an estimate of the dissipation rate, Kolmogorov length scale and the Taylor-Reynolds number. They introduce the non-dimensional forcing timescale as  $T_L^* = T_L(\epsilon^*)^{1/3}\kappa_0^{2/3}$ , with  $\kappa_0$  being the lowest forced wavenumber, which can be calculated by  $\kappa_0 = \frac{2\pi}{L}$  for domain length *L*. Then, their estimate of the non-dimensional dissipation rate is  $\epsilon_T^* = 4\epsilon^* N_f / (1 + T_L^* N_f^{\frac{1}{3}} / \beta)$ , with  $N_f$  the total number of forced wavenumbers. ( $\beta$  is a constant set to 0.8 in the reference material, and that was used here.) The estimate of the Kolmogorov length scale is given by  $\eta_T = \left(\frac{\gamma_F^3}{\epsilon_T^*}\right)^{1/4}$ . From this base, they use these parameters and their data set to generate an empirical relation for estimating the Taylor-Reynolds number as:

$$Re_{\lambda}^{T} = 8.5 / ((\eta_{T}\kappa_{0})^{5/6}N_{f}^{2/9}).$$
(5.6)

Chouippe and Uhlmann (2015) reported an overprediction by this quantity and instead proposed an alternative. They introduced the central wavenumber  $\kappa_c = (\kappa_0 + \kappa_f)/2$ , and defined a macroscopic length scale as  $L_c = 2\pi/\kappa_c$ . From this, they give their own estimate as:

$$Re_{\lambda}^{T} = \frac{20L_{c}(T_{L}\epsilon_{T}^{*})^{\frac{1}{2}}}{(3\nu_{F}))^{\frac{1}{2}}},$$
(5.7)

and consequently report an improved discrepancy (3% versus 8%) between prediction and outcome, when compared to the former paper. In the present work, the latter estimation method was also seen to be more accurate.

#### 5.2.3 Validation Case I: $Re_{\lambda} = 65$

Presented in Table 5.1 are the validation parameters used, which were derived from the non-dimensional values given in the reference paper, since they were not stated explicitly. So, the choice of  $T_L$  is somewhat arbitrary. It is the relation between the two forcing parameters that is important, and this is consistent between studies. The parameters include  $\mu$  the dynamic viscosity  $(Pa \cdot s)$ ,  $\rho$  the fluid density  $(kg \cdot m^{-3})$ ,  $\kappa_{cut}$ the cut-off wavenumber  $(m^{-1})$  normalised by the smallest wavenumber  $\kappa_0$   $(m^{-1})$ ,  $N_F$ the number of forced modes in the simulation, the non-dimensional forcing parameters, i.e.  $T_L$  the forcing timescale and  $\epsilon^*$  the forcing 'dissipation', L the box length (m), and *nelx* the number of seventh-order spectral elements in the spatial directions.

Case	μ	ρ	$\kappa_{cut}/\kappa_0$	N <sub>F</sub>	$T_L$	$oldsymbol{\epsilon}^*$	L	nelx
I	$6 \times 10^{-3}$	1.2	2.3	56	1.67389	0.00298	2π	24
11	$6 \times 10^{-3}$	1.2	2.5	80	1.67389	0.28234	2π	30



Figure 5.2: Pseudocolour non-dimensional velocity magnitude contours on the exterior surfaces shown for the turbulence box with  $Re_{\lambda} = 65$ .



Figure 5.3: Temporal convergence of the spatially averaged turbulence quantities for validation case I, Taylor-Reynolds number (top left), Taylor microscale (top right) and large eddy length scale (bottom). The black horizontal lines show the time-averaged values reported by Chouippe and Uhlmann (2015).

Figure 5.2 shows a snapshot of the  $Re_{\lambda} = 65$  turbulence velocity field, corresponding to the graphs in Figure 5.3 which display the evolution of the measured turbulence properties. There is an initial transient where the Taylor-Reynolds number rapidly increases to a value of 100, at which point it begins to decrease again towards a more stable value approximately equal to the desired value of 65. The mechanism for the peaking of the value is the consistent addition of forcing to the Navier-Stokes momentum equation through the source term. The eventual peaking and reduction back to the target value occurs through the development of the turbulence energy cascade, a process which is seen to occur between one and two normalised time units. It is worth noting that even though the forcing remains constant, thus increasing the energy in the domain, the evolution of the Taylor-Reynolds number does not increase indefinitely, due to the change and development of flow scales over this period through

the energy cascade. This is reflected by the changing length scales demonstrated in the plots of  $\lambda$  and l prior to the system reaching equilibrium. A key mechanism governing the overall evolution is the turbulence energy dissipation taking place at the smallest scales, and in the present set-up these smallest scales must first form for the dissipation to begin to take place at this scale. Overall, the system reaches equilibrium when the turbulence energy dissipation rate balances with the injected kinetic energy.

It can be seen that the fluctuations about the requisite value are significant, roughly  $\pm 10$  units; however, the sufficiently time-averaged value remains stable indefinitely as the simulation continues, and the magnitude of the fluctuations about the mean depends only on the stochastic process rather than the current velocity field.

#### 5.2.4 Validation Case II: $Re_{\lambda} = 143$

For the  $Re_{\lambda} = 143$  validation case, similar trends are seen in each of the plots of Figure 5.4. The turbulence energy cascade takes roughly the same amount of time to develop, but the peak of the Taylor-Reynolds number is four times higher. Proportionally, the Taylor-Reynolds number fluctuates about its mean value to a similar degree; however, in absolute terms, the fluctuations about the mean value are approximately  $\pm 20$  units.

The validation cases demonstrate a range of integral length scales can be resolved within the simulation domain, which places it as an improvement upon the linear forcing method of Lundgren (2003) which is known to be restricted by integral length scale considerations (Lucci et al., 2010).



Figure 5.4: Temporal convergence of the spatially averaged turbulence quantities for validation case II, Taylor-Reynolds number (top left), Taylor microscale (top right) and large eddy length scale (bottom). The black horizontal lines show the time-averaged values reported by Chouippe and Uhlmann (2015).

#### 5.2.5 Validation Case Comparison

Presented in Table 5.2 are the time-averaged results, obtained by averaging over the latter half of the simulation time. Following the reference material for this comparison, the parameters are named and calculated as follows: the large-eddy length scale  $l = k^{3/2} \epsilon_F$ ; the Taylor microscale  $\lambda = \left(\frac{15\nu_F u'^2}{\epsilon}\right)^2$ ; the Taylor-Reynolds number  $Re_{\lambda} = \frac{\lambda u'}{\nu_F}$ ; the Kolmogorov length scale  $\eta = \left(\frac{\nu_F^3}{\epsilon_F}\right)^{1/4}$ ; *L* is the domain length and  $\Delta x$  the grid step width (the average grid step width in the present case); and the vorticity fluctuation amplitude is given by  $\omega_{fluc} = \left(\frac{\epsilon}{\nu_F}\right)^{1/2}$ , with the large-eddy turnover timescale as  $T_e = \frac{u'^2}{\epsilon}$ . Here it is stated explicitly that  $\Delta x$  used in the final column of Table 5.2 is not obtained from the simulation and is therefore not a part of the validation but is shown

for comparative purposes. A Python script was used to automatically calculate the Kolmogorov scale in relation to the grid spacing, which ensured that the requirements of direct numerical simulation were met in each simulation throughout the thesis.

Case	Author	Reλ	l/L	$\lambda/L$	$\eta/L$	$T_e  \omega_{fluc}$	$\eta/\Delta x$
Ι	CU	65.5	0.5970	0.074	0.0047	16.92	1.20
Ι	Present	65.9	0.5832	0.072	0.0046	16.93	0.75
Ш	CU	143.0	0.5665	0.032	0.0014	37.26	0.70
11	Present	148.2	0.640	0.037	0.015	37.93	0.33

Table 5.2: Time-averaged results comparison against both validation cases of Chouippe and Uhlmann (2015), denoted CU.

The time-averaged values are matched very closely for Case I. Discrepancies are observed within a couple of percent in the worst cases, but this can be attributed to the slightly lower resolution, which notably does not compromise the validity of the method (DNS) since the Kolmogorov scales are still resolved, as described in Table 5.2. The numerical method for the Navier-Stokes solution is also different, with a seventh-order spectral element method arguably being preferable to a second-order finite difference scheme, particularly in a simple structured cubic and periodic domain as is the case here.

Case II is not reproduced as closely as Case I but there is still strong agreement for most quantities. In the convergence plots it was demonstrated that the key variables converged closely. In Table 5.2, the Kolmogorov scale and the vorticity multiplied by the turnover timescale are well matched. In Case II, the resolution used is considerably lower than the reference paper, as the Kolmogorov resolving nature of the simulations is met with twice the strictness by the reference paper compared to the present work. Despite this, it is the Kolmogorov scale that is met with the closest accuracy to the reference paper; whereas it is the Taylor microscale and large-eddy turnover length scale that deviate from the reference material the most in the second validation case.

Given that the boxes of turbulence used for the work in the present thesis are at Reynolds numbers below those used even in Case I, and not approaching those used in Case II, the validation demonstrates a successful implementation of the method with reliable results for the presently available computational resource. The forcing method was therefore taken forward with confidence.

## 5.3 Forcing Dimensional Turbulence Boxes

The turbulence domains used hereafter move away from non-dimensional boxes to dimensional ones. This is to ensure that the particle-phase properties remain exactly as required, since the particle models are dimensional. It is not straightforward to non-dimensionalise the particle-phase to a  $2\pi$  domain, since there are multiple competing length and time scales inherent to the various equations which describe the particle phase. Hence, it was reasoned that it was simpler to adjust the fluid-phase to the required scales of the particles rather than vice versa. A suggested improvement would be to solve the fluid within a non-dimensional box of turbulence, with the particle-phase appropriately non-dimensionalised. This would generalise the results more easily to a wider range of systems.

Relations exist for *a priori* estimation of the resolved Taylor-Reynolds number based on the chosen forcing parameters as shown above. However, when applied to the present dimensional boxes, the predictive relations were not useful. So, trial-and-error across around 30 single-phase simulations was utilised to arrive at the correct forcing parameters.

The present baseline box dimension is 1mm chosen to ensure the  $50\mu m$  particles remained small compared to the box dimensions. The buffer and bulk region approximating boxes are larger, at a length of 5mm, which was found to be necessary to arrive at the required turbulence conditions.

The strategy herein was to control the difference in the turbulence boxes by using two numbers, those were: the Taylor-Reynolds number, and the mean velocity fluctuation. The choice of the values for the three different boxes were taken from mean observations in channel flow simulations performed by Mortimer et al. (2020).

Whilst it is not possible to recreate precisely what is going on in a channel by using homogeneous and isotropic boxes of turbulence, the matching of these two values creates an environment where the first order statistics as experienced by the particles are similar. The three regions of turbulence chosen for approximation were the viscous sublayer, the buffer region, and the bulk flow region.

Shown below in Figure 5.5 are data accumulated from a single-phase channel flow with a Reynolds number based on the shear velocity of  $Re_{\tau} = 180$ . The data were generated as a supplementary run arising from the work of Mortimer et al. (2020) and are included here for completeness. The variables are non-dimensionalised by the channel half-height, with a full description of this process included in the referenced material.

This simulation was used to determine the required Taylor-Reynolds numbers used later in this chapter for the fully turbulent study. The figure, and the ultimate calculations of the required values, only considered the two directions perpendicular to the mean flow direction in the channel. This is because the calculation for the collision is performed in the reference frame of the collision, thereby reducing the necessity to consider the mean flow. As such, removing this component of the flow from the calculations allowed the effects of the fluctuating flow component to be isolated.



Figure 5.5: Non-dimensional Taylor microscale (left) and Taylor-Reynolds number (right) measured across a non-dimensional channel flow for  $Re_{\tau} = 180$ .



Figure 5.6: Pseudocolour velocity magnitude contours on the exterior surfaces shown for the viscous sublayer approximation turbulence box with  $Re_{\lambda} = 2.5$ .



Figure 5.7: Pseudocolour velocity magnitude contours on the exterior surfaces shown for the buffer layer region approximation turbulence box with  $Re_{\lambda} = 12.1$ .



Figure 5.8: Pseudocolour velocity magnitude contours on the exterior surfaces shown for the bulk flow region approximation turbulence box with  $Re_{\lambda} = 22.2$ .

Figures 5.6 to 5.8 visualise the velocity magnitude field in the three boxes. Qualitatively, the first two boxes appear similar, with similar flow scales in relation to the size of the domain. However, the domain size is different in the two cases and thus the absolute flow scales are different. Notably, the colour bar is also scaled differently, indicating a slightly higher peak velocity in the buffer layer box, and higher still in the bulk flow region box. Lastly, the box approximating the bulk flow region has a wider range of scales present and more closely resembles classic turbulence. The first two boxes appear to be dominated by a comparatively homogeneous distribution of flow scales. These boxes can be thought of as chaotic fluid motion, whilst the bulk box can more readily be thought of as classical fluid turbulence due to the wider scale separation.

These visualisations of the turbulence field in combination with the validation cases, demonstrates that a wider scale separation appears when increasing the Taylor-Reynolds number, which will have important influence upon non-spherical particle motion.



Figure 5.9: Temporal evolution of the viscous sublayer approximation box for the two controlled quantities, Taylor-Reynolds number (left) and velocity fluctuation (right).



Figure 5.10: Temporal evolution of the buffer layer region approximation box for the two controlled quantities, Taylor-Reynolds number (left) and velocity fluctuation (right).



Figure 5.11: Temporal evolution of the bulk flow region approximation box for the two controlled quantities, Taylor-Reynolds number (left) and velocity fluctuation (right).

As evidenced by Figures 5.9 to 5.11, the three boxes converge upon their required turbulence levels, where a transition has been achieved to a dimensional domain from the non-dimensional forcing method through the described modifications to the forcing method.

With respect to recreating the desired turbulence quantities as experienced in a canonical channel flow at  $Re_{\tau} = 180$ , the mean fluid velocity fluctuations are correctly matched, which should help impart representative fluid forces onto the particles. Similarly, the correct Taylor-Reynolds number is matched as calculated from the length scales observed in the spanwise and wall-normal directions of the channel.

Inspecting these boxes of turbulence, both in terms of the visualisations and statistics, it becomes clear the extent to which it is not possible to recreate exactly the properties of turbulence experienced in a channel flow through the proposed approach. In the bulk region of a channel flow, the flow tends to be characterised by larger scales and less chaotic motion, but that is not seen in the turbulence box approximating the bulk flow region. Instead, both the range of scales and the size of the velocity fluctuations observed are highest. This is simply what happens when trying to increase  $Re_{\lambda}$  under the constraints of the HIT system. From here, the boxes will be referred to colloquially by the regions that they were designed to represent (e.g. 'the bulk flow region box') in order to make referring to each case simple and clear, but it is important to bear in mind the limitations of this description when considering the results demonstrated. Relatedly, the terminology of 'lowest' and 'highest' levels of turbulence will be used; for clarity, this refers to the Taylor-Reynolds number of the box.

Table 5.3 describes the forcing parameters that were used in the present work to achieve the dimensional boxes of turbulence, whilst Table 5.4 details the time averaged quantities observed as a result of using these forcing values. Each of these resolved values are representative of those seen in the channel flow statistics described earlier.

Region	$\mu_F$	$\rho_{F}$	$\kappa_{cut}/\kappa_0$	N <sub>F</sub>	$T_L$	$oldsymbol{\epsilon}^*$	L	nelx
Viscous	$1 \times 10^{-3}$	997	2.3	56	0.1	0.1	1mm	24
Buffer	$1 \times 10^{-3}$	997	2.3	56	1.0	0.0001	5mm	30
Bulk	1×10-3	997	2.3	56	1.0	0.05	5mm	30

Table 5.3: Forcing parameters used to achieve the desired properties in thedimensional cases.

Table 5.4: Resultant values observed in the dimensional simulations and used for the following study, all variables are as before and u' is given in  $mms^{-1}$ .

Region	Reλ	u'	$\lambda/L$	$\eta/L$
Viscous	2.5	8	0.20	0.040
Buffer	12.12	20	0.21	0.035
Bulk	22.2	10	0.13	0.012

### 5.4 Simulation Set-Up

The overarching aim of the investigation was to further understand the influence of morphology and turbulence on particle agglomeration, as well as the dynamic processes that occur during particle interactions with DLVO effects. Morphology and turbulence introduce much variability into the system and so it was important to design Monte Carlo trials that would widely sample this variability whilst ultimately offering insights into the emergent behaviour of the system.

Pairs of disks and needles were injected into the turbulent boxes with randomised initial orientations using the same approach seen in Chapter 4. The same converged fluid file was used to start each of the simulations, but the location of injection was uniformly randomly distributed to ensure that a range of turbulence conditions were sampled. Further, the random nature of the fluid forcing ensured that no two simulations evolved the same. The particles were injected with a fixed minimal separation distance of  $2\mu m$ . This distance was chosen to ensure that a collision occurred, but also to ensure the particles were well outside of the effective range of

the DLVO forces so that the full short-range interaction was simulated. Initially placing the particles too far apart was seen to lead to much of the information of the initial conditions being lost in the chaos of the turbulence field by the time of collision - if a collision still occurred at all. At initial separations on the order of the particle diameters, collisions were seen to be infrequent, with the turbulence typically leading to the particles being carried far apart and not interacting at all. As such, it was a necessity to start them close enough to enforce a collision. The effects of the turbulence leading up to the collision are assumed to be partly included in the initial conditions of the particles, by matching the relative velocities observed at the moment of collision in the channel flow simulations of Mortimer et al. (2020). By starting the particles nearby, there was a retention of control over the velocity at the moment of collision, which is crucial for agglomeration outcomes. The particles were given their initial relative velocity corresponding to the measured averaged collision velocity for colliding particles in the respective region of the channel flow for each case. This means that the results are not one-to-one comparable with the earlier quiescent results, as the kinetic energy of the system is higher, and different in each box of turbulence, but overall trends can be compared.

The chemical parameters remained the same as for the quiescent box. The interactions were allowed to run for 4 *ms*, which is much longer than the previous study. This was necessary to allow enough time for the turbulence to significantly influence the particle velocities or for an agglomerate to form.

Table 5.5 introduces the initial conditions for the particle phase, where  $v_{rel}^p (mms^{-1})$  is the initial particle relative translational velocity,  $\omega^p (mms^{-1})$  is the initial angular velocity of each particle, and  $|\boldsymbol{d}|_{init} (\mu m)$  is the initial separation distance.

Approximated Region and Turbulence Level	$v_{rel}^p$	$\omega^p$	$ d _{init}$
Viscous (low turbulence)	4.2	0.0	2.0
Buffer (intermediate turbulence)	4.3	0.0	2.0
Bulk (high turbulence)	0.875	0.0	2.0

Table 5.5: Initial conditions used in the turbulence simulations.

# 5.5 Qualitative Analysis of Collision and Agglomeration Dynamics

The purpose of this section is to introduce some of the general behaviours seen for particle interactions in turbulence to bring more clarity to the results presented in the following subsections.

Figure 5.12 highlights three key cases of interacting disks in the viscous sublayer box. Case 1 is an example of an agglomerate forming. The mechanism through which this occurs is energy loss in successive collisions reducing the relative velocities of the particles enough that they cannot escape the potential well by the final collision. The separation distance plot shows the separation repeatedly going to zero, which is when a collision occurs. The relative velocity plots show sharp reductions in velocity due to the instantaneous hard-sphere collision, whilst the particles are able to accelerate towards one another through the action of the attractive van der Waals forces before colliding again. In the second case, Case 2, the particles remain close but not through an agglomeration event. Rather, the particles experience similar flow conditions due to being advected by the same turbulent eddy, and a small van der Waals contribution keeps them close before finally they are swept apart by an adverse velocity gradient. In the third case, Case 3, the disks collide on their edges and hence the van der Waals force is not strong enough to greatly influence the collision at the given velocities and so they move apart indefinitely.



Figure 5.12: Temporal evolution of separation distance (left) and relative velocity (right) of three illustrative interaction cases between disks.



Figure 5.13: Velocity magnitude  $(ms^{-1})$  contour plot for a lone needle intersecting a turbulence field.



Figure 5.14: Velocity magnitude  $(ms^{-1})$  contour plot for a pair of disks interacting in a *turbulence field.* 

Figure 5.13 depicts an example of a lone needle interacting with the turbulence field approximating the viscous sublayer. The turbulence is shown as a two-dimensional velocity magnitude contour plot, with a three-dimensional velocity vector plot overlaid at a slightly raised position, normal from the slice. The particle meshes are then overlaid in three dimensions and allowed to intersect the slice. At the location where the particle intersects the velocity magnitude contour slice, it can be observed that the
imposition of a no-slip boundary condition is forcing the fluid to around zero at the surface of the particle because the velocity of the particles is much lower than that of the surrounding fluid in this case. This influences the local velocity field, creating a boundary layer that moves with the particle and reduces the local velocity magnitude. Conversely, the local velocity field – which is acting uniformly in direction towards the top of the plot – contributes a force to the particle, through the pressure and viscous forces acting at the surface of the particle. This facilitates the advection of the particle and is balanced against any inertia that the particle already possesses to dictate how the particle motion evolves temporally.

Away from the particle's direct influence, a region of low velocity magnitude is observed in the plot to the left of the particle, identifiable as the blue region of the contour plot. This region of low velocity magnitude lies at the centre of a vortex, or turbulent eddy, characterised by the swirling flow that surrounds it. The size of this eddy is on the same order of magnitude as the needle, as such one could expect this fluid motion to induce a rotation in the particle's motion. For particles with low inertia relative to the fluid, this can be a region of entrainment where particles become trapped before eventually being ejected by a sweeping motion in the flow, whereas particles with high inertia are able to pass through such flow structures relatively unimpeded.

At the bottom right of the figure, visible as the orange region corresponding to greater velocity magnitude, there is a display of more directionally homogeneous fluid motion. Such regions are able to influence particles of a larger Stokes number, imparting more total force in one coherent direction over the surface of the particle and thus facilitating greater momentum transfer and greater influence over the particle's inertia.

Figure 5.14 displays a similar interaction taking place, in that the fluid velocity is significantly attenuated close to the surface of the particles. This reduces the velocity magnitude in the region where the agglomeration process takes place and thus reduces the disturbance from the flow field in this vicinity, favouring agglomeration. This raises an important point, which is that local to the DLVO interaction, the flow can be much more laminar, whilst the overall particle motion is subjected to more turbulent effects.

An interesting mechanism was observed through which agglomeration occurred between colliding non-spherical particles. Particles that collided on off-normal angles

and at off-centre collision points were found to convert their translational kinetic energy into rotational kinetic energy. This meant that, whilst the particles retained the correct total kinetic energy as determined from the prescribed coefficient of restitution, the energy was not always converted directly to translational kinetic energy (which generally results in particles moving away from one another). Rather, a spinning is induced that keeps the particles close to one another and facilitates a prolonged DLVO interaction. The rotation also makes a secondary collision much more likely, as points on the opposite side of the non-spherical shape are propelled back towards the opposing particle, depending on their relative orientation. Naturally, a secondary collision results in further energy loss and so this mechanism under the right conditions can quickly reduce the overall kinetic energy or total momentum of a pair of particles that beforehand would have looked unlikely to agglomerate based on their kinetic energy and initial surface collision points. This demonstrates that there are complex secondary effects at play which also determine agglomeration outcomes beyond just the initial collision points and kinetic energies.

This effect also works to counterbalance the distribution of DLVO forces, in that, the points on the non-spherical surfaces that induce this effect tend to be the points at which the agglomeration-inducing forces are lowest. The maximal forces are induced by collision points towards the centre of the particles for both non-spherical morphologies, which are the regions of the particle surfaces that do not lead to large torques upon collision, due to the symmetry of the moment of inertia tensor about these points. In other words, there is an inverse relationship when traversing the non-spherical surface between the size of DLVO effects and the size of the induced collisional torques.

# 5.6 Particle Interactions in HIT Approximating the Viscous Sublayer

Presented in the following three subsections are results demonstrating the behaviour of the three different particle morphologies (needle, sphere and disk) in the three different regions of turbulence. The results are shown here in terms of the primary trajectory data, i.e., the temporal evolution graphs of both the interparticle separation distances and the relative velocities between the interacting particle surfaces. In Section 5.8, this data is combined to form PDFs describing the overall system behaviour encompassing multiple instantiations of an interaction event in a Monte Carlo approach.

For brevity, comprehensive descriptions of the behaviours are given for the first region of turbulence, whilst for the second two regions of turbulence there is an emphasis on the contrasts seen with the first region, to avoid excessive repetition.



I. Spheres

Figure 5.15: Temporal evolution of separation distance (left) and relative velocity (right) between spheres interacting in the viscous sublayer turbulence box. Each line represents an individual interaction simulation.

