Imperial College London
Department of Mechanical Engineering

Improving the suitability of Immersed Boundary methods to model granular material

John C Pennefather
MA, MSc

September 2014

Supervised by Dr Berend van Wachem

Submitted in part fulfilment of the requirements for the degree of
Doctor of Philosophy in Mechanical Engineering of Imperial College London
and the Diploma of Imperial College London

Declaration of originality

I declare that the work presented in this thesis is my own and expressed in my own words except where declared otherwise. Any work used by any other author(s) in any form has been properly acknowledged where used. A list of all the references used has been included. This work has not already been submitted for any other degree and is not being submitted concurrently for any other degree.

Copyright

The copyright of this thesis rests with the author and is made available under a Creative Commons Attribution Non-Commercial No Derivatives licence. Researchers are free to copy, distribute or transmit the thesis on the condition that they attribute it, that they do not use it for commercial purposes and that they do not alter, transform or build upon it. For any reuse or redistribution, researchers must make clear to others the licence terms of this work.

Abstract

Embankment dams consisting of broadly graded materials have been found to be susceptible to internal erosion. The need to minimise both the cost and risk requires an accurate characterisation of a range of materials, in terms of their susceptibility.

Despite significant theoretical and experimental developments, there is a lack of consensus as to how to determine the susceptibility of a material to internal erosion, and what force conditions are required to initiate erosion in susceptible materials. A significant weakness in current geotechnical modelling is the limitation to the use of a large-scale averaged representation of the flow. The immersed-boundary method alternative is a small-scale resolving method capable of elucidating the flow at arbitrarily small scales, rectifying this limitation.

In this research the Immersed-Boundary Method (IBM) is used to accurately model the interaction of the particle surface and the fluid flow. The method is adapted and applied to modelling internal erosion, with the intention of understanding the micro-scale mechanisms involved. In particular, the fluid-particle coupling in the IBM model is adapted to be suitable to the case of densely packed particles, characteristic of embankment structures. In this thesis significant improvements are made to the immersed-boundary method's ability to model fluid particle flows. Improvements to the momentum exchange between the particles surface both when the particle is removed from others, and when two particles are within each others range of influence, are made. Additionally the dependence on mesh resolution of the flow field in proximity of the particle's surface is reduced, allowing more tractable simulation of the large scale problems characteristic of geotechnical problems. The removal of sporadic pressure fluctuations improves the accuracy of predictions. The resulting tool will provide insight into the science under-pinning internal erosion, guiding best practice with regards to predicting and preventing the onset and propagation of internal erosion.

Acknowledgements

It is with immense gratitude that I acknowledge the AXA Research Fund. Without their generous support, this thesis would not be.
I am also very grateful to many of the staff and students of Imperial College London, in particular my supervisor Dr. Berend van Wachem, for their help, inspiration and friendship. Not least for their impact on this work.

Contents

List of Figures **16**

List of Tables **19**

Nomenclature **20**

1. Introduction **22**

1.1. Overview and aims . . . 23
1.1.1. Research objectives . . . 23
1.1.2. Overview of thesis . . . 23
1.2. Embankment structures . . . 24
1.2.1. Internal Erosion . . . 24
1.2.1.1. Suffusion . . . 24
1.2.1.2. Suffosion . . . 24
1.2.1.3. Concentrated leak erosion . . . 25
1.2.1.4. Backward erosion . . . 25
1.2.2. Requirement for advanced numerical modelling . . . 25
1.3. Solid-fluid multiphase flows . . . 25
1.3.1. Immersed-boundaries representing particles . . . 25
1.3.2. Fluid governing equations . . . 26
1.3.3. Contact dynamics . . . 27
1.3.4. Lubrication force . . . 27
1.3.4.1. Definition . . . 27
1.3.4.2. Associated fluid dynamics . . . 27
1.3.4.3. Motivation for study . . . 28
1.3.4.4. Lubrication force in the immersed-boundary framework . . . 28
1.3.4.5. Combined lubrication force and contact . . . 28
1.3.5. Existing techniques for resolving solid-fluid flows . . . 28
1.3.5.1. Arbitrary Lagrangian-Eulerian methods . . . 28
1.3.5.2. Lattice-Boltzmann methods . . . 29
1.3.5.3. Immersed boundary methods . . . 29
1.3.6. Existing techniques for resolving fluid flows . . . 31
1.3.6.1. Direct Numerical Simulation . . . 31
1.3.7. Force coupling between fluid and solid phases . . . 32
1.3.8. Feasibility . . . 34

2. Literature review **36**
2.1. Forcing techniques . . . 36
2.1.1. Continuous forcing . . . 36
2.1.1.1. Introduction . . . 36
2.1.1.2. Methodology . . . 36
2.1.1.3. Method parameters . . . 37
2.1.2. Discrete forcing . . . 37
2.1.2.1. Direct forcing methods . . . 38
2.1.2.2. Ghost cell method . . . 39
2.1.2.3. Cut-cell method . . . 41
2.2. Lubrication force . . . 42
2.2.1. Analytical, experimental and empirical and predictions of lubrication force in simplified cases . . . 42
2.2.1.1. Analytic results for Stokes flow . . . 42
2.2.1.2. Further developments . . . 43
2.2.1.3. Experimental results . . . 43
2.2.2. Proposed methods for correcting the lubrication force . . . 43
2.2.2.1. Collision prevention . . . 44
2.2.2.2. Numerical investigation . . . 45
2.2.3. Applicability of existing methods . . . 45
2.2.4. Comparison of methods . . . 46
2.2.4.1. Accuracy . . . 46
2.2.4.2. Further interpretation . . . 47
2.2.5. Amalgamation of hydro-dynamic interaction and correction forces . . . 47
2.3. Mirroring immersed-boundary method . . . 48
2.3.1. Mirroring . . . 48
2.3.2. Force calculation . . . 50
2.3.3. Limitations . . . 50
2.3.3.1. Spurious oscillations . . . 50
2.4. Internal instability . . . 50
2.4.1. Assessment criteria for susceptibility to internal erosion . . . 50
2.4.1.1. Kedzi . . . 52
2.4.1.2. Kenney and Lau . . . 52
2.4.1.3. Wan and Fell . . . 54
2.4.1.4. Li and Fannin . . . 54
2.4.2. Self-filtration . . . 54

3. Numerical method **57**
3.1. Framework . . . 57
3.1.1. Single phase flow . . . 57
3.1.2. Applicability . . . 57
3.2. Enforcing the no-slip condition . . . 57
3.2.1. Determination of cell type . . . 57

3.2.2. Complete mirroring . . . 60
3.2.2.1. Developments to the mirroring equation . . . 61
3.2.2.2. Pressure condition . . . 67
3.2.3. Cut-cell partial mirroring . . . 67
3.2.3.1. Partial mirroring . . . 70
3.2.3.2. Continuity . . . 70
3.2.3.3. Adaptive mirroring . . . 82
3.2.3.4. Parity based incrementation . . . 83
3.2.4. Cut cell momentum equation . . . 83
3.2.4.1. Increased frequency of solid phase to resolve stiff collisions . . . 85
3.3. Force calculation . . . 86
3.3.1. Stencil used for interpolation . . . 86
3.3.2. Velocity gradient . . . 86
3.3.3. Pressure . . . 87
3.3.4. Special cases . . . 88
3.4. Contact model . . . 88
3.4.1. Problems associated with the use of multiple particles . . . 88
3.4.2. Discretisation Scheme . . . 89
3.4.2.1. Breakdown of overlap . . . 89
3.4.3. Force calculation from the discretised pieces . . . 92
3.4.3.1. Elastic component . . . 92
3.4.3.2. Viscous component . . . 93
3.4.3.3. Friction . . . 93
3.4.3.4. Re-combination to calculate force and torque . . . 94
3.4.4. Time incrementation of the solid phases . . . 95
3.5. Equation system solution . . . 95
3.5.1. Momentum weighted interpolation . . . 95
3.5.2. First order time scheme . . . 96
3.5.3. Matrix formulation for cut cells . . . 97

4. Particles settling under gravity **98**
4.1. Aims and Motivation . . . 98
4.2. Drag forces on a particle travelling through a viscous fluid . . . 98
4.2.1. Experimental setup . . . 99
4.2.2. Numerical setup . . . 100
4.2.2.1. Boundary conditions . . . 100
4.2.2.2. Gravity . . . 101
4.2.2.3. Refinement . . . 101
4.2.3. Stationary particle . . . 101
4.2.3.1. Equivalence of methods for the stationary case . . . 101
4.2.3.2. Dependence on refinement . . . 104
4.2.4. Moving particle . . . 109
4.2.4.1. Mesh dependence . . . 109
4.2.4.2. Force calculation . . . 114

4.2.4.3. Smoothness . . . 115
4.2.5. Comparing stationary and moving cases . . . 123
4.2.6. Dynamic case . . . 123
4.3. Oscillating sphere . . . 126
4.4. Comparison with other methods . . . 130
4.5. Discussion . . . 135

5. Lubrication forces **137**
5.1. Determination of lubrication force . . . 137
5.1.1. Prediction using dimensional analysis . . . 137
5.1.1.1. Low Reynolds number collisions . . . 137
5.1.1.2. General case . . . 146
5.1.2. Dynamic case . . . 155
5.1.3. Impact of force calculation method . . . 164
5.2. Collision dynamics . . . 164
5.3. Discussion and conclusions . . . 166

6. Flow through granular materials **168**
6.1. Constrained particles . . . 168
6.1.1. Simulation and objectives . . . 168
6.1.2. Model used . . . 169
6.1.3. Particle motion . . . 171
6.1.4. Hydro-dynamic interaction forces . . . 180
6.1.5. Collision forces . . . 186
6.1.6. DEM model . . . 187
6.2. Periodicity . . . 188
6.3. Discussion and conclusions . . . 190

7. Conclusions **191**
7.1. Reduction of spurious pressure oscillations . . . 191
7.1.1. Averaged mirroring . . . 191
7.1.2. Cut-cell partial mirroring . . . 192
7.2. Contact dynamics . . . 192
7.3. Lubrication force correction . . . 193
7.4. Test cases . . . 193
7.5. Future work . . . 194
7.5.1. Force calculation . . . 194
7.5.2. Lubrication force for arbitrary shaped particles . . . 194
7.5.3. Tractability . . . 195

Bibliography **195**

A. Appendix **204**
A.1. Interpolation 204
A.1.1. Interpolation coefficients 206
A.1.1.1. Choice function 206
A.1.2. Interpolation schemes used 207
A.1.2.1. Polynomial interpolation 207
A.1.2.2. Linear interpolation 207
A.2. Non sphericity 208
A.3. Permiameter type simulation results 210
A.3.1. Greater number of spheres 210

List of Figures

1.1. Lagrangian particle in a Eulerian fluid . . . 26
1.2. Immersed boundary principle . . . 30
1.3. Fluid particle interaction . . . 33
1.4. Refinement comparison . . . 35

2.1. Glowinski model . . . 46
2.2. Fluid particle interaction . . . 49
2.3. Diagram of the mirroring of fluid properties in the mirroring immersed boundary method 49
2.4. Fresh and dead cells . . . 51
2.5. Granular material stability criteria Venn diagram . . . 51
2.6. Kenney and Lau stability criterion: H and F explanation . . . 53
2.7. Kenney and Lau stability criterion example . . . 53
2.8. Burenkova's stability criterion . . . 55
2.9. Kenney and Lau and Kezdi stability criterion explanation . . . 55
2.10. Kenney and Lau and Kezdi stability criterion comparison . . . 56

3.1. Determination of cell type . . . 59
3.2. Fluid particle interaction . . . 61
3.3. Transition of stencil . . . 62
3.4. Compacted mirroring equation . . . 63
3.5. Mirroring point transition . . . 64
3.6. Neighbours used to smooth mirroring . . . 65
3.7. Averaged mirroring . . . 66
3.8. Over specification of boundary-interior cells . . . 68
3.9. Continuity transition problem . . . 69
3.10. Cut-cut partial mirroring . . . 69
3.11. Face types . . . 71
3.12. Removed or eclipsed face . . . 72
3.13. Cut-cell continuity . . . 72
3.14. Continuity equation and mirroring of a boundary-interior cell . . . 73
3.15. Possible inconsistency of the continuity equation . . . 75
3.16. Avoidance of the inconsistency of the continuity equation . . . 76
3.17. Cut-cell continuity contributions . . . 78
3.18. Cut-cell continuity contribution interpolation . . . 81
3.19. Timestep-parity dependent incrementation . . . 84
3.20. Determination of points used to interpolate fluid properties to an immersed-boundary element . . . 87

3.21. Potentially problematic contacts . . . 89
3.22. Contact decomposition . . . 90
3.23. Collision piece . . . 91
3.24. Subdivided collision piece . . . 91
3.25. Curvature calculation . . . 92
3.26. Internal pressure field . . . 97

4.1. Experimental setup of falling viscous sphere . . . 100
4.2. Experimental setup of falling viscous sphere . . . 102
4.3. Drag force on stationary particle . . . 103
4.4. Components of drag force on stationary particle . . . 104
4.5. Drag force on stationary particle . . . 105
4.6. Shear stress force on particle in a flow . . . 106
4.7. Pressure force on particle in a flow . . . 107
4.8. Combined (total) force on particle in a flow . . . 108
4.9. Pressure around particle . . . 110
4.10. Velocity around particle . . . 112
4.11. Out of plane velocity . . . 113
4.12. Drag force . . . 114
4.13. Drag force convergence . . . 115
4.14. Partial mirroring, 4 cells per radius . . . 116
4.15. Partial mirroring, 5 cells per radius . . . 117
4.16. Complete mirroring, 4 cells per radius . . . 118
4.17. Complete mirroring, 5 cells per radius . . . 119
4.18. Convergence of force for variations of the partial mirroring method . . . 120
4.19. Enlarged section of force variation . . . 121
4.20. Comparison of methods at low resolution . . . 122
4.21. Drag force dependence on the relative motion of the immersed-boundary and Eulerian mesh . . . 123
4.22. Comparison with experimental results of ten Cate *et al.* . . . 124
4.23. Force imbalance on falling particle . . . 125
4.24. Setup of oscillating sphere simulation . . . 126
4.25. Force on oscillating sphere . . . 128
4.26. Contributions to force on oscillating sphere . . . 129
4.27. Normal velocity around particle in the cut/merged method . . . 132
4.28. Tangential velocity around particle in the cut/merged method . . . 133
4.29. Streamlines around cut/merged case . . . 134

5.1. Lubrication force setup . . . 140
5.2. Lubrication flow for different Re . . . 141
5.3. Lubrication flow for different Re . . . 142
5.4. Under refinement of the interstitial fluid between two particles . . . 144
5.5. Corrected lubrication force . . . 145
5.6. Lubrication flow for different Re . . . 148

5.7. Difference in lubrication flow for different Re compared to Stokes flow 149
5.8. Lubrication force variation with different fluid density hence Re range 150
5.9. Error in un-corrected lubrication force relative to Stokes flow prediction 151
5.10. Error in un-corrected lubrication force relative to Cox and Brenner's low Re prediction 152
5.11. Error in corrected lubrication force relative to Stokes flow prediction 153
5.12. Error in corrected lubrication force relative to Cox and Brenner's low Re prediction . 154
5.13. Deceleration of colliding spheres 157
5.14. Deceleration of colliding spheres 158
5.15. Force acting on colliding spheres 159
5.16. Force acting on colliding spheres 160
5.17. Lubrication force dependence on separation (slow) 161
5.18. Lubrication force dependence on separation (fast) 162
5.19. Effect of the developed correction to lubrication force 163
5.20. Collision force 165
5.21. Interstice during contact 166
5.22. Combined collision and lubrication force 166

6.1. Setup for constrained particle motion 169
6.2. Comparison of particle types 170
6.3. Lubrication force for non-spherical particles 171
6.4. Evolution of constrained spherical particles 173
6.5. Evolution of constrained non-spherical particles 174
6.6. Key for evolution of constrained particle figures 175
6.7. Progression of spherical particles 176
6.8. Progression of non-spherical particles 177
6.9. Velocity through constraining mesh 178
6.10. Tangential velocity through constraining mesh 179
6.11. Non-dimensionalised force on spherical particles, x 180
6.12. Non-dimensionalised force on spherical particles, y 181
6.13. Non-dimensionalised force on spherical particles, z 182
6.14. Non-dimensionalised force on non-spherical particles, x 183
6.15. Non-dimensionalised force on non-spherical particles, y 184
6.16. Non-dimensionalised force on non-spherical particles, z 185
6.17. Non-dimensionalised collision force on a spherical particle 186
6.18. Non-dimensionalised collision force on an non-spherical particle 187
6.19. Constrained particles using DEM 188

A.1. Multiple constrained particles 210

List of Tables

4.1. Fluid properties for settling sphere 99
4.2. Settling sphere experimental results 100
4.3. Mesh resolutions 101
4.4. Simulation properties 102
4.5. Simulation properties 123

5.1. Low Re lubrication force simulation properties 138
5.2. Higher Re lubrication force simulation properties 147
5.3. Low Re colliding particle simulation properties 155

6.1. Asperity parameters 169
6.2. Constrained flow properties 169

Nomenclature

Vectors and tensors are either denoted in **bold** or in tensor notation by subscripts i (vector) or ij (tensor), depending on the context. All units are SI units.

Roman Letters

u_{inf}	Velocity prediction for case removed from walls	[m s^{-1}]
u_{max}	Maximum obtained velocity	[m s^{-1}]
V	Volume	[m^3]
$\mathbf{I}$	Inertia tensor	[Kg m^2]
$\mathbf{a}$	Surface vector	[m]
$\mathbf{a}^{\text{Dead}}$	Surface area vector of the section of the face that lies inside of the boundary	[m^2]
$\mathbf{a}^{\text{Face}}$	Surface area vector of the original (un-cut) face	[m^2]
$\mathbf{a}^{\text{Live}}$	Surface area vector of the section of the face that lies outside of the boundary	[m^2]
$\mathbf{e}_i$	Basis vectors	[m]
$\mathbf{s}_i$	In-IB-plane vectors	[m]
$\mathbf{u}$	Velocity	[m s^{-1}]
$\mathbf{x}$	Position	[m]
A	Area	[m^2]
A^{Dead}	Surface area of the section of the face that lies inside of the boundary	[m^2]
A^{Face}	Surface area of the original (un-cut) face	[m^2]
A^{Live}	Surface area of the section of the face that lies outside of the boundary	[m^2]
b	Harmonic average of radius of curvature	[m]
B, C	Constants used for non-dimensionalised lubrication force prediction (Low Re)	[−]
C	Curvature	[m]
c_v	Centre of volume	[m]
d_p	Particle diameter	[m]
E	Young's modulus	[N m^{-2}]
E^*	Effective Young's modulus	[N m^{-2}]
F	Volume flux	[m^3 s^{-1}]
f	Force	[N]
f_c	Sub grid lubrication force contribution	[N]
F_d	Volume flux through the 'dead' face	[m^3 s^{-1}]
F_l	Volume flux through the 'live' face	[m^3 s^{-1}]
F_t	Volume flux through the entire (total) face	[m^3 s^{-1}]
g	Distribution function	[−]
h	Separation between surfaces	[m]
L	Length scale of the large scale eddies	[m]
m	Mass	[Kg]

N	Component of distance in normal direction	[m]
n_c	Contact direction/normal	[m]
p	Pressure	[N m^{-2}]
p, q	Functions used for non-dimensionalised lubrication force prediction (Any Re)	[–]
r	Radius of sphere	[m]
r^*	Effective particle radius	[m]
r^m	Radius of inscribed sphere of non-spherical particle	[m]
S	Forcing source term	[N]
T	Component of distance in tangential direction	[m]
t	Time	[s]
u	Rate of approach of surfaces	[m s^{-1}]
u_p	Particle velocity	[m s^{-1}]
w^{Dead}	Numerical weight of the 'dead' face	[–]
w^{Live}	Numerical weight of the 'live' face	[–]

Greek Letters

η	Length scale of the smallest scale eddies	[m]
Γ	Particle surface	[m^2]
μ	dynamic viscosity	[N s m^{-2}]
μ_f	Fluid viscosity	[N s m^{-2}]
ϕ	Interpolation coefficient	[–]
ϕ_i^{Dead}	(Direct) interpolation coefficient to the centre of area of the 'dead' section of the face of the i^{th} cell in the stencil	[–]
ϕ_i^{Live}	(Direct) interpolation coefficient to the centre of area of the 'live' section of the face of the i^{th} cell in the stencil	[–]
ρ	Density	[Kg m^{-3}]
ρ_f	Fluid density	[Kg m^{-3}]
ρ_p	Particle density	[Kg m^{-3}]
τ	Stress tensor	[N m^{-2}]

Superscripts

i	Direction	[–]
j	Direction	[–]
k	Index of neighbouring cell	[–]

Non-Dimensional Numbers

α	Ratio of distance to IB of mirroring and mirrored points	[–]
ϵ	Non-dimensional geometric separation parameter	[–]
e	Relative error	[–]
n	Fluid cells per particle radius	[–]
Re	Reynolds number	[–]
Re_p	Particle Reynolds number	[–]
Re_L	Reynolds number of large-scale flow structures	[–]

1. Introduction

Embankment dams and flood embankments benefit society by providing flood protection, water supply, power generation, etc. However, collapse of these structures can be catastrophic causing loss of life and significant damage to property and infrastructure. ICOLD (International Commission on Large Dams) estimates that there are over 45,000 large dams globally. Although the risk of a failure of an individual dam is low, the high consequence of any failure and the enormous number of dams globally requires careful management of risks associated with dams. Defra estimates that there are 7500 km of flood embankments in the UK (some of these are 500 years old). The ASCE estimates that there are 100,000 miles of levees in the US. The vulnerability of these structures was highlighted by the floods in the UK in 2007 as well as hurricanes Katrina and Rita in the US in 2005.

Charles and Boden (1985) stated that internal erosion is probably the greatest and the least well understood hazard posed to embankment dams – there has been little progress in its understanding since. A serious incident at WAC Bennett dam in British Columbia, Canada, in the late 1990s highlighted the vulnerability of embankment dams to internal erosion. Foster *et al.* [1] conducted an international study of 11,192 dams, of the 136 failures identified, about 50% related to internal erosion. The need to consider seepage/internal erosion in flood embankment design is well established *e.g.* Dyer [2].

The vulnerability of Europe's embankment infrastructure to internal erosion in particular prompted the formation in 1993 of the ICOLD European Working Group (EWG) on Internal Erosion in Embankment Dams. The UK environment agency recently commissioned WS Atkins to review the Interim Guide to Quantitative Risk Assessment at UK Reservoirs, which addresses the risks related to internal erosion in British dams. The American Federal Emergency Management Agency (FEMA) also recognises the need for research into internal erosion.

Typically the Computational Fluid Dynamics (CFD) and particle dynamics used to simulate sections of embankment structures is a sub-grid scale approach. Is such models the complex flow amongst the particles is not elucidated. Consequently the developed models are dependent on empirical relations, and do not capture the complex behaviour of granular systems spontaneously.

In this thesis the Immersed-Boundary Method (IBM), where by individual particles in a Lagrangian framework are added to a Eulerian fluid, is modified to be suitable for the required modelling. In particular the Mirroring Immersed-Boundary Method (MIBM) is adapted to improve its capability to model granular materials. The developed method addresses some of the previous limitations of immersed-boundary methods including:

- Difficulty handling particles in close proximity to one another. In particular the ability of the method to capture the hydro-dynamic interaction forces acting on particle in close proximity to one another.
- Limitation of contact between particles, especially for non-spherical particles.

- The large number of Eulerian fluid cells required to accurately model a Lagrangian particle, with affects the feasibility of modelling granular materials.
- Oscillations in the pressure field surrounding the particles (or immersed-boundaries) that have limited other implementations of the immersed-boundary method.

1.1. Overview and aims

1.1.1. Research objectives

The overall aim for this investigation is to develop a numerical method to accurately model granular materials. This will allow the identification of the micro-mechanical basis for current guidelines, and potential improvement. Of particular interest is the relationship between susceptibility to internal erosion and particle size distribution (PSD). This aim will be achieved using coupled fluid particle simulations and the research will examine the algorithms used in these coupled simulations. This project has a number of objectives in relation to this overall aim, namely:

Objective 1 To develop a computationally efficient means to incorporate local flow heterogeneities in coupled particle-fluid systems using an immersed boundary method. In particular to develop modifications to the method that have previously limited its applicability to modelling densely packed granular systems. Most notably are the methods susceptibility to pressure oscillations as particles within the system move and the inability of the method to resolve the forces between particles in close proximity, *i.e.* lubrication forces.

Objective 2 To quantify the sensitivity of the response of the fluid and particle phases, observed in coupled fluid-particle CFD codes to the level of refinement adopted in the modelling of the fluid phase by simulating laboratory experiments.

Objective 3 To develop insight into the extent to which the local flow heterogeneities influence the behaviour of granular systems.

It is hoped this understanding will validate and improve the current best practice guidelines.

1.1.2. Overview of thesis

Chapter 1 Gives an introduction to embankment structures and the current state of numerical tools available to model them.

Chapter 2 Reviews the current literature for the relevant numerical methods and experimental data.

Chapter 3 Details the numerical method developed to improve the modelling of embankment structures.

Chapter 4 Discusses the simulations relating to single particles in free stream, in particular relating to the replication of experimental results and the influence of numerical parameters.

Chapter 5 The development and testing of a modification to the calculation of hydro-dynamic interaction forces acting on particles in close proximity is presented. The modification is based on a sub-grid scale contribution from the interstitial flow, known as lubrication force, which is under predicted due to the limited resolution of the fluid separating particles in close proximity.

Chapter 6 Details simulations performed on arrays of particles which are undertaken to investigate the effects of resolving the flow around individual particles over the more commonly used methods, such as discrete element method (DEM), in which a the flow is averaged over a number of particles.

Chapter 7 Reviews the major findings of the thesis and outlines potential future work.

1.2. Embankment structures

1.2.1. Internal Erosion

Internal erosion of a granular material is the detachment, and subsequent transportation, of individual particles from the intact parent material. Internal erosion is categorised into two types; erosion caused as a result of a significant meso-scale defect, often refereed to as concentrated leak erosion, and erosion that takes place with meso-scale homogeneous (defect free) bulk material. Erosion that takes place within bulk material in homogeneous material can be sub-categorised into suffusion and suffosion, and will be the focus of this work. In order for erosion to occur in a homogeneous, compacted granular material (in a percolating, stress transmitting environment), not only is there required to be sufficient drag and buoyancy forces to initiate and maintain movement of the particle but also sufficiently large voids in the parent material through which the particle can travel. While the relative sizes of the largest voids and the smallest particles are in part dependent on the construction of the material it is predominantly determined by the particle-size distribution, or PSD. Materials with a particle size distribution that permits this type of erosion under sufficient hydraulic loads are said to be internally unstable.

1.2.1.1. Suffusion

Suffusion is characterised by the redistribution of grains over a large volume. It occurs in matrix supported soils when the fine content of the soils is selectively eroded downstream, leaving an intact soil skeleton. A matrix supported soils is a soil in which larger particles are supported by the fine particles, and hence are involved in force transmission. Despite leaving an intact load-bearing skeleton the resultant change in porosity can can cause large hydraulic gradients (due to locally low permeability) which may be the precursor to mass moving erosion Garner and Sobkowicz [3]. Although the time development for for suffusion is slow Fell *et al.* [4] relatively little motion is required to significantly alter permeability.

Li [5] links the susceptibility of a material to suffusion with grain or particle size distribution of the material. Garner and Fannin [6] link the occurrence of suffusion in susceptible materials with the stress and hydraulic gradient to which the material is subjected.

1.2.1.2. Suffosion

Gardner and Sobkowicz [3] distinguish suffosion from suffusion if the it does not leave an intact load bearing skeleton. This usually occurs in matrix supported soils in which the fines are required to support the structure. Suffosion is considered significantly more damaging than the comparatively benign suffusion as if it propagates it can cause backwards erosion, see section 1.2.1.4, associated with a with significant mass movement on a macroscopic scale.

1.2.1.3. Concentrated leak erosion

Concentrated leak erosion occurs around a significant meso-scale defect in the material, resulting in a continuous permeable zone (perhaps due to an irregularity in the grading, poor compaction, cracking or as a result of suffosion). The large resultant flow can cause a sheer stress sufficient to break of particles from the wall of the leak-path, perpetuating the leak. The severity of this type of erosion depends on the ability of the downstream filter to seal the leak and withstand the resultant stress. While concentrated leak erosion will not be investigated directly, its implications for the reliability of predictions and its interaction with suffusion and suffosion require its understanding and inclusion.

1.2.1.4. Backward erosion

The term backward erosion can apply to a number of mechanisms of erosion, whereby the noticeable effects originate upstream and propagate downstream. It indicates that there is not only a redistribution of particles but a net flux of particles downstream. This is considered particularly damaging as mass moving erosion is often perpetuated and can lead to structural as well as hydrodynamic failure.

1.2.2. Requirement for advanced numerical modelling

Numerous physical experiments have been carried out in order to determine the micro-dynamics involved in internal instability. However physical experiments can only reveal a limited amount of information about the process. One of the principal problems with experimental investigation is that the micro-dynamics of most of the phenomena, such as individual particle motion, are difficult to determine. This is principally due to the difficulties of tracking individual particles. This is possible near the surface of a granular material however this is not helpful as the interesting phenomenon happen inside the material. Tracking particles inside the granular material is currently unfeasible. Additionally the it is also currently unfeasible to elucidate forces between particles and the momentum exchange between the fluid and particle phases.

Accurate numerical modelling would address these issues as micro-mechanical data is available. While numerical methods have been applied to the modelling of embankment structures this has mainly been limited to discrete element modelling, in which not all of the scales are resolved. It is hoped the development of an accurate, fully-resolved flow solver would provide the insight necessary to generate consensus and justification regarding the dynamics and hence stability criterion.

1.3. Solid-fluid multiphase flows

1.3.1. Immersed-boundaries representing particles

The concept of using a Lagrangian reference for representing a solid particle in a Eulerian framework is discussed extensively in this work. The basis of the method is to include a representation of the surface of a particle and adjust the equations governing the flow near this surface. This allows the inclusion of particles larger than the fluid cells. Section 1.3.4.5 introduces methods used to enforce the no-slip condition at the boundaries of these surfaces, moving in a Eulerian framework. Section 1.3.5.3 discusses this in the context of immersed-boundary methods. Section 2.3 discusses the current literature with regards to the enforcement of the immersed-boundary condition in the

mirroring immersed-boundary method framework which forms the basis of the method developed in this thesis.

Figure 1.1 shows an example of a such a particle and its surface. The figure also shows how the fluid near the particle is modified by the presence of the particle. In particular the component of velocity tangential to the velocity of the particle is shown along with streamlines of the fluid velocity relative to that of the particle.

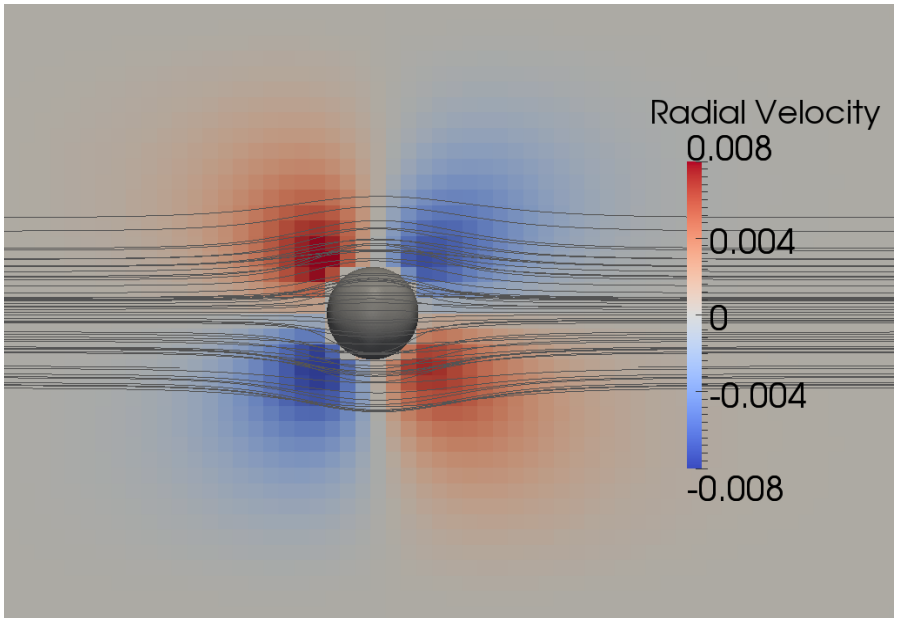


Figure 1.1.: An example of the surface used to represent a Lagrangian particle in a Eulerian fluid. The component of velocity tangential to the velocity of the particle is shown along with streamlines of the fluid velocity relative to that of the particle. The particle is moving with velocity 0.0361 m s^{-2} in the left direction of the image.

1.3.2. Fluid governing equations

A solution to the flow of an incompressible Newtonian fluid can be obtained by solving the Navier-Stokes equations given by:

$$\rho\frac{\partial u^i}{\partial t} + \rho\frac{\partial}{\partial x^j}\left(u^i u^j\right) = -\frac{\partial}{\partial x^i}p + \frac{\partial}{\partial x^j}\tau^{ij} \tag{1.1a}$$

$$\frac{\partial}{\partial x^i}u^i = 0 \tag{1.1b}$$

Satisfying the conservation of momentum (1.1a) and mass (1.1b). This can be extended to fluids containing a particle phase with appropriate boundary conditions at the solid phases' surface (Γ):

$$u^i = u^i_\Gamma \quad \text{at any point on } \Gamma \tag{1.2}$$

Newtonian mechanics is applied to the solid phase along with any constitutive laws required, for example particle-particle collision.

1.3.3. Contact dynamics

Contact modelling is a necessary component of modelling a granular material. The literature regarding contact modelling in the immersed boundary framework is discussed in section 2.2 however most of the work is limited to sphere-sphere contacts. Elastic Hertzian contact often forms the basis of these methods and is readily adapted to non-spherical particles by considering the reduced radius of curvature of the two points in contact. Models of visco-elastic spheres are given in [7, 8, 9, 10]. Some modifications must be made for complex collisions, for example with multiple contact points. It is possible to adopt more general contact dynamics methods such as the finite element method to resolve the contacts however the associated computation cost would be large. A compromise is proposed and validated in this work, see sections 3.4 and 5.2.

1.3.4. Lubrication force

Also important to the modelling of particles in close proximity is the repulsive force imparted by the thin interstitial fluid layer, known as the lubrication force. While the lubrication force is theoretically resolved by the fluid phase, the inevitable under-refinement of the discretisation of the fluid phase near contact often results in under-prediction for the force. Hence it is often required that the lubrication force be modified with a sub-grid scale contribution [11, 12, 13]. Literature regarding the lubrication force is discussed in section 2.2 and the topic is discussed in detail in chapter 5. As some implementations of the immersed-boundary method require that the particles do not overlap, due to numerical restrictions [14] it is often common that this correction is used to prevent overlap. However this is not feasible for modelling densely packed particle arrays as contact is an important aspect of the model. Discussions of models for spheres and their performance in viscous fluids can be found [15, 16].

1.3.4.1. Definition

The lubrication force between two surfaces separated by a small gap, filled with incompressible, or negligibly compressible, fluid is the repulsive force exerted on the surface pair. It is of significance that this force can be much greater than the hydro-dynamic interaction of the same surfaces with the same relative velocity if the surfaces are not in close proximity.

1.3.4.2. Associated fluid dynamics

The magnitude of the lubrication force arises from the region separating two close parallel surfaces, moving with respect to each other along their mutual normals, the sheer stress and hence energy dissipation is great. This is because a large volume of fluid must be moved relative through a thin opening resulting in very large velocity gradients over a short space. The thinness of the opening is also the reason behind the difficulty in accurately modelling the fluid and hence the resultant force. In order to resolve the force the flow through the opening that gives rise to them must also be modelled. However this requires the resolution of flow at a very small length scale, *i.e.* that of the opening. Further more the constriction is liable to occur without prior knowledge of when and where it will occur hence local refinement is not always possible.

1.3.4.3. Motivation for study

In the numerical modelling of immersed particles, forming a stable structure, the accurate treatment of lubrication force is critical. It determines the particle particle interactions both in the formation of an equilibrium when stable, and in all but the most energetic of collisions at the onset of motion. Hence it is of particular importance to the stability analysis of granular materials. It is noted in [11] that the precise nature of the interaction of the solids phases is significant, even at very small scales, typically up to 1% of the particles length-scale, for spherical particles.

1.3.4.4. Lubrication force in the immersed-boundary framework

While the lubrication force is theoretically resolved by the fluid phase, the inevitable under-refinement of near contact often results in under-prediction for the force. Hence it is often required that the lubrication force be corrected [11, 12, 13]. As some implementations of the immersed-boundary method require that the particles do not overlap, due to numerical restrictions [14] it is often common that this correction be used to prevent overlap. If this is not required a huge number of particle-particle near contact forcing models are available. In this thesis work is undertaken to accurately account for the under-prediction of lubrication force. This is presented in chapter 5.

1.3.4.5. Combined lubrication force and contact

Discussions of models for spheres and their performance in viscous fluids are discussed [15, 16]. Elastic Hertzian contact often forms the basis of these methods and is readily adapted to non-spherical particles by considering the reduced radius of curvature of the two points in contact. Joseph and Hunt [17] investigated the combined effect of wall collisions and lubrication force

1.3.5. Existing techniques for resolving solid-fluid flows

A number of techniques for fluid-solid flow simulation are available. A brief outline of the most common approaches are given.

1.3.5.1. Arbitrary Lagrangian-Eulerian methods

Arbitrary Lagrangian-Eulerian (ALE) methods resolve the presence of a submerged solid phase using body-fitted grids [18]. By using a body fitted grid, the particles body is represented by a domain boundary. Hence the boundary condition is enforced at the existing cell faces. Techniques for this discussed in the literature. [19] For individual time steps, the motion of a rigid body is captured by enforcing the correct velocity at the boundary. ALE methods have the advantage that:

- The method enforces the boundary condition very accurately.
- The method requires little empirical input.

However its drawbacks are:

- In order to maintain a body fitted mesh the method requires frequent re-meshing. This makes the method computationally expensive [20, 21].

- The computational mesh must be generated automatically and is required to fit potentially complex geometry. Hence mesh quality is likely to affect results [22, 23, 24].
- The re-meshed computational grid results in new integration points. As these new integration points did not exist at the previous time instance the is no readily accessible 'old time instance' properties at these points. These properties must be interpolated, this comes with an associated cost and decrease in accuracy.

There have been numerous applications of the arbitrary Lagrangian-Eulerian method to simulate solid-fluid flows [25, 26].

1.3.5.2. Lattice-Boltzmann methods

Lattice-Boltzmann methods (LBM) significantly differ from other methods described in this section. Instead of solving fluid equations using the finite volume, finite difference approach of finite element method, the Lattice Boltzmann method solves the Boltzmann equation on a discrete system of fictitious fluid 'elements'. The fluid 'elements' are positioned on a regular lattice. They are allowed to move with discrete steps and to collide with other 'elements' in accordance with a number of collision rules. If appropriate lattice topology and collision rules are chosen, the resulting solution satisfies the Navier-Stokes equations and boundary conditions [27].

The Lattice-Boltzmann method avoids the re-meshing issue encountered by body-fitted grid methods. The main disadvantage of this approach is the time-step limitation, which is much more strict than finite volume discretisation methods. LBM was used for the direct numerical simulation of particulate flows in [28, 29]. Feng and Michaelides [30] introduce a method that combines the Lattice-Boltzmann method with an Immersed boundary method (see section 1.3.5.3). The method was used by Holzer and Sommerfeld [31] in the study of a flow past non-spherical particles.

1.3.5.3. Immersed boundary methods

Introduced by Peskin [32] as a method of simulating the complex moving boundaries in the circulatory system. Immersed-Boundary Methods (IBM) resolve both solid and fluid phases by treating the solid phase as a boundary within the fluid domain. This immersed boundary is represented by a discretised mesh of the solids surface. The boundary is enforced by modifying the fluid equations in the regions around these surfaces. Figure 1.2 illustrates the schematic used in two dimensions.