Displayed in Figure 5.15 are the time evolution plots for the closest separation between particle surfaces dx and the relative velocities at these points dv for all of the spherical particle interaction simulations. The most general trend is for the particles to move apart and accelerate away from one another post-collision. Evidently, these are the rebound events. At the same time, there is agglomeration occurring where particles stick together and thus remain nearby with low relative velocity. Naturally, the extent to which this occurs is less visible from the separation plot but can be seen slightly more clearly from the velocity plot.

For spheres in the lowest turbulence box, particles that separate post-collision continue to separate indefinitely over the studied time interval, as opposed to experiencing forces that bring them back together. At very short time intervals after a collision, DLVO forces remain effective and are thus the prevailing mechanism through which particles move back towards each other – but, outside of this short effective

range, the relative motion of the particles is a result of the background fluid velocity field, as well as the inertia that the particles possess. As such, energetically unfavourable (with respect to agglomeration) collision events occur over very short time periods, whilst agglomeration events continue to occur for extended periods, and indefinitely if the turbulence is not strong enough to separate them. From a computational modelling perspective this places a requirement on the numerical interaction models to always be active. This can be challenging as there is a need to achieve a balance, or quasi-equilibrium, in the energy that is exchanged as a result of the two separate methods – one that does not diverge over time.

The trend in this region of turbulence (of continued separation) is instructive on the background fluid motion. Specifically, the resolved fluid dynamics are such that the particles experience longer and more uniform sweeping motions in their frame of reference. This means that in the simulation time studied, particles tend not to experience wide variability in the background fluid force.

The tendency for the particles to move apart after collision creates a strong separation of scales in the simulations between cases where agglomeration is observed and where it is not. This underscores the challenge with studying these systems. It becomes abundantly clear that the DLVO interactions occur on very short timescales compared to the turbulence, or fluid motion. This creates challenges for lower-fidelity modelling approaches, like those seen in industry, because capturing both elements of the physics in its entirety requires very low timesteps and thus large compute times. Hence why an instantaneous model is typically used for agglomeration physics. The existing models, as described in the literature review, are primitive in the case of nonspherical particle interactions and thus this area of study has the added benefit of informing any further model developments that may take place.

The continued separation tendency of the particles is compounded by an acceleration, seen in the relative velocity plot, wherein particle relative velocities continue to increase. The general trend appears to be an approximately quadratic increase in the particle separation for a linear increase in the particle relative velocities.

Since the particles are injected at very close separations to ensure collision, there is an obvious accumulation of collision behaviours seen in the first 0.5 ms of the simulations, best seen in the relative velocity plots represented by a sharp and instantaneous decrease in velocity. Each of these lines, for the individual simulations, almost lie atop one another. As the respective simulations advance, further collisions are seen infrequently in cases where either the turbulence, the DLVO forces, or both acted to ensure particles remained close by. The short-time collision behaviour is therefore predictable, but the chaos of the turbulence means that the further time advances from the initial condition, the less predictable the behaviour of the system and, specifically, the less predictable the occurrence of a further collision. Some collisions happen late into the simulation time but, by this point, their frequence is random.





Figure 5.16: Temporal evolution of separation distance (left) and relative velocity (right) between needles interacting in the viscous sublayer turbulence box.

Figure 5.16 shows that the general behaviour of the needles possesses greater variability than the spheres in the same turbulence box. The maximum separations experienced are similar between the two morphologies, but the maximum relative velocities achieved are marginally greater in the case of the needles. This indicates that the needles are slightly more susceptible to the fluid forces in the extreme cases, likely owing to the greater surface area of the needles. When this surface area aligns in such a way that the symmetry axes lie perpendicular to the local mean flow direction, the forces will have their greatest effect, and a greater total force will be imparted than in the equivalent scenario with a sphere. When the end of the needle faces the local flow direction, in a parallel alignment, the total force imparted will be lower than that of a sphere. (However, this configuration is more susceptible to perturbation, meaning it

does not sustain itself.) This point is relevant to the total behaviour over morphologies discussed in Section 5.9.

As with the spheres, particle pairs that separate quickly post-collision continue to separate generally, but two simulation cases are seen to begin to move back towards each other, which is only possible if the fluid force experienced towards the end of the simulation opposes the direction of the original fluid force. In general, this effect relates to the length and time scales of the fluid, which are in theory constant between the different morphological cases studied. So, the difference between the two scenarios can be explained by the morphology of the shape: the longest axis of the needle is greater than that of the spheres. Therefore, the extent of the velocity field probed by the needles is greater and this allows the needle-like particles to experience the fluid at a slightly different length scale. The particles will also alter the turbulence field itself differently based on morphology in a coupled way.

Most clearly seen in the relative velocity plots, the needles experience far more collisions, extending further into the total simulation time with sustained frequency. By interrogating the specific cases where this happened, it was observed that one mechanism for this was due to secondary collisions induced by the specific configuration of the first collision. In such cases, the orientational configuration is generally such that the collision points lie away from the centre of the particle's principal axes. This means that the force acts off-normal with respect to the centre of mass and thus there is a torque induced which leads to the kinetic energy of the particles being redistributed into predominantly rotational motion as described earlier. When this occurs, the particles remain nearby, since the translational motion is the key mechanism through which the particles separate. With the particles rotating close to one another, a further collision is inevitable; thus, without considering DLVO forces at all, this leads to repeated collisions and thus repeated losses of energy through the inelastic collisions (e = 0.4). This is therefore observed to be a secondary mechanism keeping particles close together with a reduced energy besides the attractive van der Waals force. These two effects can also work in tandem to further increase agglomeration propensity, since the lack of translational separation allows the DLVO forces to act for longer in the case of non-spherical shapes. Further to this, the edgebased collisions inducing a spin causes the orientational configuration often to pass

through face-face or long-edge-long-edge alignment, thus promoting more maximal van der Waals forces in this window.

It can be seen from the relative velocity plots that the amount of velocity lost in a collision is much greater than 60% as prescribed by the coefficient of restitution. This indicates a strong influence from the DLVO forces in reducing the kinetic energy of the system.

III. Disks



Figure 5.17: Temporal evolution of separation distance (left) and relative velocity (right) between disks interacting in the viscous sublayer turbulence box.

The disks, shown in Figure 5.17, are similar to the needles, however they reach greater maximum values of separation and of relative velocity than both the spheres and needles. Again, with different principal lengths compared to the spheres, there is the opportunity for disks to sample the flow differently due to morphology. As a result, cases can be seen where the separation begins to reverse, this time much quicker than the needles, showcasing the unique acceleratory effect the flow is able to have on the disks.

The results for the agglomeration versus bouncing outcomes are displayed in Table 5.6, helping to give weight to the prior observations. The spheres and needles are shown to give comparable rates of agglomeration, with the disks showing the lowest rate of agglomeration over the full parameter space.

Result	Needle (5:1)	Disk (5:1)	Sphere	Total
Bounced	65	73	35	173
Agglomerated	32	23	15	70
Agglomeration Rate	32%	23%	30%	28%

Table 5.6: Agglomeration results for the lowest turbulence level box.

# 5.7 Particle Interactions in HIT Approximating the Buffer Layer Region

#### I. Spheres



Figure 5.18: Temporal evolution of separation distance (left) and relative velocity (right) between spheres interacting in the buffer layer turbulence box.





Figure 5.19: Temporal evolution of separation distance (left) and relative velocity (right) between needles interacting in the buffer layer turbulence box.

III. Disks



Figure 5.20: Temporal evolution of separation distance (left) and relative velocity (right) between disks interacting in the buffer layer turbulence box.

The general trends are repeated in the case of all morphologies in the buffer layer box when compared with the viscous sublayer results, this can be seen from Figures 5.18 to 5.20. One key difference is that all three morphologies are accessing higher velocities and separations in the extreme cases, which is indicative of the overall higher available kinetic energy budget of the fluid. Clearly, in the studied interval there is enough time for this difference in turbulence level to influence the outcomes of the particles. Secondly, the relative velocity plots are beginning to look less smooth for all three morphological cases. The slightly wider range of scales present in the intermediate turbulence box increases the variability of the velocities.

Presented in Table 5.7 are the overall breakdowns for agglomeration rate. Again, the spheres are seen to be comparable with the needles with the disks now falling even further behind. The trend for the needles to behave similarly to the spheres is retained, and both morphologies maintain similar values of agglomeration rate when mildly increasing the turbulence level. Interestingly, the overall agglomeration rate of the box is the same (28%) as the lowest turbulence box.

Table 5.7: Agglomeration results for the intermediate turbulence level box.

Result	Needle (5:1)	Disk (5:1)	Sphere	Total
Bounced	64	79	34	177
Agglomerated	34	18	16	68
Agglomeration Rate	34%	18%	32%	28%

### 5.8 Particle Interactions in HIT Approximating the Bulk Flow Region

I. Spheres



Figure 5.21: Temporal evolution of separation distance (left) and relative velocity (right) between spheres interacting in the bulk flow turbulence box.

#### II. Needles



Figure 5.22: Temporal evolution of separation distance (left) and relative velocity (right) between needles interacting in the bulk flow turbulence box.

III. Disks



Figure 5.23: Temporal evolution of separation distance (left) and relative velocity (right) between disks interacting in the bulk flow turbulence box.

Observing Figures 5.21 to 5.23, the contrast with the previous box is similar again in that the separation distances and relative velocities reached in both the extreme cases and in general have greatly increased. The variability in the relative velocity plots is also much more pronounced in the highest turbulence level box. Collisions, seen most clearly by sharp reductions in the relative velocity, are much less frequent later into the simulation time. This is because the particles are more quickly and reliably transported to regions in the simulation domain away from one another where they cannot collide. This shows agreement with what is generally observed in a channel flow, where the

bulk flow region experiences collisions least frequently (Mortimer et al., 2020), however the underpinning mechanism for this outcome is not the same.

For all three boxes of turbulence, the separation distance evolution for all three morphologies remains smooth over the studied time interval, indicating that in the timeframe under investigation the turbulence levels are too low to affect particles of this size in a more chaotic way, which would typically be synonymous with particle advection due to turbulence. As shown, more nonuniform and chaotic particle motions are beginning to be observed in the highest turbulence level box, approximating the bulk flow region. This runs in contrast to the general theory associated with turbulent particle motion across the channel flow (Mortimer et al., 2019).

Finally, it was demonstrated in the present simulations that the background turbulence field was not strong enough to break agglomerates; however, there is no reason this should not be possible under the DLVO potential model. The fluid must supply enough energy to allow the particles to escape the potential well, but this is not realised at such a low Reynolds number. The inclusion of many particles is the next avenue for study, and a tertiary particle collision could also supply enough energy to break an agglomerate.

The overall statistics for agglomeration rate are presented in Table 5.8. The same trends are repeated. The spheres and needles are again similar in agglomeration rate: this time the spheres have the same agglomeration rate as the needles, despite lagging very slightly behind in previous cases. Both agglomeration rates are being attenuated by the turbulence compared with lower values of the Taylor-Reynolds number. The disks now have less than half the agglomeration rate of the other two morphologies. Lastly, the overall agglomeration rate has gone down significantly compared to the previous two turbulence levels.

Result	Needle (5:1)	Disk (5:1)	Sphere	Total
Bounced	60	87	40	187
Agglomerated	22	13	15	50
Agglomeration Rate	27%	13%	27%	21%

Table 5.8: Agglomeration results for the highest turbulence level box.

# 5.9 Influence of Turbulence Characteristics on Particle Agglomeration

A wide study has taken place into the role of fluid turbulence properties on the agglomeration and collision behaviours of interacting pairs of spheres, needles and disks with an aspect ratio of 5:1 and particle volume kept fixed. With three turbulence boxes and three morphologies, there is a large data set accumulated with a total of 725 independent simulations taking place. The following subsection aims to analyse these results with a focus on the influence of turbulence strength.

Starting with a direct comparison of the overall data from the three boxes, it was observed that the rate of agglomeration reduced as the turbulence strength increased. Specifically, as the Taylor-Reynolds number of the box increased from 2.5 to 12.1 to 22.2, the overall agglomeration rate inclusive of all morphologies reduced from 40% to 38% to 21%, respectively. This implies that agglomeration is negatively impacted by increasing turbulence levels.

Whilst this conclusion is true for the present system set-up, there are two key points to consider in order to view this result in context. Firstly, collisions are guaranteed in the present set-up, whereas in industrial or natural systems, collisions between particles are not guaranteed (in the precise sense). In industrial cases or in natural flows, increased turbulence can therefore help to promote collisions through increased collision frequency which in turn promotes agglomeration, even if the agglomeration efficiency is unchanged. In the present system, that frequency-based argument does not apply. Secondly, in industrial or natural systems, there is likely an inflection point where the turbulence becomes too high so that almost all collisions are energetically unfavourable for agglomeration. Similarly, the collision frequency may become so high that tertiary particle collisions are repeatedly breaking agglomerates into smaller constituents. These kinds of effects can only be studied by moving to multi-particle systems, as opposed to binary. What the present results do indicate is that where a collision occurs, increased fluid forces are overall disruptive of agglomeration, and thus increasing the turbulence level decreases the agglomeration level. There were instances where the fluid acted in a direction such that it decelerated the particles enough to promote agglomeration, but in general the fluid effects accelerated the particles enough to make the interaction energetically unfavourable. Due to the nonuniform and turbulent nature of the flow, particles experienced different velocities to one another which on average increased the relative velocity magnitude between them. The acceleration was also partly down to the initial prescribed velocities of the particles being below that of the convergent/mean particle velocity should they have been free to move around indefinitely, thus the particles relaxing to the background flow conditions was also an acceleratory effect in general.

Another explanation for the difference in agglomeration outcomes relates to the initial separation distance of the particles. This distance was chosen as a matter of necessity to ensure that the collisional conditions were similar to those determined in the complementary channel simulations. In boxes where  $Re_{\lambda} = 2.5$  or  $Re_{\lambda} = 12.1$ , the closing of the initial separation distance (which is facilitated mostly by the initial relative velocity of the particles) is not long enough for the fluid to greatly impact the overall dynamics as the typical force is not high enough to impart a large acceleration. Whereas, once the turbulence level reaches  $Re_{\lambda} = 22.2$  the fluid is able to disrupt the agglomeration process by transferring momentum to the particles quickly enough to prevent agglomeration. In both cases, the fluid is having an effect, however this effect is not prolonged enough for the particles to relax away from their initial conditions towards the exact conditions of the turbulence field. This is a limitation in the sense that the particle dynamics at the moment of collision are not fully representative of the turbulence field they lie in; however, this also helps to keep the results in line with the complementary channel flow simulations, since the initial conditions contain the relevant information about the mean flow they theoretically exist within. The initial velocity of the particle-phase is lower in the bulk box, giving a greater window for the fluid, which will not act in the direction of the particles' approach in general, to change the direction slightly and sweep the particles apart in some cases.

It was observed throughout the present simulations that none of the turbulence levels possessed the means to separate the particles once an agglomerate had formed. This is because once the particles fall into the potential well, the attractive force dominates over the fluid force, in a strong separation of scales. The overall statistics indicate that the turbulence conditions in the lesser two boxes of turbulence were not different enough from one another to greatly change the agglomeration outcomes despite the different Taylor-Reynolds numbers. It can therefore be inferred that in these two boxes

the orientationally dependent DLVO forces are more influential than the fluid forces in determining an interaction outcome.

Looking at the trajectory data in the previous sections, the disturbance to smooth particle motion, which would arise from the background turbulence field, is not visible on the timescales studied for these two boxes, whereas noise begins to become more visible in the higher turbulence box. Such velocity fluctuations as seen in the highest turbulence case may be large enough to disrupt the agglomeration process as it happens, as the particles try to reach their computational equilibrium, and thus prevent an agglomerate forming, either by imparting enough kinetic energy into the interaction to make agglomeration unfavourable, or to move the surfaces away from one another before the agglomeration process has enough time to act. The chosen Taylor-Reynolds numbers were low, with only the bulk box showing the scale separations associated with turbulence. It is this scale separation that leads to the noise in the velocity plots, as the fluid can affect the particles on multiple length scales. The smallest of these scales act on a length scale lower than that of the particles, which allows gradients to form between the particles more often, making it more likely for them to be swept apart. The lower two turbulence boxes also resolve this effect, but to a lesser extent and so they tend to demonstrate similar local flow conditions more frequently, in comparison. It is this coherence between the two particles induced motion that therefore reduces relative velocity and allows an agglomerate to form. Once the agglomerate has formed, a subsequent increase in turbulence can occur to unfavourable levels, and the particles still remain connected.

These smaller and faster fluid motions also create velocity gradients across the respective surfaces of the particles, promoting a different kind of motion, wherein they rotate about their centre of mass, as opposed to the larger sweeping motions observed in the other boxes which more readily affect the translational component of motion with less rotation induced.

Both the lowest  $Re_{\lambda}$  turbulence boxes started with virtually the same particle-phase kinetic energy: this was prescribed through their initial velocities, as given in Table 5.5 and those velocities were chosen based on results from the reference channel flow simulations. With no difference in initial velocity, the 2% reduction in agglomeration rate could feasibly be attributed to the changing turbulence level. However, a sample

size of O(100) is small enough to contribute error of a similar order to this computed difference, despite the convergence of the PDFs. The reduction in agglomeration is seen to be much more significant when contrasting the former two turbulence boxes to the highest turbulence level box. This is even though the initial kinetic energy condition of the particle-phase was lower in this highest box of turbulence. This suggests that the differences in agglomeration between the lower two boxes are reflective of turbulence, rather than intrinsic error in the simulations.

Relating these findings back to the theory of particle agglomeration in channel flows, the results are contrary to current understanding, but can be remedied by the following. The picture in full-scale particle transport flows is that, in the bulk flow region of a channel or pipe, the relative velocities at collision events are low because particles share similar trajectories, advected by larger coherent motions than those typically observed nearer to the wall. Hence, the relative velocities of the particles are lower on average despite their absolute velocities being higher. In a channel, therefore, the agglomeration rate (or probability of an agglomeration given a collision) observed in the bulk region is highest, with particles colliding at low relative velocity and with low collision angles. The collision frequency is also lowest in this region (Mortimer et al., 2020) because particle paths are not frequently crossing the trajectories of others. This outcome is not replicated in the boxes of homogeneous and isotropic turbulence discussed presently. Removing the large-scale coherent structures from the flow to satisfy the conditions of homogeneity and isotropy means a loss of significant information to the behaviour of particles in realistic flows. Rather than particles in the bulk region box experiencing lesser flow gradients, they experience higher flow gradients because all aspects of the flow are forced to be higher to satisfy the HIT requirements. For this reason, it is sensible to not infer anything about specific channel flow scenarios from the present results in a one-to-one comparison, but rather information can be taken about the interaction between turbulence and particle agglomeration.

Similar to the approach taken throughout Chapter 4, the trajectory data of the particles can now be recast in terms of PDFs that describe the overall system behaviour. From this, it is possible to further identify general trends with respect to the turbulence level in the respective boxes.

205



Figure 5.24: PDFs demonstrating the overall tendencies of the spheres with respect to separation distance (top left), relative velocity (top right) and van der Waals force (bottom) in the three turbulences boxes.



Figure 5.25: PDFs demonstrating the overall tendencies of the needles with respect to separation distance (top left), relative velocity (top right) and van der Waals force (bottom) in the three turbulences boxes.



Figure 5.26: PDFs demonstrating the overall tendencies of the disks with respect to separation distance (top left), relative velocity (top right) and van der Waals force (bottom) in the three turbulences boxes.

Figures 5.24 to 5.26 are the PDFs of separation distance, relative velocity and maximum achieved van der Waals forces for the spheres, needles and disks, respectively. These graphs indicate similar relationships between the plotted metrics, and the turbulence box level, consistent across all three morphologies. In particular, the viscous box creates the closest mean separations between the particles due to the accompanying lower velocity states reached by the particles.

The variance of the plots also follows the same trend, where the most variance is seen in the bulk flow region, followed by the buffer and then viscous sublayer regions. Again, this is consistent across morphologies. The increase in variance is due to two main factors. Firstly, a wider range of flow scales allows for a wider range of motions to be experienced by the particles that separate after collision, giving a wider variability in both separation and velocity. Secondly, the greater turbulence level facilitates greater maximum velocities and thus the tail of the distribution can span a wider range. The van der Waals force plots show that greater maximum attractive forces were achieved in the low turbulence level box, for the two non-spherical shapes. In theory, this maximum value should only be dependent on the orientation experienced at the point of collision because collisions are guaranteed - collisions imply touching surfaces and separations approaching zero. However, to access the very lowest possible separations, and thus maximal forces, it is required that the particles repeatedly collide and lose relative velocity through the previously described 'vibration effect' due to the hard-sphere collision scheme. This allows the dynamic collision detection distance to relax and closer separations to be achieved. That means the distributions of maximum van der Waals forces shown above are reflective of the number of agglomeration events in the respective boxes and so only in the box with the lowest background turbulence was this effect allowed to predominate such that the particles could regularly achieve their maximal forces. This effect could also be influenced by the effect of turbulence on particle orientation or alignment, which will be investigated in the next section. Relating this back to physical scenarios, even though the initial agglomeration process occurs on very short timescales, the turbulence of the fluid is able to disrupt the strengthening of the bonds between the particles which occurs over a slightly longer period under the present model.

The spherical van der Waals distributions are bimodal with a peak at near zero force and a peak at the maximum achievable force. Due to the lack of orientational effects, there is not much variance between simulations and the space in the distribution between these peaks is mostly populated as the initial agglomeration process takes place.

### 5.10 Influence of Morphology on Particle Agglomeration in Turbulence

It is also of interest to analyse the observed effects based upon the morphology of the particles, the other key variable in the simulations, to understand the role non-sphericity has on overall agglomeration propensities and particle alignment.

In Figure 5.24, for the spheres, the peaks of the relative velocity distributions are very similar between the two lowest turbulence level cases ( $Re_{\lambda} = 2.5$  and  $Re_{\lambda} = 12.1$ ). This trend is not replicated by the disk morphology. This indicates that the non-sphericity of

the other two particles gives rise to a set of effects that creates a disparity in relative velocities not seen in the spheres. For non-spherical particles, unique motions are induced as the flow gradients interact with the anisotropic inertia tensors, which has a coupled and dynamic behaviour. Different sized flow structures and flow gradients exist in these two boxes as shown in Section 5.3, which could thus create a difference that is not seen for the spheres, because only non-spherical particles have access to this effect. Secondary to this, the principal radii of the non-spherical particles are of course different to the spheres as well as to each other, which means different characteristic lengths. Thus, the way the particles experience the turbulent flow from a particle response time point of view is different based upon the morphology.

Below, in Figures 5.27 to 5.29, the data are recast as a morphological comparison, now with the symmetry axes included in the third subfigure. These plots further demonstrate the tension between the acceleration due to the fluid and the orientational attractive forces. Needles were demonstrated in the previous chapter to have the highest agglomeration propensity in a quiescent box but increasing the turbulence levels means that the dynamic measurements of this propensity, i.e., separation and relative velocity, begin to favour spherical particles more and more. In the viscous box, the needles retain their agglomeration predominance but as the turbulence increases, they are overtaken by the spheres in some measures. Even in this lowest turbulence level box, the mean of the relative velocity is higher for the needles despite the agglomeration rate also being higher. This is since the tails of the needles' distribution extends further, influenced by the non-agglomerating cases.

The suggested reason for this disparity is due to the reduced drag on the spherical particles. Conversely, the surface area of the disks for the hydrodynamic forces to act on is greatest, hence drag on the disk is likely to be stronger. This is supported by the disks routinely having the worst agglomeration characteristics, although this cannot be decoupled from the disks also showing less agglomeration favourability over their total orientational parameter space, as detailed in Chapter 4.

It is known that non-spherical particles in turbulence will tend to orientate themselves locally within the flow in such a way that maximises their drag (van Wachem et al., 2015). The injection of the particles and their initial orientation has no relation towards the local flow, however as the simulations develop, the particles experience this dragrelated effect, and they align in this way (an effect which is not present in the quiescent case). As such, the difference in surface area between the particles shows in the results here also.

Relatedly, as the particles act to orientate themselves in a maximal drag configuration, the characteristic length of the particles lying normal to the flow increases. Again, this leads to differences in the ratio between the particles' characteristic length and the turbulent length scales, which can influence the types of motion experienced by the particles. As the particles agglomerate and form larger structures, this length scale can increase again and so the relative effects of the turbulence will change over time, for a fixed Reynolds number.