There are a number immersed-boundary methods which are determined by the nature of these modifications to the fluid equations in the regions around the immersed surfaces. Since its introduction by Peskin, the immersed boundary method has been applied by numerous researchers in various fields. Numerous variants of the approach are available. Mittal and Iaccarino [34] review a number of current methods for coupling Eulerian and Lagrangian phases, for an immersed boundary. They proposed the following the classification of the currently available approaches:

Continuous forcing methods A number of the methods discussed involve a (smoothed) constitutive relation, in which the interface is allowed to deform. A constitutive relation between momentum transfer and this displacement is imposed which models the coupling. Such methods are not appropriate for the simulation of flows dominated by a high pressure gradient, such as those through densely packed solid phases, as a sharp boundary is not maintained. For such flows it

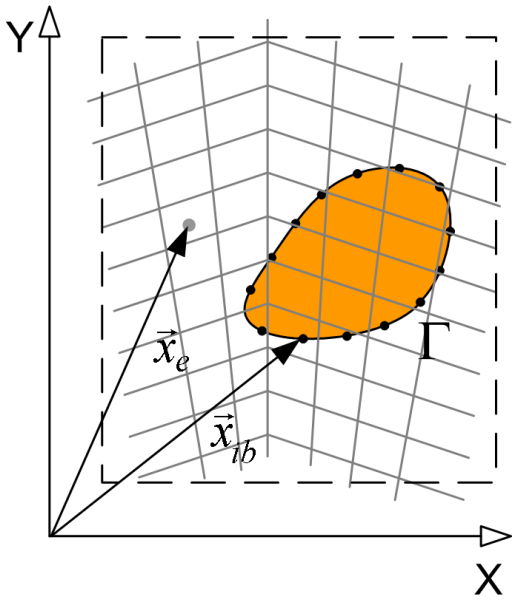


Figure 1.2.: Schematic of the immersed-boundary method showing an arbitrary Lagrangian particle mesh (Γ) over a static Eulerian fluid mesh. Reproduced from [33].

is necessary to maintain a sharp interface in order to resolve the flows accurately. Such methods where intended to model interfaces for which this was appropriate, such as the vascular system (Peskin), or situations in which as small error of the fluid relative to the interface is permissible, such as low Stokes number particles. However, with an appropriately small time step, stiff systems can be modelled successfully with this approach. See section 2.1.1.

Direct forcing methods Such methods are also referred to as sharp interface resolution methods. These methods attempt to directly enforce the no slip condition for the ridged immersed-boundary by only modifying the equation system of cells that are cut by, or inside of, an immersed boundary. Section 2.1.2 details a number of such methods available.

Immersed boundary methods have a number of advantages, both in general and in the application to embankment structures:

- The fluid phase domain of an embankment structure is highly complex and varies over time. This is well suited to the adaptability of the immersed-boundary method to vary the position of boundaries.

- As the solid phase is discretised separately from the fluid phase, their properties can be adjusted without having to account for the additional impact on the fluid phase. For example arbitrary shaped or deforming particles, particle-particle interaction and particle constraint can be implemented without additional modification to the coupling of the fluid and solid phases. Additionally the method can be applied to most existing fluid flow solvers without interference to the fluid phase.

- Contrary to the ALE approach, in immersed-boundary methods there is no need for the Eulerian grid re-meshing associated with particle movement. Not only does this save computational expense, it prevents the loss of accuracy related to the interpolation of the flow field to the new mesh. Additionally the permanent mesh aids meshing. Furthermore as the complexity

of embankment structure cases often arises from only the immersed phase, often very simple meshes can be employed for complex problems. Indeed the IBM often permits the use of a Cartesian mesh to model complex cases.

- The method is well suited to parallelisation which can further improve its feasibility for large simulations.

A sharp interface resolution IBM has been chosen for this project as its advantages are particularly relevant. In particular the Mirroring Immersed Boundary method (MIBM) was chosen as the basis of the method developed. Its primary advantage being its not dependence on force to impose the boundary conditions, hence its ability to resolve the no-slip condition flows dominated by high pressure gradient flows. Further details are presented in sections 2.1, 2.3 and 3.2. In addition to Mittal and Iaccarino [34], a review of the available immersed-boundary methods is given by Haeri and Shrimpton [35].

1.3.6. Existing techniques for resolving fluid flows

Flows through fine granular materials such as the core of embankment structures, which are of interest to this research, are invariably laminar. In such cases, the simulations of the fluid phase can be considered as Direct Numerical Simulation (DNS) and no turbulence modelling is required. Hence for this work only DNS will be considered see section 1.3.6.1.

However, numerous techniques are available for the solution of the Navier-Stokes equations. Many of which can be integrated with a MIBM. It is hoped that with the increase in computational power over time, the use of immersed-boundary methods will expand to more engineering applications. In light of this it is worth noting the MIBM may potentially be adapted to be used with the following common techniques in addition to DNS:

- Large Eddy Simulation (LES)
- Reynolds Averaged Navier-Stokes (RANS)

However, some modification to the technique must be made to account for the more complex boundary conditions. Balaras [36] makes such a modification to allow the use of LES in conjunction with the continuous forcing method. This method is developed by Yang and Balaras [37]. The mirroring immersed boundary presents a greater challenge as the wall shear stresses are not directly controllable, but the product of the mirroring equation. However, using a modification developed in this thesis, the cut-cell variation of the momentum equation described in section 3.2.4, the wall shear stress can be set directly again. This could be used in conjunction with the development made by Balaras [36] to introduce LES to the mirroring immersed boundary method.

1.3.6.1. Direct Numerical Simulation

Direct Numerical Simulation (DNS) is the most accurate way to simulate fluid flow. DNS computes the fluid flow at a resolution of the smallest spatial and temporal scales present in the flow. Hence the Navier-Stokes equations can be discretised directly, with no requirement for additional turbulence modelling.

Most flows involve velocity and pressure fluctuations on a range of scales. In such cases the computational grid require an enormous number of cells and many very small time-steps. This

is because it must be fine enough in order to resolve the smallest oscillations and eddies while maintaining the scope of the large scale structures of the flow. This can make DNS impractical for simulation of large, complex or highly turbulent cases.

The required grid spacing shrinks relative to the large scale structures of the flow as Reynolds number grows. In particular in turbulent flow: $\frac{L}{\eta} \propto Re_L^{0.75}$ [38]. Where:

- Re_L is the Reynolds number of the large scale structures.
- L is the length scale of the large scale eddies
- η is the length scale of the smallest scale eddies

Clearly grid spacing $\propto \eta$ hence the number of points grows $\propto Re_L^{2.25}$ per particle. In the case of low Re flow, when η and L are equivalent, the solution is more tractable. This limits DNS to flows with low Re. Hence most cases of engineering interest are not suitable for modelling with DNS.

Although DNS is limited to analyses of simple problems at low Reynolds numbers, it is still a very useful technique. It allow to understand the physics governing the observed processes. DNS can be used to quantify various flow-particle interaction properties, along with the mechanisms of turbulence production and dissipation. However the limit imposed by Reynolds number is not such a severe constraint for dense particle flows such as those found in embankment structures as Re is small in these cases.

1.3.7. Force coupling between fluid and solid phases

The force between the solid and fluid components of a multi-phase flow is a critical component of its modelling. It is particularly important in high particle density situations, where significant lubrication forces (section 1.3.4) are expected, as it is critical to determining the dynamics of the system. This is especially so if there are large stresses or significant relative motion between the particles.

There are many methods of evaluating this force and hydro-dynamic interaction forces in general section 2.3.2 discusses some of the methods used. Section 3.3 discusses the method developed in this thesis. The choice of method which is used to account for this force is often limited by the numerical framework. Additionally, the suitability of the method will depend on the Reynolds number of the flow, the Stokes number of the solid particles and numerical scheme used. In discrete forcing immersed-boundary methods the calculation of the force a fluid exerts on an immersed boundary is not a part of the immersed-boundary enforcement method. In these cases the force is often taken to be balance momentum in the Navier-Stokes equations as calculated for the fluid phase near the boundary. These forces are subsequently integrated over the surface, see section 3.3. This type of method is employed in this project, the details of which are also given in section 3.3. The discussion in this section should be applicable to all such methods. However some methods described in section 1.3.5.3 of enforcing the immersed-boundary do not use this approach. In such cases where the force is directly involved in the coupling of the fluid phases, such as continuous forcing methods, it is possible to transfer momentum between solid phases directly. Nguyen and Ladd [11] give an example of this. It is hoped the discussion in this section could be adapted for use in such methods, however focus is given to the methods employed in this work.

Methods that integrate stress over the immersed-boundary, to determine the momentum exchange between the immersed-boundary and the fluid phase, require an accurate description of the flow around the particle. A sufficiently accurate description may not always be available. For example if the fluid is not resolved sufficiently or if the calculation of the flow is inaccurate for numerical reasons such as the difficulty of accounting for the presence of the immersed-boundary. This lack of resolution is particularly relevant when two solid phases are in close proximity *i.e.* lubrication force, see section 1.3.4, because:

- When two particles are in close proximity a small number of cell lengths, perhaps less than one, separate the two solids.
- Within this gap the fluid flow can vary significantly, and is not well resolved without a number of cells.
- The correct force in such a scenario is potentially quite large.

In these cases an empirical model may be applied however this is often only accurate in a limited number of special cases. Common restrictions include limitation to spherical bodies and restriction to laminar, or even Stokes flow regimes [14].

For some immersed-boundary methods, the force acting on the particle from the fluid phase requires an integral over the particle surface of the force terms of the Navier-Stokes equations. Hence a surface integral method calculation based on the flow around the immersed-boundary must be employed see figure 1.3. This calculation is complicated by the movement of the boundary and that the boundary does not usually apply to a fluid cell face. Hence it is necessary to employ a number of interpolation and extrapolation equations to accurately predict what the fluid properties at the surface are. This is discussed in detail in section 3.3.

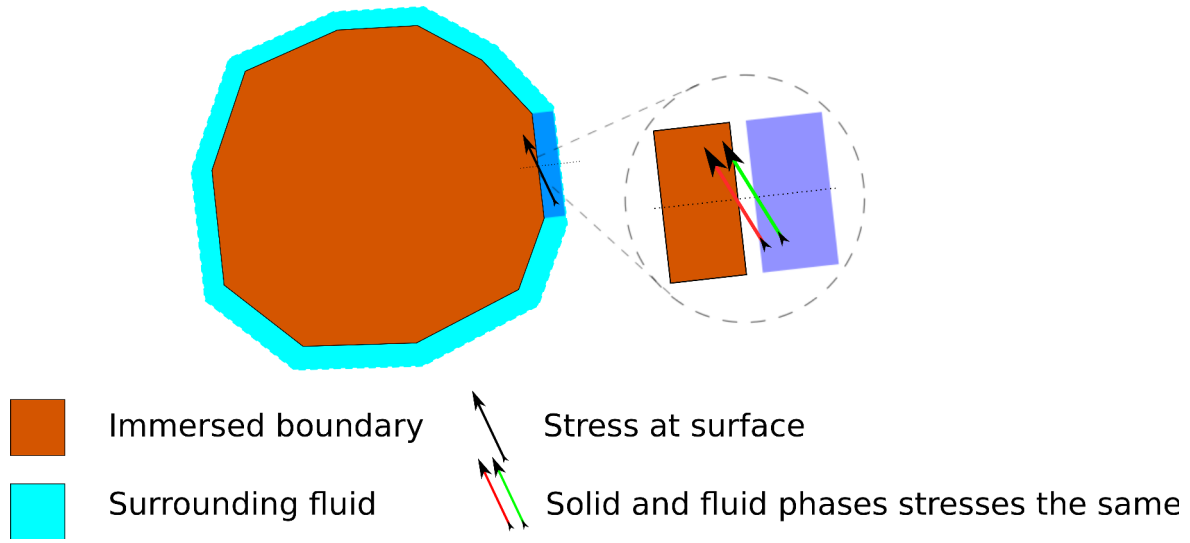


Figure 1.3.: A diagram of an arbitrary immersed body in two dimensions, showing the fluid immediately surrounding the body. The surface stresses on the particle are taken as equal to the fluid stress in the adjacent fluid, via conservation of momentum. The stresses in the fluid phase are calculated from the Navier-Stokes equations. The force acting on the body is taken as a surface integral over these stresses.

If the problems associated with evaluating the fluid properties at an immersed-boundary can be addressed, the force acting on an immersed-boundary may be obtained via evaluating the stress terms

in the Navier-Stokes equations across the boundary. This analysis gives that the i^{th} component of force acting on an immersed object (f_i) is given by:

$$-\int_{IB}\left(p\cdot\delta_{ij}+\mu\left[\frac{\partial u_i}{\partial x_j}+\frac{\partial u_j}{\partial x_i}\right]\right)a_j dS \tag{1.3}$$

It is notable that this is only possible to evaluate f_i if it is possible to accurately evaluate both the pressure field and the velocity gradient field across the surface of the immersed object. This is non-trivial as the varying length scale of variation of flow around the particle makes extrapolating properties numerically difficult. This is exacerbated by the collisions or near collisions of the particles. Indeed in a near collision of two similarly sized particle separated by fraction (x) of their radii the volume ahead of one of the particles, determining the interaction force, in which the fluid flow deviates from the particle free flow only by the interaction with that particle, will be a fraction of x itself. However in low Reynolds number flow of a particle through an otherwise undisturbed flow may have a region many times the particles radius, ahead of which the variation is representative of the fluid particle coupling.
This is rectifiable in principle, either by determining the appropriate range or by taking sufficient data to ensure both the range and density to have the necessary data in all cases. However these solutions are numerically and computationally prohibitive. Additionally the density of data points is limited by the resolution of the mesh. Section 3.3 addresses the issue of interpolation of fluid properties to the surface and section 3.2 discusses the implementation of the boundary conditions at immersed-boundaries.

1.3.8. Feasibility

The main limitation of the fully resolved solid-fluid coupled flow is feasibility. While in sub-grid scale methods a single fluid cell may contain a large number of particles a typical fully resolved case will take a number of fluid cells across a single particle. This means a significantly greater number of fluid cells are required in order to utilise a fully resolved method, associated with a large decrease in feasibility. An example is given in figure 1.4. While the relative required refinement is subject to both the case and accuracy and fidelity requirements, feasibility is generally an issue. This is exacerbated for embankment structures where typically a large number of particles are required. As a result, feasibility is considered a central issue for fully resolved methods.

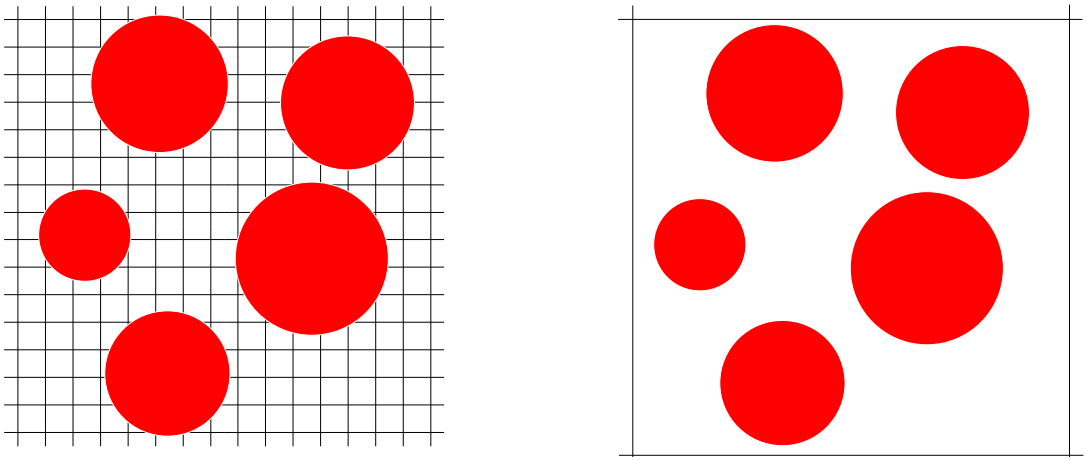


Figure 1.4.: A two dimensional schematic comparison of the number of fluid cells required to resolve a fixed number of particles. The left image shows an example ofa simulation refined to see the flow around individual particles. The right image shows an example ofan under-refined simulation in which only the averaged flow is seen. In this example the relative spacial refinement is 15. In three dimensions this would require $15^3 = 3375$ times as many cells as well as 15 times the number of time-steps to achieve the same CFL number.
Even in the refined case, the flow is not particularly well resolved. In order to have good resolution a further spacial refinement of 5 is required. The overall cell count for this resolution in three dimensions would be 421875, more then five orders of magnitude more than the original.

2. Literature review

2.1. Forcing techniques

This section will cover both forcing techniques as classified by Mittal and Iaccarino [34]:

- Continuous forcing approach
- Discrete forcing approach

2.1.1. Continuous forcing

2.1.1.1. Introduction

The continuous forcing method, the original immersed-boundary method, Peskin [32], was developed to simulate immersed-boundaries with elastic properties. Inparticular , the vascular system was studied. The method is well suited to simulating any immersed-boundary with constitutive force-displacement laws such as linear-elastic materials. The major problem with this method is numerical instability if this relationship is too stiff. This is a common problem for many solid-fluid flows where the solid phase does not deform significantly as a result of fluid forces.

2.1.1.2. Methodology

The continuous forcing technique is implemented by modifying the momentum equation 1.1a with the inclusion of a forcing term (S). I.e. equation 2.1:

$$\rho\frac{\partial u^i}{\partial t} + \rho\frac{\partial}{\partial x^j}\left(u^i u^j\right) = -\frac{\partial}{\partial x^i}p + \frac{\partial}{\partial x^j}\tau^{ij} + S^i \tag{2.1}$$

Here S^i is based on a the required constitutive forcing of the surrounding cells. I.e.:

$$S^i = \sum_k S_k^{\text{local } i}\, g_k \tag{2.2}$$

Where:

- k ranges over nearby cells
- S_k^{local} is the constitutive force at cell k
- g_k is a distribution function, centred at the current cell, evaluated at cell k

2.1.1.3. Method parameters

The main variation within continuous forcing methods is the choice of distribution functions (g). Numerous distribution functions have been proposed. These include:

- A Dirac delta function [32]
- A compact three-point distribution function [39]
- A two-point hat distribution function [40]
- A 4-point distribution function [41]
- A 6-point distribution function [42]

Haeri and Shrimpton [35] study the influence of a number of proposed distribution functions. They conclude that functions that attenuate smoothly, *i.e.* that proposed by [42], tend to produce more reliable and smoother results than those with abrupt changes, *i.e.* [32]. However they also conclude that a large number of points are required for a practically smooth distribution function. This can significantly increase the computational cost associated with the method, as it reduces diagonal dominance of the solution matrix. It also leads to a very smoothed representation of the interface. This usually renders the method first order in space and requires a very refined mesh to remain physical [43].

2.1.2. Discrete forcing

Discrete forcing methods attempt to address the problems inherent to the continuous forcing approach. The discrete forcing approach modifies the flow equations only in the immediate vicinity of the interface. Hence in discrete forcing methods the forcing is applied directly to the fluid volume from which the force arises. There are three main ways of doing this, currently in the literature:

Direct forcing methods Discussed in further detail in section 2.1.2.1. There are two ways in which this method can be used to enforce the no slip condition at the interface:

- By adapting the momentum equation in the boundary cells.
  In this method a force is added to the momentum equations in a similar way to that of the continuous forcing method. However the forcing term is usually evaluated iteratively on the pre-solved flow in order to ensure the no slip condition is satisfied. Additionally, the force is only applied to the particular cell in which the boundary is being enforced.
- By forcing the velocity directly.
  This is achieved by forcing the fluid velocity of the boundary fluid cells such that the imposed boundary condition coincides with the exterior flow. This is usually achieved by extrapolating from the flow in the vicinity of the boundary.

Ghost cell methods Discussed in further detail in section 2.1.2.2. The ghost cell method enforces the velocities in cells inside of the body, named 'ghost cells'. This is done such that the interpolated velocity at the surface, from the surrounding fluid and ghost cells, is equal to the required velocity of the surface. This has the advantage that only the Navier-Stokes equations are solved outside of the body.

Cut-cell methods Discussed in further detail in section 2.1.2.3. The cut-cell method attempts to solve the Navier-Stokes equations, momentum and the continuity, without the necessity of using cells inside the immersed-boundary. It does this by considering a new set of cells formed by the intersection of the immersed-boundary and the fluid mesh.

2.1.2.1. Direct forcing methods

Numerous variations on the direct forcing methods exist in the literature.

Mohd-Yusof [44] calculate the forcing term imposed by the boundary using the previous time instance flow field. This corrective force is applied to interior cells, and enforces the boundary using the resultant non-physical flow inside the immersed-boundary. A diffusion step is applied to diminish discontinuities over the boundary. The main disadvantages of this method are:

- It's dependence on the interior, non-physical, flow.
- Potential error in the enforcement of the immersed-boundary due to the way it is enforced, and the diffusion step.

One of its advantages is the interior flow, near the boundary, is useful in reconstructing the fluid properties at the boundary. This helps ascertain the momentum transfer between the phases.

Fadlun et al. [45] furthered the direct forcing method of Mohd-Yusof by studying various approaches for relating the immersed-boundary velocity, the fluid velocity in the surrounding fluid cells and the forcing technique. They also experimented with the distribution of the force within the cell. One particular beneficial development was modifying the forces applied to each cell depending on the volume fraction of the cell currently inside the body. Their major conclusions are:

- The interpolation scheme used to find an appropriate values of the immersed-boundary, at the necessary points strongly influences the behaviour of the method.
- Of the interpolation techniques tested, a linear interpolation of the surrounding fluid and immersed-boundary was optimum.
- Within the framework, the flow deep inside the body had little affect on the exterior flow.

Pressure is set by enforcing a Poisson equation, in relation to velocity. As the Velocity is forced to be linear around this this cell, this equation reduces to a zero normal pressure gradient condition. This is enforced by the same interpolation and stencil used for velocity equations.

Balaras [36] noted problems can occur for arbitrarily shaped, highly curved boundaries. In particular that when using the method of Fadlun *et al.* the reconstruction at the boundary is dependent on the direction of interpolation used. This choice is largely dependent on the surrounding cell's type, but is often an ambiguous choice. Hence the method was altered to make it more compatible with a wider range of geometries. In particular, the method proposes using a interpolation stencil removed from the immersed-boundary and skewed in the direction normal to the plane. Doing so removes some of the problems of local concavity of the immersed-boundary as translating in the normal direction makes it more likely that a full set of exterior fluid points are available. This is achieved used an interpolated 'virtual' point, slightly removed from the exterior boundary cells (in the direction normal to the immersed-boundary). A linear

relation is enforced between the 'virtual' point, the forced cell and the boundary, in terms of velocity. The 'virtual' point itself has velocity set from tri-linear interpolation from its surrounding eight cells. An additional advantage of this method is the is no discrete transition in the normal directions taken.

Uhlmann [46] resolves the implicit force at the boundary over a distribution function, similar to those used in continuous forcing models. This smooths the representation of the interface resulting in smoother flow profiles however losses fidelity, necessary to resolve some interactions.

Gilmanov [47, 48, 49] implemented a similar model to that of Balaras [36].

- Slight differences include the projection distance of the 'virtual' point, to the first cell face
- Upwind scheme for convection

In addition to these modifications, a number of simulations were carried out, testing a number of variants of the model. Of particular note is the use of a quadratic interpolation and fit through the points used to enforce the boundary. This method was used to simulate relatively complex cases with some success.

Choi et al. [50] adapted available models by forcing the fluid on either side of the interface, taking advantage of the observation of Fadlun *et al.* that the unphysical interior flow does not heavily influence the exterior flow. Further improvement was made by the use of a weighted interpolation, were points close to the 'virtual' point where considered more heavily. This aided in a smooth transition as the immersed-boundary moved. See section A.1 for more details of interpolation techniques.

Huang and Sung [51] added mass source terms based on the mass loss errors encountered in the above framework. Their method is based in ideas introduced by Kim *et al.* [52].

2.1.2.2. Ghost cell method

Ghost cell methods enforce the necessary boundary conditions of the immersed-boundary by modifying flow variables inside the immersed-boundary. This has the advantage that outside the body only the Navier-Stokes equations are solved. The ghost cell method was first introduced in two dimensions by Majumdar [53]. The method involves setting the velocity field at interior cells, near the boundary, such that the flow is reconstructed accurately at the immersed-boundary. These interior cells are referred to as 'ghost' cells. An implicit relation between the velocities in the 'ghost' cells, the velocity at the surface of the body and the surrounding fluid cells is set. The choice of this relationship is such that the no-slip boundary condition is satisfied. The principle is much the same as is employed in 'ghost' cells used in Eulerian grids to set boundary conditions for CFD methods which do not employ the use of an immersed-boundary. However unlike in CFD methods in which the boundaries are static, the boundaries cannot be made to always correspond nicely to a cell face. Hence a form of interpolation is needed to ascertain the correct relations between the ghost cells and the exterior cells. Additionally, as the boundaries are not guaranteed to be static relative to the Eulerian mesh over time, this interpolation must be updated each timestep. Similarly to otherdiscrete forcingapproaches, the choice of the interpolation method is critical.

The ghost-cell finite difference approach involves modifying internal, boundary-adjacent cells and using an interpolation scheme to incorporate the presence of the immersed boundary directly into the

equations of the fluid cells. The method is heavily dependent on the interpolation method, of which many are available. Section A.1 discusses the issue of interpolation with regard to CFD. Ghias *et al.* [54] outlines a number of possible interpolation methods specific to modifying internal variables in the ghost-cell method. In particular a bi-linear interpolation scheme is found to produce good results. Ghias *et al.* [54] outlines a number of possible interpolation methods for use in the ghost cell method. Their method has been used to simulate flows with Reynold's numbers up to 3900 using a large eddie simulation technique. However the mesh resolutions used where substantial and required local refinement in the vicinity of the immersed-boundary. Also the results did not include the simulations of moving bodies. Both of these are substantial limitations if a large number of moving bodies need to be simulated.
It is notable however that in regions of rapidly varying flow, such as at high Reynolds numbers, the interpolated value is particularly sensitive to the choice of interpolation. Furthermore a poor choice of interpolation scheme can lead to erroneous evaluations of fluid properties at non-cell-centre points. Majumdar *et al.* [53] have had some success at addressing the issue of interpolation for the ghost cell method.
Majumdar investigated the effects of:

- Linear interpolation
- Quadratic interpolation
- Bi-linear interpolation

One observation of Majumdar *et al.* is that some interpolation coefficients give rise to an equation system that is difficult to solve for numerical reasons. In particular coefficients that are too large or negative, or non-diagonally-dominant coefficients can cause problems. In order to prevent the interpolation scheme from creating interpolation coefficients that present difficulties to the solver, the following scheme was adopted:

- For each ghost cell, p_g, a reflection point, p_r, is created outside the immersed-boundary. p_r is chosen to be in the direction of the surface normal near p_g.
- Interpolation coefficients from the surrounding cells are assigned to this point using an interpolation scheme that fits pseudo-quadratically. That is the velocity field is fitted to a polynomial with zeroth-order, first-order and two second-order terms. Hence $u = a_0 + a_1 N + a_2 T + a_3 NT + a_4 N^2$ where N and T are components of position in the normal and tangential direction respectively.
- Interpolation between p_g and p_r is used to enforce the boundary.

The restriction on interpolation coefficients may in some cases prohibit an ideal interpolation scheme. However it was found to notably improve the stability of the solution to the system's equations.

It is notable that to extend the fit used here to a three dimensional system would require a large stencil (seven points), to resolve the large number of cross terms. However this is fewer than the ten points required for a general quadratic fit. In particular, this allows the use of the 'box' of *eight* integration points surrounding the interpolation point for a three dimensional Cartesian case. However, in a similar way to the problem that occurs when two interpolation points lie on a line in two dimensions, this is not always sufficient.

The a three dimensional analog of the ghost cell method of Majumdar was introduced by Tseng and Ferziger [55]. In addition Tseng and Ferziger proposed an alternative to the reflection point method for controlling interpolation coefficients. The idea behind the method is to treat exterior fluid cells, close to an immersed-boundary, to control the boundary in lieu of the 'ghost' cells. In such a case these cells would directly take the boundary velocity. This occurs when an immersed-boundary is within a fixed fraction of the grid spacing to the cell centre of an exterior point.

The ghost cell method was first implemented in conjunction with compressible flows by Ghias *et al.* [56]. In their implementation, the use of a reflection point was abandoned if the interpolation to the boundary yielded sufficiently 'nice' integration coefficients. In this case direct interpolation to the boundary is used instead.

A number of contributions where made by Mittal *et al.* [57]. Including:

- Numerical techniques to handle complex bodies more efficiently.
- Implementation of a zero normal pressure gradient condition for the 'ghost' cells.
- Handling some of the problems associated with a moving immersed-boundary. In particular, the treatment of transitional cells, in which previous time instance values o flow properties may be unreliable.

Mark and van Wachem [58] observed that a non-physical mass flow through the body interface may be present in some ghost cell methods. The lack of local mass conservation is particularly problematic in the case of a moving body as the mass error fluctuates. The fluctuation in mass error leads to spurious pressure oscillations in the exterior fluid. This has been observed in the simulations of moving particles. The significance of this phenomenon has motivated the methodology used in the current study, see section 3.2.3 for further details.

The method proposed by Mark and van Wachem [58], referred to as the Mirroring Immersed Boundary Method, addresses this issue. In order to ensure a local mass conservation, the *ghost cell* velocities were excluded from the continuity equation. The no-slip boundary condition at the particle surface is enforced using the 'ghost' and reflection points. Work on similar approaches has been conducted by Shinn *et al.* [59] and [60]. Developments to this method have been made [61] [62].

The mirroring immersed-boundary method forms the base of the numerical method developed and used in the current study. Detailed discussion of the method's implementation is presented in chapter 3.

2.1.2.3. Cut-cell method

The cut-cell finite volume approach has the advantage that it is based around satisfying the underlying conservation laws for the cells in the neighbourhood of the immersed solid phase. To achieve this the control volume in the vicinity of the immersed-boundary is modified along the existing grid, in order to conform to the boundary. A number of implementations based on this work is available in the literature. Common variations include differences regarding: how to discretise the cut cell continuity equation and how to treat small cut cells and the viscous sheer stress contribution from the immersed-boundary. Ye *et al.* [63] have shown promising results using this method to describe isolated particles. However for the specific implementation, the need to have a number of cells around the immersed boundary makes this infeasible to describe nearly colliding or colliding particles.

Seo and Mittal [64] introduce virtual cell merging. A technique where the a cut-cell version of the continuity equation is considered along with the un-modified method. The difference between the cut-cell version of the continuity equation and the un-modified version is then 'added' to an adjacent cell. In this way continuity is conserved overall over these two 'merged' cells. There is a similar problem for this method for closely pack particles that the stencil required may be too large. Indeed it would not always be feasible to 'merge' a cut-cell if there are no nearby fluid cells to merge with.

2.2. Lubrication force

As introduced in section 1.3.4 the lubrication force, an attractive or repulsive force imparted by the thin interstitial fluid layer, is a critical component of particle-particle interaction however is often poorly modelled by immersed-boundary methods. The inclusion of particle-particle contact, and near contact, in the immersed-boundary framework is discussed extensively in the literature. This section covers work in the literature covering sphere-sphere dry contact, the lubrication force during wet contact and the amalgamation of the two in the immersed-boundary framework.

2.2.1. Analytical, experimental and empirical and predictions of lubrication force in simplified cases

In order to amend the calculated lubrication force to the correct one in the case of near collisions, the correct result for near collisions must be known. While in general this will be highly dependent on the exact case the major determinant of the lubrication force is the thin film of fluid separating the two bodies, which can be simply parametrised. A number of approaches to finding the appropriate parametrisation and determining the appropriate resultant force are discussed in this section.

2.2.1.1. Analytic results for Stokes flow

Brenner [65] gives an analytical solution to the flow around a sphere approaching a flat surface for Reynolds number much less than one ($Re \ll 1$) (Stokes flow). This gives a prediction for force between the sphere and wall of:

$$\begin{aligned} f &= 6\pi\mu r\lambda u_p \\ \lambda &= \frac{3}{4}\sinh(\beta)\sum_{n=1}^{\infty}\frac{n(n+1)}{(2n-1)(2n+3)} \\ &\quad \left(\frac{2\sinh((2n+1)\beta)+(2n+1)\sinh(2\beta)}{4\sinh^2((n+1/2)\beta)-(2n+1)^2\sinh^2(\beta)}-1\right) \end{aligned} \tag{2.3}$$

Where:

u_p = rate of approach of the sphere to the wall $(m \cdot s^{-1})$

$\beta = \cosh^{-1}(h/r)$ (dimensionless)

The singular term in this equation, which forms the major contribution to the force when the two are in close proximity reduces to $f = \frac{6\pi\mu r^2 u_p}{l}$.

2.2.1.2. Further developments

Cox and Brenner Cox and Brenner [66] improve upon the model proposed by Brenner [65] to take into account some inertial forces. This results in:

$$f = \pi\mu r u_p \frac{1}{\epsilon}\left[1 + \frac{1}{5}\left(1 + \frac{Re_p}{2}\right)\epsilon \ln\frac{1}{\epsilon}\right] \tag{2.4}$$

Where:

$$\epsilon = (h)/r \qquad \text{(a non-dimensional geometrical factor)}$$
$$Re_p = \rho_f u_p r/\mu \qquad \text{(the particle Reynold's number)}$$

Adams and Perchard Adams and Perchard [67] makes a prediction of the interstitial pressure forces in the case of a sphere-wall collision, namely:

$$f = \frac{6\pi\mu r^2 u_p}{h} \tag{2.5}$$

Which is the asymptotic limit of the result of Brenner [65].

2.2.1.3. Experimental results

Experimental data does not elucidate the lubrication force directly, as a collision cannot easily be separated into solid-solid and solid-fluid interaction. However, in conjunction with a solid-solid interaction model (see section 3.4), experimental results can be used to corroborate an overall model. Joseph *et al.* [68] and Joseph and Hunt [17] present data on particle wall collisions. Zhang *et al.* [16] compared simulation data of sphere-sphere collisions with experimental results and theory in two dimensions using a lattice-Boltzmann method. A drag correction factor is determined to predict the effects of close range interactions between two particles at higher Reynolds numbers. This factor is determined as a function of a series of non-dimensional quantities, to fit with experimental results, shown in equation 2.6.

$$f = \frac{6\pi\mu r u_p}{h}(1 + 0.15Re_p^{0.687})exp\left(\frac{Re_p}{1.7}^{0.44}\frac{\rho_p}{\rho_l}^{0.19}\frac{r}{h}^{Re_p^{0.47}}\right) \tag{2.6}$$

They account for the non-resolved flow in the interstice using this drag correction factor. With the incorporation of this factor, the method shows very good consistency with the experimental results. In particular, there is a strong improvement over the unmodified method. It is noteworthy that the increased force is dependent on the unmodified force, hence its incorporation into a general framework requires careful evaluation of the error in unmodified prediction of force.

2.2.2. Proposed methods for correcting the lubrication force

Lubrication force between a two surfaces is a form of hydro-dynamic interaction and hence potentially modelled by any accurate fluid particle coupling model. However, unmodified, a general hydro-dynamic interaction model, *i.e.* one that does not directly take account of the proximity of the two

surfaces, is invariably an inaccurate approach. This is the case because in the near field of two particles in close proximity, flow with a characteristic length much smaller than that of two distal particles becomes significant. This has been noted by many authors [14, 11, 69, 30, 13, 12].

The determination of the correct force can be achieved by resolving the fluid and boundary with a sufficient resolution to capture the flow even at this length scale. However addressing the problem in this way is not currently feasible using current methods, as it is prohibitively expensive due to the necessary mesh refinement. Indeed resolving the flow sufficiently to capture behaviour at a separation of 0.01 of the particle radius would require $(\frac{2}{0.01} \times \text{minimum refinement})^3$ cells per smallest particle. Even with only three cells to resolve the gap this would be6.4e7 cells per particle. Eight cells across a gap would require in excess of4e9 cells per particle. cells per particle. Such an approach would require a non-feasible number of fluid cells even for very small case. This approach is however significantly more tractable for a two dimensional model.

The problem is particularly relevant in some frame works, in which overlap of particles is prohibited for numerical reasons. In such cases it is common to add a repulsive force in order to prevent this. It has been noted [11] that the nature of this force is not significant in many cases. However in cases when particles come into persistent contact these forces play an important role in the determination of the physics of the system. Benyahia *et al.* [70] noted that even in dilute systems, such interactions can be significant.

A number of propositions have been made as to how to improve the treatment of lubrication force and general hydro-dynamic interaction when particles are in close proximity. Many of these proposals involve adding force terms based on a simplified system. There are a number of problems associated with this:

- The correct results for the lubrication force for an arbitrary shape and arbitrary relative velocity in any flow is not known.
- The lubrication force is a part of the hydro-dynamic interaction and does not replace it. It therefore remains to ascertain how to correct the un-modified hydro-dynamic interaction to account for any additions made.

2.2.2.1. Collision prevention

Glowinski Glowinski *et al.* [14] stipulate a functional relationship of position and lubrication force to account for the under prediction of the lubrication force for the general hydrodynamic interaction directly, namely:

$$f_{ij} = \begin{cases} 0 & \text{if } d_{ij} \geq R_i + R_j + \kappa \\ \frac{c_{ij}}{\epsilon} \left(\frac{d_{ij} - R_i - R_j - \kappa}{\kappa} \right)^{-2} \frac{\vec{G_i G_j}}{d_{ij}} & \text{if } d_{ij} < R_i + R_j + \kappa \end{cases} \tag{2.7}$$

Where:

$\kappa,\ \epsilon,\ c_{ij}$	are fixed constants
$\vec{G_i G_j}$	is the vector joining the centres of the particles in close proximity
$R_i + R_j$	are the radii of the particles
d_{ij}	the distance between the two particles centre's *i.e.* $\lvert\vec{G_i G_j}\rvert$

A number of authors, including Feng and Michealides [30] also use this approach.

Nguyen and Ladd Within the lattice-Boltzmann method, Nguyen and Ladd [11] investigate the effect of incorporating analytical Stokes-flow results. The method shows promise and demonstrates accurate forcing even when particles are in very close proximity. Force as a result of sliding and relative rotation is also considered. However the method is only valid in cases of Stokes-flow in limited geometry cases. There is also the potential for particles to overlap, which while physical may present problems for some methods.

2.2.2.2. Numerical investigation

Feng and Michealides [30] investigate hydro-dynamic interaction in an immersed boundary-lattice Boltzmann framework. They pay particular attention to the case of particle collision. Their model includes allowances for particle collision and deformation. The additional contribution from non-captured interstitial repulsive force is added. The functional form of this force is that introduced by Glowinski *et al.* [14], as discussed in section 2.2.2.1. They also introduce a wall collision model in which a fictitious additional particle is placed on the other side of the wall, to replicate a non-deformable wall.

Numerical investigations using this method, in order to validate it, were conducted. A number of conclusions are made. Of interest to this work is the observation of influence of the choice of stiffness parameters and interaction length scale. For the case of two falling spheres, which subsequently collide the following observations where made.

- The method was shown not to have heavy dependence on these parameters, provided they where in a certain range.
- Some dependence was observed all variations.
- If the range of influence was increased to be too large, a significant influence on the results was observed.

2.2.3. Applicability of existing methods

There are a number of implementations of the immersed-boundary method in which close range solids-solid interaction is modelled by a non-physical repulsive force. In situations where such interactions are short lived and largely repulsive, such as dilute flows, this is not problematic and does not result in unphysical behaviour. Indeed a simple parametrisation of a collision, in terms of its impulse, often suffices to capture the interaction between two particles accurately. However, some fluid-solid flows are more strongly dependent on close range solid-solid interaction. One set of examples in which this is the case are granular beds, where particles are consistently in contact. In such cases the collision model is significant. The possibility of a significant impact of the collision model, even in a simple case, is supported by the results of Feng and Michealides [30] where the influence of the treatment of a single interaction is shown to have an affect on two sedimenting spheres. It is also noteworthy that often only spherical particles are considered, given the simplicity with which their collisions are described and the collision model implemented. In situations where the (non)sphericity is a determining factor in the system's dynamics, an alternative in which arbitrary shapes are modelled

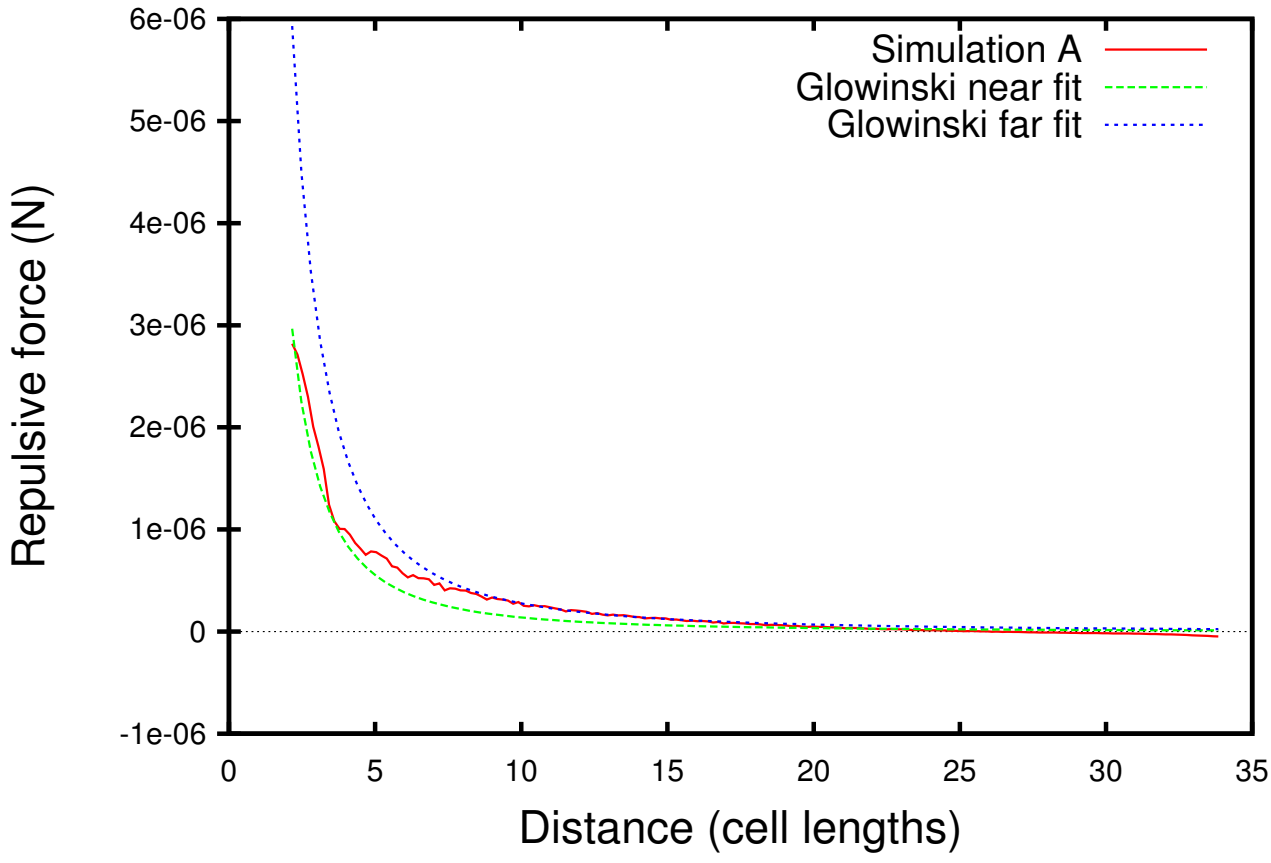
is necessary. There are numerous challenges to accurately modelling the short range interaction of arbitrary shaped particles in addition to the problems of modelling spherical-spherical interactions. Namely:

- Arbitrary contacts are difficult to describe in a numerical frame work.
- Large number of possible collision scenarios each require a force model able to account for them.
- This requires a large amount of parametrisation either though refined simulations or experiments.