These differences in experience of the flow are indicated by the symmetry axis plots, which show that: in the viscous sublayer box, disks have a higher tendency to be aligned; in the buffer layer box needles have greater alignment tendency; and in the bulk flow box the alignment is very similar, with disks showing slightly less variance the opposite of that which was observed in the quiescent case. In other words, depending upon the local flow scales the extent to which the disks and needles align appears to change. A deeper analysis focused on particle alignment could investigate how the ratio of the particles' principal axes to the flow scales alters alignment characteristics in HIT. Such an analysis could be compared to the distributions shown presently to understand the extent of the attractive forces on particle alignment and how they balance with the influence of the turbulence on this same property. The surfaces of the particles are also able to introduce shear to the flow, thus introducing a mechanism to generate turbulence, and increasing the local energy dissipation. It would be of interest to study how the agglomeration process effects this phenomenon. Lastly, there have been numerous studies on spheroidal particles and their alignment with turbulent flow (Voth and Sodati, 2017), an extension of this existing work could involve examining the alignment of agglomerated particles with the turbulent flow, which will likely be much less stable and predictable, given the intricacies of the shapes formed.



Figure 5.27: PDFs demonstrating the overall tendencies of the viscous sublayer box with respect to separation distance (top left), relative velocity (top right) and van der Waals (bottom) for the three morphologies. Vertical lines indicate the statistical mean.



Figure 5.28: PDFs demonstrating the overall tendencies of the buffer layer box with respect to separation distance (top left), relative velocity (top right) and van der Waals (bottom) for the three morphologies. Vertical lines indicate the statistical mean.



Figure 5.29: PDFs demonstrating the overall tendencies of the bulk flow box with respect to separation distance (top left), relative velocity (top right) and van der Waals (bottom) for the three morphologies. Vertical lines indicate the statistical mean.

In the buffer layer box, the difference between spherical and non-spherical particles begins to become much more pronounced in terms of the agglomeration characteristics, compared to the viscous sublayer box where small differences were present, but the means and distributions occupied similar areas of the parameter space. In the highest two turbulence boxes, there is a clear division of spheres and non-spheres into two separate and distinct categories. In that, the difference between spheres and non-spheres is greater than the difference between the non-spherical particles themselves, an effect owing to the increasing turbulence levels. In these two boxes, the drag effects clearly begin to dominate over the attraction effects. The conclusion therefore is that to create a system where weakly attractive particles remain close by, the best choice is spheres due to their optimal drag characteristics, so long as the turbulence level is high. As the turbulence level reduces towards  $Re_{\lambda} \sim O(1)$ , then needles become the preferred choice due to their greater attractive force

characteristics over the orientational parameter space. Interestingly, the gap has closed between disks and needles in the highest turbulence level box, and in fact this is seen sequentially moving up the turbulence levels. Rather than the highest surface area shape, the disk, becoming more and more separated from the other two shapes, in terms of agglomeration propensity, the results demonstrate needles and disks becoming similar in these metrics and distinct from the spheres. Despite this, the needles still retained the highest integer number of agglomeration events and this is due to the closeness of the separation. The turbulence effects discussed presently, seen in the PDFs, are largely accumulated at later simulation times than the moment that agglomeration typically occurred, which was very early into the simulation time. Since it is not possible to separate the particles widely and still ensure the correct collisional conditions, an improvement here would be to have unique initial velocities for the different morphologies, more representative of their different behaviours in the same turbulence field.

For both non-spherical morphologies, the alignment of the particle's symmetry axes appears to have shifted to the left, or to more parallel configurations. In the quiescent case shown in the previous chapter, the total symmetry axes distributions tended to be roughly equally distributed about the  $\pi/4$  half-way point; whereas, in turbulence, each of the distributions collect slightly to the left of this mark. This indicates that the turbulence is promoting alignment between the particles, as well as the DLVO effects. It could also be the case that the longer simulation time is giving the DLVO torque longer to enact, which also favours parallel alignments.

In slower moving systems, and with the agglomeration forces allowed to act over a significant period, a significant DLVO torque is induced that ultimately acts to align particles with one another, as implied by Schiller et al. (2011). In the present model, this torque is resolved, but it was observed to be a very small contribution to the overall dynamics of the particles, with the turbulence and collision mechanics dominating. Two particles agglomerated and allowed to interact for extended periods could change their orientation towards the maximal configuration through this torque. The DLVO force is calculated at a single point on each surface in the present model, which means that there is *not* a force resolved coming from regions where the van der Waals force is higher, pulling the particles into a stronger alignment case. However, it is hypothesised here that any perturbations (which naturally arise due to the flow) to the orientational

configuration will on average favour the more maximal van der Waals configurations which over time should align the particles further.

Having the DLVO forces acting at a single point per particle surface could be seen as a limitation, due to the above. However, since these interaction forces are so close range, they can be assumed to be acting very locally, especially on a convex body where neighbouring points to the interaction always fall away in distance. This is underscored by the Derjaguin approximation described in Chapter 3.

On the present timescales, agglomerated particles were observed to stick in their initial alignment with strong stability. In fact, they were seen to maintain their orientational configuration with strict coherence even when the combined agglomerate was still moving rapidly and chaotically in the turbulence. For lower fidelity models, this implies that it would be a reasonable assumption to agglomerate particles into a fixed configuration based upon the initial collisional state.

Lastly, the increased alignment could also be due to the fact that the particles were injected very close to one another: Voth and Soldati (2017) relay that the orientation of particles is strongly dependent upon the local velocity gradient tensor. Hence, two nearby particles experiencing similar local flow conditions are more likely to become aligned with one another.



Figure 5.30: Symmetry axis alignment (left) and collision point location (right) for needle interactions in the viscous sublayer box.



Figure 5.31: Symmetry axis alignment (left) and collision point location (right) for disk interactions in the viscous sublayer box.



Figure 5.32: Symmetry axis alignment (left) and collision point location (right) for needle interactions in the buffer layer box.



Figure 5.33: Symmetry axis alignment (left) and collision point location (right) for disk interactions in the buffer layer box.



Figure 5.34: Symmetry axis alignment (left) and collision point location (right) for needle interactions in the bulk flow box.



Figure 5.35: Symmetry axis alignment (left) and collision point location (right) for disk interactions in the bulk flow box.

To understand further the orientations that lead to agglomeration in turbulence, and hence the kind of structures most likely to form, data were again collected on the collision points of disks and needles and compared in cases with and without agglomeration, these results are presented in tandem with the symmetry axis configurations at the point of collision.

It can be observed from Figure 5.30, which shows particle behaviour in the viscous sublayer box, that the particle pairs which agglomerate tended to have their symmetry axes aligned for the needles, as was the case in the quiescent system analysed in Chapter 4. In comparison to the quiescent box, the influence of the turbulence field, and higher particle initial velocity, is that an alignment of the symmetry axes has become more pronounced in the case of agglomeration. A first interpretation of these results indicated that introducing turbulence meant a stricter requirement on alignment

in order to see agglomeration. However, this is not replicated in the increasing turbulence level boxes. Instead, it appears that alignment becomes less important, implying that the kinetic energy of the particles is the most important metric and alignment runs secondary as the turbulence level increases.

Both distributions, that of the agglomerated particles and the rebounded ones, appear to have shifted towards lower relative orientations, indicating that there is a promotion of parallel alignment irrespective of agglomeration or rebound. The behaviour reverses in the case of the disks in Figure 5.31. The particles that agglomerated are evenly distributed across the parameter space, whilst the rebounding particles show a bias towards parallel alignment. The 'agglomerated' distribution may be anomalous for the disks since so few particles agglomerated and hence the sample size is low. In fact, that is true of each of the cases for the disks. Since the disks were shown to be more susceptible to the forces from the flow field, the local flow dominates the interaction to a greater extent, and it is more difficult to extract the DLVO effects in the analysis stage.

This underscores the fact that as the turbulence level increases, the effect of particle alignment diminishes in determining agglomeration outcomes, or at least becomes less visible in the present PDFs. This is best seen by observing the needles' PDFs, the morphology for which the sample size for both categories of agglomeration and bouncing is large enough to support a conclusion. From the viscous sublayer box to the buffer layer box, the requirement for alignment appears diminished, but the peak of the agglomerated distribution still shows a favouring towards alignment more than bouncing events. Moving to the bulk flow box, this requirement is seen to diminish almost entirely, with the distributions between agglomerating and bouncing needles becoming very similar. Only in the region of perpendicular alignment is there a significant difference between the two PDFs. This again shows how the increasing turbulence begins to predominate. In other words, the interacting particles are relying much more on their lower kinetic energy, having not been advected by a strong fluid motion, as opposed to relying on parallel alignment, to facilitate agglomeration.

It is also possible to repeat the analysis shown in the previous chapter wherein the collision points were measured on the particles' surfaces, allowing understanding to be generated on the importance of the particle surface characteristics. The presented

results omit the  $\theta$  distributions, which were shown to be unimportant. Similarly, the absolute value of  $\phi$  is taken in order to double the sample size for a given angle, this helps when interpreting the PDFs since the turbulence introduces a variability. The operation is valid since the interval the angles are defined on is symmetric.

Now in the case of turbulence, it can be seen from the rightmost plots in Figures 5.30 to 5.35 that the overall trends are retained, however the extent changes. Over the full parameter space, the disks in particular are shown to be less reliant on surface collision points when determining agglomeration outcomes. As described in the case of symmetry axes, this is due to the increasing reliance on the kinetic energy state as determined by the flow experienced rather than specific orientation or surface conditions. There is still an overrepresentation of agglomeration observed at collision points where the surface interaction is maximised. However, this diminishes for the disks and in the high turbulence box the trend reverses, with an overrepresentation of off-normal collisions. Here, the sample size is low and so it is not appropriate to draw strong conclusions on this particular result. Nevertheless, it is possible that the offnormal initial collisions actually become preferable for agglomeration when the overall system is not favouring agglomeration. This is because of the off-normal collisions inducing rotation and leaving the particles nearby in a lower velocity state. In some cases, this may lead the particles towards a favourable configuration that ultimately facilitates agglomeration on a secondary or tertiary interaction when the kinetic energy is lower. However, this mechanism is reliant on the particles remaining close by in the state induced by the collision. In many cases, the stronger turbulence will not allow this, sweeping the particles apart. So, it is not clear the extent to which this effect is present. Lastly, the total distribution of collision points (i.e., not considering agglomeration or rebound) has flattened to become a much more uniform distribution, which indicates that turbulence in some way prevents disks from colliding on their edges, or promotes face-face collisions, much more than observed in quiescent conditions.

For the needles, the peaks have remained in the same position for the collision points but the distributions appear to be less coherent away from the peak. Across the three boxes, the very stark overrepresentation of agglomeration at the favourable configuration is seen, much more so than the equivalent points on the disks. This implies that even in turbulence the types of structures that are likely to form will have needles agglomerating along their long edges since it is such a starkly favoured state under the present model. Conversely, the disks may take on different shapes but are less likely to agglomerate at all based on the assumptions in this model. Naturally, there are many effects not captured by binary interactions that will be pertinent to larger structures.

#### 5.11 Conclusions

The work presented in this chapter represents the successful implementation of a stochastic forcing method which was able to resolve both non-dimensional ( $Re_{\lambda} = 65, 143$ ) and dimensional ( $Re_{\lambda} = 2.5, 12.1, 22.2$ ) homogeneous and isotropic turbulence boxes. A validation was performed against two references cases with strong agreement demonstrated. Furthermore, the developed agglomeration methodology in the previous chapters was successfully coupled to the fully resolved turbulent fluid. This is the first time non-spherical immersed boundaries with orientationally resolved DLVO forces have been coupled to such a turbulence field, highlighting the novelty of the work.

The results indicated a clear dependency on both turbulence and morphological characteristics for agglomeration outcomes, alignment, separation trajectories, and relative velocities between particles. It was found that increasing turbulence levels reduced the agglomeration rates between all three morphologies and that introducing turbulence to the domain promoted alignment between particles. The viscous sublayer box exhibited the highest agglomeration rate, followed by the buffer layer and then the bulk flow boxes. This finding was contrary to the general theory of wall-bounded flows; however, this was understood as a limitation of the HIT approach, given that increasing the Taylor-Reynolds number to the measured channel flow value leads to very different turbulence conditions to that which is being approximated. Despite this, the results gave insight into the relationship between turbulence and agglomeration. The increasing turbulence levels were observed to accelerate the particles to relative velocities where agglomeration became energetically unfavourable due to the kinetic energy possessed by the particles. The net effect of acceleration is due to the fact that the pertinent velocity quantity is relative velocity, and the turbulence in general creates different flow conditions for the different particles as a result of the scales and thus gradients present in the box. The disparity in flow conditions between particles is an increasing function of turbulence level due to the increasing separation of scales.

The overall reduction in agglomeration rate as a function of turbulence strength is due to multiple reasons. The first is the additional energy imparted by the fluid onto the particles by the increased turbulence levels. In particular, it was observed that increased mean fluid velocity fluctuations, means an increased turbulence kinetic energy budget of the fluid and thus more of this is imparted onto the particles. Secondly, this increased energy allowed particles to be carried further from one another in the same period of time, thus removing the particles more quickly from the effective DLVO range and moving the particles to locations in the box where they are unlikely to meet again for the rest of the time simulated. Thirdly, increasing turbulence levels was found to increase the presence of different flow scales and structures which were able to disrupt the agglomeration process via increased noise in particle motion, although this noise still remained relatively low for the studied turbulence levels.

The non-spherical differences in the particles were shown to change the way that they interacted with the different sized flow structures, which was hypothesised to be due to the different characteristic lengths of the particles in their principal directions. This morphological difference also gave rise to differences in the drag effects experienced. There is a coupled relationship between the motion of the particles and the turbulence field local to them, and this is something that could be investigated further in future work, i.e., to what extent do certain collision and agglomeration events create certain flow conditions? For example, if two particles collide such that they induce significant rotational motion, does this therefore induce a certain kind of vortical structure on average – and does this relate to the axes of the particle, since this geometry dictates the size of the wakes following the rotating particle? The coupled interaction here will lead to interesting changes in the local turbulence field, and the role that collisions and agglomeration will have on this turbulence modulation would offer new insight to the field.

There was seen to be a tension between the influence of the turbulence field and the influence of relative orientation: as the turbulence level increased, the fluid dominated, and orientation became less important to determining interaction outcomes in general, which was evidenced by the distributions of particle symmetry axes and collision points

compared over increasing turbulence levels. Particles became reliant on having low kinetic energy in order to experience agglomeration. Needles were observed to still demonstrate a massive overrepresentation towards their maximal configuration, whereas the disks were seen to reduce their reliance on maximal configuration. This was suggested to be because the reduced sample size of the agglomerated distribution in comparison to the bouncing distribution, lowering the confidence of overall trends particularly in the higher turbulence level boxes where the effects were observed.

Over the full parameter space, needles were shown to favour agglomeration the most followed by spheres and then disks, however the extent of this difference reduced as the turbulence level increased. Increasing the turbulence level was seen to allow the spheres to take the lead in some measures of agglomeration propensity, this was due to their reduced drag keeping the kinetic energy of the particles in low enough regimes to sustain agglomeration. In particular, it was shown that in the highest turbulence level box, the agglomeration rate of spheres was equal to that of needles, which indicates the advantage the spheres had in terms of drag was significant enough to reach an inflection point where the strength of the needles' van der Waals contribution can be overtaken.

Returning the fluid domain to a non-dimensional box of length  $2\pi$  could be the next step. This would further generalise the results and allow for easier comparison with other systems, as well as being a more general framework that can be quickly adapted and reproduced. In achieving this aim, however, there is the need to non-dimensionalise the particle-phase such that the relationship between the particle- and fluid-phases remains the same. Another clear next point of study is to introduce many particles into the system which will generate insight on the larger scale structures formed. This would remove the need to start the particles very close by, which ultimately reduced the quality of the results obtained. Nevertheless, a detailed analysis was able to draw out some key effects on the relationship between agglomeration and turbulence. Throughout this chapter, suggestions were made on effects that could be studied in more detail. For example, the effects of agglomerated particles on the alteration to the turbulence field as a result of the dynamic collision and agglomeration interactions. Also, the alignment and dynamic motion characteristics of agglomerated

shapes sitting within a turbulence field could be further investigated and contrasted to see how these quantities evolve over time, as agglomerate sizes grow.

### 6 Multi-Particle Soft-Sphere Agglomeration and Behavioural Modification Effects

#### 6.1 Introduction

The start of this chapter continues to build on the story so far, with the goal of bringing the thesis to a logical and satisfactory conclusion. This is realised by delivering insight to the nuclear systems underpinning the work, in the form of non-spherical behavioural modification studies. As such, the present chapter begins with use of the established methodology (i.e., binary hard-sphere collisions) but later expands on this, with the significant development process and achievements discussed. In particular, the second part of the chapter demonstrates the extension of the binary particle models to be capable of robustly treating multi-particle agglomeration and collisions in turbulence. This is achieved through:

- 1. The development and implementation of a non-spherical soft-sphere collision scheme;
- 2. The extension of the code and algorithms to a multi-particle framework;
- 3. MPI parallelisation of the most expensive particle routines within the code;
- 4. A variable time-stepping scheme that allows resolution of the van der Waals forces without making computational requirements prohibitively high.

Before this is introduced, the chapter begins by presenting results generated using the original immersed boundary-direct numerical simulation framework, which is used here to study binary interactions between pairs of spherical and spheroidal particles. Now, an in-depth sensitivity analysis takes place to investigate the relative importance of the identified key parameters.

As introduced and motivated in Chapter 1, the idea behind such sensitivity analyses is to understand the ways in which systems of agglomerating particles in turbulence can be controlled, with a view to improved process and equipment design. Exploiting
the behaviours of the system to achieve a desired flow result is termed a *behavioural modification* and it is one of the primary goals of this section to generate insight in this direction. This has great value for the nuclear waste processing industry, where certain processes involving turbulent multiphase flows operate with extreme caution due to a lack of understanding of the complex interactions therein. It is argued that complete understanding of these systems begins at the smallest and most fundamental scales, and that is the contribution of the present work.

# 6.2 Behavioural Modifications for Binary Particle Interactions

At the University of Leeds, there has been significant investigation into particle-laden flows, with a recent emphasis placed on behavioural modifications (Mortimer et al., 2021; Wolde, 2023), delivering valuable insight for the nuclear industry. This modelling work has typically been approached using a particle-unresolved Lagrangian particle tracking (LPT) approach for spherical particles, or DNS-IBM calculations between spheres. Two key objectives therefore arise from this. The first is to continue to generate fundamental insight into the agglomeration process at the finest scale. The outputs can be used as detailed information for the development and improvement of the macroscopic models inherent to the LPT code – this objective has already been partly realised in Chapters 4 and 5, and it is furthered here. The second objective is to understand the impact of non-sphericity on the conclusions previously made about behavioural modifications. As motivated, the spherical assumption predominates the particle-laden flow literature, including the agglomeration literature, and so the findings are wider reaching beyond these two objectives.

To give an overview of the recent findings, Mortimer et al. (2021) reported that the coefficient of restitution and the level of turbulence were the most influential areas when studying binary spherical interactions with structure-resolved particles in turbulence. The Hamaker constant was seen to mildly increase the system's relative velocity, but it was not seen to greatly increase agglomeration despite this parameter governing the attractive strength of the interactions.

In structure-unresolved LPT simulations, with millions of interacting particles in a turbulent pipe-flow, Wolde (2023) reported the coefficient of restitution had the most

significant influence on agglomeration, noting in particular that the case of e = 0.2 was the most agglomeration-favouring configuration of those studied. This is consistent with the structure-resolved simulations. However, the two studies differed for the case of Hamaker constant. The LPT studies were shown to be largely dependent on Hamaker constant, with agglomerated clusters of a greater particle number forming earlier into the simulations as the Hamaker constant was increased. It will be of interest here to try and reconcile this difference through careful investigation of the present simulations. The studies of Wolde (2023) also demonstrated a strong reliance on Reynolds number, comparable with variations in the level of turbulence as performed by Mortimer et al. (2021).

Table 6.1: Simulation matrix for the behavioural modification studies of Hamaker constant *A*, Taylor Reynolds number  $Re_{\lambda}$ , and coefficient of restitution *e*. Each column represents a unique set of ~50 simulations.

	Collisional Energy		Attractive Strength		Turbulence Strength				
	Low	Mid	High	Low	Mid	High	Low	Mid	High
Re <sub>λ</sub>	2.5	2.5	2.5	2.5	2.5	2.5	2.5	12.1	22.2
A (zJ)	22.3	22.3	22.3	7.84	22.3	36.37	22.3	22.3	22.3
е	0.2	0.4	0.6	0.4	0.4	0.4	0.4	0.4	0.4

For this reason, the chosen parameters to vary in the present work are the Hamaker constant and the coefficient of restitution, with a brief word also on turbulence level. The results aim to bring out the interaction behaviours and dynamics, through comparison of the system probability distributions for all three morphologies. This is performed in the lowest turbulence level box introduced in Chapter 5 because, pragmatically, this reduces the overall influence of the turbulence field and thus makes it easier to extract the particle-phase parameter's effects, reducing the number of required simulations. The simulation set-up remains exactly the same as the previous chapter (Table 5.5), but the aforementioned parameters are varied. For completeness, the turbulence levels of the box are discussed briefly in the context of behavioural modification but no further simulations were conducted varying the turbulence levels. The simulation parameters varied are shown below in Table 6.1. The chosen values

are consistent with Mortimer et al. (2021) for comparative purposes, except for the turbulence level. This is because the comparison across turbulence analyses the work done in Chapter 5.

#### 6.2.1 Behavioural Modification I: Collisional Energy

One way that particle interactions can be effectively mediated is through the amount of energy, or momentum, retained post-collision – which Mortimer et al. (2021) and Wolde (2023) demonstrated the efficacy of in the context of spheres. In industry, such a parameter could be controlled, for example, through the coating of particles to decrease the amount of energy retained, through the addition of salt to the fluid (Shahidzadeh et al., 1998), through altering the system temperature (Bergström, 1997) or modifying the viscosity of the carrier fluid, and so it is a viable control option.

This is controlled in the simulation code through the parameter e, the coefficient of restitution. This indirectly mediates the kinetic energy retained by the pair of colliding particles by enforcing that the particles retain a proportion e of their pre-collision relative velocity. The default value for calcite particles is e = 0.4 or 40%, which is equivalent to the particles retaining 16% of their kinetic energy post-collision.

It has been demonstrated in the cited work that collisions are essential in promoting agglomeration in the present flows. Further established in Chapter 5, collisions between particles generate a much more favourable energy state for agglomeration, with particles remaining close by, and with lower velocity. It was shown that off-normal and off-centre collisions between non-spherical particles were sometimes effective in promoting agglomeration even though the van der Waals forces were not maximised in the orientational configurations of the initial collision, since these interactions converted much of the total kinetic energy to rotational motion, inadvertently causing the particles to stay relatively nearby and thus the likelihood of a secondary and tertiary collision is increased at more favourable surface locations. In such scenarios, the collisional energy loss can compound. It will be of interest to investigate how this effect changes with different values of e.

Presented below are the probability density functions for three values of the coefficient of restitution,  $e \in \{0.2, 0.4, 0.6\}$ , or equivalently a 20%, 40% and 60% retention of the collisional velocity state. It is worth noting that this is enforced at the direct point of collision on both surfaces, which is not the centre of the particle as in spherical LPT

codes for example. This allows the transfer of energy from translational to rotational motion as described, and offers more detailed insight on the full collisional interplay.

The results are broken down into the respective morphologies for comparative purposes. Firstly, the spheres are presented in Figure 6.1, followed by needles in Figure 6.3 and the disks in Figure 6.4. Exactly as with the simulations presented in Chapter 5, Monte Carlo sampling takes place over uniformly randomly distributed orientations and injection locations of particles in a statistically stationary box of homogeneous, isotropic turbulence.

For the collisional energy study here, the case of e = 0.4 was already investigated in Chapter 5 and so simulations only needed to be conducted for the higher and lower cases either side of this. The same is true of the attractive strength in the second behavioural modification study, where A = 22.3 zJ was already simulated, leaving the other two cases to be investigated. Around 50 simulations were conducted to populate the sample space for these cases, for each of the three morphologies, leading to around 600 further simulations. Each simulation used 24 processors for approximately 12 hours.

I.I: Spheres



Figure 6.1: Probability distributions of separation distance (left) and relative velocity (right) between spheres interacting in the viscous sublayer turbulence box, contrasted for three values of the coefficient of restitution.

The results displaying the separation and relative velocity tendencies of the spherical particle pairs are shown in Figure 6.1 contrasting the three coefficient of restitution values. The spherical system naturally has fewer degrees of freedom and so the trend

would be expected to be straightforward compared to the non-sphericals, as long as enough instantiations (individual simulations) are taken to fully sample the turbulence – which was ensured by waiting for the PDFs to remain smooth and stable under addition of new data (approximately 40-50 simulations).