In addition to these problems, it is necessary to have a lubrication force model that can be integrated into an existing hydro-dynamic interaction model, which allows particles to overlap. Hence any lubrication model must be able to account for the existing forces hydro-dynamic interaction and collision forces.

2.2.4. Comparison of methods

2.2.4.1. Accuracy



(a) Simulation A

Figure 2.1.: Graphs showing comparisons of the lubrication force as simulated and as predicted by Glowinski *et al.* [14]
2.1a shows a low Reynolds number simulation ($Re = 0.25$) and 2.1b shows a higher Reynolds number simulation ($Re = 5.0$)

The functional form stipulated by Glowinski *et al.* has a rapidly increasing repulsive force, as a function of decreasing minimum separation. While this prevents overlap which may be difficult to model, it makes it impossible to calibrate the model for a range of minimum separation distances. For example, two low resolution simulations calibrated at different separations give rise to two separate repulsive forces, as shown in figures 2.1a and 2.1b. In both cases the two calibrations while well calibrated for a particular separation differ by a factor of two. An additional limitation is that it does not allow for dependency on the rate of approach, which is in stark contrast to the theoretical results in which the rate of approach is critical.

2.2.4.2. Further interpretation

It is of significance that equations 2.3, 2.4 and 2.5 are only applicable for low Reynolds numbers *i.e.* $Re_p \ll 1$. In this case the Reynolds number is of the order of:

$$Re_p = \frac{\rho_f u_p r}{\mu}$$

Additionally the equations predicting lubrication force apply to the collision of a sphere and a wall, and not the more general case of spheres of potentially different radii. Note that sphere-sphere collisions encompass sphere-wall collisions as they can be treated as the limiting case when one sphere's radius tends to infinity. The restriction to low Reynolds number, while restrictive to other applications, is acceptable to the current work as low Reynolds numbers are expected. However the limitation to the form of sphere-wall collisions is substantial, as a wide variety of shapes and sizes of particle are present in embankment structures. Further more, as size and shape are significant factors in determining the dynamics of the system [71, 72, 4], it is important to have a model that accurately models a wide range of particles.

2.2.5. Amalgamation of hydro-dynamic interaction and correction forces

The repulsive force models in section 2.2.1 give predictions for the hydro-dynamic interaction force in an undisturbed flow. Incorporating this force into an existing hydro-dynamic interaction force model is potentially problematic. The problem arises as the repulsive force models do not incorporate the force as a result of any additional flow. Further more the hydro-dynamic interaction force already *inaccurately* accounts for the lubrication force. Hence neither of the forces or their sum is accurate.

Mammoli [69, 12] discusses the merits and implementation of a method for overlaying an empirical relation based on exclusively spherical particles, in a boundary integral equation framework. The method involves shrinking any two spheres in close proximity (close enough to have significant error in the prediction of lubrication force) and calculating the hydro-dynamic interaction based on this reduced radius. The close sphere correction is then added directly to the calculated force. This is justified by verifying:

- The only significant change to the force acting on any particle, of reducing the radius a small amount, is when the particles are in close proximity.
- In this case, the calculated hydro-dynamic interaction force acting on the shrunken particle is negligible compared with the force predicted by the force models.

This method is potentially adaptable to fit non-spherical bodies. However the following difficulties would need to be addressed:

- 'Shrinking' an arbitrary particle is potentially more complex than shrinking a sphere.
  For example a thin sheet in the vicinity of a sphere cannot easily be adjusted to fit the requirements.

- More complex shape amending algorithms may be computationally expensive and have unintended effects on the curvature of the particle's surface.

This method requires the accurate calculation of the correct lubrication force in the range of cases considered in advance.

The model proposed by Glowinski *et al.* [14] also has the advantage that the force can be added directly. Indeed the parameter κ is chosen such that the model does not contribute to particles within the region the un-modified method is considered accurate without modification.

The results from dimensional analysis could also be potentially adjusted to be incorporated into the hydro-dynamic interaction by limiting the region over which the force is computed using dimensional analysis. To do this the equation would have to be adapted to:

$$\phi\left(\left[\frac{r}{h}\right], \left[\frac{\mu}{vr}\right], \left[\frac{d}{r}\right]\right) = \frac{f}{\kappa \cdot \mu \cdot u \cdot r} \tag{2.8}$$

Where d is the distance over which the force is computed using the described method.

2.3. Mirroring immersed-boundary method

2.3.1. Mirroring

While a number of immersed-boundary methods are available, the Mirroring Immersed-Boundary Method (MIBM) has been chosen for this project, see section 1.3.5.3. The mirroring immersed-boundary method used is based on the ideas introduced in [58]. It is fully implicit and resolves the immersed boundaries, in the form of a triangulated surface meshes, by modifying the Navier-Stokes equations inside the boundary. In particular, the mirroring immersed-boundary method mirrors the velocity field through the surface triangles by replacing the Navier-Stokes equations immediately inside the body. Thus creating a fictitious flow in a layer inside the particle see figure 2.2.

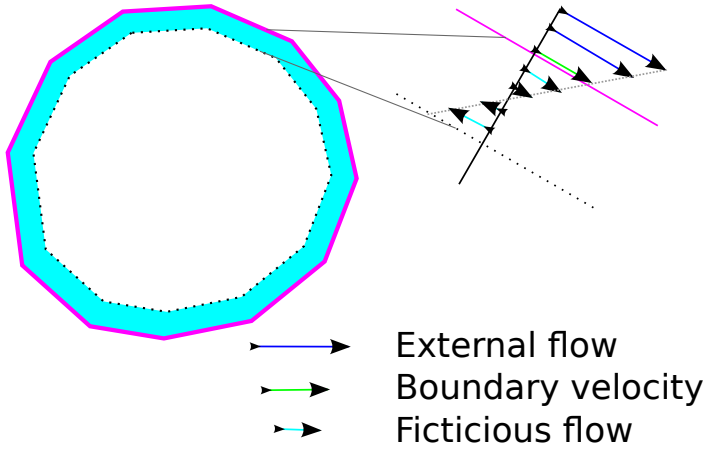


Figure 2.2.: An arbitrary immersed body showing the immediate interior of the body (blue). The blue, fictitious, region of the fluid is considered to represent both the fluid and particle velocities. The force acting on the body is calculated as the opposite of the force acting on the blue region of the fluid, from the exterior fluid. This force is calculated from the Navier-Stokes equations as described in sections 1.3.7 and 2.3.2.

This is done for each interior cell near the boundary, by 'mirroring' the fluid through the immersed-boundary. For each of the interior points ($\vec{x}_{iib}$) an exterior 'mirroring point' ($\vec{x}_e$) is constructed by using the closest point on the immersed-boundary to $\vec{x}_{iib}$ ($\vec{x}_b$). This arrangement is shown diagrammatically in figure 2.3.

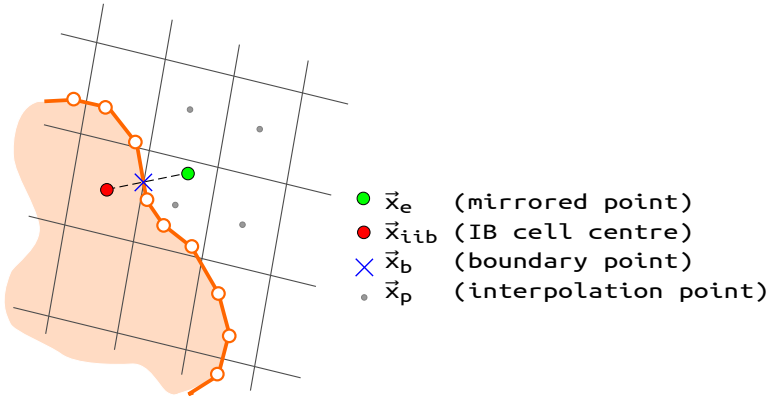


Figure 2.3.: Diagram showing the mirroring point in the flow domain ($\vec{x}_{\mathrm{e}}$) arising from the interior cell near the boundary ($\vec{x}_{\mathrm{iib}}$) and the nearest point on the immersed-boundary ($\vec{x}_{\mathrm{b}}$).

In particular, for a given interior cell ($\vec{x}_{iib}$), an exterior 'mirroring point' ($\vec{x}_e$) is constructed along the line joining $\vec{x}_{iib}$ and the closest point on the surface to $\vec{x}_{iib}$ ($\vec{x}_b$). Equivalently, the mirrored point's location is given by:

$$\vec{x}_e = \vec{x}_{iib} + (1+\alpha)\mathbf{dn} \tag{2.9}$$

$$= \mathbf{x}_b + \alpha\mathbf{dn} \tag{2.10}$$

where $\mathbf{dn}$ is the vector from the $\mathbf{x}_{iib}$ to $\mathbf{x}_b$, the closest point on the immersed-boundary, and α a fixed parameter of the method. If $\alpha = 2$, as in [58], $\mathbf{x}_{iib}$ is at an equal distance to the immersed-boundary as $\mathbf{x}_e$.

The boundary velocity is then enforced by imposing a linear variation in the velocity field between

$\mathbf{x}_{iib}$, $\mathbf{x}_b$ and $\mathbf{x}_e$ at each time-instance. This is achieved implicitly by setting a relation between the velocity $\mathbf{x}_{iib}$ and implicit interpolation coefficients of $\mathbf{x}_e$. Namely:

$$\mathbf{u}_b = \frac{(\alpha \mathbf{u}_{iib}) + (\sum^i \phi^{\mathbf{x}_{p^i}} \mathbf{u}_{p^i})}{1+\alpha} \tag{2.11}$$

where $\phi_i^{\mathbf{x}_p}$ is the interpolation coefficient of the i^{th} cell centre to in the interpolation stencil p_i, to $\mathbf{x}_e$. In [58] the interpolation coefficients are given by Shepard's interpolation. Literature on the available interpolation schemes is discussed in the appendix (section A.1).

2.3.2. Force calculation

As discussed in section 1.3.7, the mirroring immersed-boundary method does not explicitly give a value for the force required for the coupling. Force is reconstructed by integrating the fluid stress over the surface *i.e.* equation 1.3. Torque is calculated from a similar surface integral. The quantities involved in this calculation are not explicitly known at points at which they are required. Possibilities for the estimation of these quantities are to interpolate the quantities to boundary or approximate the boundary by a surface over which these quantities are known. Examples are given by Zastawny [61]. A more detailed description of how this interpolation is implemented in this project is given in section 3.3.

2.3.3. Limitations

2.3.3.1. Spurious oscillations

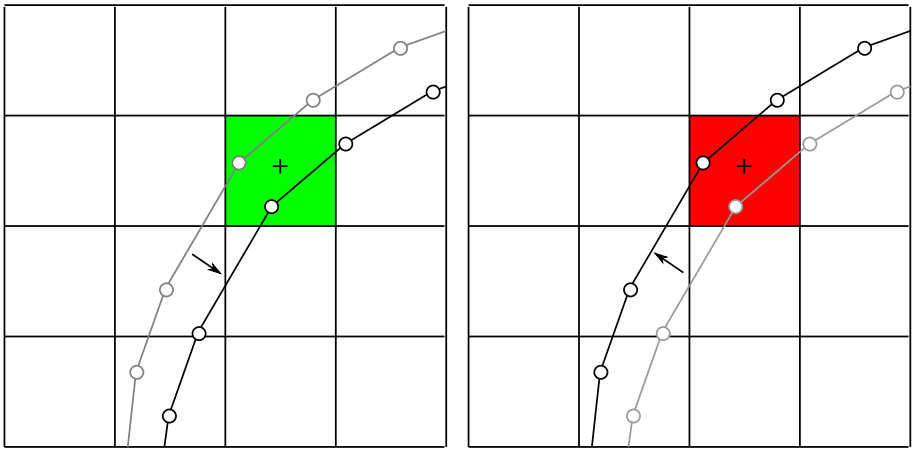
One of the principle limitations of sharp interface resolving methods is their susceptibility to sporadic oscillations as particles move [64, 73, 43]. Seo and Mittal [64] identify the source of these oscillations as the transition of Eulerian grid cells and their role in enforcing the immersed-boundary. In particular, when a cell transitions from being a boundary-interior to boundary-exterior creating a 'fresh' cell, or conversely from a boundary-exterior to boundary-interior creating a 'dead' cell. These transformations are illustrated in figure 2.4.

In this case a discrete change in the continuity and momentum equations occurs over a single timestep. The rapidly changing flow field required to accommodate such transitions often results in large fluctuations in the pressure field around the immersed-boundary. The magnitude of these oscillations is heavily dependent on the spacial and temporal resolution of the simulation. While increasing the spacial resolution decreases the necessary jump to the fluid phase, increasing the temporal resolution decreases the time over which the change takes place, increasing the magnitude of the oscillations [64].

2.4. Internal instability

2.4.1. Assessment criteria for susceptibility to internal erosion

Referring to figure 2.5, Garner and Fannin [6] propose three conditions that determine the instability of granular materials these are; the stress conditions, the hydraulic loading and the material's susceptibility. While there are problems that can arise as a result of a combination of stress conditions and hydraulic load - including heave, seepage and arching - these are predominately well understood



(a) Fresh cell (b) Dead cell

Figure 2.4.: Schematic of 'fresh' and 'dead' cell transition. The cell marked '+' changes between interior and exterior as the immersed-boundary moves. Figure 2.4a shows the point transitioning from being interior to being exterior, creating a 'fresh' cell. Figure 2.4b shows the point transitioning from being exterior to being interior, creating a 'dead' cell.

and designed against. Material susceptibility however is the subject of ongoing research, as the relationship between material properties, such as particle size distribution and frictional and resitutive laws, and the stability of the material are not well understood. There is however consensus amongst the literature to suggest that the principle determinant of the material susceptibility is their particle size distribution (Kenney and Lau [72], Kezdi [71], Fell *et al.* [4], Rönnqvist [74] and USACE [75]). The mechanical basis for this, and the effects of hydraulic conditions on which gradings are stable and vice versa, are poorly understood.

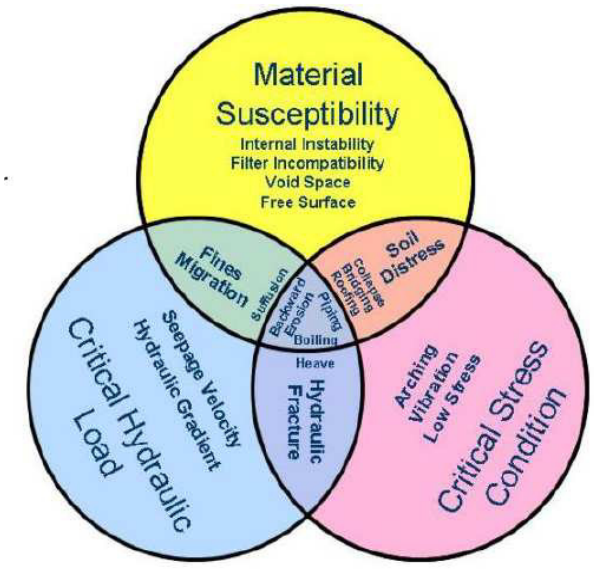


Figure 2.5.: A Venn diagram of the criteria for failure and possible consequences as proposed by Garner and Fannin [6].

2.4.1.1. Kedzi

Kedzi made the prediction that significant suffusion would not occur in a granular material if the material was capable of self-filtering, the condition for which is based on Terzaghi's filter criterion, see equation 2.15. This is a non-conservative approach as it over looks the problems associated with a self filtered material, namely significantly varying porosities and hence permeability, leading to significant hydraulic gradients. See figure 2.9.

2.4.1.2. Kenney and Lau

Kenney and Lau [72] made an evaluation of the susceptibility to internal erosion of a material of a certain particle size distribution, based on the interactions of particles of differing length scales. The principle assumptions are:

- That the material under evaluation has a matrix support structure (or under-filled packing Vaughan [76]).
- That any particle retained by this clast is retained in such a way as to retain smaller particles as if they where themselves part of the clast.
- The primary fabric constricts particles greater than 0.25 times the size of themselves.

These assumptions are then exploited via $H : F$ curve analysis, where:

$$F(x) \text{ is the weight, or \% weight, of particles of size } x \text{ or less} \tag{2.12}$$

$$H(x) = F(4x) - F(x) \tag{2.13}$$

The predicted requirement for stability is then:

$$H > F \tag{2.14a}$$

or

$$F(4x) > 2F(x) \tag{2.14b}$$

Figure 2.6 shows an example of how H and F are determined. Figure 2.7 show an example of how this criterion is applied to a range of soil gradings.

The justification for this is that for each particle size (x), if the the soil contains a sufficient relative volume of particles that will retain the that size of particle that size will itself be retained. This assumes the total volume of voids of the soil, less grains of size less than $4 \times x$, is related to volume of particles of size x.

The model does not evaluate what the length scale/size of particles is in the principle fabric, as opposed to merely being retained by it. Indeed no distinction is made between the size that can't wash out of a load bearing structure and the size that would be sufficient to contribute to the structure. This may result in non-conservative estimates of internal stability as the throat size of a retained but non-stressed array of particles is determined by the size of the voids in the restricting fabric, not the size of voids the array would make if stressed.

In some cases the model may be overly conservative as it does not account for the ability of a soil to remain stable via a matrix supported skeleton and hence may predict a the case of a stable,

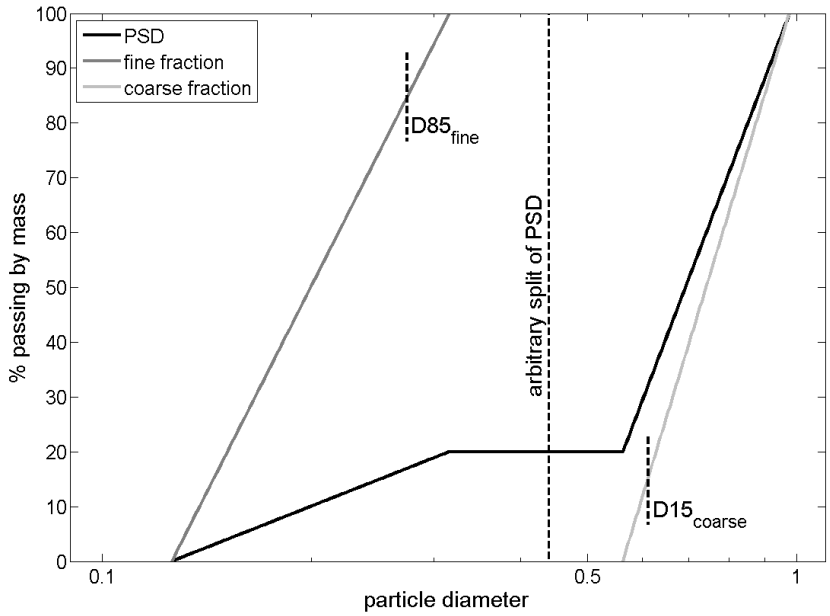


Figure 2.6.: An explanation of the H and F evaluations used in the Kenney and Lau method of determining stability of soils. Reproduced from the work of Shire [77].