By observing the location of the peaks of the distributions, it is seen that the expected trend is followed. Namely, the lowest coefficient of restitution leads to a peak closest to zero separation. This is very closely followed by the moderate coefficient of restitution, which has almost identical shape as the lowest case, but with the peak shifted slightly. Finally, the higher coefficient case leads to a noticeably shifted distribution to the right. The lowest two cases peak much closer to zero, in the region of separation that implies agglomeration; whereas, the higher coefficient case has a peak away from zero and outside of the van der Waals effective range, implying fewer instances of agglomeration.

The agglomeration outcomes are enumerated in Table 6.2 where the spheres are seen to have a very high agglomeration rate for the lowest coefficient of restitution. In line with Wolde (2023) and Mortimer et al. (2021), agglomeration outcomes are strongly dependent on the retained velocity in the collision. As was described in Chapter 4, the agglomeration process tends to be a function of the post-collision velocity, with particles only becoming agglomerated, or trapped in the DLVO potential well, *after* experiencing a collision.

Result	e = 0.2	e = 0.4	e = 0.6
Bounced	19	35	38
Agglomerated	21	15	12
Agglomeration Rate	52%	30%	24%

Table 6.2: Agglomeration data for the spherical particles, with varied coefficient of restitution.

A secondary effect that is visible in the PDFs and should be stated is the following: in the hard-sphere model, it was established that agglomerating particles are accelerated towards each other at very narrow separations, they then lose that accumulated energy in the collision, before regaining some of it from the DLVO contribution – this

occurs continually and cyclically. Simply, the particles gain kinetic energy from the DLVO potential and then lose it again via a collision. This was shown to create a very small vibrational state of quasi equilibrium where particles became 'agglomerated', trapped in the potential well, but constantly exchanging energy. In this state of agglomeration, the velocity and the separation take on the form of a decaying oscillation, and that will be shown again in Section 6.3.4. It can be noted that increasing the coefficient of restitution would allow particles to separate their surfaces further, before being recaptured by the van der Waals forces. Firstly, this naturally gives them much greater chance of escaping agglomeration, which is consistent with the physical expectations. But even in cases where particles agglomerate, it will slightly increase the mean separation distances observed in the plots. The same argument leads to the same outcome in the relative velocity graphs. What this means is that the minimum of the energy potential resolved in the simulations, or the mean position that the particles sit within that well, could be thought of as shifting based on coefficient of restitution, which is a computational effect, not a physical one, but this will have an effect on agglomeration resolution.

The plot for relative velocity tells a very similar story to the separation plot, with the two lowest coefficient of restitution (COR) cases gravitating heavily towards zero relative velocity, and the higher COR case seeing a peak shifted away from zero, again implying less agglomeration in the latter case. These two plots give the impression that for a given set of parameters there is a cut-off value of the COR, just as there was with velocity, where the system moves from strong agglomeration favourability to the alternative. This stands to reason, since it was established in earlier chapters that the post-collision velocity is a key metric determining agglomeration. However, it does mean that varying COR linearly will not necessarily result in linear changes to agglomeration variables like dx, dv or agglomeration rate, as the system at some point appears to move from one attractive state to another less attractive state, in a somewhat binary way. That is implied by the PDF plots, however the table of results does not necessarily support this, since there is a seemingly gradual decrease in agglomeration rate. Thus, the cases of needles and disks can be analysed with this hypothesis in mind.



Figure 6.2 Probability distributions of van der Waals force measured between spheres interacting in the viscous sublayer turbulence box, contrasted for three values of the coefficient of restitution.

The graph of van der Waals forces, shown in Figure 6.2, does not depend on the coefficient of restitution. Different agglomeration outcomes allow slightly different samplings of the two main states but the shape of the distribution remains the same, because from the point of view of the full physical system, the van der Waals forces are quite binary for spheres. Either particles agglomerate and they experience in-and-around the maximal force. Or, they rebound and separate, leading to zero van der Waals contribution quickly. One conclusion here may be that a binary 'on-off' agglomeration model may be a good approximation for spheres, but not for non-spheres.



Figure 6.3: Probability distributions of separation distance (left), relative velocity (right) and van der Waals force (bottom), measured between needles interacting in the viscous sublayer turbulence box, contrasted for three values of the coefficient of restitution.

The PDFs of the needle interactions shown in Figure 6.3 follow the same trends as the spheres; in particular, a very strong agglomeration favourability is seen for the lowest COR case, which falls away as COR is increased. The relative velocity graph implies a similarity between the two highest COR cases, and this is supported by the agglomeration rate data shown in Table 6.3 where the two highest COR cases have almost identical agglomeration rates. The lowest COR case separates itself with a sharp increase in agglomeration rate, which gives credence to the idea that beyond a certain value, the systems become similar.

Lastly, the graphs of the van der Waals forces experienced remains similar for all values, as was the case with the spheres. Noticeably, the graphs for the needles are more uniform, without two clear peaks, because there is a wider spectrum of van der

Waals forces that can be experienced by needles in an agglomerated state due to the orientational considerations in the force model. It would appear this space of possible van der Waals forces is sampled quite evenly under uniformly random orientations.

Result	e = 0.2	e = 0.4	e = 0.6
Bounced	29	65	33
Agglomerated	21	32	15
Agglomeration Rate	42%	33%	32%

Table 6.3: Agglomeration data for the needles, with varied coefficient of restitution.

I.III: Disks



Figure 6.4: Probability distributions of separation distance (left), relative velocity (right) and van der Waals force (bottom), measured between disks interacting in the viscous sublayer turbulence box, contrasted for three values of the coefficient of restitution.

Figure 6.4 presents the results for the disks, which do not follow the same obvious pattern. The lower two COR cases resolve very similar PDFs of separation and of velocity, with the highest COR case showing a mildly reduced peak of both separation and velocity, but a similar trend overall is observed. This is backed up by the agglomeration results presented in Table 6.4, which do not indicate a strong reliance on COR.

One interesting difference in the relative velocity plot is that different velocity regimes are dominated by different COR values. The range  $dv \in [0, 5] mms^{-1}$  is most sampled by e = 0.2, the range  $dv \in [5, 10] mms^{-1}$  is most sampled by e = 0.4, and e = 0.6predominates well beyond that, implying that the relative velocity dependence breaks itself into regimes. In other words, to reliably access the lowest values, one requires a low COR and to reliably access the highest velocities in the simulation one requires a higher COR. This is to be expected. Strangely, this does not translate particularly well to the agglomeration outcomes shown in Table 6.4.

Result	e = 0.2	e = 0.4	e = 0.6
Bounced	40	73	49
Agglomerated	10	23	11
Agglomeration Rate	20%	23%	19%

Table 6.4: Agglomeration data for the disks, with varied coefficient of restitution.

To round off the analysis, the relative orientations of the particle symmetry axes were again measured for the two non-spherical morphologies, shown in Figure 6.5. Both the needles and disks appeared to show very little dependence on COR when it came to observing the relative orientations of the particles. The distributions for both morphologies seemed to be almost identical for the two lower COR values, and similar also for the highest case. In the case of the needles, the highest COR seems to very marginally favour parallel alignments compared to the other two COR values, whereas the disks appear to show a slight widening of the variability when increasing COR to the highest value. A higher coefficient of restitution, as established, allows for more energy to be retained post-collision. This also means more rotational kinetic energy can be retained post collision (and more can also be induced by the collisions).

Therefore, this could facilitate a wider range of orientations being traversed over the same simulation period. This can possibly explain the wider variability in the highest coefficient of restitution case. Higher energy collisions perhaps create certain orientational states more readily than low energy collisions.



Figure 6.5: Probability distributions of relative orientation, measured between needles (left) and disks (right) interacting in the viscous sublayer turbulence box, contrasted for three values of the coefficient of restitution.

The statistical breakdowns for the agglomeration efficacy were presented in the tables, reflecting roughly what was observed through the PDFs. The key takeaway is that as the COR is increased, more energy is retained in the collision and thus, over the sampled space, the proportion of agglomeration events is reduced. Between the three morphologies there are different disparities when jumping between different values of e, which perhaps relates to different cut-off values depending on morphology, beyond which the systems move from agglomeration favouring to not favouring.

At a low COR, i.e., 0.2, the spheres agglomerate at a very significant rate, despite being immersed in turbulence, and the needles also showed a strong favourability towards this value. This implies a strongly agglomeration-favouring system for both morphologies. The trend was not repeated for the disks, which exhibited similar agglomeration outcomes over the full range of COR values. Perhaps because the agglomeration rate is lower for the disks in general, the differences are not quite so pronounced.

#### 6.2.2 Behavioural Modification II: Strength of Attraction

The second parameter under investigation is the Hamaker constant, which mediates the strength of the attractive van der Waals contribution. Wolde (2023) showed this to

be the most influential parameter within the LPT framework adopted in that thesis. The work of Mortimer et al. (2021), which was the starting point for the present work, was inconclusive on this, and so it is of interest to investigate the parameter further.

There is a wide array of parameters that can be varied within the DLVO framework, particularly for the many parameters that make up the electric double layer contribution, but those were shown to be less significant in previous works. Particularly given the baseline parameters used here, the electric double layer contribution is negligible compared to the van der Waals contribution so it is sensible that the Hamaker constant should be the most influential parameter in the DLVO model, and it would be expected to mediate agglomeration outcomes. Practically, modifying the ionic strength of the carrier fluid can be used to affect changes in the van der Waals forces, e.g., the addition of salt reduces such forces.

As with the collisional energy, three values are chosen in accordance with Mortimer et al. (2021), which includes the base case stemming from the viscous sublayer results in Chapter 5. The values investigated are  $A = \{7.84, 22.3, 36.37\} zJ$ .

#### II.I Spheres

Figure 6.6 shows the effect of varying the Hamaker constant on the resolved separations, velocities and van der Waals forces. Clearly, observing the bottom plot in Figure 6.6, the van der Waals forces resolved in the simulation are strongly dependent on the varied parameter, as would be expected. The peaks of the van der Waals distributions follow the changes in Hamaker constant in a relatively simple and linear way.

The distributions of separation and relative velocity follow a similar pattern to one another, for the spherical case. The highest Hamaker constant scenario, A = 36.76 zJ, separates itself from the lower value of Hamaker constant, A = 7.84 zJ, and the moderate value of A = 22.3 zJ. The lowest two cases are very similar, peaking at the same value of separation distance and almost lying on top of one another. However, the highest Hamaker constant scenario has a greater concentration of low velocities and of low separation distances, indicating a greater agglomeration propensity.



Figure 6.6: Probability distributions of separation distance (left), relative velocity (right) and van der Waals force (bottom), measured between spheres interacting in the viscous sublayer turbulence box, contrasted for three values of the Hamaker constant.

This implies again that there may be a threshold value beyond which agglomeration propensity sharply increases, as the lower two cases seem to fit into a category separate from the higher case, as shown by the results in Table 6.5. This stands to reason, as the spheres are prescribed the same initial velocity in all simulations, and there is a cut-off value for velocity (as shown in Chapter 4) beyond which agglomeration no longer occurs. As the Hamaker constant changes, this cut-off value would be expected to change due to the DLVO potential changing (more kinetic energy is required to escape the well). The turbulence alters the particle velocity at collision somewhat and so there is a distribution of collisional velocities about the initial value. However, given the present set-up, it still centres on a value strongly related to the initial velocity given to the particles. Therefore, if the Hamaker constant is to change enough between two cases, such that the initial velocities go above or below a

threshold cut-off velocity, then it would be expected to see a drastic change in simulation tendencies from one case to another.

Result	<i>A</i> = 7.84 <i>zJ</i>	A=22.3 zJ	A=36.76 zJ
Bounced	30	35	26
Agglomerated	12	15	22
Agglomeration Rate	29%	30%	46%

Table 6.5: Agglomeration data for the spheres, with varied Hamaker constant.

II.II Needles



Figure 6.7: Probability distributions of separation distance (left), relative velocity (right) and van der Waals force (bottom), measured between needles interacting in the viscous sublayer turbulence box, contrasted for three values of the Hamaker constant.

The separation and relative velocity distributions shown in Figure 6.7 again demonstrate a binary split in the behaviour of the respective cases. The higher two

Hamaker cases strongly favour agglomeration whilst the lower case lags behind slightly. In Table 6.6, the agglomeration outcomes show a fairly consistent increase in agglomeration rate with Hamaker constant, rather than the binary nature suggested by the plots.

Result	A=7.84 zJ	A=22.3 zJ	A=36.76 zJ
Bounced	36	65	28
Agglomerated	12	32	20
Agglomeration Rate	25%	33%	42%

Table 6.6: Agglomeration data for the needles, with varied Hamaker constant.

II.III Disks



Figure 6.8: Probability distributions of separation distance (left), relative velocity (right) and van der Waals force (bottom), measured between disks interacting in the viscous sublayer turbulence box, contrasted for three values of the Hamaker constant.

The separation of the disks follows a very similar pattern for all three values of the Hamaker constant, as displayed in the left graph of Figure 6.8. In the other two plots, relative velocity and van der Waals force, comparatively little agglomeration propensity is shown for the lowest value of the Hamaker constant. The sampling of non-zero van der Waals forces is particular low compared to both the needles and the spheres. In the lower plot of Figure 6.8, displaying the van der Waals forces for the disks, the distributions no longer follow the same shape as the needles. Rather, the distribution appears more bimodal, particularly for the lowest Hamaker case. This implies that agglomeration is being achieved only by a narrow band of van der Waals force cases, which will correspond to the maximal configurations. In the other cases, very little van der Waals interaction takes place. The peak is greater at the lower end of the force distribution for A = 7.84zJ since fewer agglomeration events can be achieved and thus the state of van der Waals attraction remains low throughout. This is supported by the lower agglomeration rate reported in Table 6.7. In these lower Hamaker systems, the effect of DLVO forces is very much reduced, and the system is moving towards one where exclusion of such forces may be a reasonable modelling assumption. Notably, this was not the case for the spheres or needles at the same values of attractive strength.

Result	A=7.84 zJ	A=22.3 zJ	A=36.76 zJ
Bounced	42	73	38
Agglomerated	6	23	12
Agglomeration Rate	12.5%	23%	24%

Table 6.7: Agglomeration data for the disks, with varied Hamaker constant.

Finally, the relative orientations of the particle symmetry axes are compared for the needles and disks in Figure 6.9. The needles (left) are shown not to change their preferred relative orientation based on the system's Hamaker constant. In contrast, the disks (right) appear to increase their variability as Hamaker constant is reduced. A similar argument can be made here as with the coefficient of restitution study where a similar trend was observed: the lower Hamaker constant will generally equate to more kinetic energy retained post-collision, which means more rotational kinetic energy and

a wider sampling of the rotational sample space in a given simulation. Why this would only affect disks to this extent is not clear.

There is an interesting distinction drawn here between disks and needles, where the relative orientations of the disks appears to be more sensitive to the system parameters in general than the relative orientations of the needles.



Figure 6.9: Probability distributions of relative orientation, measured between needles (left) and disks (right) interacting in the viscous sublayer turbulence box, contrasted for three values of the Hamaker constant.

#### 6.2.3 Behavioural Modification III: Strength of Turbulence

For completeness, turbulence level should be commented on in the context of behavioural modification, since it was clearly shown to be influential in the previous chapter, and the work of Mortimer et al. (2021) and of Wolde (2023) highlighted this variable as having strong influence. The turbulence level relates quite straightforwardly to system control as well, since it can be controlled directly through the pumping force, or choice of pipe geometry, for the specific application of waste processing, for example.

The previous chapter highlighted that the separation and velocity tendencies of the disks, needles, and spheres were directly correlated to the turbulence level. Lower levels of turbulence led to lower particle kinetic energy and thus agglomeration favourability. The way this changed with morphology was quite stark. In particular, the spheres were able to handle the increasing turbulence much better than the non-spherical particles, retaining similar agglomeration levels as turbulence increased. This was said to be because of the reduced drag characteristics of a sphere, reducing acceleration due to the flow. This is supported by the results of Table 6.8 where the

overall agglomeration rate is seen to be relatively independent of the turbulence. A similar trend is seen for the needles, although the full PDFs indicated that the needles were generally experiencing less agglomeration-favouring conditions in contrast to the spheres. Lastly, the disks are seen to reduce their agglomeration rate significantly as a function of turbulence. Thus, from the point of view of process design, understanding the morphology of the typical particle shape (through characterisation) can help to inform what choices to make on pumping/extraction conditions.

 against increasing turbulence level.

 Result
 Needle
 Disk
 Sphere

Table 6.8: Collated data showing the agglomeration rate for the three morphologies

Result	Needle	Disk	Sphere
$Re_{\lambda} = 2.5$ Agglomeration Rate	32%	23%	30%
$Re_{\lambda} = 12.1$ Agglomeration Rate	34%	18%	32%
<i>Re</i> $\lambda$ =22.2 Agglomeration Rate	27%	13%	27%

#### 6.2.4 Behavioural Modification Summary

In summary, the system was seen to be dependent on coefficient of restitution for all three morphologies, which could be addressed in industry through the coating of particles or modifying the viscosity of the carrier fluid. The effects were less pronounced for the disks. The resolved van der Waals forces were seen to be independent of COR and so it was the collisional energy exchange that caused particle agglomeration to increase for reduced COR. In particular, particles were able to remain close by for longer, giving the DLVO interaction more opportunity to form an agglomerate.

The Hamaker constant was also found to be influential on the system, particularly in terms of the van der Waals forces experienced over the simulation, as would be expected. Agglomeration rate increased with Hamaker constant for all three morphologies, but each system displayed a binary behaviour where two of the three cases would have similar values of agglomeration rate, and the other case would separate itself on this metric. A wider sampling of the different values of Hamaker constant could help understand if this binary nature is a trend over the full span of values.

Confirmation of the dependence on Hamaker constant is consistent with Wolde (2023) and builds upon the work of Mortimer et al. (2021) where the findings were inconclusive. In particular, the Taylor-Reynolds number studied in the base case ( $Re_{\lambda} = 2.5$ ) here is much lower than the latter mentioned work. Hence, the conditions in the vicinity of the particle interaction are less chaotic and perhaps lead to longer more stable interactions such that the Hamaker constant can have a greater effect – which also feeds back into the discussion on turbulence.

Applying this insight to industrial considerations, the system is clearly sensitive to the Hamaker constant and so experimental measurements of the industrial system should be geared towards characterising this parameter to get a prediction of the system behaviour. If it could be confirmed that there is a cut-off value across which the agglomeration propensity of the system changes widely, as indicated by the present results, this could be a useful finding for understanding the likely behaviour of the system.

For both mentioned parameters, the distributions of relative orientations of the disks appeared to increase their variance for the cases that implied reduced energy loss, an effect which was not present for the needles.

Lastly, turbulence was discussed as influential but its relative effect was seen to be strongly dependent on morphology, with disks exhibiting the greatest dependence on this parameter. This is in contrast to the other two behavioural modifications, upon which disks were seen to have the least dependence. This was attributed to the disks overall having lower agglomeration rates, thus trends were more difficult to extract.

### 6.3 Development of a Multi-Particle Agglomeration Framework for Non-Spherical Particles

#### 6.3.1 Background and Motivation

Throughout the thesis, the story has been developing towards the notion that more particles are needed to be studied at once to further expand on previous conclusions. Detailed insight has been gained on the pair interactions of non-spherical particles, which has established the vocabulary and the lens through which to assess systems of many particle-particle and particle-fluid interactions, but to extrapolate findings to a larger scale system there needs to be the multitude of simultaneous competing forces whose coupled interactions will undoubtably give rise to interesting and unexpected non-linear phenomena. Emergent behaviours can be observed for example in the fluidisation of particles (Gan et al., 2016a), the flow's changing deposition or sedimentation characteristics (Wolde, 2023), or the formation of crystal-like structures in the flow (Ochsenbein et al., 2015). Moving to a multi-particle framework is not as simple as just changing the number of particles, however, and there is a myriad of challenges to overcome. A few are outlined below.

The general observation can be made that by introducing agglomeration effects into the multiphase model, one naturally increases the particle loading in localised regions of the flow. So, even with what might be termed a 'dilute' suspension, there may still be many locations where many multiples of particles interact at once, if the suspension is self-attracting. It is well documented within the literature that soft-sphere collision models are better able to deal with the type of physical scenario (Ma et al., 2022) that involves high particle loading where collisions are much more frequent. This is because soft-sphere models can handle multiple colliding particles at once, where hard-sphere models struggle. Multiple collisions involving the same particle means that the way the velocity is updated breaks down for the standard hard-sphere implementation, where collisions necessarily are assumed binary and infinitesimal. The present simulation code, as with many discrete element method (DEM) and LPT solvers, visits particle pairs sequentially, meaning there is an implicit assumption around the ordering of collisions, which may be untrue and erroneously alter the physics. A similar limitation is true of methods that visit particle pairs through more complex frameworks like tree searches or multigrid searches (which are used to reduce the compute time for high particle numbers). Such an assumption may even prevent some particles colliding, as the velocity vectors are updated by a former collision. This was observed in the development process herein, where the first attempt at multi-particle interactions was made using a hard-sphere approach.

There are ways to order the collision pairs to circumvent this. One example would be working out the order of collisions based on velocity and looping through the pertaining list. Alternatively, one could design an algorithm that iteratively reduces the particle timestep separately from the fluid timestep until no two collision partners meet the collision condition in the same step. Such workarounds exist but can become unviable when the number of interacting particles is high. There, the edge case being treated becomes the norm. There is the possibility that the simulation timestep eventually needs to be reduced to below the collision duration, at which point the model assumptions are weakened. It is typical in the literature to ensure a collision lasts for at least O(10) timesteps (Costa et al., 2015) such that the application of the force is smooth, ensuring the correct energy transfer.

Arguably more important than the above is the strict no overlap condition that is placed on the simulation to ensure stability. If there is no means to compute overlap depth, then the slightest overlap can cause problems in the simulation, and this was also observed first-hand. To reiterate a previously described problem, the collision detection (common normal) implementation strictly requires that surfaces do not overlap, otherwise the in-built assumption about the convexity between interacting non-spherical surfaces is no longer true, and the optimisation algorithm cannot converge. There is no longer one unique point on the respective surfaces where the 'common normal' criteria is satisfied. When this happens, the closest distance vector is not uniquely defined and so the DLVO and collisional algorithms present unphysical and sometimes seemingly random values, and in some instances diverge when not treated.

It was shown in the previous chapter that even for two particles this poses a problem when one also considers attractive forces like those arising from the DLVO potentials because the attractive forces lead to many repeated collisions in a narrow band of timesteps, and small numerical errors were able to overcome the collision detection distance in rare instances, compromising the robustness of the solver. This is worsened because the minimum of the DLVO potential (Chapter 3.15) lies at zero separation, or in other words the interparticle forces *want* to make the surfaces touch, and the collisional implementation must act in opposition to ensure they never do. One can imagine then that a small perturbation to this balance can lead to overlap in rare cases when particles are agglomerated and are thus 'colliding' hundreds of times. As mentioned, this was observed in rare instances in the binary particle case, and it was further observed with high particle loading in the first efforts at a multi-particle framework.

Everything described up to this point in the thesis, in terms of the methodology, is theoretically applicable to a multi-particle framework; however, when the particle

number increased, the non-overlap problem led to a soft-sphere approach being sought. The other consideration here is that the advantage of the hard-sphere method in reducing the number of computations (as it is an instantaneous model) is almost made redundant in the case of agglomeration, where collisions become so frequent that they need to be computed every timestep anyway.

An alternative to the soft-sphere approach not pursued here is the potential to identify agglomeration as a change of state (*from a simulation perspective, not in terms of matter*) and to fix the agglomerating particle meshes together for the remainder of the simulation, or until enough energy is supplied to allow their separation. This would take place once the first hard-sphere-DLVO exchange has completed, and resulted in an agglomeration event. This would rely on a dependable way of determining an agglomerated particle pair, which the results of Chapter 5 lay a strong foundation for achieving. This idea is similar to that of the types of models discussed extensively by Almohammed (2018) where particles either agglomerates would also be resolved, as would the initial time-dependent DLVO interaction that leads to agglomeration.

This previous discussion motivates the undertaking of the work in this chapter, where the overlap problem is successfully addressed. In fact, the depth of overlap is used to define the strength of the penalty force incurred and this acts as a robust way of treating particle collisions. This is the essence of the soft-sphere approach introduced in Chapter 3: Section 2.10. A spring model which is a function of the overlap depth  $\delta$ , acts to separate particles over time, whilst the energy loss is achieved as before through a coefficient of restitution, which is implemented through a dashpot model, operating as a function of relative velocity dv. This forms the spring-dashpot implementation.

Tuning the model coefficients (Equation (3.67)) correctly can be a fine art, as this will affect whether the correct dissipation of energy is resolved, as well as whether particle forces remain stable and realistic. The main parameter to control is  $N_c$ , the number of timesteps per collision, but this needs to be weighed against the characteristic velocities of the particles and the simulation timestep to ensure collisions are being resolved appropriately.