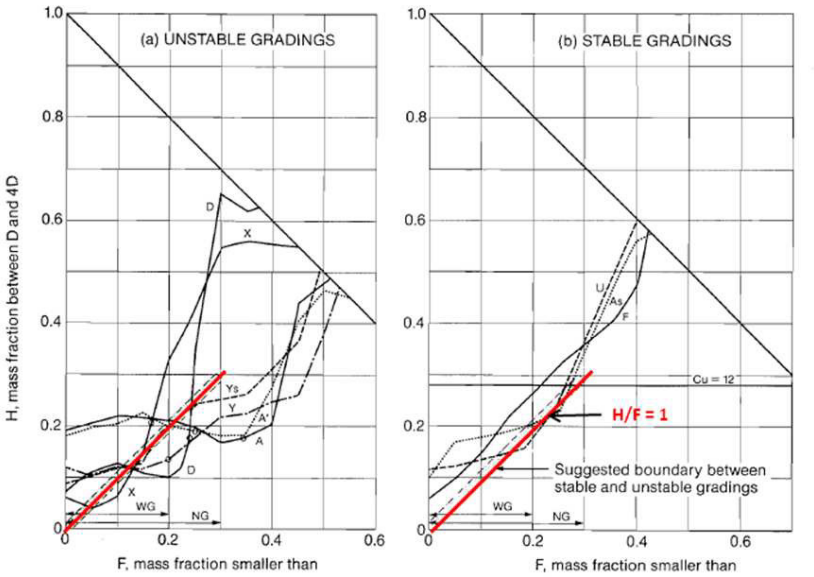


Figure 2.7.: Example and validation of the Kenney and Lau method of determining stability of soils using experimental results [77].

matrix supported, soil as unstable. This is not necessarily problematic as matrix supported soils are prone to suffosion, and the removed fines content would make the soil prone to instability. An under conservative nature is also inherent as the method assumes that if the soil is stable at the arbitrary division of the PSD curve the entire soil will be stable. This may be an accurate assumption in the case of a typical soil however this may not always be the case. Indeed Li and Fannin [78] find Kenney and Lau's method less accurate at predicting internal instability in gap-graded soils.

2.4.1.3. Wan and Fell

Wan and Fell [79] investigated the factors that influence the internal stability in widely graded soils, using upwards and downwards seepage tests. They concluded that gaps within the particle size density distribution within soils to be the primary cause of susceptibility to instability (*i.e.* gap graded soils are prone to internal erosion). A dependence on the fines content, and density of the soils was also identified however no clear relationship was established.

Wan and Fell [79] also evaluated the impact of the likelihood of internal instability based on the predictors of the method of Burenkova [80] (shown in figure 2.8). This method considers discrete points along the material's particle-size cumulative probability function, as well as the method proposed by Kenney and Lau [72]. It was found that whilst variations on Kenney and Lau's predictions showed good correlation with probability of internal instability the inclusion of Burenkova's condition was found advantageous.

In addition Wan and Fell [79] showed that not more than 40% of any material is erodible. The implication being that soils with more than 40% fines must be matrix supported, and hence whilst these soils are not inherently unstable, they are prone to backwards erosion (see section 1.2.1.4) and unlikely to be capable of effective self filtration, this is supported by Skempton and Brogan [81].

2.4.1.4. Li and Fannin

Li and Fannin [78] undertook an evaluation of the two criteria, for internal stability, in common use; that of Kezdi and that of Kenney and Lau. Figures 2.9 and 2.10 show the comparison of the methods. It is concluded that the Kezdi's method is more successful at predicting instability in gap-graded materials whereas Kenney and Lau's method is moremsuccessful at predicting instability in well graded material.

2.4.2. Self-filtration

The above sections discuss the potential for preferential movement of particles within a given homogeneous structure, defined by the PSD, and packing density. However in some cases the movement may be stable, *i.e.* not perpetuated, and hence benign. Indeed if a limited amount of particle migration, induced by hydraulic action, yields a stable soil then the soil is unlikely to undergo internal erosion. The potential for this to occur will be considered. Self-filtration may occur via a number of mechanisms. Lafleur *et al.* [82], Chapuis [83], Skempton and Brogan[81] and Li [5] present elaborations upon the Terzahgi criterion (2.15), which is widely accepted as valid.

$$4 \times d_{15} < D_{15} < 4 \times d_{85} \tag{2.15}$$

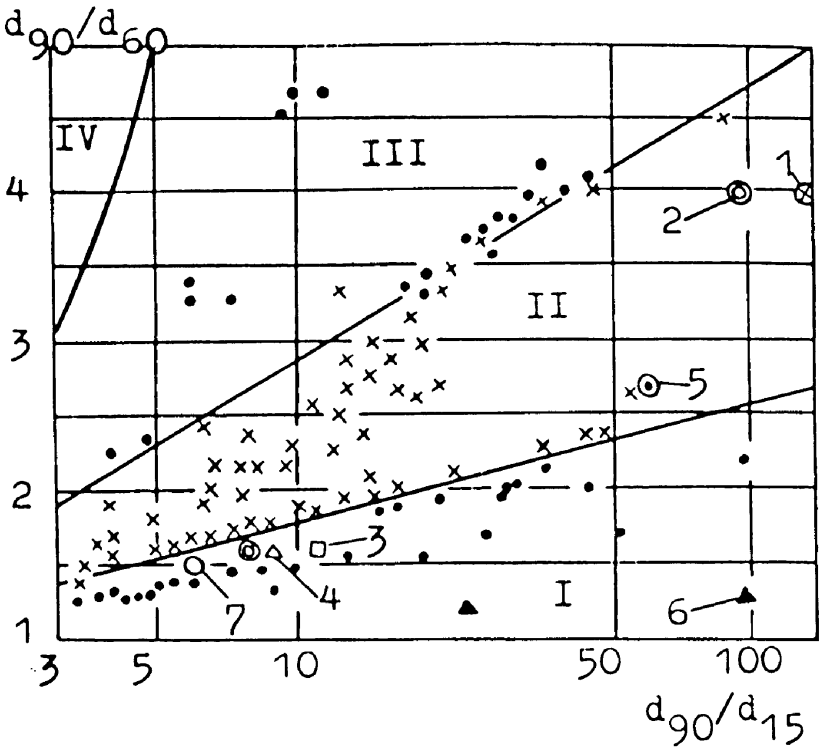


Figure 2.8.: Shows the boundaries between suffusive (internally unstable) and non-suffusive (stable) soils as proposed by Burenkova [80].
I Zone of suffusive soils.
II Zone of non-suffusive soils.
III Zone of suffusive soils.
IV Zone of artificial soils.
· Suffusive soil.
x Non-suffusive soil.

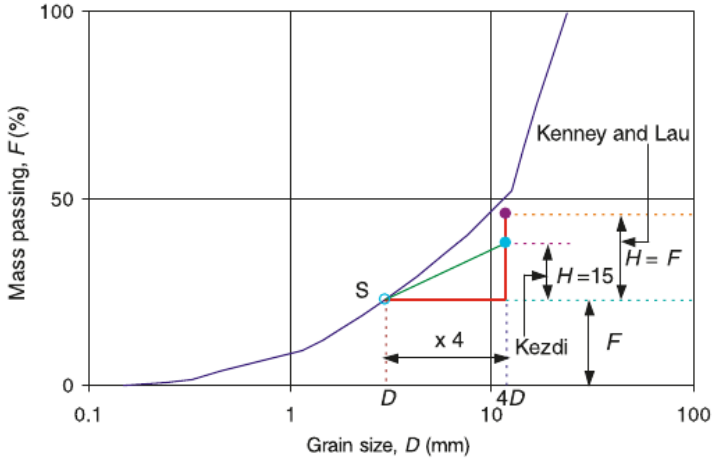


Figure 2.9.: An explanation of the Kenney and Lau and Kezdi methods of determining stability of soils.

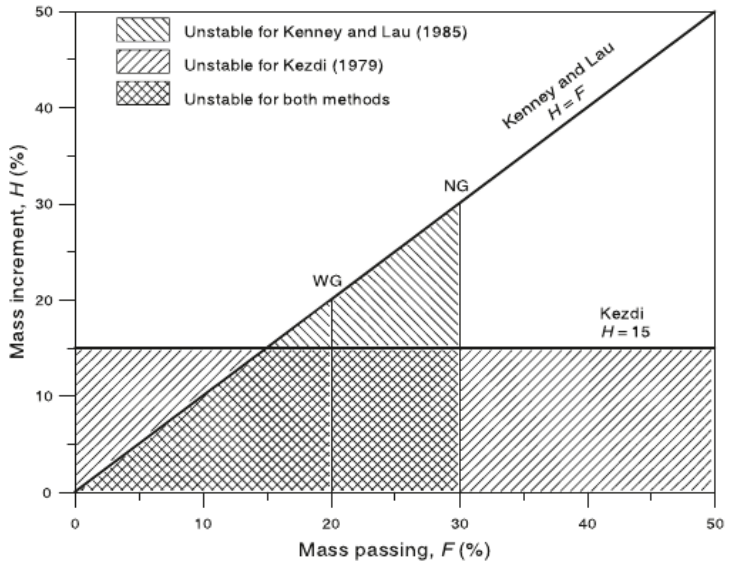


Figure 2.10.: A comparison of the Kenney and Lau and Kezdi methods of determining stability of soils.

However Rönnqvist [74] notes widely graded moraines (glacial tills), common in embankment dams, are unlikely to satisfy the Terzahgi criterion. In such cases the criteria for initiation of suffusion is no longer overly conservative in the prediction of it's propagation.

3. Numerical method

3.1. Framework

3.1.1. Single phase flow

The current work considers incompressible Newtonian fluids, governed by the Navier-Stokes equations 1.1, consisting of the continuity equation, equation 3.1, and the momentum equation, equation 3.2, repeated below. The method employed to solve these equations that given by Oliveira and van Wachem [60].
The equations are discretised by integration over the control volumes *i.e.* cells. Second order terms become the product of coefficients and old time-instance values.
The continuity equation 3.1 is solved using momentum weighted interpolation, to prevent pressure velocity de-coupling. The method used is that given by Rhie and Chow [84]. In particular the CFD code, MultiFlow, developed by the Berend van Wachem research group is employed.

$$\frac{\partial}{\partial x^i}\left(u^i\right) = 0 \tag{3.1}$$

$$\rho\frac{\partial u^j}{\partial t} + \rho\frac{\partial}{\partial x^i}\left(u^i u^j\right) = -\frac{\partial}{\partial x^j}p + \frac{\partial}{\partial x^i}\tau_{ij} - \left(R^j u^j + S^j\right) \tag{3.2}$$

where R represents the source contributions linear in velocity components, S represents flow field independent sources. An example of a momentum source term (S) is gravity. Source terms linear in velocity (R), while not used in this work, can be used to represent the effects of an additional immersed phase. τ_{ij} represents the stress tensor, given in equation 3.3:

$$\tau^{ij} = \mu\left(\frac{\partial u^i}{\partial x^j} + \frac{\partial u^j}{\partial x^i}\right) + \delta^{ij}\mu\left(\frac{\partial u^k}{\partial x^k}\right) \tag{3.3}$$

3.1.2. Applicability

A number of improvements have been made to the Mirroring Immersed Boundary Method (MIBM), introduced in section 2.3. In all cases studied the underlying discretisation scheme is that given in this section. However it is hoped the subsequent developments will translate and apply to other discretisation schemes.

3.2. Enforcing the no-slip condition

3.2.1. Determination of cell type

In themirroring immersed-boundary method the no-slip boundary condition in enforced by modifying the equations of the cells that are in or near an immersed-boundary. Indeed, through out the

mirroring immersed-boundary method the equation system for each cell depends on it's location relative to theimmersed-boundaries. Eulerian grid cells in the method are characterised into one of the following categories:

Complete-exterior Cells in this category are those that are completely outside of any immersed boundary and are treated in an unmodified way. *I.e.* the Navier-Stokes equations are solved, assuming the surrounding cells to have reliable fluid data.

Boundary-exterior Cells in this category are those that have a cell centre outside of any immersed boundary however are close and may be cut by the immersed-boundary. In some variations of the mirroring immersed-boundary method, these cells are treated differently than the complete-exterior cells however a variation the premise remains the solution of the Navier-Stokes equations. Modifications are typically in the form of a different calculation of terms in the Navier-Stokes equations to account for unreliable fluid properties in boundary-interior cells, or the presence of an immersed-boundary changing the cell's boundary.

Boundary-interior Cells in this category are those that have a cell centre inside of the immersed boundary however are close and affect the solution of the exterior flow field. These cells form the boundary and account for the presence of the immersed-boundary. Numerous propositions or the handling of these cells have been made. The principle behind the governing equations of these cells is the continue/extend the fluid domain described by complete-exterior and boundary-exterior cells. This is done in such a way as to satisfy the boundary conditions imposed by the immersed-boundary.

Complete-interior Cells completely inside the immersed-boundary are excluded from influencing the flow in the other cells. Flow values and coefficients of these cells are not calculated or used for other cells. In order to simplify the numerical framework an equations for these cells are set however are chosen only to provide solid phase values for post-processing.

In addition to being used to determine the treatment of the cell itself, the category of the cell is also taken into consideration when determining interpolation weights for surface force calculation.

The process of determining the categorisation of cells is as follows:

1. Cells with a cell centre inside an immersed-boundary are determined by ray tracing.
2. Interior cells which share a face with an exterior cell, or an interior cell of a different body, are tagged boundary-interior.
3. Exterior cells which share a vertex with an interior cell are tagged boundary-exterior.
4. Remaining cells are tagged in the obvious way.

Figure 3.1 shows the cell type determination in two dimensions. The motivation for this choice of determination of cell type is as follows:

- Boundary-interior cells predominant effect is via there contribution to the continuity equation of adjacent cells. Such contributions are made over faces and hence are only necessary for face neighbours of cells with continuity equations

- Boundary-exterior cells are chosen to be the minimum set of cells that might be cut by the immersed-boundary or share a face with a boundary-interior cell. Specifying a larger set of boundary-exterior, as opposed to complete-exterior cells is not problematic as the resultant equation system would be identical. However due to the increased cost of checking for the influence of the immersed-boundary means cells not expected to be influenced by the boundary should not be tagged.

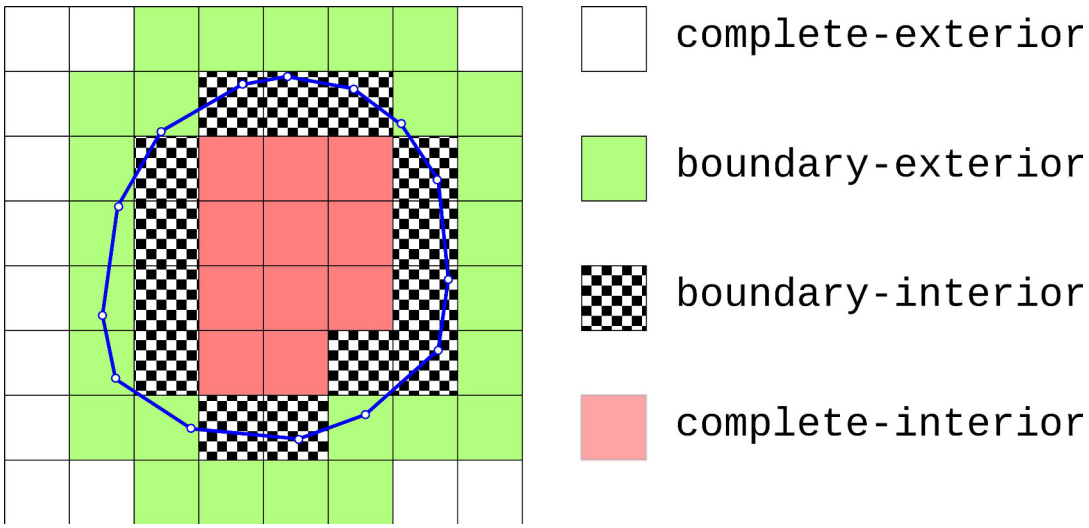


Figure 3.1.: A schematic of a two dimensional Eulerian grid with an immersed-boundary and the resulting cell-type as determined by the method given in section 3.2.1.

It is assumed that no complete-exterior cells are intersected by an immersed-boundary. *I.e.* that an immersed-boundary cannot intersect a cell that is more than two cells away from the nearest interior cell.
It is possible, for some combinations of immersed-boundary geometry and Eulerian mesh, that this condition would not hold. However the assumption that this condition aways holds is reasonable. Indeed, for a contraditction, the immersed-boundary would have to contain significant sub-cell-scale features, *i.e.* the immersed-boundary must contain features more detailed than the Eulerian mesh. Where this to be the case, the flow around the body would necessarily be under resolved.
If desired, although not recommended, an additional layer, or layers, of cells around the body could be tagged as boundary-exterior. The use of additional boundary-exterior cell tags, instead of complete-exterior cell tags is safe, and would not otherwise effect the method, as to treat an un-intersected cell as boundary-exterior cell does not change the equation system. In the developed method, as only smooth bodies (relative to the grid spacing) are considered, this modification is not used.

A number of methods of treating the various cell types are investigated in this research. As complete-exterior cells are required to be governed by the Navier-Stokes equations and the flow inside the body, *i.e.* complete-interior cells, is not physical the focus is on the treatment of boundary-interior and boundary-exterior cells. The methods used are described in this section.

3.2.2. Complete mirroring

The principle of the complete mirroring method is based on the mirroring immersed boundary method introduced by Mark and van Wachem [58]. Introduced in section 2.3, the mirroring principle is to enforce a relationship between the velocity inside and outside of the immersed-boundary to reconstruct the correct velocity at the interface. To achieve this, for each boundary-interior cell, an exterior point at a position mirroring the boundary-interior cell centre through the immersed-boundary is created. The fluid properties at the boundary-interior cell are then set to be the mirror of the fluid properties of this 'mirror point' outside of the immersed-boundary. The location of the mirror point of a boundary-interior is determined using the point on the immersed-boundary, nearest to the boundary-interior cell centre. The mirror plane is then taken as the plane with normal in same direction as the surface of the immersed-boundary at the closest point. This mirroring is analogous to the technique used to enforce a boundary condition at a cell face using a ghost cell, as shown in figure 3.2.

The situation is slightly different however, as there is not an obvious choice of how to interpolate fluid variables to the face from the cell-centres. The two main reasons for this are:

- The domain boundary face separates two cells, the cell on the boundary and the ghost, these two cells form the basis of the interpolation to the face. This does not always hold for the mirroring immersed-boundary method.

- As the mirroring point moves, it in necessary to update the interpolation to the point, unlike a fixed mesh for which the interpolation requirements are constant. This adds the additional complexity of ensuring a smooth interpolation to the immersed-boundary as the immersed-boundary moves.

Section 3.2.2.1 discusses the specific implementations used in the current work.

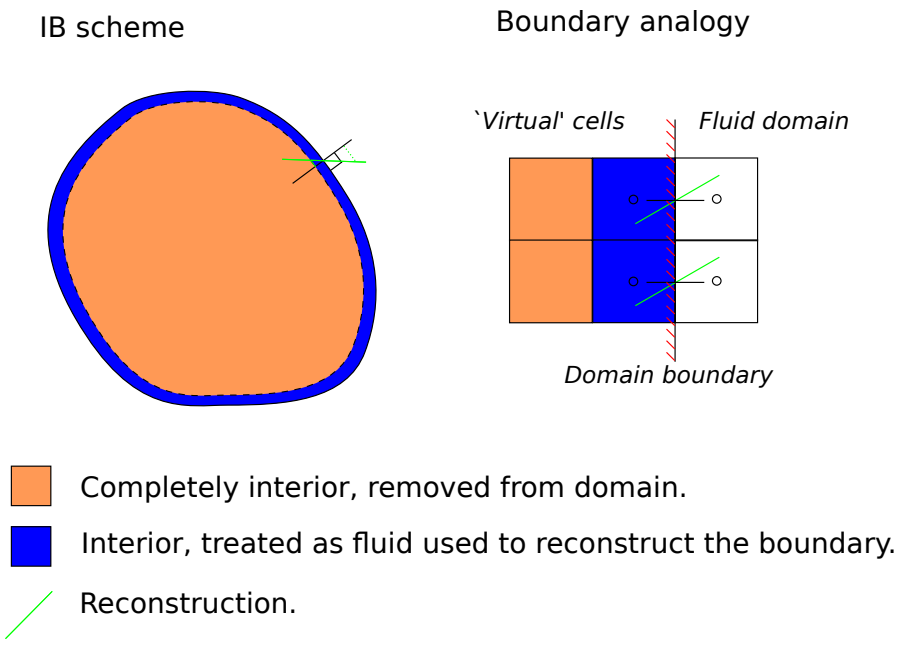


Figure 3.2.: Analogy between an arbitrary immersed body and a domain boundary. Left: immersed-boundary, right: domain boundary.
In both cases the region outside but near to the fluid (blue) is considered to represent the continuation of the exterior flow field such that the boundary velocity is reconstructed. The deep interior (orange) is discounted.