The literature on non-spherical soft-sphere algorithms and theories is sparse in the context of multiphase systems, and thus an area where there is scope for contribution. Most implementations opt to convert or reduce the problem to a spherical one, which is done in a variety of ways (Ardekani et al., 2016). As outlined in the literature review, the spherical case allows for much simpler implementations, so trying to convert the problem to a spherical one is a sensible approach. Some efforts use a multi-sphere method (van Wachem et al., 2015), where the algorithm fills non-spherical particles with an inscribed number of spheres which roughly define the surface to compute collisions. This removes the need to complexify the mathematics; instead, the computational time increases, with multiple theoretical spherical particles now approximating one non-spherical particle. Ardekani et al.'s (2016) approach detects the collision through 'moving balls' and uses ghost spheres with a scaled radius based on the local curvature to approximate the collision physics. Wynn (2008) published a useful paper that attempts to make some of the non-spherical treatments of such local curvature scaling robust, as did Zheng et al. (2013) who utilised full simulations of structural finite elements to give credence to their proposals.

As highlighted in Chapter 2, some more detailed approaches are also seen in the literature that use a geometric potential algorithm (Dziugys and Peters, 2001). This is sometimes done in DEM frameworks, which provides a means to locate the colliding surface points by utilising a potential function. The physical accuracy of this model is contested, since it does not satisfy a common normal constraint and is known to deviate from common normal approaches (Ma et al., 2022).

The key problem that arises in non-spherical soft-sphere models is that a relatively simple set of assumptions can become complex quite quickly due to the many degrees of freedom involved in defining the overlap problem continuously and rigorously for arbitrary orientations. For example, a reasonable assumption for computing the force at a single point (using a single point is a standard in-built assumption) would be to say that the *deepest points of overlap* (of the penetrating particle past the opposing surface) are the requisite points for defining the strength of the collisional force. Defining the 'deepest' overlap points is not especially rigorous for complex shapes, and once one converges on those points by some chosen definition, one is left with a pair of arbitrary points on the respective surfaces that must then be used to define a single shared 'collision point' or 'contact point', somewhere between these two surface

points. Again, there are a multitude of ways to do this, and different implementations for each. One option is to select the midpoint between the surface points. This provides a simple means to define a contact point but for two particles of very different sizes and shape in the vicinity of the contact point, it is questionable how accurate this treatment is. It could be more appropriate to use the centre of mass of the overlapping region, employing a calculation based on the shared overlap volume (which is difficult to calculate precisely and quickly). Then there is the problem of defining the shared normal vector along which the forces act. This can be approximated by surface normal vectors computed at the collision point, or by forming a vector between the two colliding surface points on the non-spherical bodies, but again there is more than one approach, with different connections to the overall physics of the problem.

It could also be appropriate to return to the initial assumptions. Recall, the idea of the soft-sphere model is partly to approximate the time dependent loss of energy through the surface deformation of colliding particles. In terms of the physics, when two arbitrarily orientated ellipsoids present themselves, the collision points on the respective surfaces will also mediate the amount of deformation of the body, i.e. the stiffness will be a function of the surface curvature that presents itself in the collision, which is a function of orientation, and will be different on the respective bodies in general. This was addressed in Equation (3.67a).

There is a further complication that simply having an algorithm that finds the deepest overlap points does not fully address the problem: a test is also required to confirm overlap, as it is important in a dynamic simulation to have a metric that tracks the overlap state of each collision partner. Depending on the mathematical formulation it is possible to get both results at once (Dziugys and Peters, 2001), i.e., overlap state and overlap depth, or to not (Podlozhnyuk et al., 2017). If one does not, there is an associated computational cost, possibly of solving a high-order polynomial (Gan et al., 2016a) to determine the overlap state. This will be calculated prior to conducting an algorithm for defining the collision point, so as to reduce the overall computational requirements of the simulation, as this step will be less costly than solving the entire problem to determine overlap state.

As evidenced then, there is no single accepted solution to the problem and it is an area still requiring much investigation, with many open questions, perhaps more than

it would seem on the surface. The presently available approaches each have their limitations, and the challenges outlined in detail above are done so because the same considerations and limitations were met during the development process of a soft-sphere framework for the present code. It is the first goal of this section to propose a novel approach to the problem that could circumvent some of the computational cost whilst still offering a robust framework for computing particle collisions. Ultimately, two ideas were put forward and tested and the results are discussed here. Upon later finding a robust solution in the literature, this was implemented further to the two proposed models, contrasted with the novel approaches, and ultimately taken forward for the final presented results as a matter of pragmatism.

#### 6.3.2 Development of an Overlap Depth Calculation Scheme

Both novel ideas were introduced in Chapter 3: Section 2.9. They try to make use of the available tools within the code, namely the common normal algorithm and rigidbody framework, paired with assumptions or reformulations that provide a means to the calculation of overlap depth. The first approach, *Method 1*, involved detecting the closest distance points up until the previous definition of a contact, and rather than calculating the hard-sphere collision physics upon detection, the colliding points on the surface were stored in a separate vector  $T_0$ . The coordinates  $T_0$  are updated independently, which effectively tracks the specific location on the surface of the particles that first collided. This relies on the assumption that the initial contact points remain the deepest points of overlap throughout the collision which is a significant assumption but is theoretically well suited to collisions involving small overlaps and convex shapes like present. In simple testing and development, this was observed to be a satisfactory assumption for a wide range of trialled normal collisions. This facilitated a computation of overlap and allowed a soft-sphere force to be implemented for two particles which accurately resolved the collision physics. The first result of this was demonstrated in Figure 3.18.

There were two limitations observed with this approach. The first is that the deepest overlap point assumption becomes less valid when significant torques are induced over the collision period or when the fluid is able to appreciably deflect the orientational configuration between particles as they interact – as both effects disrupted the location of the true deepest overlap points, in relation to the assumed ones. Thus, it can be said that this model is more suited to systems where the collisional timescale is

significantly less than the fluid timescale (which is most multiphase systems studied using CFD-DEM approaches). The quality of the assumption is also a function of the curvature of the body close to the collision point, since for certain configurations, small perturbations to orientation can have a more significant effect on the error.

As motivated, the key objective was to allow for multiple particle collisions to be computed simultaneously. Adding many simultaneous collisions to the simulations unfortunately further reduced the strength of the above assumption, as multiple competing directional granular forces complicate the evolution of the overlaps, which in turn reduces the strength of the idea. Overall, it can be expected that this method works well for binary particle interactions and small overlaps and can be used as a simple approach in dilute systems or to quickly set up a framework for studying softsphere interactions and agglomeration events.

On the plus side, this is an extremely computationally efficient way of calculating overlap in comparison to existing methods. The collision points  $T_0$  are converged upon by the optimisation algorithm (any suitable algorithm can be used), and once calculated they never have to be calculated again for the remainder of the soft-sphere collision, meaning that computational resource is reduced for overlapping particles. The updating of the  $T_0$  vectors can be slotted into the existing implementation for updating positions and orientations, which is already a computationally inconsequential procedure in comparison to the various algorithms that must converge for particle interactions.

Because of this very small cost, this could also be an effective way to generate the initial guesses for other algorithms that use a trial solution to begin their search or optimisation scheme. This could be preferable over other initial guess approaches (e.g., starting with the points intersected by the line joining particle centroids) because it is based in a form of physical reality – at least initially. In codes where the overlap computation dominates the computational load, then this may offer a speed up as a supplementary calculation rather than the main computation of overlap depth.

Further to this approach, there is the second proposed idea, *Method 2*, where the particle surfaces are iteratively shrunk until the common normal algorithm converges, thus satisfying a relaxed version of the common normal constraint. This method is at the opposite end of the spectrum. In theory, it should offer a similar accuracy to the

geometric potential method described earlier and is no less precise of a solution than other typically used methods. However, it requires multiple iterations of an optimisation scheme, so there is the potential for a high computational cost penalty, the deeper the allowed overlap, which will also scale with particle number as more particles generally means deeper overlap – as well as this, particle interactions scale like  $O(N_p^2)$  using a sequential ordering of pairs, so it is of importance to not have overly costly routines for calculating particle interactions. There were also concerns raised in the testing surrounding robustness, given the lack of a complementary criteria to identify overlap in situations where the common normal algorithm erroneously converges. This edge case requires special treatment.

Lastly, there is a third described approach, *Method 3*, which takes an existing method of Podlozhnyuk et al. (2017) and adapts it to the special case of spheroids. Given initial testing, all methods showed promise but it was Method 3 that was taken forward as a matter of pragmatism, given the time constraints on this final piece of work: method 3 has already been proven to be robust. Results comparing all three methods are now presented, illustrating the described findings.

#### 6.3.3 Comparative Study of Soft-Sphere Collisions

Results are presented contrasting the performance of the different soft-sphere schemes for two configurations of aspect ratio 5:1 needle particles, shown below in Figure 6.10. The simulation set-ups were very basic, with the particles directed at each other with a velocity of  $1 mms^{-1}$ , and all other forces turned off. This allowed a timestep of  $dt = 1 \times 10^{-3} s$  to be used. The first configuration is the standard long edge-to-long edge needle configuration, used many times up to this point, and the second configuration is an arbitrary case achieved by setting the first particle's Euler angles to be  $e_{\theta} = \left(\frac{\pi}{4}, \frac{\pi}{4}, \frac{\pi}{4}\right)$  and translating it 0.25r in the *z*-direction. The approach vectors for the initial velocity prescription remained unchanged, indicated by the orange arrows. This ensures an off-normal and off-centre collision is also studied, such that the curvature scaling in the spring constant, and the particle torques, can have a dynamic effect on the interaction.



Figure 6.10: The two orientational configurations used for the illustrative studies.



Figure 6.11: Distance and overlap calculation between particles colliding in configuration 1. (Top left: Method 1; top right: Method 2; bottom left: Method 3.) The non-zero metrics are then combined and compared for the three methods in the bottom right.

As can be seen in Figure 6.11, the three methods result in very similar predictions of the separation evolution between surfaces for the first configuration, and all proposed

methods are seen to be capable of handling a soft-sphere collision. Methods 2 and 3 perform virtually identically but there is a small discrepancy with the first method, where once again the transfer between the distance detection scheme and the overlap scheme is not perfect. This was also highlighted in Chapter 3 Figure 3.18.



Figure 6.12: Comparison of the overlap computation normalised by *r* (left) and the resulting relative velocity evolution normalised by initial velocity (right), for configuration 1.

In Figure 6.12, the results of the first configuration are further interrogated and contrasted. A simulation has been conducted using the hard-sphere method as a basis for comparison also. As already indicated by Figure 6.11, the predictions of Methods 2 and 3 are identical: confirmed by Figure 6.12, the overlap depth lines lie directly atop one another, leading to the same evolution of velocity. The discrepancy with the first method is further seen, both when transferring to the overlap regime and when exiting it. This leads to the initial reduction in the velocity happening sooner, as well as an overall deviation in the velocities that appears to compound over time. The hard-sphere method precisely retards the velocity to the specified amount at the collision points, whereas the soft-sphere schemes each underpredict the post-collision velocity. This underscores the strong advantage of the hard-sphere scheme in dilute suspensions and binary particle interactions, justifying its use in previous chapters. It allows tighter control over the conservation of energy of the system which is generally a considerable benefit to any physical modelling. In contrast, Method 1 does not offer this same level of control.



Figure 6.13: Comparison of the overlap computation normalised by *r* (left) and the resulting relative velocity evolution normalised by initial velocity (right), for configuration 2.

The same plots are generated for configuration 2. Again, Methods 2 and 3 are very similar, with the lines almost lying atop one another again. However, the resolved post-collision velocity now differs between the two methods. Method 2 overpredicts the post-collision velocity very slightly but is the closest of the three. Method 3 underpredicts slightly, consistent with its performance in the previous configuration, and Method 1 largely overpredicts the resolved velocity.

The overprediction by Method 1 can be related to the overlap depth prediction, where a very unsmooth transition back to non-overlap is seen. Towards the end of the overlap evolution, Method 1 begins to deviate from the two more accurate methods, until the overlap check algorithm abruptly realises the  $T_0$  points are no longer in a state of overlap, terminating the application of the soft-sphere forces. However, this means that the final bit of energy left to be dissipated by the dashpot element of the model is unable to act and so the particle interaction finishes with the particles in a higher energy state than the model intends to enforce theoretically – this has the potential to destabilise simulations.

The implication of this study is that Methods 2 and 3 offer the most promise. For the more complicated configuration, Method 2 offered the closest prediction to the desired restitution coefficient, however it did so at significantly more computational cost. Method 3 took around 60 iterations to converge on average, where Method 2 took around 500 for the deepest penetrations and 100 for the smallest (the required iterations varies as a function of overlap depth). This was done using very small

convergence tolerances, i.e.,  $\varepsilon \sim O(10^{-9})$ , so this could be relaxed. Meanwhile, Method 1 requires zero iterations to converge on a solution after the initial collision points have been identified, which is a significant advantage, but perhaps not enough to overcome the accuracy concerns.

If there was a means to recalibrate the tracked points mid-collision, this could be a promising way forward. Perhaps, there is the possibility of combining Method 1 with either Method 2 or 3, so that the precise points do not have to be converged upon every timestep through a complicated optimisation scheme, but rather the detailed scheme (Method 2 or 3) is run every few timesteps to recalibrate the collision location, and these points are advected using the framework of Method 1 in the time period where no optimisation scheme is run. Like a means of interpolation through time. This could be a good way to offset some of the computational costs whilst maintaining accuracy, however it was beyond the time constraints of the present work to implement this idea.

#### 6.3.4 Demonstration of Agglomeration Capability

To demonstrate the capability of the three approaches in achieving agglomeration between non-spherical particles, Figure 6.14 highlights the same simulation case for the three approaches with the DLVO force model reintroduced. The orientation was the same as Configuration 1 in the last study. The simulation timestep was reduced to  $dt = 2 \times 10^{-7} s$  and the particle velocities were set to  $0.5 mms^{-1}$  as this value was shown to facilitate agglomeration for this configuration in Chapter 4.

As can be seen from Figure 6.14, all three methods are capable of resolving agglomeration. Interestingly, agglomeration is achieved in much the same way as the hard-sphere model, which was not expected. That is, the particles repeatedly collide with one another to successively reduce velocity. The reason this happens is because the van der Waals forces are turned off during overlap and thus the two forces are never acting simultaneously, preventing them smoothly reaching a stationary equilibrium. This is demonstrated explicitly below in Figure 6.15. The same decaying oscillation is seen as the particles move towards a quasi-equilibrium state as with the hard-sphere calculations. Interestingly, the three methods predict slightly different frequencies of oscillation as the interactions evolve. Methods 1 and 3 predict longer slower transfers between overlap and DLVO, whilst Method 2 operates at a slightly

higher frequency. It is not known why this is the case. Methods 2 and 3 predict smoother transitions than Method 1 which has small spikes in its separation.



Figure 6.14: Separation distance evolution between particles agglomerating in configuration 1. (Top left: Method 1; top right: Method 2; and bottom: Method 3.)

On the left of Figure 6.15, the normalised velocity evolution is also shown for the case of Method 3. Since the needles are in the maximal configuration, a very significant acceleration is experienced momentarily before collision as a result of van der Waals attraction, but this is dissipated by the collision, and a similar exchange occurs successively as the velocity and kinetic energy decay towards zero.

Outlined earlier by Figure 6.11, there is a hybrid approach to determining distance metrics in the solver. The separation distance is tracked by a separate algorithm to the overlap depth. There is the same distinction drawn in the application of the forces, which is seen in the right-hand plot in Figure 6.15: the van der Waals force is only non-zero when the soft-sphere force is zero, and vice versa. Perhaps a smoother transition to agglomeration between particles could be achieved by allowing some, possibly reduced, component of the van der Waals force to stay active during overlap. This was

not recommended in the literature, in particular by Mihajlovic et al. (2020), due to an overprediction of agglomeration observed in DEM simulations. However, it may be appropriate to reintroduce this attractive overlap force in the context of turbulence, which will be further commented upon in Section 6.6. It should be noted the relative time that the two force schemes are active for as the agglomerate forms as this will be commented on and prove crucial to findings in the turbulent box.



Figure 6.15: Relative velocity evolution (left) and force interaction (right) between particles agglomerating in Configuration 1 under Method 3.

Overall, in the present state of all three methods, Methods 2 and 3 present strong approaches to the problem. The results of the previous investigations indicate either could be chosen, but given the uncertainty around Method 2, given its novel and untested nature, Method 3 was taken forward when weighed against the time limitations of the project. As future work, the other two methods could be combined to create a strong approach to the problem as discussed.

# 6.4 Validation Against Hard-Sphere Agglomeration Results

As a means of validating the implementation for the purposes of agglomeration, and comparing the previous hard-sphere to the proposed soft-sphere method, key results from Chapter 4.3 were recreated, wherein needle particles were collided in their minimal and maximal configurations at a fixed range of velocities to determine the so-called *cut-off point* for agglomeration.

The spherical case is first compared as a baseline, which validates the soft-sphere model before non-spherical effects are included; however, the non-spherical distance metrics are active for the computation of forces. After this, the minimal and maximal cases for needles are compared.

#### I. Spherical Case



Figure 6.16: Separation distance (left) and relative velocity (right) time evolution for colliding spheres under the soft-sphere model. Distinct lines are colour-coded by initial velocity and highlighted in the figure legend.

The spherical case, depicted in Figure 6.16, overall displays a very similar result to the use of a hard-sphere collision method. The cut-off velocity is seen to change very marginally. In particular, for the  $0.700 \ mms^{-1}$  velocity case, an agglomerate is not formed, and the particles rebound. In the hard-sphere case, shown in Figure 4.2 (c), the particles were observed to nearly escape the potential well, but ultimately be recaptured by the attractive forces, leading to agglomeration. This represented the very edge case of cut-off given how far the particles were able to separate before this took place. In the soft-sphere case, the particles are just able to escape, which implies they retained very slightly more energy post-collision as compared with the hard-sphere. This small difference can be attributed to the imperfection of the overlap distance which is not perfectly accurate due to the assumptions of the model.

#### II. Maximal and Minimal Needle Configurations

The same validation was performed for the needle configurations, with the separation evolutions shown in Figure 6.16. The minimal force inducing configuration has a cutoff velocity of  $0.350s^{-1}$  and the maximal configuration  $1.225 mms^{-1}$ . These are the same two values resolved by the hard-sphere model, demonstrating consistency in the resolved agglomeration outcomes and implying that the new model will be consistent with the findings of the previous chapters.



Figure 6.16: Separation distance evolutions for colliding needles under the softsphere model, demonstrating the cut-off initial velocity for agglomeration in the case of minimal configuration (left) and maximal configuration (right) needles.

### 6.5 Multi-Particle Collisions and Agglomeration

#### 6.5.1 Simulation Initial Conditions and Challenges

To test the robustness of the particle interaction solver, a test-case needed to be created which would place the numerical models under significant stress. This aided the development process and allowed problems to be found sooner.

The chosen case was injection into a quiescent box of fluid at random locations throughout the domain. The particles were given an initial velocity proportional to their normalised coordinate location in the box, which ensured every particle's velocity vector pointed at the box centre. This was done typically with a proportionality constant of  $C = -u_B$ , where  $u_B = 0.875 mms^{-1}$  is half of the measured average collision velocity in the bulk flow region from the complementary channel flow work (Mortimer et al., 2020). However, many values were experimented with during the development process to test a range of conditions.

This initial condition meant that the particles travelled back towards the centre of the box, due to the symmetry of the domain. The centres of the particles would all meet back at the origin after the same number of steps, if not for the non-sphericity of the

shapes and their finite size, as well as the crossing of their trajectories – which all lead to the particles also colliding en route, but mostly in the vicinity of the origin.

Further to this, a slight variability was given to the initial velocities to begin to introduce the variability that would be present in a physical system and to try and maintain a continuous stream of collisions over a wider span of time for testing purposes. The exact definition of this was to scale the velocity magnitude by a factor that allowed a small degree of uniform variability to enter, i.e.,

$$C_i = -u_B \cdot (1.0 + 0.25 \cdot (1 - 2 \cdot X_i)) \tag{6.1}$$

for particle *i*, where  $X_i \sim U(0, 1)$ , which equates to varying the velocity above and below the mean by at most 25%. For the purposes of testing, this value could be decreased to try to increase the stress on solver at a localised point in time, or increased to ensure a continual stream of binary collisions for the purpose of recreating edge cases more quickly.

By injecting the particles uniformly randomly into the domain it was found that often overlaps would be present at the zeroth timestep which introduced numerical problems. The first reason for this was that the initial overlap depth may have been large, leading to a large repulsive penalty force, quickly increasing the kinetic energy above the amount prescribed by the initial conditions. To circumvent this, the softsphere force was capped through the overlap depth,  $\delta = \min(\delta, b/3)$  for minimal principal axis b, which also treats other such edge cases. In a robust framework, the physics in the collision model would ensure that overlap depths do not reach anywhere near the order of the particle size - this is the idea behind the penalty component opposing overlap – but the allowable depth can be controlled by the user/developer depending on the 'softness' of the collisions, using  $N_c$ . However, there are two frequent cases that can overcome this treatment. The first is the injection of the particles as described. The second is when there are many interacting, and thus overlapping, particles creating a granular force effect that overcomes the soft-sphere penalty for some particle pairs in the matrix and allows further penetration than typical. Both must be considered.

A second reason compounded the instabilities first observed, which is the fact that the relationship between the velocity state and the overlap depth are not defined by any
physical process at the point of injection – in the model these two quantities need to be related as they are coupled.

The biggest issues were observed when particles are overlapping to an extent whereby the deepest point penetrates beyond one of the semi-axes of the particle, i.e. it goes over half-way. This happened in injection but was also seen to occur when the collisions were too 'soft' or multiple interactions pushed an overlap too far. This leads to confusion in the spring-dashpot model because there is a relative sign change between the overlap depth and the relative velocity defined at the contact point. Instead of dissipating energy, the particles were observed to accelerate via the dashpot component, and this led to large spikes in the kinetic energy of the system and a violation of conservation of energy (effectively a coefficient of restitution greater than unity is erroneously implemented). Correct tuning of the soft-sphere parameters counteracts this issue, which involves making sure  $N_c$  is low enough to resolve the collision properly (based on the velocity of the colliding particles) as well as choosing sensible values for the parameters outlined in Equation (3.67).

The described mathematical effects caused a problem in the injection step but it was also partly why the first methods had to be discarded – when many particles are involved, the definition of the contact point has to be highly robust. This requirement increases with aspect ratio of the particles since the distance to travel to pass the smallest semi-axis decreases. It can be noted here that 'robust' does not necessarily mean the most physically precise. There just needs to be a consistent mathematical rule that is followed to define the contact point, which also can be converged upon over the full range of orientational and thus geometrically overlapping degrees of freedom.

### 6.5.2 Particle Injection Algorithm

To avoid the injection issue, an algorithm was developed for the injection step, which ensured no initial overlap could be achieved. Firstly, the cubic domain is broken down into smaller cubic sub-domains. The length of these domains is defined by the extent of the longest axis of the spheroid, so that the spheroids could be lined up tip-to-tip without breaching the boundaries of the domain, nor the boundaries of the subdomains. Since the domain is a cube, this can be simplified to one dimension: the number of subdomains in each direction is the length of the box divided by the longest edge of the spheroid, but naturally this will not give an integer in general and so a small buffer is left at the edges of the domain where no particles will be injected.

Now in three dimensions, the algorithm calculates the midpoint of each of these subdomains and this forms a candidate location where a spheroid centre can be injected, whilst also ascribing a unique integer ID in m to each location. The number of allowable subdomains is thus an integer M (for this to work, the number of particles must be lower than M). The injection algorithm generates an integer at random,  $m_i \in \{1, ..., M\}$ , and removes  $m_i$  from the set of allowable integers. By construction of the array of midpoints, the chosen integer corresponds to a unique volume within the box for a particle to be injected. This was seen to work quickly and well.



Figure 6.17: Demonstration of random injection (left) and the main collisions at the centre of the domain (right).

Figure 6.17 demonstrates the injection step and the timestep at which the particles theoretically reach the origin based on  $u_{bulk}$ . It can be seen that the particles are nicely spread throughout the domain at random and that they subsequently collide at a similar time, as required.

### 6.5.3 Initial Results

Initially this was a useful test case for understanding the limitations of the collision implementations, but taking measurements of the system was also able to give some physical insight.



Figure 6.18: Cumulative system separation between particles (left) and cumulative relative velocity (right), measured over time.



Figure 6.19: Kinetic energy of the particle phase (left) and cumulative soft-sphere force (right), measured over time.



Figure 6.20: Mean coordination number of the particles over time.

The results of this test simulation are shown in Figures 6.18 to 6.20. The quantities on display are the interparticle surface separation, relative velocity, kinetic energy, soft-sphere force and mean coordination number. These quantities are *total* quantities, measured pairwise between particles and then appended to a cumulative sum to give a rough metric on the overall system behaviour at each timestep.

The separation of the particles decreases linearly as the particles move towards the centre of the box, representing their constant velocity, the system separation then reaches a minimum whilst the collisions take place. This separation quantity does not reach zero because, even when all surfaces are touching, particles at opposite sides of the congestion remain a significant distance apart. The minimum value obtained, which seems somewhat arbitrary, does actually give insight into the system as it roughly represents the density of the structure formed when all particles are interacting. Lastly, the separation between the particles begins to increase again as the collisions complete, but the gradient of the line is seen to be reduced, representing a decrease in overall velocity state and a loss of energy, in much the same way that binary particle interactions were observed to behave; in fact, the graphs are noticeably very similar.

The relative velocity plot in Figure 6.18 tells a similar story to the tracking of the kinetic energy shown in Figure 6.19, which is that the collisions are dissipative. Robustness is demonstrated by these lines monotonically decreasing, ensuring that conservation of energy is not violated. There can be very small increases in the relative velocity because this value is a magnitude and particles will temporarily reach zero relative velocity magnitude after deceleration has occurred in the collision, before they are again accelerated outside of one another.

Lastly, Figure 6.20 shows the mean coordination number of the system which is defined here as the number of particle pairs that are undergoing a soft-sphere collision divided by the total number of particles. This places a metric on how many collision partners the average particle has. Naturally, every step change in this graph is by a fixed amount  $(1/N_p)$  and so the graph is not especially smooth. The coordination number starts at zero and ends at zero, indicating that all collisions were successful in separating particle surfaces, giving indication to the robustness of the method. As

more particles collide, the coordination number grows, indicating multiple interacting particles, with certain particles having multiple collision partners.

#### 6.5.4 Application to Behavioural Modifications

In the spirit of the first part of the chapter, it seemed relevant to assess the system in terms of a behavioural modification and so the coefficient of restitution was varied through the same three values as earlier,  $e \in \{0.2, 0.4, 0.6\}$ . The results are shown below in Figures 6.21 to 6.22.



Figure 6.21: Cumulative system separation between particles (left) and cumulative relative velocity (right), measured over time for three values of the coefficient of restitution.



Figure 6.22: Particle-phase kinetic energy (left) and coordination number (right), measured over time for three values of the coefficient of restitution

Interestingly, the difference in the two key metrics, i.e., separation and relative velocity, is not as large as would maybe be expected, shown in Figure 6.21. The systems with the lower two CORs behave very similarly in terms of the amount of energy lost. This

may be due to the limitations of the dashpot model when there are many interacting collision pairs with competing forces. In such cases, the spring-dashpot model no longer has the same level of *control* over overlap depth as the interaction progresses – there are other competing forces. Thus the assumption that a fixed amount of energy will be dissipated is not always met with full accuracy, as the overlap evolution has moved away from the ideal scenario in which the model constants were derived. Therefore, a recommendation is made here that a combination of the hard- and softsphere models could be employed where the hard-sphere model is used as the default, thus allowing the correct dissipation of energy to be ensured in simple binary collision cases, and the soft-sphere model is used as a back-up model whenever the limitations of the hard-sphere model are encountered or agglomerates begin to form (as will be the case when DLVO forces are re-introduced).

The coordination number is significantly higher for the lowest coefficient of restitution case, meaning that the average particle has more collision partners. Therefore, a lower coefficient of restitution has the potential to make the simulation less accurate from a energy dissipation standpoint.



### 6.6 Reintroduction of Turbulence and DLVO Forces

Figure 6.23: Cumulative system separation between particles (left) and cumulative relative velocity (right), measured over time.

Optimistically, the highest turbulence box ( $Re_{\lambda} = 22.2$ ) was first chosen to study multiparticle agglomeration. The same initial conditions were used as the previous study, for 24 particles, however the turbulence was quickly able to overcome the initial conditions as seen in the relative velocity plot of Figure 6.23. The chaos of the turbulence field can be seen in the separation of the particles. It is noted that sharp spikes are not indicative of instabilities, but just represent discontinuities in the separation between particles as the particles pass through the periodic boundaries and are injected at the opposite side. There is not currently a treatment to make this calculation continuous in the code and can also lead to agglomerates breaking as one particle passes through the boundary first.



Figure 6.24: Coordination number (left) and force magnitude (right), measured over time.

The coordination number is plotted again on the left of Figure 6.24, and this time it is the absolute value shown rather than being scaled by the number of particles. In retrospect this gives a clearer picture of what is going on: a value of 2.0 means that two soft-sphere interactions are taking place at once, and so on. Over the simulation period plotted, a few interactions take place, climbing towards a maximum value of four simultaneous interactions. For this application, the coordination number should ideally also take into account the number of particles in the DLVO potential, to give an impression of the number of formed agglomerates. This should be addressed in future work.

The right plot of Figure 6.24 displays the van der Waals and soft-sphere competition for the multi-particle case. As expected, the two forces seem to counter one another. The soft-sphere forces reach a maximum whilst the van der Waals forces are not active. Agglomeration would be seen where the van der Waals forces are active and the soft-sphere forces are very minimal, because the magnitude of the force is a function of overlap depth. As the particle agglomerates, the successive overlap depths reduce, and so the soft-sphere force reduces also, as was shown in Figure 6.15. This was not really observed in the present simulation.

Under close inspection of the simulations, an effect is observed where for particles that ought to agglomerate, they spend a lot of their interaction time in the state of overlap, for which there is no attractive force resolved between them. This gives a window for other competing forces to act to accelerate the particles apart with no reply from the van der Waals force. By the time the van der Waals forces reactivate, the particles may have already been accelerated to a degree that allows them to escape. This actually creates a scenario where agglomerated particles are too easy to break apart. This is in direct contrast to the findings relating to the hard-sphere model in Chapter 5, where agglomerates were unable to be broken by the turbulence field. There, the van der Waals forces are always active, and this underscores the difference in compatibility between models. For future work, it is recommended to implement some consideration of agglomerative force when the particles are in a state of overlap, such that agglomerate structures are more stable in the flow – and more representative of the physical reality.

Since the van der Waals forces are so short range, they tend appear sharp, but they are resolved over a number of timesteps and 'zooming in' on those spikes reveals a smooth application of force. It is this requirement that massively restricts the simulation time. Just to simulate this simple interaction took two days of compute time on 48 cores. Since there is no splitting of the timesteps between particles and the fluid, a significant chunk of this compute time is taken redundantly calculating a high-order accuracy velocity field which has changed imperceptibly between steps.

To try and achieve agglomeration without changing the code, higher and higher numbers of particles were assessed, however turbulence at this level ( $Re_{\lambda} = 22.2$ ) was mostly found to separate the particles and disperse them quite evenly as shown in Figure 6.25 below which represents the final state of a simulation with the same simulation conditions as previous, but with 50 particles included. The final timestep is shown here with no sustained agglomerates formed. Again, this is probably also a function of the particle collision and DLVO interaction between one another leading to agglomerates which are overly susceptible to breakage.



Figure 6.25: Final state of the particles after a turbulent simulation with  $N_p = 50$ .

In all of the studied cases there was a significant limitation resulting from the low timestep requirements of the DLVO forces. Either the initial velocities, or the turbulence strength were set high, thus allowing many interactions to take place, but almost always too strong for agglomerates to form; or, the velocities and turbulence strength were set low, but the simulation time to allow for many interactions was so prohibitively high that few interactions could be studied, even though agglomeration was observed. It was possible to reduce the maximal strength of the attractive forces and thus reduce the simulation timestep without fear of divergence, but then agglomeration similarly became unlikely because the forces are reduced. There is also very fine resolution required in the particle motion, such that the amount particles move relative to one another gives the chance for the separation distances to enter the potential well of the DLVO effects, giving them enough chance to contribute to the pair's motion. If the movement of the particles in a single step is too large, the particles can easily bypass the potential well, even if they collided.

# 6.7 Solver Improvements: Parallelisation, Time-Stepping and Speed-Up

As described, long compute times were required in each of the simulations presented in Section 6.6, even though particle numbers of 24 and 50 are relatively low. As such, there was a clear need to speed the simulations up for this methodology to be

practically viable on a larger scale. To bring real-world applicability would involve being able to resolve a much larger simulation time period, which means far more time-steps. The challenge here is the vastly different timescales of the fluid and the DLVO forces. The boxes of turbulence require  $dt \sim O(10^{-3} - 10^{-5}) s$  depending on the resolved Taylor Reynolds number, whereas the DLVO forces are requiring  $dt \sim 0(10^{-6} 10^{-7}$ ) s depending on morphology. For particles to move with realistic velocities based on the channel-flow analogue (or nuclear waste systems in general) and to fully sample the related turbulence field as well as interact sufficiently with each other, an inordinate number of timesteps is required (hundreds of millions) even in millimetre domains. However, since there is such a separation of scales, it is a reasonable assumption to allow the update of interparticle forces to decouple from the bulk fluid motion at the time-level. The main source of particle motion is the background fluid and naturally the time-stepping involved in this coupling needs to be on the order of said fluid, but for very close range interactions it is possible to vastly reduce the timestep of just the particle phase, such that there are a number of internal time loops run within one global timestep for just the calculation of the particle motion.

The first element of speed up was achieved by creating a subloop for the advection of the particles due to DLVO forces. In the main simulation loop a check is applied that identifies particle interactions that are found to be entering into a DLVO potential well. This is presently identified as interparticle separations falling below two microns, consistent with Chapter 5. These particle interactions are flagged and then placed into a subloop of their own which takes place before the main particle advection loop and gives sufficient resolution to just the particle pairs that require it. A reduced timestep is applied,  $dt_{vdW} = dt/100$ , over which the force calculation and consequent particle advection is calculated for particles experiencing agglomeration. If particle surfaces are found to overlap then their flag is removed and they are exited early from this subloop awaiting overlap attention in the next iteration of the soft-sphere forces (the next timestep). This allowed the overall simulation timestep to be increased by two orders of magnitude, giving incredible speed ups, and allowing simulations that were previously taking many days to complete, to now complete in hours. This meant a much fuller/longer turbulent agglomeration interaction could feasibly be simulated.

The second speed up was achieved by distributing particle interactions over the multiple cores used by the fluid solver. Nek5000 uses an MPI parallelisation technique,

and within this framework the user has access to the node ID that a calculation is taking place on. Thus, a simple computation takes place at the start of the simulation to identify the available number of processors, and any particle loops that are independent of one another (i.e., the particle interactions, which are computed pairwise) are evenly distributed across the available cores, which also significantly sped up the code. This is done by looping over an integer range that is defined locally each timestep according to the identified local node ID.

One nice factor here was that the main bottleneck is the DLVO forces, and since the expensive subloop only contains particles that are undergoing considerable DLVO forces, there are no redundant calculations taking place within this framework. All of this allowed the results presented in the next section.

### 6.8 Soft-Sphere Agglomeration in Turbulence

A method has now been developed which gives the real possibility of studying nonspherical agglomeration in turbulence at a very fine level of detail, for the case of many particles interacting for extended periods of time. Unfortunately, as this was achieved, the remaining time on the project was short so the presented results outline a capability but fall short of delivering new insight into the topic, which is now certainly possible.



Figure 6.26: Three simultaneous binary particle interactions (highlighted in red, blue, and green) take place, advected by a homogeneous and isotropic turbulence field ( $Re_{\lambda} = 2.5$ ). The pictures are separated by 0.01ms and the blue arrow indicates direction of time.

To demonstrate the capability, three pair interaction cases are shown in the turbulent multiphase system in Figure 6.26. To aid the agglomeration process, the higher turbulence level was replaced by the lower turbulence level box.

In the interaction that takes place between the particles highlighted in blue, at the first snapshot the particles begin to collide in an almost long-edge to long-edge configuration. In the second snapshot, the first soft-sphere collision has completed and the particles are seen to have reorientated themselves into the precise maximal force configuration. This was helped by the natural torques induced by the collision, but was held in place by the attraction towards this particular configuration which is favoured by the DLVO forces. In this case, the absolute velocity of the particles is quite low in comparison to most of the rest of the particles shown and the relative velocity between the particles is also relatively low. At the same time, the needles had this perfect configuration, where the van der Waals forces were shown to be increased in Chapter 4, compared to a spherical interaction. These favourable conditions lead to an agglomerate forming which was seen to last for the remainder of the simulation. In the third snapshot this is demonstrated as the particles are advected away as a coherent pair. This demonstrates the capability of the methodology to facilitate softsphere agglomeration events in turbulence.

In the particle interaction highlighted in green, the two particles possess considerably more kinetic energy. The particle that starts to the left in the first snapshot has considerable energy in the form of rotational kinetic energy, and even though this particle pair moves through the same maximal attractive state as the blue particle pair, which occurs in the second snapshot, this higher energy ensures that the particles do not agglomerate. By the third snapshot, the particles are swept apart by the local velocity gradients and the interaction is much shorter overall.

In the third case, where the particles are highlighted in red, an unusually lengthy interaction was seen to take place due to a region where the background turbulence field happens to be relatively calm. This allows two particles to undergo a soft-sphere collision that results in an agglomerate forming for a short while, this is seen to last for the full time shown by the snapshots. However, as was shown in previous chapters the configuration where needles are connected by their extremities is not particularly

strong and this agglomerate was seen to break up as soon as their local flow became stronger.

# 6.9 Agglomerate Structures Formed by Interacting Non-Spherical Particles

Finally, an avenue for investigation and validation of the solver and subsequent findings is outlined. Ochsenbein et al. (2015) studied a system of interacting and agglomerating needles of a similar aspect ratio to present (5:1). Needle particles became agglomerated in a reactor of monosodium salt water through a stirring mechanism. Considerable measurements were taken which can be compared to, as a means of validation of the types of structures formed by agglomerating non-spherical particles. Some examples of the formed structures are displayed below in Figure 6.27, which have similarity to those observed in the final simulations conducted as part of the thesis.



Figure 6.27: [Taken from Ochsenbein et al. (2015)] image recognition technique used to characterise needle agglomerates of a ~5:1 aspect ratio.



Figure 6.28: Comparison with Ochsenbein et al. (2015) (left) shows a similar triplet needle agglomerate formed in the flow with the present methodology (right).

Figure 6.28 shows a typical agglomerate formed between needles of a 5:1 aspect ratio in the reference material. Like that which was typically observed in the present simulations, both the experiment (left) and the simulation (right) show clustering around a central point for multi-particle agglomerates. This is shown in most examples above in Figure 6.27, which was also seen in present simulations when triplet agglomerates formed. Further analysis should look to conduct many simulations and attempt to quantify this effect over a statistically significant sample, whilst also comparing to experiment.

## 6.10 Conclusions

In conclusion, the final chapter has covered wide ground. To begin, a comprehensive behavioural modification study took place to analyse the relative importance of Hamaker constant and coefficient of restitution across morphologies. It was found that both parameters played a significant role in agglomeration, building upon the work of Mortimer et al. (2021) and Wolde (2023), confirming hypotheses presented therein and extending those ideas to non-spherical particles. In general, this extension showed similar conclusions across morphologies, although interesting behaviours were seen with the disks. The disks seemed to be more independent of the Hamaker constant and COR, compared to the other morphologies in terms of the resolved agglomeration rate, and also exhibited interesting alignment tendencies for both mentioned behavioural modifications. The needles and spheres were found to consistently agglomerate at higher rates than the disks and they showed more

sensitivity to the behavioural modifications of COR and Hamaker constant, whereas they showed higher sensitivity to turbulence level.

Further to this, an extensive developmental journey was outlined for the simulation of multi-particle agglomerates arising due to a turbulence field. This is the first time this has been attempted for structure resolved non-spherical particles, with all competing forces modelled explicitly, and the full process was outlined with consideration given to the challenges and solutions. Overall, a powerful simulation methodology has now been developed that facilitates a multitude of competing physical phenomena, with wide applicability – not least to the nuclear industry and the challenges faced by Sellafield Ltd.

The recommendations from this final piece of work are quite wide since there are so many potential avenues to go next. The novel overlap detection schemes could be made robust and compared to the findings of other schemes in the literature. The newly developed methodology should be applied at scale to investigate the agglomerate structures formed in turbulence by non-spherical particles. This in and of itself is a very wide topic and it is unchartered territory in the literature, especially at this level of fidelity. The application of the described methodology can offer insight into a lot of different multiphase systems, and represents the cutting-edge. It would be interesting to compare the types of structures formed by different flow conditions, and this can be controlled through the prescribed turbulence conditions of the box.

# 7 Conclusions and Suggestions for Further Work

### 7.1 Summary of Main Developments and Findings

Over the course of this project, significant ground has been covered with novel contributions to the general understanding of turbulent multiphase flows, non-spherical colloidal particle pair interactions, and a wide range of computational modelling techniques and methods. It is the goal of this final chapter to highlight these key achievements of the project and give suggestions for future work.

The thesis began with an introduction that motivated the work undertaken, both in terms of its general significance to developing scientific understanding, as well as its particular importance to the nuclear industry and Sellafield Ltd. To concretise this, a series of ambitious but realistic objectives for the project were outlined. Each of these stated goals were ultimately met, which will be demonstrated throughout this section, highlighting how the attainment of these goals directly contributes to, and moves forward, the understanding and capability surrounding the relevant systems.

In support of this, the literature review covered a breadth of fields, underscoring the wide array of physics that feed into this modelling problem. After introducing the key background literature and concepts relating to turbulent flows, computational fluid dynamics, and particle modelling, the current state-of-the-art was relayed in terms of the more specific areas of particle-laden flows, non-spherical modelling techniques, and particle agglomeration in turbulent systems. Stemming from this, there were several key areas identified where there was scope for novel contribution. In particular, it was highlighted that very few studies exist that combine particle-resolved multiphase fluid dynamics with agglomeration effects, and no existing studies were found where particle non-sphericity was explicitly captured in this context. Moreover, there were no studies identified that delve into the effect of turbulent fluid dynamics on the agglomeration process between non-spherical particles.

The multiphase agglomeration literature itself was seen to be quite sparse, despite this being shown to be a key emergent physical process that can dominate the overall behaviour of particle-laden systems, which appear in a wide range of situations, from the nuclear industry to environmental flows. This means that a complete picture of the process of agglomeration in turbulence is still to be developed, and there is significant room to develop theoretical understanding. At the macroscopic level, it was shown that there are studies in the multiphase flow literature that address spherical particle agglomeration, and similar studies were found addressing fluidisation. Separate to this, it was shown that there are robust physical theories about particle interactions at the microscopic scale and there are successful efforts to extend this to non-spherical particles under a theoretical framework. However, there are very few attempts to bridge the gap between these two scales to understand the interplay in full detail, explicitly capturing the emergent behaviours as they take place – this is where the present simulations fit in.

There are reasons for the lack of study in the stated areas in the literature, but two key points are as follows. Firstly, there is the difficulty of including non-sphericity into fourway coupled flows, due to the challenge of robustly treating collisions and accurately capturing orientationally-dependent forces – this poses a strong mathematical challenge, with many separate treatments and models that must be devised and developed of a significantly higher complexity than the equivalent spherical models. Notably, there are attempts to overcome the specific challenges associated with nonspherical collision modelling, but the different approaches do not necessarily align with one another and thus this is an active and evolving area of research. Beyond this, there are very few attempts to include agglomerative forces into the models, and none that consider non-spherical particles due to the orientationally-dependent complications this brings. Secondly, there is the much-discussed challenge associated with the vast range of scales present in these flows. There are typically two to three orders of magnitude separating the scales associated with the turbulence field which drives the bulk particle motion, and the DLVO effects which lead to particle agglomeration events. As such, it is typically unfeasible to capture the physics at both scales simultaneously. By restricting the scale of investigation, the present work was able to capture both effects at a very high level of detail to inform the typical simulations of large-scale systems that cannot realistically model both.

Before any results were generated, a significant development process first needed to take place to address the above challenges. In pursuit of this, a multiscale and multiphysics solver was developed to address the problem of particle agglomeration in turbulent flows for systems of non-spherical particles. The project started with an inherited methodology (Mortimer, 2019) which laid the foundations for the work undertaken in this thesis. The inherited code had successfully incorporated DLVO forces between pairs of spherical immersed boundaries but was not yet seen to reliably resolve agglomeration. One of the key features of this code was its ability to read particle meshes into a data structure that could then be used to accurately couple the particle-phase to the fluid-phase, such that fluid forces could be interpolated directly from the high-fidelity fluid calculation in situ. This proved to be useful in the process of extending to non-spherical particles as the mesh reader could handle any kind of icosphere mesh and the treatment for particle advection due to the fluid was second-order accurate. Building upon this, extensive adaptations, refinements and model developments took place during the present work such that agglomeration became robust, and that this could be achieved also between non-spherical particles. In particular, models were either identified in the literature or developed from first principles to address the problems of: particle injection, boundary conditions, forces and torques, rotation and tracking, robust distance metrics between surfaces, orientationally-dependent forces derived from the local surface curvature, collision detection and physically-accurate treatment, and robust turbulence forcing. This methodology represents the state-of-the-art and there are no other known codes that combine all of these models into a high-fidelity fluid simulation capable of direct numerical simulation (DNS) of all fluid time and length scales, which was achieved using Nek5000 throughout. This methodology was used to generate the results presented in Chapters 4, 5 and 6.

On top of this, the findings of the mentioned chapters made clear what the next developments to the solver needed to be, and as such a second stage of developmental work took place to devise algorithms for *multi-particle* agglomeration in turbulence, with the extensions to the code being non-trivial. In particular, two novel methods were presented addressing the challenge of soft-sphere collisions and agglomeration between non-spherical particles, and a further method was also implemented which ultimately allowed the final goal to be achieved, demonstrated at

the end of Chapter 6. Further adaptations to the code, namely a new proposed technique for the variable timestepping of DLVO forces and parallelisation of the particle module, led to very significant speed-ups which made a new kind of simulation possible wherein multi-particle agglomeration could be captured in its full detail, over comparatively significant lengths of time. This is a significant achievement and gives the University of Leeds a powerful and unique simulation technique for investigating turbulent agglomeration between non-spherical particles. After the first stage of the methodology gave detailed insight into the pair interactions between particles, this new described framework has that which is required to take the next step, vastly increasing the scope of the simulations in terms of the length and time scales that can be studied, which will allow detailed study of multiparticle agglomeration into crystal-like structures.

Utilising the described techniques, several studies were designed and undertaken to probe the systems of interest. By systematically investigating the interactions between particle agglomeration events and turbulent fluid dynamics, several novel insights have been gained that offer theoretical and practical application, particularly in industrial contexts such as nuclear waste management.

Chapter 4 was the first results chapter of the thesis, which first centred on demonstrating the robust modelling of spherical and non-spherical colloidal particles interacting under the influence of van der Waals and electric double layer forces within the DLVO framework. Such were the chosen system parameters, the electric double layer repulsion was seen to have very little influence in comparison to the van der Waals effect, meaning that the DLVO interaction was dominated by the attractive van der Waals forces for the entire set of simulations investigated throughout the thesis. The parameters were chosen in line with calcite particles in water, a realistic nuclear waste analogue.

One of the first findings was the identification of a 'cut-off velocity' for spherical particles – approximately  $0.7 mms^{-1}$  under the present conditions – which marked the threshold for particle agglomeration. As described, this result was arrived at through fundamental models that capture the microscale effects, and thus provided a validation to previous studies operating at a coarser level of modelling (Mortimer et al., 2020), which used

an energy-based deterministic ('on/off') method. This result was able to confirm the validity of the chosen value in those simulations.

Further, the behaviours and dynamics of non-spherical particles, specifically of disks and needles with a 5:1 aspect ratio, were explored in detail. The cut-off velocities for either morphology was seen to vary dramatically as a function of relative orientation, a result which pairs with later studies over the orientational parameter space to underscore the importance of considering particle morphology on any study that attempts to model agglomeration in multiphase systems. Notably, the disks when orientated face-to-face continued to agglomerate at much higher velocities, with maximal van der Waals forces three times those observed in needle configurations. Needles, conversely, demonstrated the lowest possible attractive forces at their endpoints, and therefore agglomeration was unable to occur in all but one studied velocity case in this configuration. This set of simulations also provided a means to understand the timestep requirements of the solver, with the stable values being  $dt = 1 \times 10^{-6} s$  for spheres and  $dt = 2 \times 10^{-7} s$  for the most extreme disk configurations, and similar for the needles, which was thus taken forward for all subsequent non-spherical studies.