3.2.2.1. Developments to the mirroring equation

Interpolation to the mirroring point The stencil used for interpolation is chosen to be eight points surrounding the mirrored point ($\mathbf{x}_e$). These eight points form the verticies of a cube of connected cells around the mirrored point. The interpolation method used, for determining the interpolation coefficients to these points is a weighted linear interpolation. On a Cartesian grid, this is continuous even under the transition of stencil, which occurs as the point passes one of the cubes faces. This is because the four points not on the 'crossed' face have zero contribution, *i.e.* coefficient of interpolation equal to zero. If the Eulerian mesh is skewed, continuity can be maintained by interpolating linearly from a reduced stencil. The reduced stencil is comprised of the cell centre closest to $\mathbf{x}_e$ and three of is neighbours, from the original stencil. The three neighbours are chosen such that the interpolation stencil forms the smallest enclosing tetrahedron. Figure 3.3 shows how this scheme maintains continuity. During the transition shown, as $\mathbf{x}_e$ passes one of the faces of the tetrahedron, the removed and added cells do not contribute, *i.e.* have coefficient of interpolation equal to zero. An additional advantage of this reduced interpolation stencil is that all of the weights are guaranteed to be positive, which is advantagious for numerical reasons.

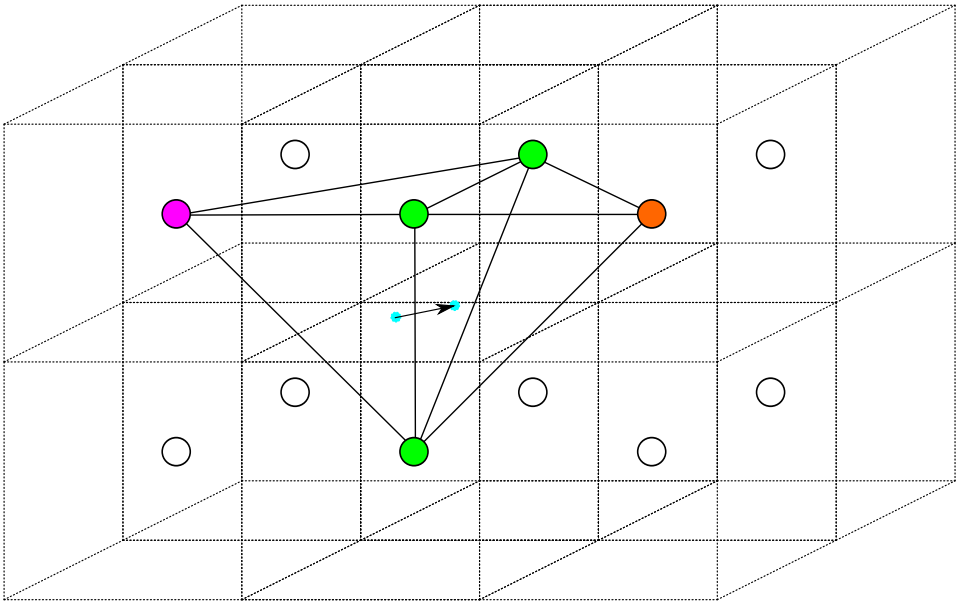


Figure 3.3.: Schematic of a changing stencil for the mirroring point. As the mirroring point (small blue dot) crosses the plane of points used for interpolation (green stencil points) the interpolation stencil changes. The purple stencil point is removed and the orange one added. However during this process only the green stencil points actively contribute as the interpolation point lies on the face bounded by these points. Hence the interpolation coefficients do not 'jump'.

Location of the mirroring point The location of each mirror point uses the method introduced in section 2.3 and described in section 3.2.2. In particular as given by in equation 2.9, with $\alpha = 0.5$. The mirroring equation with $\alpha = 0.5$ reduces to equation 3.4, given below.

$$\mathbf{u}_b = ((1/3)\mathbf{u}_{iib}) + (2/3)(\sum^{i} \phi^{\mathbf{x}_{p^i}} \mathbf{u}_{p^i}) \tag{3.4}$$

$\alpha = 0.5$ was chosen as the smaller the value of α the more closely the equation enforces that the immersed-boundary velocity matches the interpolated velocity from a compact stencil, *i.e.* from only it's fluid neighbours. Additionally the interpolation stencil is more compact, hence less far-removed cells from the immersed-boundary point and the boundary-interior are used to enforce the relationship. This principle is shown in figure 3.4. Indeed if $\alpha = 0$ equation 2.9 reduces to interpolating directly to the immersed-boundary point, *i.e.* $\mathbf{u}_b = \sum^{i}(\phi^{\mathbf{x}_{p^i}} \mathbf{u}_{p^i})$. However if α is too low the equation system does not retain the desirable numerics associated with the diagonal dominance guaranteed when α is large. As α becomes small, the diagonals for the velocity (momentum) equations at the boundary-interior cell reduce. Indeed unless the boundary-interior cell forms part of the stencil, which is not guaranteed, the diagonal coefficient is $\frac{\alpha}{1+\alpha}$. In the case $\alpha = 0.5$ the diagonal remains at least a third, where the off-diagonal coefficients sum to less than two thirds. An additional advantage is the mirroring point moves more smoothly as the immersed-boundary moves. Indeed the movement in the mirror point as a cell transitions from boundary-interior to boundary-exterior is proportional to α. This slower movement in beneficial as the 'snap' at a 'fresh' or 'dead' transition is smaller.

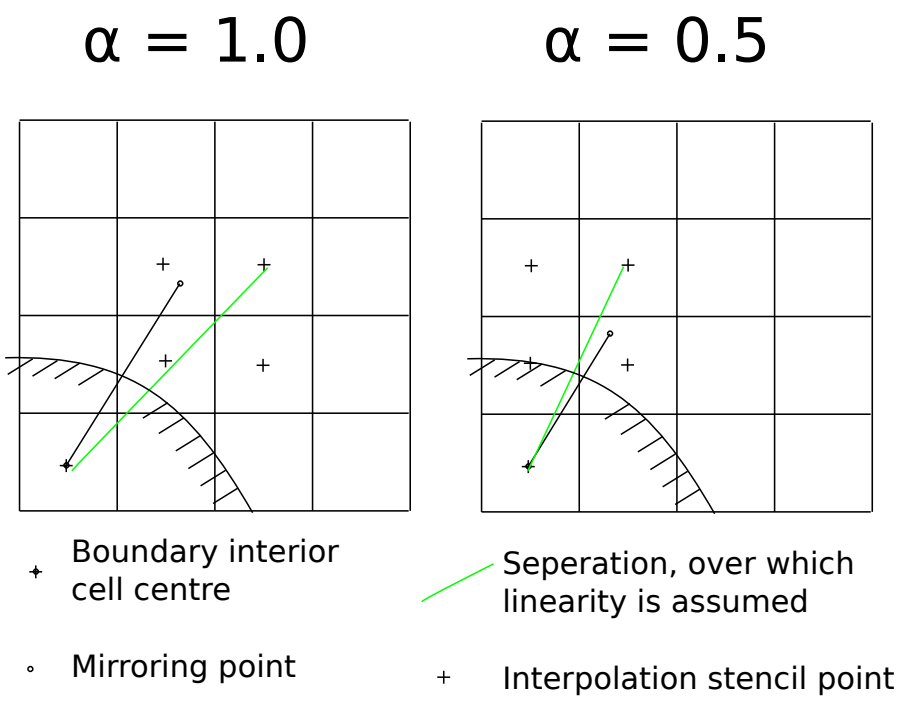


Figure 3.4.: A reduction in α decreases the size of the stencil used to enforce the mirroring condition and hence does so more accurately and reliably.

Modification to the mirroring equation to ensure smoothness Spurious oscillations in flow field, pressure and hence forcing arise from the change in stencil used to enforce the immersed boundary. A potential contributing factor to this is the transition in position, and hence governing equations, of the mirroring point. This can be problematic if the solution at the old timestep is not close to the solution of the current equation system. In the case of the mirroring equation the potential discontinuity arises when a boundary-exterior cell becomes a boundary-interior cell (a 'dead' cell). After such a transition, the mirroring equation will enforce the immersed-boundary velocity on the new boundary-interior cell. However, at the previous time instance, when the cell was boundary-exterior, no such equation was present. The Navier-Stokes equations, combined with the adjacent immersed-boundary cells will provide a prediction that, depending on refinement, interpolation scheme *etc*, should give a fluid velocity close to the value before the transition. However the accuracy of this prediction is not directly enforced, further more there is no guarantee the velocities converge as the time-step separating the transition tends to zero. Hence when this transition occurs there may be a discontinuity. Figure 3.5 shows such a transition and how the boundary-exterior cell may have fluid properties different to that of an adjacent immersed-boundary.

The problem may be seen as the inability of the method to recognise when a boundary-exterior cell is close to the immersed-boundary. This is improved (*i.e.* smoothed) by enforcing the immersed-boundary mirroring condition in a way that enforces that exterior cells, close to the boundary, have velocity close to that of the immersed-boundary. These are the cells that may form part of the immersed-boundary stencil in the near future or, less significantly, have been excluded in the near past. The principle of the 'averaged-mirroring' modification is to incorporate a mirroring about the boundary-exterior cells into the existing mirroring (about the boundary-interior cells). This allows the correct enforcement of velocity at boundary-exterior cells close to the boundary. This is achieved by modifying the mirroring equation for each of the boundary-interior cells with a contribution from

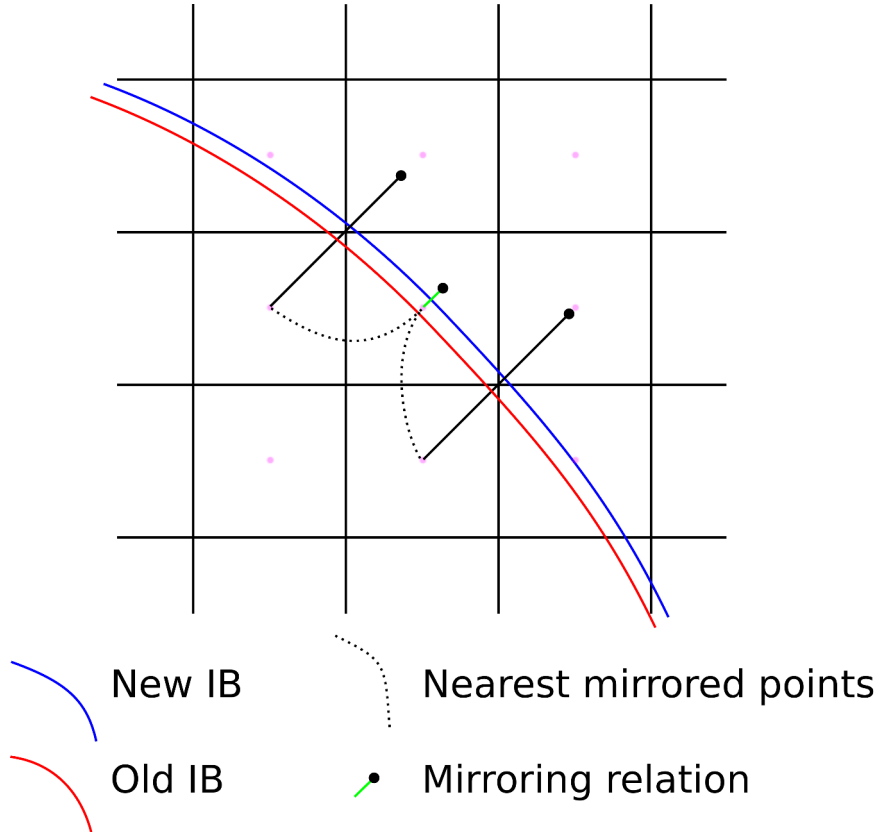


Figure 3.5.: A two dimensional schematic of an immersed-boundary at two time instances and the resulting mirroring equations. In particular, the transition of a mirroring point and the new mirroring equation arising from a newly boundary-interior cell. The previous time instance only loosely enforces the boundary velocity at the current time instance. *I.e.* with error $\mathcal{O}(\delta x)$.

each of its boundary-exterior face neighbours. The contribution from these neighbours is derived from the following procedure:

- For each of the boundary-exterior neighbours, the closest point to the boundary is computed.
- An internal mirroring point (M) is calculated in the same way as it is for boundary-interior cells through their closest surface points. These points will be inside the immersed-boundary.
- Each internal mirroring point is interpolated to via the same stencil as is used for the exterior mirroring point of the boundary-interior cell.
- The resultant equations are combined using a weighting average.

Figure 3.6 shows the configuration of the two types of mirroring point, the original mirroring point [58] and the additional 'interior' mirroring point (M).

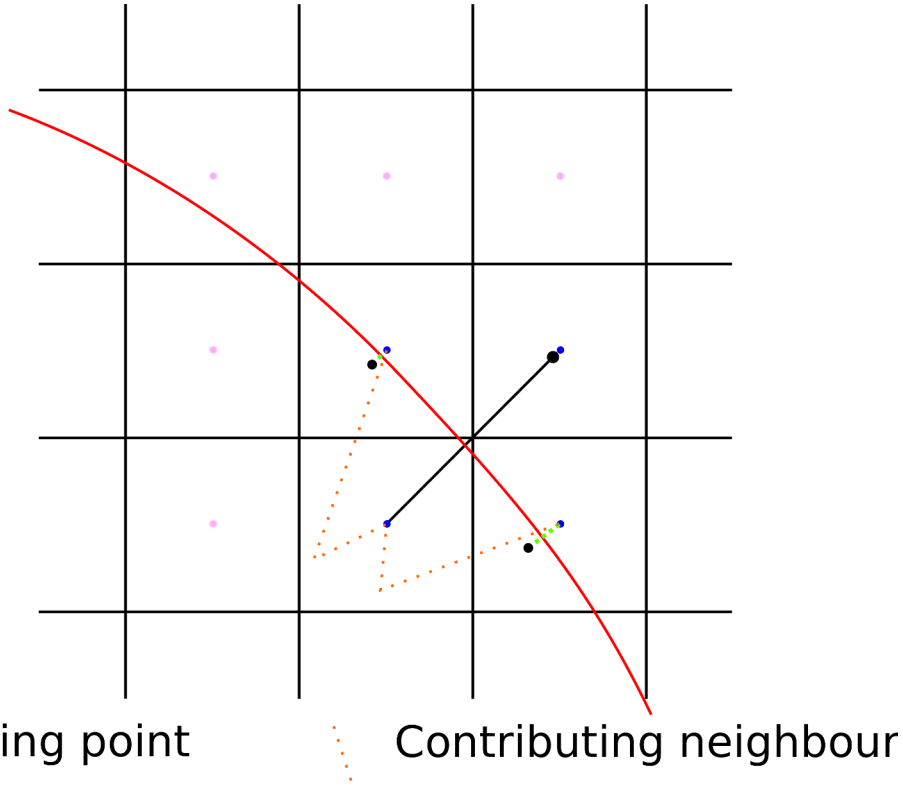


Figure 3.6.: Diagram showing the vicinity of an immersed-boundary (shown with a red line) with the old and new mirroring points (blue dots), giving an example of how boundary-exterior neighbours are used to smooth mirroring. One central mirroring equation (black line) is shown, along with contributions made from two boundary-exterior cells and the corresponding internal mirroring points (dotted green lines). The introduction of the boundary-exterior mirroring equations, which replicate the potential mirroring equation after the immersed-boundary crosses the cell centre, allows a smooth transition of equation systems.

The combination of the mirroring equations (weighting) is based on how close the mirroring point is from the immersed-boundary. The requirement being that if the mirroring point is very close to the surface it should be weighted heavily. As a result points close to the boundary are enforced to have the correct velocity. The overall mirroring equation becomes:

$$\psi_I \mathbf{u}_b + \sum^{j}(\psi_{Ej}\mathbf{u}_{bj}) = \tag{3.5}$$

$$((1/3)\psi_I \mathbf{u}_{iib} \sum^{j}(\psi_{Ej}\mathbf{u}_{Ej}))$$

$$+$$

$$(2/3)(\sum^{i}(\psi_I \phi^{\mathbf{x}_{p^i}} \mathbf{u}_{p^i} + \sum^{j}(\psi_{Ej}\phi^{\mathbf{x}_{p^i}Mj}\mathbf{u}_{p^i})))$$

Where:
$\phi^{\mathbf{x}_{p^i}Mj}$ is the interpolation coefficient to Mj (the mirroring point of $\mathbf{x}_{Ej}$) from $\mathbf{x}_{p^i}$. ψ_I is the weight

of the boundary-interior mirroring equation ψ_{Ej} is the weight of the j^{th} boundary-exterior mirroring equation. The mirroring of the boundary-interior cell is weighted more heavily to reproduce more accurately the original scheme when no boundary-exterior cell centres are close to the boundary. Hence the weights, ψ_I and ψ_{Ej} are given by:

$$\psi_I = \frac{1}{||\mathbf{x}_{iib} - \mathbf{x}_b||} \tag{3.6}$$

$$\psi_{Ej} = \frac{k}{||\mathbf{x}_{Ej} - \mathbf{x}_{bj||}} \tag{3.7}$$

Where:
$\mathbf{x}_{Ej}$ is the position of the j^{th} boundary-exterior neighbour and $\mathbf{x}_{bj}$ is the corresponding closest point on the immersed-boundary. k is a constant of the method, the decreased weighting factor of the external mirroring equations for this work $k = 0.05$. $k = 0$ would correspond to the original method. The effect of the scheme during the transition of cell types (creating a 'fresh' or 'dead' cell) is illustrated in figure 3.7.

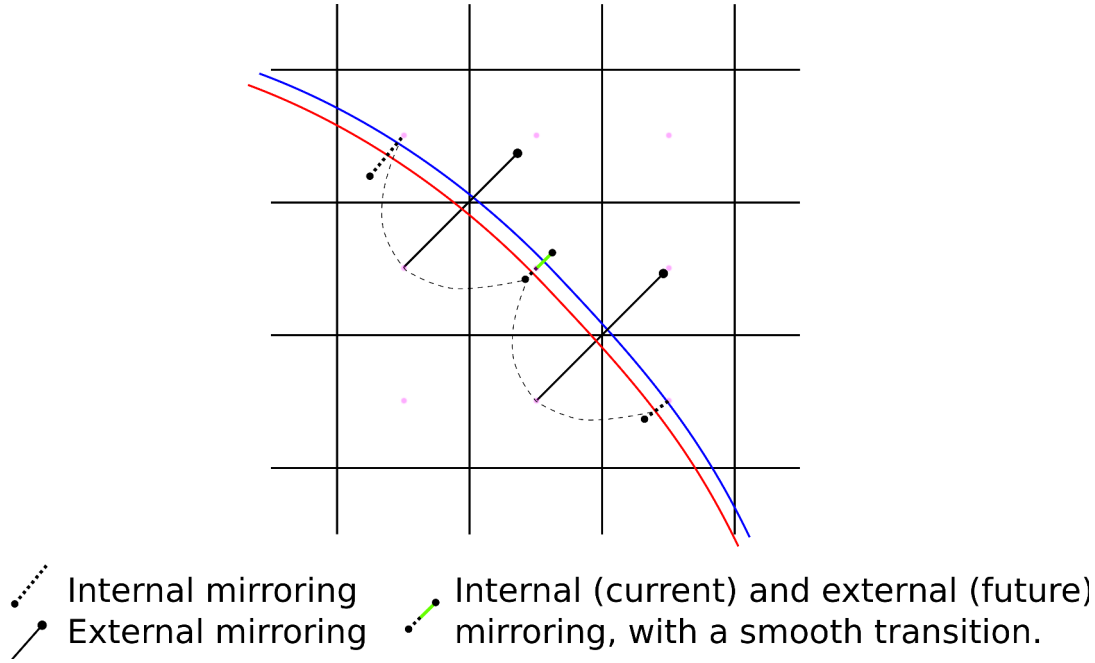


Figure 3.7.: Diagram showing an example, in two dimensions, of a Eulerian mesh with the mirroring points used in the averaged mirroring scheme, for two positions of an immersed-boundary. The new boundary-interior cells are governed by an equation similar to that when they where boundary-exterior cells. Using equation 3.5 with weights given by equation 3.6 the transition is continuous. This scheme is used to ensure smooth transition of the mirroring equations as during 'fresh' and 'dead' cell transitions.