The next study looked to investigate the role of orientation in determining agglomeration outcomes, which was achieved through uniform sampling of the orientational parameter space and a Monte Carlo approach. Orientation was consequently seen to be a crucial variable, alongside the energy of the particles, in determining interaction outcomes. Surprisingly, despite having the capacity for the highest available DLVO forces, disks were seen to agglomerate in far few instances than needles. To understand this result, both the relative orientations between the nonspherical particles symmetry axes and the points on the surface that experienced the collision were interrogated. This demonstrated an interesting result: the distributions of collision points over the surface of the shapes is different between the two morphologies under uniform sampling of the initial orientations. This was understood through a number of described geometric effects and highlighted as an interesting area for further study. The relationship between this result and the distribution of surface curvature (which mediates the strength of the DLVO interaction) in combination gave a detailed picture of the precise mechanism that alters the agglomeration propensity as particle morphology is changed. It was the *distribution* of

281

the curvature across the shape which was seen to be the most important factor, due to the fact the two spheroidal shapes had the same 'total curvature', defined in the mathematical sense. More specifically, the analysis of collision point distributions was able to be related back to the particle geometry which emphasised that while disks tended to collide at their edge regions (with higher curvature and weaker attractive forces), needles more frequently collided in areas away from their curvature maxima, which facilitated stronger interactions over the generated sample space. Overall, the study concluded that agglomeration in non-spherical particles is highly sensitive to particle morphology, and orientation, and the spatial distribution of surface curvature was highlighted as the factor that drives the differences between morphologies.

Extrapolating from these findings to a physical hypothesis: needles with their stronger interactions over the full orientational parameter space and thus higher agglomeration rates could be expected to experience more agglomeration and thus faster sedimentation compared to systems of disks; however, this hypothesis needs to be tested in a multi particle system to capture the full range of relevant effects. It could be expected that systems of disk-like particles remain more stable (i.e., separated and resistant to aggregation) compared to needle-like suspensions under the same flow conditions. It is not clear to what extent uniform sampling of orientation is the appropriate analogue to physical systems, since the flow conditions will naturally influence the typical relative orientation at the moment of collision, however it is a good first probing of the system that has generated an unexpected result of value. Placing the particles in turbulence will help to more naturally bias their configurations.

Chapter 5 therefore looked to expand the study of particle agglomeration by introducing turbulent fluid dynamics. Firstly, to achieve this, the successful implementation of a robust stochastic turbulence forcing field, coupled to the flow solver Nek5000 through a source term, was demonstrated – with the validation showing strong agreement with the reference cases studied. Using adjacent channel flow calculations performed at a macroscopic scale, the turbulence boxes' parameters were matched to specific regions of the channel flow. The attempt to match turbulence conditions in this way was ultimately shown to be problematic due to the limitations of the idealised case of homogeneous isotropic turbulence (HIT) which forces turbulence differently from that observed in channel flows. Matching the velocity fluctuations of the flow between the channel and the HIT boxes allowed representative fluid forces to

282

be imparted on the particles, which mediated the agglomeration process at the correct level of magnitude as desired, but ultimately it led to overall flow conditions different to those experienced by particles in a channel, even when considering only the reference frame of the collision (thus theoretically avoiding mean flow effects). Despite this, the altered levels of turbulence between the boxes ensured that the results were meaningful and contributed to the understanding of the interplay between turbulence and agglomeration.

Key findings from this chapter included the observation that increasing turbulence in general reduced agglomeration rates across morphologies. Over all simulations, needles were seen to favour agglomeration the most, followed by spheres and then by disks. This was in keeping with the findings of Chapter 4. In turbulent environments, particles became more energetically agitated through increases to their relative velocities, making agglomeration increasingly energetically unfavourable as the turbulence level was raised. However, this effect was more pronounced for disks due to their greater surface area for the fluid forces to act upon, as well as their tendency to align with the flow in such a way that maximised their drag. To a lesser extent, this also affected the needles, and ultimately the spheres were able to overtake the needles on some agglomeration metrics in the box with the highest levels of turbulence, which was attributed to the reduced drag experienced by spheres, despite their lesser DLVO attraction. This highlighted a crucial interplay between turbulence and morphology as it relates to agglomeration. It is not the case that one morphology agglomerates more than another for all systems, but rather the relative strength of certain system parameters will favour some morphologies over others, an idea which is further investigated in Chapter 6.

The highest level of turbulence was seen to most closely resemble classical turbulence, whereas the lower two turbulence boxes were of such low Taylor-Reynolds numbers that they could be thought of more readily as 'chaotic flows', since they did not display the same separation of scales of classical turbulence. Even so, the flow scales were different in the three boxes, and this contributed to the alignment statistics of the particles. There was no clear morphology that aligned more than the other across all three studied boxes, however. In fact, the candidate morphology favouring alignment seemed to change between boxes. This was hypothesised to be due to the different principal lengths of the particles, which will thus experience the flow in unique

ways, having a different characteristic length compared to the average flow scales. As those flow scales change between boxes, different motions will be induced by the different shapes. This is something that could be looked into in further detail.

In summary, this chapter highlighted that turbulence generally reduced agglomeration rates and confirmed that the findings of the quiescent case were roughly transferrable to the more complicated system. As before, particle morphology was a key determinant of the observed agglomeration statistics, but an added variable of surface area became important due to its significant influence over particle acceleration and the overall dynamics.

The final results chapter, Chapter 6, delved into the importance of key parameters in determining particle agglomeration, extending previous spherical studies to nonspherical particles, aiming to understand the difference when introducing differing morphologies and aiming to validate macroscale studies performed within the research group. Both parameters studied, the Hamaker constant and the coefficient of restitution (COR), were found to have a significant impact on agglomeration, with needles and spheres showing a higher sensitivity to both of these behavioural modification parameters when compared to disks, which appeared relatively independent of these factors in terms of their resolved agglomeration rate. Disks however exhibited notable alignment tendencies under varying conditions of COR and Hamaker constant, which contrasted with the behaviours observed for needles and spheres. The disks were also shown to be most affected by the turbulence strength as was noticed in Chapter 5. This indicated that different behavioural modification techniques were more suitable for certain morphologies, and thus the morphological characteristics of a suspension should be characterised before applying any such modifications.

In addition, the simulation of multi-particle agglomerates in a turbulence field was explored for the first time with structure-resolved non-spherical particles. This approach, which accounted for all competing forces in a direct way, marked a significant step forward in agglomeration modelling and places the research group in a position to generate significant understanding of the structures formed by agglomerating non-spherical particles in turbulence, which will be of significance to many areas of research. The method was successfully applied to simulate the dynamics of agglomerate formation and provided a robust framework for future studies of use to the nuclear industry, particularly for applications such as waste treatment at Sellafield. To reach this point, three methods were compared for overlap detection, two of which were novel, however the finally chosen method was a pre-existing technique chosen as a matter of pragmatism. The method was shown to be capable of resolving agglomerates between many binary particles at once, and was thus far observed to resolve triplet agglomerates in the present turbulent conditions. The interplay between the soft-sphere scheme and the DLVO force model was seen to be capable of resolving agglomeration, but agglomerated particles were seen to spend much more simulation time in the state of mild overlap, which meant that the DLVO forces were often not active, allowing a significant window for tertiary forces to disrupt the state of agglomeration. It was suggested that both force models should be active to increase the stability of the formed agglomerates and thus allow the coordination number of the agglomerate structures to increase, to the expected levels.

### 7.2 Suggestions for Further Work

There has been significant progress achieved, attacking what is a difficult modelling problem, but there remain several exciting avenues for further exploration, particularly given the tools developed. First and foremost, however, whilst the simulations presented here offer theoretical insights, experimental validation would strengthen confidence in the findings. In particular, experiments with non-spherical particles in turbulent flow could provide valuable benchmarks for further refining these models. Experimental investigators face the same problems as computational modellers for these systems due to the difficulty in assessing all scales, of which there are many. The most practical way to cross-check modelling predictions would therefore be through comparison to bulk flow measurements. The capability to measure these turbulent agglomerating systems is presently being developed at the University of Leeds' *MULTIForm* facility and so steps are underway to achieve this. This would provide validation at the macroscopic level and could be assessed in terms of agglomeration rate, sedimentation rate, or by qualitative and quantitative comparison of the structures formed in the flow.

Equally as important would be increased understanding at a microscopic level, without necessarily considering the fluid dynamics, to further calibrate the DLVO models used

in the present work. There is uncertainty throughout the literature into the correct values to set certain parameters, such as the cut-off separation for the van der Waals force (given its mathematical representation is a  $1/r^2$  law it thus diverges, so an arbitrary value must be chosen to limit the force). Throughout the agglomeration literature, this parameter is chosen as a way of making simulations stable rather than being directly tied to a physical effect but it has significant implications for the maximal allowed interaction strength. Further investigation into this would be of great benefit to the DLVO-induced agglomeration literature, allowing greater consistency between studies if a physical and consistent way of determining this parameter was developed.

Stemming from the work in the quiescent conditions, a potentially rich mathematical problem was identified where it would appear there are a number of interesting and complicated effects occurring between the geometric nature of the particles and the resultant interaction behaviours under random sampling. Spheroidal particles are created either by 'stretching' or 'squashing' a sphere with volume kept fixed, and yet these two similar shapes have very different properties when analysing the regions of the surface that accumulate collisions under uniform sampling, which then has importance for the magnitudes and distributions of the forces in pair interactions.

The random initial conditions for the orientations of the two bodies can each be defined in terms of a single unit vector which is randomly uniformly distributed on the surface of the sphere in  $\mathbb{R}^3$ . The dot product of these two independent random vectors forms a new probability distribution which can provide insight into the underlying distribution formed between two randomly orientated particles – which will not necessarily lead to a uniform random sampling of relative orientation, which needs to be understood to analytically understand the results of Chapter 4 in full detail, as well as to generate fundamental understanding of interaction modelling. Beyond this, when moving from a unit vector to a three-dimensional shape, the way that the distribution of the surface curvature then affects the sampling of collision points creates a very complicated multiparameter probability distribution which would be extremely difficult to quantify analytically. At this stage, an approach like machine learning may be able to drive towards the underlying model. If fed enough training data on the interaction outcomes obtained at this high level of detail for pair interactions, then a trained model could potentially characterise all of the interesting behaviours into something that could be applied at a macroscopic level saving significant time and resources. In a macroscale situation, the model could be given parameters like the relative orientation and kinetic energy to quickly determine the agglomeration state post-collision for pairs of non-spherical particles, without having to explicitly model the small-scale phenomena that places a stress on compute time. Such a model could be applied to point particle Lagrangian particle tracking simulations. The analysis framework developed throughout Chapters 4 and 5 showed that it was possible through these types of simulations to get a sense of which collisional conditions, particularly in terms of orientation and energy, lead to agglomerates – a machine learning approach could potentially get to a suitable model much quicker.

This still leaves the problem of knowing the correct shape to replace agglomerating particles in the point particle scenario, and this becomes even more complicated for non-spherical particles. Thankfully, the capability has now been developed to explicitly resolve the structures formed as particles interact in turbulence. So, a main recommendation for future work involves utilising the multi-particle soft-sphere technique to gain insight into this area. The different types of structures that form will lead to particles of unique shapes, which will have interesting drag characteristics and interactions with the turbulence field. Fortunately, the capability is there to study this motion in a turbulence field, and machine learning could perhaps even be applied to the problem of identifying the correct drag relationships to be used in point particle models given the agglomeration conditions of the particles that formed the agglomerate.

Beyond just improving point particle modelling approaches, quantifying the structures formed by agglomerating non-spherical particles is of direct relevance to the rheological challenges faced in the nuclear industry as complex shapes interact to form sludges, which are poorly characterised and understood due to the safety risks with assessing waste material directly. Further still, agglomerated shapes will have an interesting and coupled relationship with the turbulence field and will likely result in different levels of turbulence modulation based upon the type of structure formed and the number of constituent particles, which in turn can develop new insight into the relationship between particles and turbulence. Characterising these effects would begin to shed light on the relationship between agglomeration and turbulence, from the point of view of the fluid, i.e., how do certain collision and agglomeration events create certain flow conditions – and what kind of vortical structures are formed by

agglomerates in the flow, and how does this changing turbulence field further influence the particle dynamics and thus alter the agglomerates formed later in the flow. It would similarly be of interest to understand whether these complex agglomerates are able to stabilise and align with the flow, or whether their complex shape prohibits this as well as at what point they deposit out of the flow and how this changes with structure, as opposed to just considering mass or volume to predict deposition.

Now that significant computational improvements have been made, perhaps there is the scope to return the developed models to a larger-scale and more realistic physical scenario, such as a channel or a pipe, to begin to model agglomeration in these scenarios and contrast those results with those obtained in the idealised flow scenarios. In pursuit of this, there are further adaptations that could be made to the code to speed it up: the immersed boundary solver could be further parallelised, and the particle interaction code could incorporate a search mechanism that is faster than a sequential visiting of pairs. For this new soft-sphere multi-particle technique, the proposed overlap detection methods could be taken further to see whether they can be made as robust as the presently implemented method, and there were suggestions made in Chapter 6 surrounding the ways the models could be combined to theoretically obtain speed-ups. The so called 'Method 2' showed real promise, and perhaps offered a strong mathematical solution to the problem given that it satisfies a relaxed version of the common normal constraint, which is the ideal way of defining distance metrics between convex particles.

# References

- Abbasfard, H., Evans, G., & Moreno-Atanasio, R. (2016). Effect of van der Waals force cut-off distance on adhesive collision parameters in DEM simulation. *Powder Technology, 299,* 9-18.
- Abbasov, M. (2015). New optimization algorithm for finding distance between two convex sets. 2015 International Conference "Stability and Control Processes" in Memory of V.I. Zubov (SCP), 293–294.
- Agarwal, R., & Lewis, J. (1992). Computational fluid dynamics on parallel processors. *Computing Systems in Engineering, 3*(1-4), 251-259.
- Almohammed, N. (2018). *Modeling and Simulation of Particle Agglomeration, Droplet Coalescence and Particle-Wall Adhesion in Turbulent Multiphase Flows.* Helmut Schmidt University. PhD Thesis.
- Alvelius, K. (1999). Random forcing of three-dimensional homogeneous turbulence. *Physics of Fluids*, *11*, 1880–1889.
- Ardekani, M. N., Costa, P., Breugem, W. P., & Brandt, L. (2016). Numerical study of the sedimentation of spheroidal particles. *International Journal of Multiphase Flow, 87*, 16-34.
- Aref, H., & Balachandar, S. (2018). *A First Course in Computational Fluid Dynamics*. Cambridge University Press.
- Bailey, S., Hultmark, M., Schumacher, J., Yakhot, V., & Smits, A. (2009). Measurement of Local Dissipation Scales in Turbulent Pipe Flow. *103*(1:014502).
- Bailly, C., & Juve, D. (1999). A Stochastic Approach to Compute Subsonic Noise Using Linearized Euler's Equations. 5th AIAA/CEAS Aeroacoustics Conference and Exhibit. American Institute of Aeronautics and Astronautics.
- Balachandar, S., & Eaton, J. (2010). Turbulent Dispersed Multiphase Flow. *Annual Review of Fluid Mechanics*, *42*(1), 111-133.
- Bassenne, M., Urzay, J., Park, G., & Moin, P. (2016). Constant-energetics physical-space forcing methods for improved convergence to homogeneous-isotropic turbulence with application to particle-laden flows. *Physics of Fluids, 28*, 035114.
- Bechara, W., Bailly, C., Lafon, P., & Candel, S. M. (1994). Stochastic Approach to Noise Modeling for Free Turbulent Flows. *American Institute of Aeronautics and Astronautics Journal*, 32(3), 455-463.
- Bellani, G., Byron, M., Collignon, A., Meyer, C., & Variano, E. (2012). Shape effects on turbulent modulation by large nearly neutrally buoyant particles. *Journal of Fluid Mechanics*, 712, 41-60.
- Bennett, R., Fischer, K., Lavoie, D., Bryant, W., & Rezak, R. (1989). Porometry and fabric of marine clay and carbonate sediments: Determinants of permeability. *Marine Geology*, 89, 127-152.

- Bergström, L. (1997). Hamaker constants of inorganic materials. *Advances in Colloid and Interface Science*. 70, pp.125-169.
- Best, J. (1992). On the entrainment of sediment and initiation of bed defects: insights from recent developments within turbulent boundary layer research. *Sedimentology*, 39, 797-811.
- Bin, H., Yang, Y., Cai, L., Zhulin, Y., Roszak, S., & Linjun, Y. (2018). Experimental study on particles agglomeration by chemical and turbulent agglomeration before electrostatic precipitators. *Powder Technology*, *335*, 186-194.
- Borée, J., & Caraman, N. (2005). Dilute bidispersed tube flow: Role of interclass collisions at increased loadings. *Physics of Fluids*, *17*(5), 055108.
- Boyd, J. P. (2001). Chebyshev and Fourier Spectral Methods. Dover Books.
- Bretherton, F. P. (1962). The motion of rigid particles in a shear flow at low Reynolds number. Journal of Fluid Mechanics, 14(2), 284-304.
- Breuer, M., & Almohammed, N. (2015). Modeling and simulation of particle agglomeration in turbulent flows using a hard-sphere model with deterministic collision detection and enhanced structure models. *International Journal of Multiphase Flow, 73*, 171-206.
- Breuer, M., & Khalifa, A. (2019). Revisiting and improving models for the breakup of compact dry powder agglomerates in turbulent flows within Eulerian–Lagrangian simulations. *Powder Technology, 348*, 105-125.
- Breuer, M., Alletto, M., & Flow, F. (2012). Efficient simulation of particle-laden turbulent flows with high mass loadings using LES. *International Journal of Heat and Fluid Flow, 35*, 2-12.
- Buist, K., Seelen, L., Deen, N., Padding, J., & Kuipers, J. (2016). On an efficient hybrid soft and hard sphere collision integration scheme for DEM. *Chemical Engineering Science*, *153*, 363-373.
- Caraman, N., Borée, J., & Simonin, O. (2003). Effect of collisions on the dispersed phase fluctuation in a dilute tube flow: Experimental and theoretical analysis. *Physics of Fluids*, *15*(12), 3602-3612.
- Carlsson, A., Lundell, F., & Söderberg, D. (2006). The wall effect on the orientation of fibres in a shear flow. *Annual Transactions of the Nordic Rheology Society, 14*.
- Carroll, P., & Blanquart, G. (2013). A proposed modification to Lundgren's physical space velocity forcing method for isotropic turbulence. *Physics of Fluids, 25*, 105114.
- Chen, S., & Doolen, G. D. (1998). Lattice Boltzmann Method for Fluid Flows. *Annual Review of Fluid Mechanics*, 30, 329-364.
- Chen, H., Grey, C.P. (2008). Molten Salt Synthesis and High Rate Performance of the "Desert-Rose" form of LiCoO<sub>2</sub>. *Advanced Materials*. 20(11), pp.2206-2210.
- Chen, S., & Li, S. (2020). Collision-induced breakage of agglomerates in homogenous isotropic turbulence laden with adhesive particles. *Journal of Fluid Mechanics*, 902, A28.

- Chen, S., Li, S., & Marshall, J. S. (2019). Exponential scaling in early-stage agglomeration of adhesive particles in turbulence. *Physical Review Fluids*, *4*, 024304.
- Chouippe, A., & Uhlmann, M. (2015). Forcing homogeneous turbulence in direct numerical simulation of particulate flow with interface resolution and gravity. *Physics of Fluids, 27*, 123301.
- Costa, P., Boersma, B. J., Westerweel, J., & Breugem, W.-P. (2015). Collision model for fully resolved simulations of flows laden with finite-size particles. *Physical Review E, 92*, 053012.
- Cundall, P. A., & Strack, O. D. (1979). A discrete numerical model for granular assemblies. *Géotechnique*, 29(1), 47-65.
- Dearing, S., Campolo, M., & Capone, A. (2013). Phase discrimination and object fitting to measure fibers distribution and orientation in turbulent pipe flows. *Experiments in Fluids*, *54*, 1419.
- Den Toonder, J., & Nieuwstadt, F. (1997). Reynolds number effects in a turbulent pipe flow for low to moderate Re. *Physics of Fluids*, 9(11), 3398-3409.
- Derjaguin, B., & Landau, L. (1941). Theory of the stability of strongly charged lyophobic sols and of the adhesion of strongly charged particles in solutions of electrolytes. *Acta Physicochim. URSS.*, *14*(6), 633-662.
- Derksen, J. J. (2012). Direct numerical simulations of aggregation of monosized spherical particles in homogeneous isotropic turbulence. *AIChE Journal*, *58*, 2589-2600.
- Donev, A., Connelly, R., Stillinger, F. H., & Torquato, S. (2007). Underconstrained jammed packings of nonspherical hard particles: Ellipses and ellipsoids. *Physical Review E*, 75(5), 051304.
- Donev, A., Torquato, S., & Stillinger, F. H. (2005). Neighbor list collision-driven molecular dynamics simulation for nonspherical hard particles. *Journal of Computational Physics*, 202, 765–793.
- Dubrulle, B. (2019). Beyond Kolmogorov cascades. Journal of Fluid Mechanics, 867 P1.
- Durst, F., Miloievic, D., & Schönung, B. (1984). Eulerian and Lagrangian predictions of particulate two-phase flows: a numerical study. *Applied Mathematical Modelling*, 8(2), 101-115.
- Džiugys, A., & Peters, B. (2001). An approach to simulate the motion of spherical and nonspherical fuel particles in combustion chambers. *Granular Matter, 3*, 231-266.
- Eaton, J. (2009). Two-way coupled turbulence simulations of gas-particle flows using pointparticle tracking. *International Journal of Multiphase Flow, 35*, 792–800.
- Eaton, J., & Fessler, J. (1994). Preferential concentration of particles by turbulence. International Journal of Multiphase Flow, 20(Supplement 1), 169-209.
- Elghobashi, S. (1991). Particle-laden turbulent flows: direct simulation and closure models. *Applied Scientific Research, 48*(3-4), 301-314.

- Elghobashi, S. (1994). On predicting particle-laden turbulent flows. *Applied Scientific Research*, 52, 309-329.
- Elghobashi, S. (2007). An Updated Classification Map of Particle-Laden Turbulent Flows. *IUTAM* Symposium on Computational Approaches to Multiphase Flow: Proceedings of an *IUTAM Symposium held at Argonne National Laboratory, October 4-7, 2004:*, 3.
- Elghobashi, S. E., & Truesdell, G. C. (1989a). Direct simulation of particle dispersion in grid turbulence and homogeneous shear flows. *The Bulletin of the American Physical Society, 34*, 2311.
- Elghobashi, S. E., & Truesdell, G. C. (1989b). Direct simulation of particle dispersion in a decaying isotropic turbulence. *Seventh Symp. on Turbulent Shear Flows, M*, 121–122.
- Elghobashi, S. E., & Truesdell, G. C. (1991). On the interaction between solid particles and decaying turbulence. *Eighth Symp. on Turbulent Shear Flows, 1*, 731–736.
- Elghobashi, S. E., & Truesdell, G. C. (1993). On the two-way interaction between homogeneous turbulence and dispersed solid particles. *29th Joint Propulsion Conference and Exhibit*.
- Español, P., & Warren, P. B. (2017). Perspective: Dissipative particle dynamics. *Journal of Chemical Physics*, *146*, 150901.
- Eswaran, V., & Pope, S. (1988). An examination of forcing in direct numerical simulations of turbulence. *Computers and Fluids*, *16*(3), 257-278.
- Evans, D., & Murad, S. (1977). Singularity free algorithm for molecular dynamics simulation of rigid polyatomics. *An International Journal at the Interface Between Chemistry and Physics*, *34*(2), 327-331.
- Everaers, R., & Ejtehadi, M. R. (2003). Interaction potentials for soft and hard ellipsoids. *Physical Review E*, 67, 041710.
- Fan, F., & Ahmadi, G. (1995). A sublayer model for wall deposition of ellipsoidal particles in turbulent streams. *Journal of Aerosol Science*, *26*(5), 813-840.
- Fischer, P., Lottes, J., & Kerkemeier, S. (2008). Nek5000 [Online: Accessed 1st December 2024] Available from: http://nek5000.mcs.anl.gov .
- Fröhlich, K., Schneiders, L., Meinke, M., & Schröder, W. (2018). Validation of Lagrangian Two-Way Coupled Point-Particle Models in Large-Eddy Simulations. *Flow, Turbulence and Combustion*, 101, 317–341.
- Fröhlich, J., Hafemann, T.E., Jain, R. (2022). Phase-resolving direct numerical simulations of particle transport in liquids—from microfluidics to sediment. *GAMM Mitteilungen*, 45
- Fujita, M., & Yamaguchi, Y. (2007). Multiscale simulation method for self-organization of nanoparticles in dense suspension. *Journal of Computational Physics, 223*(1), 108-120.
- Gan, J., Yu, A., & Zhou, Z. (2016b). DEM simulation on the packing of fine ellipsoids. *Chemical Engineering Science*, *156*, 64-76.
- Gan, J., Zhou, Z., & Yu, A. (2016a). Particle scale study of heat transfer in packed and fluidized beds of ellipsoidal particles. *Chemical Engineering Science*, *144*, 201-215.

George, W. K. (2013). Lectures in Turbulence for the 21st Century. Technical Report.

- Gilbert, E. G., Johnson, D. W., & Keerthi, S. S. (1988). A fast procedure for computing the distance between complex objects in three-dimensional space. *IEEE Journal on Robotics and Automation*, 4(2), 193-203.
- Girault, I., Chadil, M.-A., & Vincent, S. (2022). Comparison of methods computing the distance between two ellipsoids. *Journal of Computational Physics*, *458*, 111100.
- Goldstein, H. (1950). Classical Mechanics. Boston: Addison-Wesley.
- Gouesbet, G., & Berlemont, A. (1999). Eulerian and Lagrangian approaches for predicting the behaviour of discrete particles in turbulent flows. *Progress in Energy and Combustion Science*, *25*(2), 133-159.
- Government Office for Science. (2021). Large-scale computing: the case for greater UK coordination. Gov.Uk. Retrieved 12 10, 2023, from https://www.gov.uk/government/publications/large-scale-computing-the-case-forgreater-uk-coordination
- Gregson, C.R., Goddard, D.T., Sarsfield, M.J., Taylor, R. J. (2011). Combined electron microscopy and vibrational spectroscopy study of corroded Magnox sludge from a legacy spent nuclear fuel storage pond. *Journal of Nuclear Materials*. 412, 45-156.
- Gu, Y., Ozel, A., & Sundaresan, S. (2016). A modified cohesion model for CFD–DEM simulations of fluidization. *Powder Technology*, 296, 17-28.
- Håkansson, K. K., Wittberg, L. P., & Söderberg, D. (2013). Measurement of width and intensity of particle streaks in turbulent flows. *Experiments in Fluids*, *54*, 1555.
- Hamaker, H. (1937). The London—van der Waals attraction between spherical particles. *Physica, 4*(10), 1058-1072.
- Hiller, R. (1981). Der Einfluss von Partikelstoß und Partikel-Haftung auf die Abscheidung von Partikeln in Faserfiltern. Universität Fridericiana Karlsruhe, Germany. PhD Thesis.
- Ho, C. A., & Sommerfeld, M. (2002). Modelling of micro-particle agglomeration in turbulent flows. *Chemical Engineering Science*, *57*(15), 3073-3084.
- Houlsby, G. (2009). Potential particles: a method for modelling non-circular particles in DEM. *Computational Geotechnology*, *3*6, 953-959.
- Hoomans, B., Kuipers, J., Briels, W., & van Swaaij, W. (1996). Discrete particle simulation of bubble and slug formation in a two-dimensional gas-fluidised bed: A hard-sphere approach. *Chemical Engineering Science*, *51*(1), 99-118.
- Hu, H. (1996). Direct simulation of flows of solid-liquid mixtures. *International Journal of Multiphase Flow, 22*(2), 335-352.
- Hu, H., Patankar, N., & Zhu, M. (2001). Direct numerical simulation of fluid–solid systems using the arbitrary Lagrangian–Eulerian technique. *Journal of Computational Physics*, 169, 427–462.
- Huang, H., Yang, X., Krafczyk, M., & Lu, X.-Y. (2012). Rotation of spheroidal particles in Couette flows. *Journal of Fluid Mechanics*, 692, 369-394.

- Infrastructure and Projects Authority. (2023). Setting up major projects for success: A site visit to Sellafield. [Online: Accessed 06/11/23]. Available from: <u>https://ipa.blog.gov.uk/2023/02/21/setting-up-major-projects-for-success-a-site-visit-to-sellafield/</u>
- Ishimoto, K. (2023). Jeffery's Orbits and Microswimmers in Flows: A Theoretical Review. *Journal* of the Physical Society of Japan, 92, 6.
- Israelachvili, J. N. (1992). Intermolecular and Surface Forces. London: Academic Press Inc.
- Jain, R., Tschisgale, S., & Fröhlich, J. (2019). A collision model for DNS with ellipsoidal particles in viscous fluid. *International Journal of Multiphase Flow*, *120*, 103087.
- Jain, R., Tschisgale, S., & Fröhlich, J. (2021). Impact of shape: DNS of sediment transport with non-spherical particles. *Journal of Fluid Mechanics*, 916, A38.
- Jain, R., Tschisgale, S., & Frohlich, J. (2022). A collision model for DNS with ellipsoidal particles in viscous fluid [Corrigendum]. *International Journal of Multiphase Flow, 150*, 104009.
- Jamshed, S. (2015). Using HPC for Computational Fluid Dynamics. Elsevier Science.
- Janin, J., Duval, F., Friess, C., & Sagaut, P. (2021). A new linear forcing method for isotropic turbulence with controlled integral length scale. *Physics of Fluids*, 33(4), 045127.
- Jeffery, G. B. (1922). The Motion of Ellipsoidal Particles Immersed in a Viscous Fluid. Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character, 102(715), 161-179.
- Jiang, L., Wang, C., Liu, S., Sun, C., & Calzavarini, E. (2022). Dynamics of finite-size spheroids in turbulent flow: the roles of flow structures and particle boundary layers. *Journal of Fluid Mechanics*, 939, A22.
- Kim, D., & Choi, H. (2006). Immersed boundary method for flow around an arbitrarily moving body. *Journal of Computational Physics, 212*, 662–680.
- Kloss, C., Goniva, C., Hager, A., Amberger, S., & Pirker, S. (2012). Models, algorithms and validation for opensource DEM and CFD-DEM. *Progress in Computational Fluid Dynamics*, *12*, 140-152.
- Kobayashi, T., Tanaka, T., Shimada, N., & Kawaguchi, T. (2013). DEM–CFD analysis of fluidization behavior of geldart group a particles using a dynamic adhesion force model. *Powder Technology, 248*, 143–152.
- Kolmogorov, A. (1941). The Local Structure of Turbulence in Incompressible Viscous Fluid for Very Large Reynolds Numbers. *Doklady Akademiia Nauk SSSR, 30*, 301-305.
- Kraichnan, R. H. (1970). Diffusion by a Random Velocity Field. Physics of Fluids, 13(1), 22-31.
- Kuerten, J. (2016). Point-Particle DNS and LES of Particle-Laden Turbulent flow a state-of-theart review. *Flow, Turbulence and Combustion, 97*, 689–713.
- Kuerten, J., & Vreman, A. (2005). Can turbophoresis be predicted by large-eddy simulation? *Physics of Fluids, 17*(1), 011701-1/4.

- Kulick, J. F. (1994). Particle response and turbulence modification in fully developed channel flow. *Journal of Fluid Mechanics*, *277*, 109–134.
- Langston, P., Ai, J., & Yu, H. (2013). Simple shear in 3D DEM polyhedral particles and in a simplified 2D continuum model. *Granular Matter, 15*, 595–606.
- Lee, S., & Durst, F. (1982). On the motion of particles in turbulent duct flows. *International Journal of Multiphase Flow,* 33, 125-146.
- Li, Y., Yu, Q., Gao, S., & Flemming, B. (2020). Settling velocity and drag coefficient of platy shell fragments. *Sedimentology*, *67*, 2095-2110.
- Lin, A., & Han, S.-P. (2022). On the distance between two ellipsoids. *SIAM Journal on Optimization*, 13(1), 298-308.
- Löffler, F., & Muhr, W. (1972). Die Abscheidung von Feststoffteilchen und Tropfen an Kreiszylindern infolge von Tr<sup>°</sup> agheitskr<sup>°</sup> aften. *Chemie Ingenieur Technik, 44*(8), 510-514.
- London, F. (1930). Zur Theorie und Systematik der Molekularkräfte. *Zeitschrift für Physik*, 63(3-4), 245.
- Loth, E. (2023). *Fluid Dynamics of Particles, Drops, and Bubbles*. Cambridge: Cambridge University Press.
- Lu, G., Third, J., & Müller, C. (2015). Discrete element models for non-spherical particle systems: From theoretical developments to applications. *Chemical Engineering Science*, *127*, 425-465.
- Lucci, F., Ferrante, A., & Elghobashi, S. (2010). Modulation of isotropic turbulence by particles of Taylor length-scale size. *Journal of Fluid Mechanics*, 650, 5-55.
- Lundgren, T. S. (2003). *Linearly Forced Isotropic Turbulence*. Stanford University, Department of Mechanical Engineering. Centre for Turbulence Research.
- Ma, H., Zhou, L., Liu, Z., Chen, M., Xia, X., & Zhao, Y. (2022). A review of recent development for the CFD-DEM investigations of non-spherical particles. *Powder Technology, 412,* 117972.
- Machado, M., Moreira, P., Flores, P., & Lankarani, H. M. (2012). Compliant contact force models in multibody dynamics: Evolution of the Hertz contact theory. *Mechanism and Machine Theory*, 53, 99-121.
- Magan, R., & Sureshkumar, R. (2006). Multiscale-linking simulation of irreversible colloidal deposition in the presence of DLVO interactions. *Journal of Colloid Interface Science*, 297(2), 389-406.
- Mallouppas, G., George, W. K., & van Wachem, B. G. (2013). New forcing scheme to sustain particle-laden homogeneous and isotropic turbulence. *Physics of Fluids*, 083304.
- Mandø, M., Yin, C., Sørensen, H., & Rosendahl, L. (2007). On the modelling of motion of nonspherical particles in two-phase flow. *6th International Conference on Multiphase Flow ICMF 2007*.

- Marchioli, C., & Soldati, A. (2013). Rotation statistics of fibers in wall shear turbulence. *Acta Mechanica*, *224*, 2311–2329.
- Marchioli, C., Giusti, A., Salvetti, M., & Soldati, A. (2003). Direct numerical simulation of particle wall transfer and deposition in upward turbulent pipe flow. *International Journal of Multiphase Flow, 29*(6), 1017-1038.
- Marchioli, C., Soldati, A., Kuerten, J., Arcen, B., Taniere, A., Goldensoph, G., . . . Portela, L. (2008). Statistics of particle dispersion in direct numerical simulations of wall-bounded turbulence: Results of an international collaborative benchmark test. *International Journal of Multiphase Flow, 34*(9), 879-893.
- Marchioli, C., Zhao, L., & Andersson, H. I. (2016). On the relative rotational motion between rigid fibers and fluid in turbulent channel flow. *Physics of Fluids, 28*, 013301.
- Mauleon Amieva, A. (2020). *The phase behaviour of active colloidal rollers*. PhD Thesis. University of Bristol.
- Maxey, M. (1987). The motion of small spherical particles in a cellular flow field. *Physics of Fluids*, *30*(7), 1915-1928.
- Maxey, M., & Riley, J. (1983). Equation of motion for a small rigid sphere in a nonuniform flow. *Physics of Fluids*, *26*(4), 883-889.
- McConkey, R., Yee, E., & Lien, F.-S. (2021). A curated dataset for data-driven turbulence modelling. *Scientific Data*, *8*, 255.
- Mihajlovic, M., Roghair, I., & Annaland, M. (2020). On the numerical implementation of the Van der Waals force in Soft-Sphere Discrete Element Models for gas-solid fluidization. *Chemical Engineering Science*, *226*, 115794.
- Mindlin, R. D., & Deresiewicz, H. (1953). Elastic Spheres in Contact Under Varying Oblique Forces. *Journal of Applied Mechanics, 20*(3), 327-344.
- Mittal, R. (1999). Planar Symmetry in the Unsteady Wake of a Sphere. *American Institute of Aeronautics and Astronautics Journal*, *37*(3), 388-390.
- Mochizuki, S., & Nieuwstadt, F. (1996). Reynolds-number-dependence of the maximum in the streamwise velocity fluctuations in wall turbulence. *Experiments in Fluids, 21*, 218–226.
- Mohd-Yusof, J. (1997). Combined Immersed Boundaries/B-Splines Methods for Simulations of Flows in Complex Geometries. *Center for Turbulence Research, Annual Research Briefs*.
- Mohd-Yusof, J. (1998). Development of immersed boundary methods for complex geometries. *Center for Turbulence Research, Annual Research Briefs*.
- Monaghan, J. J. (1992). Smoothed Particle Hydrodynamics. *Annual Review of Astronomy and Astrophysics*, *30*, 543-574.
- Mortensen, P. H., Andersson, H. I., Gillissen, J. J., & Boersma, B. J. (2008). Dynamics of prolate ellipsoidal particles in a turbulent channel flow. *Physics of Fluids*, *20*(9), 093302.
- Mortimer, L. F. (2019). *Particle transport, interaction and agglomeration for nuclear reactor and nuclear waste flow applications*. University of Leeds. PhD Thesis.

- Mortimer, L. F., & Fairweather, M. (2021). Assessment of behavioural modification techniques through immersed boundary method simulation of binary particle interactions in isotropic turbulence. *Physics of Fluids*, *33*(7), 073307.
- Mortimer, L. F., Njobuenwu, D. O., & Fairweather, M. (2019). Near-wall dynamics of inertial particles in dilute turbulent. *Physics of Fluids*, *31*(6), 063302-1.
- Mortimer, L. F., Njobuenwu, D. O., & Fairweather, M. (2020). Agglomeration dynamics in liquid– solid particle-laden turbulent channel flows using an energy-based deterministic approach. *Physics of Fluids*, *32*(4), 043301.
- Nagata, T., Hosaka, M., Takahashi, S., Shimizu, K., Fukuda, K., & Obayashi, S. (2020). A simple collision algorithm for arbitrarily shaped objects in particle-resolved flow simulation using an immersed boundary method. *International Journal for Numerical Methods in Fluids*, 92, 1256–1273.
- Njobuenwu, D., & Fairweather, M. (2013a). Inertial particle behaviour in turbulent fluid flow. *Computer Aided Chemical Engineering*, *32*, 253-258.
- Njobuenwu, D., & Fairweather, M. (2013b). Non-spherical particle translation and orientation in wall-bounded turbulence. *11th International Conference of Numerical Analysis and Applied Mathematics*, 132-135.
- Njobuenwu, D., & Fairweather, M. (2014). Effect of Shape on Inertial Particle Dynamics in a Channel Flow. *Flow, Turbulence and Combustion*, 92, 83-101.
- Njobuenwu, D., & Fairweather, M. (2015). Dynamics of single, non-spherical ellipsoidal particles in a turbulent channel flow. *Chemical Engineering Science*, *123*, 265-282.
- Obukhov, A. (1941). On the distribution of energy in the spectrum of turbulent flow. *Dokl. Akad. Nauk SSSR.*
- Ochsenbein, D., Vetter, T., Schorsch, S., Morari, M., & Mazzotti, M. (2015). Agglomeration of Needle-like Crystals in Suspension: I. Measurements. *Crystal Growth and Design*, *15*.
- Office for Nuclear Regulation. (2023). Sellafield, Decommissioning, Fuel and Waste (SDFW). [Online: Accessed 06/11/2023]. Available from: <u>https://www.onr.org.uk/sellafield-dfw.htm</u>
- Oseen, C. (1927). Neuere Methoden und Ergebnisse in der Hydrodynamik. *Akademische Verlagsgesellschaft, Leipzig.*
- Pazouki, A., Mazhar, H., & Negrut, D. (2012). Parallel collision detection of ellipsoids with applications in large scale multibody dynamics. *Mathematics and Computers in Simulation, 82*, 879-894.
- Peskin, C. (1977). Numerical analysis of blood flow in the heart. *Journal of Computational Physics, 25*(3), 220-252.
- Peskin, C. (2002). The immersed boundary method. Acta Numerica, 11, 479-517.
- Podlozhnyuk, A., Pirker, S., & Kloss, C. (2017). Efficient implementation of superquadric particles in Discrete Element Method within an open-source framework. *Computational Particle Mechanics, 4, 101-118*.
Pope, S. (2000). Turbulent Flows. Cambridge University Press.

- Pumir, A., & Wilkinson, M. (2011). Orientation statistics of small particles in turbulence. *New Journal of Physics*, *13*, 093030.
- Qi, D., & Luo, L. (2002). Transitions in rotations of a nonspherical particle in a three-dimensional moderate Reynolds number Couette flow. *Physics of Fluids*, *14*(12), 4440-4443.
- Qi, D., & Luo, L.-S. (2003). Rotational and orientational behaviour of three-dimensional spheroidal particles in Couette flows. *Journal of Fluid Mechanics*, 477, 201-213.
- Rani, S. L., Winkler, C., & Vanka, S. (2004). Numerical simulations of turbulence modulation by dense particles in a fully developed pipe flow. *Powder Technology*, 141(1-2), 80-99.
- Rani, S., & Vanka, S. (2000). Direct numerical simulation of two-way coupling effects in a particle-laden turbulent pipe flow. *In: 36th AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit*.
- Reeks, M. (1983). The transport of discrete particles in inhomogeneous turbulence. *Journal of Aerosol Science, 14*, 729-739.
- Richardson, L. F. (1922). *Weather Prediction by Numerical Processes*. Cambridge University Press.
- Rietema, K., Cottaar, E., & Piepers, H. (1993). The effects of interparticle forces on the stability of gas-fluidized beds—II. Theoretical derivation of bed elasticity on the basis of van der Waals forces between powder particles. *Chemical Engineering Science*, 48(9), 1687-1697.
- Rosales, C., & Meneveau, C. (2005). Linear Forcing in Numerical Simulations of isotropic Turbulence: Physical Space Implementations and Convergence Properties. *Physics of Fluids, 17*(9), 095106.
- Ruan, X., Chen, S., & Li, S. (2021). Effect of long-range Coulomb repulsion on adhesive particle agglomeration in homogeneous isotropic turbulence. *Journal of Fluid Mechanics*, 915, A131.
- Saffman, P., & Turner, J. (1956). On the collision of drops in turbulent clouds. *Journal of Fluid Mechanics, 1*(1), 16-30.
- Schiller, P., Krüger, S., & Wahab, M. M. (2011). Interactions between Spheroidal Colloidal Particles. *Langmuir*, *27*(17), 10429–10437.
- Schmidt, S., & Breuer, M. (2017). Source term based synthetic turbulence inflow generator for eddy-resolving predictions of an airfoil flow including a laminar separation bubble. *Computers and Fluids, 146*, 1-22.
- Schneiders, L., Fröhlich, K., Meinke, M., & Schröder, W. (2019). The decay of isotropic turbulence carrying non-spherical finite-size particles. *Journal of Fluid Mechanic*, 875, 520-542.
- Sciacchitano, A., & Wieneke, B. (2016). PIV uncertainty propagation. *Measurement Science and Technology*, *27*, 084006.

- Shahidzadeh, N., Bonn, D., Ragil, K., Broseta, D. and Meunier, J. (1998). Sequence of Two Wetting Transitions Induced by Tuning the Hamaker Constant. *Physical Review Letters*, 80, 3992-3995.
- Shardt, O., & Derksen, J. J. (2012). Direct simulations of dense suspensions of non-spherical particles. *International Journal of Multiphase Flow, 47*, 25-36.
- Shirolkar, J., Coimbra, C., & McQuay M.Q. (1996). Fundamental aspects of modeling turbulent particle dispersion in dilute flows. *Progress in Energy and Combustion Science, 22*(4), 363-399.
- Shokri, R., Ghaemi, S., Nobes, D., & Sanders, R. (2017). Investigation of particle-laden turbulent pipe flow at high-Reynolds-number using particle image/tracking velocimetry (PIV/PTV). *International Journal of Multiphase Flow*, 89, 136-149.
- Slootman, A., de Kruijf, M., Glatz, G., Eggenhuisen, J., Zuhlke, R., & Reijmer, J. (2023). Shapedependent settling velocity of skeletal carbonate grains: implications for calciturbidites. *Sedimentology*, 70, 1683-1722.
- Snyder, W. H., & Lumley, J. L. (1971). Some measurements of particle velocity autocorrelation functions in a turbulent flow. *Journal of Fluid Mechanics*, 48(1), 41-71.
- Spoelstra, J. (1989). The Modelling of Oil Agglomeration of Coal Fines. *Journal of Computational and Applied Mathematics, 28,* 359-366.
- Squires, K., & Eaton, J. (1991). Preferential concentration of particles by turbulence. *Physics of Fluids*, *3*, 1169-1178.
- Taylor, G. I. (1935). Statistical theory of turbulence. *Proceedings of the Royal Society of London,* Series A: Mathematical, Physical and Engineering Sciences, 421-444.
- Tennekes, H., & Lumley, J. (1972). A First Course in Turbulence. The MIT Press.
- Trefethen, L. (2000). Spectral Methods in MATLAB. Philadelphia: SIAM.
- Tseng, Y.-H., & Ferzinger, J. H. (2003). A ghost-cell immersed boundary method for flow in complex geometry. *Journal of Computational Physics*, *192*, 593-623.
- Tsuji, Y., Kawaguchi, T., & Tanaka, T. (1993). Discrete particle simulation of two-dimensional fluidized bed. *Powder Technology*, 77(1), 79-87.
- Uhlmann, M. (2005). An immersed boundary method with direct forcing for the simulation of particulate flows. *Journal of Computational Physics*, 209(2), 448-476.
- Uteshev, A., & Goncharova, M. (2008). Computation of the distance from an ellipsoid to a linear surface and a quadric in ℝ<sup>^</sup>n. *Doklady Mathematics*, *77*, 269-272.
- van Wachem, B. G., Zastawny, M., Zhao, F., & Mallouppas, G. (2015). Modelling of gas–solid turbulent channel flow with non-spherical particles with large Stokes numbers. *Multiphase Flow,* 68, 80-92.
- Varaksin, A., & Polezhaev, Y. P. (2000). Effect of particle concentration on fluctuating velocity of the disperse phase for turbulent pipe flow. *International Journal of Heat and Fluid Flow,* 21(5), 562-567.

- Verwey, E., & Overbeek, J. (1955). Theory of the stability of lyophobic colloids. *Journal of Colloid Science*, *10*(2), 224-225.
- Voth, G. A., & Soldati, A. (2017). Anisotropic Particles in Turbulence. *Annual Review of Fluid Mechanics*, 49(1), 249-276.
- Vreman, A. (2007). Turbulence characteristics of particle-laden pipe flow. *Journal of Fluid Mechanics*, 584, 235-279.
- Weber, M., Hoffman, D., & Hrenya, C. (2004). Discrete-particle simulations of cohesive granular flow using a square-well potential. *Granular Matter*, 6, pages239–254.
- Wellmann, C., Lillie, C., & Wriggers, P. (2008). A contact detection algorithm for superellipsoids based on the common-normal concept. *Engineering Computations*, *25*(5), 432 442.
- Westerweel, J. (1997). Fundamentals of digital particle image velocimetry. *Measurement Science and Technology*, *8*, 1379.
- White, L. R. (1983). On the deryaguin approximation for the interaction of macrobodies. *Journal* of Colloid and Interface Science, 95(1), 286-288.
- Williams, J., & O'Connor, R. (1999). Discrete element simulation and the contact problem. Archives of Computational Methods in Engineering, 6, 279–304.
- Winterton, R. (1970). Van der Waals forces. Contemporary Physics, 11(6), 559-574.
- Wolde, B. (2023). Simulation of Behavioural Modification Effects and Deposition in Suspension Waste Pipe Flows. PhD Thesis. University of Leeds.
- Wynn, E. (2009). Simulations of rebound of an elastic ellipsoid colliding with a plane. *Powder Technology*, *1*96(1), 62-73.
- Yamamoto, Y., Potthoff, M., Tanaka, T., Kajishima, T., & Tsuji, Y. (2001). Large-eddy simulation of turbulent gas particle flow in a vertical channel: effect of considering inter-particle collisions. *Journal of Fluid Mechanics, 442*(1), 303-334.
- Ye, M., van der Hoef, M., & Kuipers, J. (2004). A numerical study of fluidization behavior of geldart a particles using a discrete particle model. *Powder Technology*, *139*, 129–139.
- Zhao, F., & van Wachem, B. G. (2013). A novel quaternion integration approach for describing the behaviour of non-spherical particles. *Acta Mechanica*, *224*, 3091–3109.
- Zhao, L., Challabotla, N. R., Andersson, H. I., & Variano, E. A. (2015). Rotation of Nonspherical Particles in Turbulent Channel Flow. *Physical Review Letters*, *115*, 244501.
- Zheng, G., Zhou, Z., & Yu, A. (2013). Contact forces between viscoelastic ellipsoidal particles. *Powder Technology, 248*, 25-33.
- Zhou, Z., Pinson, D., Zou, R., & Yu, A. (2011). Discrete particle simulation of gas fluidization of ellipsoidal particles. *Chemical Engineering Science*, 66(23), 6128–6145.