The advantage of this method is that the boundary-exterior cells close to the immersed-boundary have fluid velocities close to that of the immersed-boundary. Hence when a transition is made, the velocity is already close to that required, *i.e.* the immersed-boundary velocity. This results in a smooth transition of stencils when a fluid cell becomes in interior cell. Equivalently the same equation is enforced just prior to the transition as after for a boundary-interior to boundary-exterior transition.

3.2.2.2. Pressure condition

As mentioned previously, a pressure condition is set for the boundary-interior cells, to enforce continuity. In particular it is set such that derivative of pressure in the direction normal to the surface of the immersed-boundary is the same as it is slightly removed from the cell centre, again in the normal direction. This is done fully implicitly to prevent any dependence on the pressure from a previous time-instance.

This is achieved with the use of the mirroring point of the boundary-interior cell in question. A linear coefficient interpolation, from the surrounding cells, is used to determine the gradient and value of pressure at the mirroring point implicitly. The pressure at the cell centre is then set such that pressure varies linearly in the direction normal to the immersed-boundary, as described in equation 3.8.

$$p_c = p_m + [\bigtriangledown p] \cdot (\mathbf{x_c} - \mathbf{x_m}) \tag{3.8}$$

Hence pressure at the boundary-interior cells is at most $\mathcal{O}(\delta x^2)$ inaccurate. More specifically the error is proportional to the second derivative of pressure in the normal direction and the square of the distance to the closest point on the immersed-boundary to the cell. This 'pressure continuity' equation enforces the spacial continuity, *i.e.* smoothness, of the pressure field.

3.2.3. Cut-cell partial mirroring

In a three dimensional system, it is not feasible to solve three mirroring equations and continuity in the same cell. In such a system pressure is not governed and the velocity field is over-specified. This is because the mirroring of the three components of velocity and the continuity equation (eqn 1.1b) are equations related solely to velocity. Figure 3.8 shows an example in which the out of plane component of velocity of a face is largely determined by both the mirroring equation and the continuity equation. Indeed if the face and the immersed-boundary where to coincide, the equations would be identical.

This is not necessarily and over-specification in the equation system in the strict sense as it is possible to incorporate a pressure condition into the continuity equation, and indeed is common place (via use of the Rhie and Chow pressure equation). However in this case the pressure terms are numerical correction terms, unsuitable as the only governor of pressure. As a result, as well as being unphysical, there is no guarantee of continuity in pressure.

The necessity for a consistent continuity equation is two-fold. Firstly it allows effective control of mass loss errors. Secondly it allows control over the problems that occur as the result of the transition of cell types. Figure 3.9 shows such a transition and the potential problem that the flow field prior to the transition may not satisfy the continuity equation of system afterwards. To see this consider the flux over the marked 'green' faces (these fluxes will be determined by interpolation from the cell centre to its neighbours). If the cell centre is considered boundary-interior, the cell centre values are determined by the velocity of the immersed-boundary close to the transition point. If the cell centre is considered boundary-exterior, the fluxes will be determined by flux over the remaining faces, via continuity. However the value of these fluxes is determined by all the cells neighbours. There is no guarantee the two values of flux, that result from the ywo cases, will be the same. Enforcing that they are the same would be highly complex due to the large stencil involved. The change that results leads to pressure oscillations, as discussed earlier.

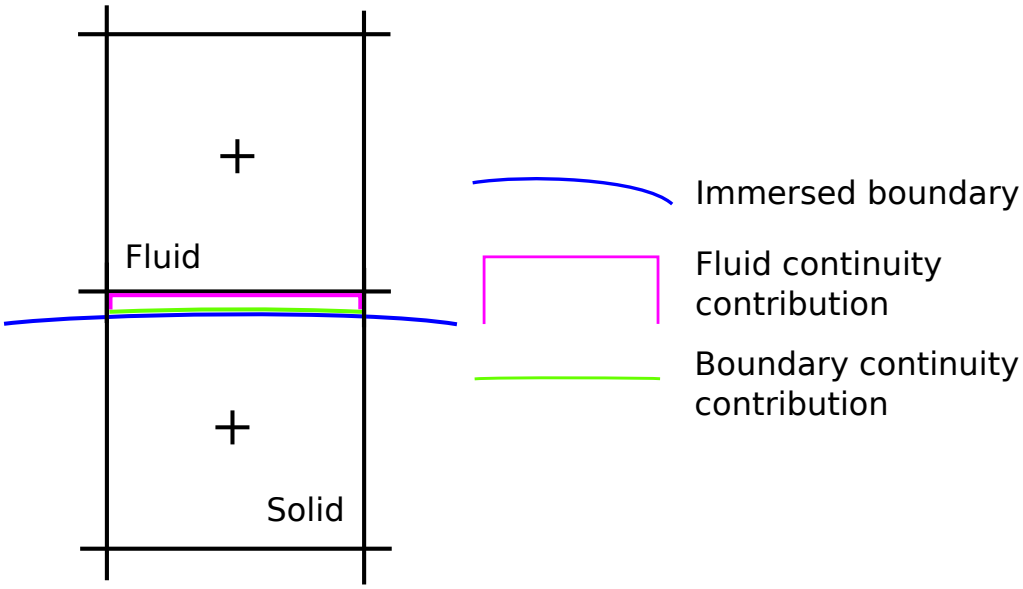


Figure 3.8.: Example of a boundary-interior cell with a redundant mirroring equation. In this case both the continuity equation and the mirroring equation approximately reduce to determining the velocity at the face *i.e.* $u_{iib} = u_s + u_n$. One of these equations is redundant, and the inclusion of both effectively over-specifies the equation system.

Methods have been proposed where the continuity equation is calculated added as a correction term to another (adjacent) cell. This removes the problem of over-specification as two continuity equations are merged into one. However, the problem of a discontinuous set of equations remain as the introduction and replacing of these 'links' occur instantaneously. In addition to this the merged cells create problems for estimating velocity gradients, necessary for force calculations. This is due to the small jumps in velocity field, close to the immersed-boundary, that result.

In this section a novel solution is presented where this over-specification is resolved. The principle of the method is reformulating the mirroring principle by 'dropping' one of the mirroring equations. In particular the out of plane component of velocity in the mirroring equation.

The advantage of such a method it that it provides a consistent continuity equation of all cells (except complete-interior which are removed from the method). This prevents mass loss errors locally and hence globally. Additionally the spurious oscillations that result from the abrupt changes to the continuity equations of exterior cells are diminished, as there are no longer such changes.

Hence for cut boundary-interior cells: two mirroring equations, continuity, and the 'pressure continuity' equation, described in section 3.2.2.2 are solved. The imposed continuity equation reconstructs the interior flow in the out of plane direction, while the new (remaining) mirroring equations reconstruct the remaining components of the flow. The result of which is a fluid velocity field that interpolates to the correct value on the boundary, a smooth pressure field and no abrupt changes to the equation systems as cells transition. The scheme is described schematically in figure 3.10.

It is possible that some the boundary-interior cells are completely inside the immersed-boundary, and hence a non-governable by a cut-cell continuity equation. Depending on the implementation, velocities at these points may be required, for example for interpolation in other aspects of the methodology. In such a case the complete mirroring method (section 3.2.2), of mirroring all three components of velocity, is resumed. Care is taken such that this transition occurs smoothly.

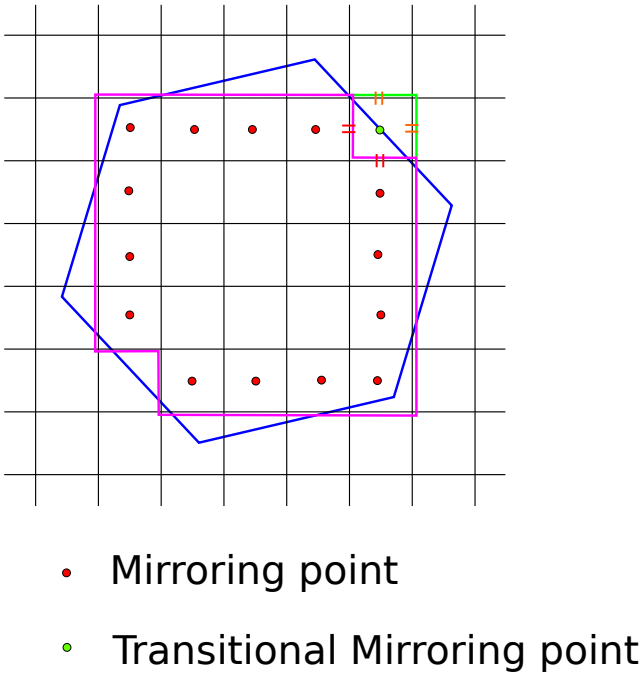


Figure 3.9.: Schematic of the limiting case of a cell centre lying on an immersed-boundary, with different continuity equations used depending on which 'side' is taken. If no continuity equation is solved for interior cells the following situation can occur:
If the transitional point is taken as outside, and hence governed by a continuity equation, there is no guarantee this continuity equation will be satisfied by the velocity that field that would occur if the point had been considered inside and mirroring used. Hence during a transition, a discontinuity is likely to occur.

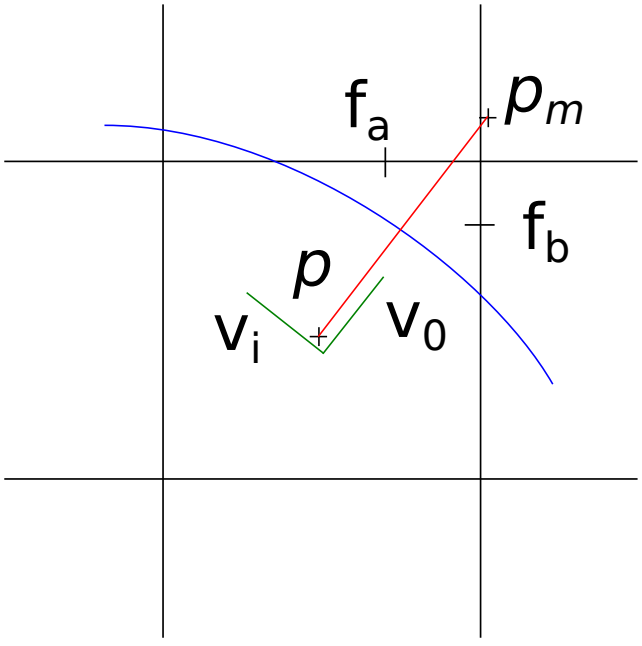


Figure 3.10.: A two dimensional schematic of the cut-cell partial mirroring method for a cell cut by an immersed-boundary (blue line).
Key:
Cell centre p, and mirroring point p_m. Out-of-plane component of velocity v_0, in plane components of velocity v_i.
f_a and f_b are the fluxes across the 'active' faces.

3.2.3.1. Partial mirroring

As stated in the introduction, in the developed method, mirroring, used to enforce the no slip condition over the immersed-boundary, is only used for the in-plane components of the velocity. Here the surface-plane is taken as the plane normal to the vector connecting the mirror point and the cell centre. The mirroring scheme is similar to that which is described in section 2.3. The mirroring can be adapted in a similar manner to that described in section 3.2.2.1, the details of which are given in section 3.2.3.3.

The partial mirroring is achieved using the Gram-Schmidt process to create vectors $\mathbf{s}_i$ and $\mathbf{n}$, where $\mathbf{s}_i$ are two in-plane vectors and $\mathbf{n}$ is the surface normal an (out-of-plane vector). The mirroring equations are then set so that the in-plane components of the velocity, determined by a dot product with $\mathbf{s}_i$, are mirrored. Depending on whether the mirroring equation being employed is an 'averaged' mirroring (*i.e.* subject to the modifications given in section 3.2.2.1, the exact form of the equation will vary. However, given a set of immersed-boundary velocities and corresponding interpolation stencil and coefficients from the complete mirroring method (equation 3.9) the modified equation system is given by equation 3.10. Complete mirroring:

$$u^b = \sum_{s\ \epsilon\ \text{stencil}} [\phi^s u^s] \tag{3.9a}$$

$$v^b = \sum_{s\ \epsilon\ \text{stencil}} [\phi^s v^s] \tag{3.9b}$$

$$w^b = \sum_{s\ \epsilon\ \text{stencil}} [\phi^s w^s] \tag{3.9c}$$

Partial mirroring:

$$\mathbf{u}^b \cdot \mathbf{s}_1 = \sum_{s\ \epsilon\ \text{stencil}} [\phi^s \mathbf{u}^s \cdot \mathbf{s}_1] \tag{3.10a}$$

$$\mathbf{u}^b \cdot \mathbf{s}_2 = \sum_{s\ \epsilon\ \text{stencil}} [\phi^s \mathbf{u}^s \cdot \mathbf{s}_2] \tag{3.10b}$$

3.2.3.2. Continuity

Treatment of individual cells A cut-cell variation of the continuity equation is solved for all boundary-interior and boundary-exterior cells, as defined in section 3.2.1. In the case of boundary-interior cells, this is in addition to the two mirroring equations and a pressure equation. In the case of boundary-exterior cells, the cut-cell continuity equation replaces the original continuity equation.

The premise is to replace the original control volume with only the exterior (or fluid) section. The surface of this new control volume is bounded by two types of surface:

- The immersed-boundary.
- Fluid sections of the original control volume surfaces. These can be further sub-divided into:
  - Cut faces
  - Un-cut faces

Here cut and un-cut faces are with respect to the presence of an immersed-boundary in particular if any part of a face lies inside an immersed-boundary it is deemed to be cut. Figure 3.11 gives an example of the classification cut and un-cut faces. In some cases a face may lie completely inside an immersed-boundary, in such a case the face does not make a contribution to the continuity equation. This case is illustrated in figure 3.12.

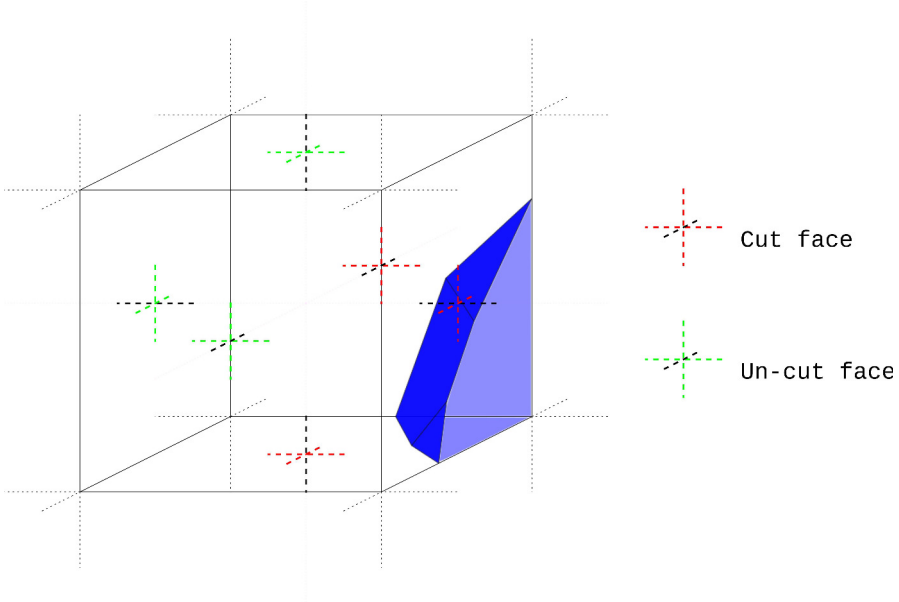


Figure 3.11.: Example of a Eulerian mesh cell intersected by an immersed-boundary. Faces are classified according to their type; cut or un-cut.

The continuity equation then becomes a sum over the contributions from the immersed-boundary and the fluid directly. Hence the new continuity equation for a cut-cell becomes:

$$\begin{aligned} &\sum_{\text{Intersecting IB elements}} [\text{boundary area} \times (\text{boundary normal} \cdot \text{boundary velocity})] \\ &\qquad + \\ &\sum_{\text{Fluid section of faces}} [\text{face area} \times (\text{face normal} \cdot \text{face velocity})] \\ &\qquad = 0 \end{aligned} \tag{3.11}$$

The separation of these contributions is depicted in figure 3.13 in three dimensions. It is also shown in two dimensions, in conjunction with the mirroring equations for velocity, in figure 3.14.

In the case a boundary-interior cell is completely inside, the immersed-boundary the continuity equation, as described, is no longer valid. This is because all of the faces are removed, and the continuity equation would reduce to a tautology (*i.e.* $0 = 0$). Such a cell is still considered boundary-interior, as it is adjacent to an exterior cell, and cannot be removed from the system. Hence it is necessary to set values for the velocity and pressure at such cells.

This is done using the following equations:

Pressure The same pressure condition as other boundary-interior cells.

Velocity Mirroring equations are used for all three components of velocity.

i

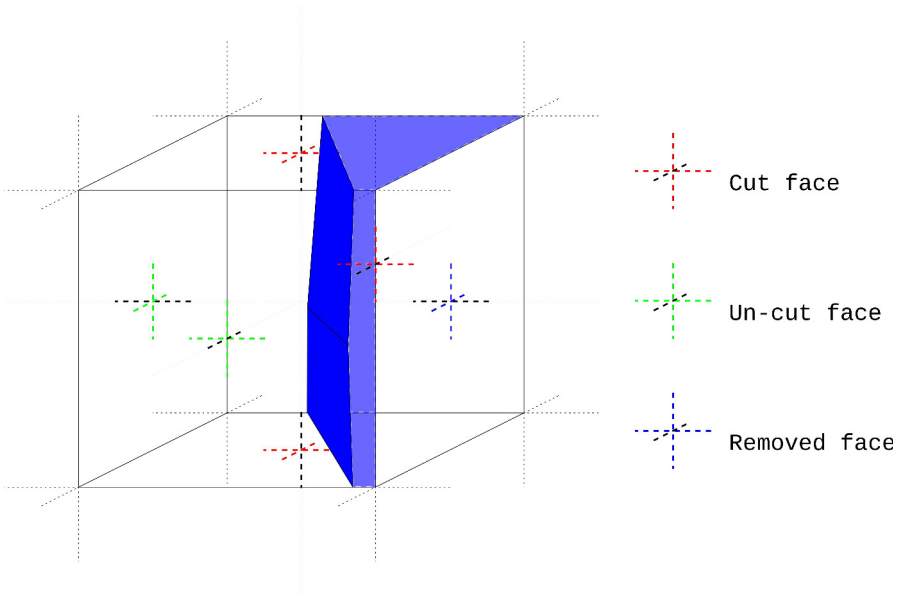


Figure 3.12.: Example of a Eulerian mesh cell intersected significantly by an immersed-boundary. In the example one of the faces (the right) is completely inside the immersed-boundary and hence will be removed from the continuity equation. Faces are classified according to their type; cut, un-cut or removed.

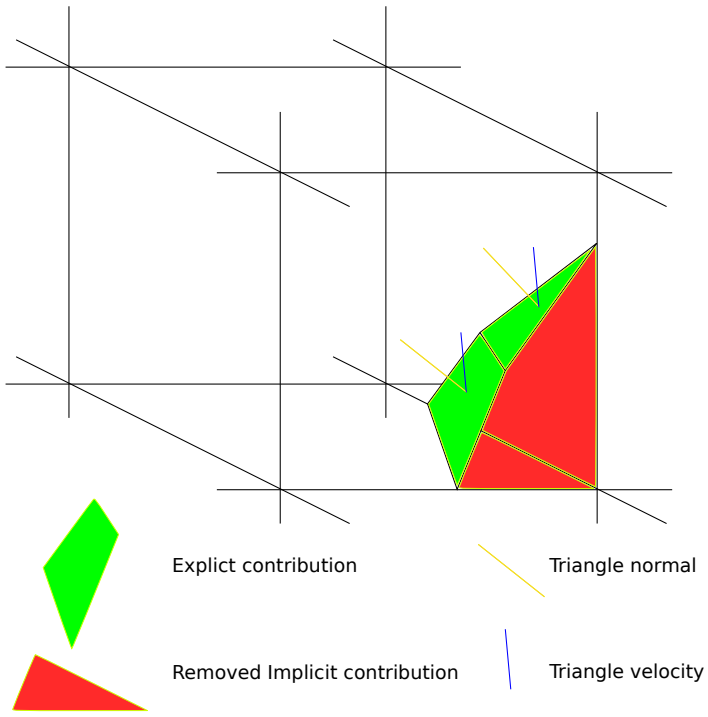


Figure 3.13.: Example of a cut-cell, a Eulerian mesh cell intersected by an immersed-boundary. Each cell face is sub-divided into sections that are inside and outside of the immersed-boundary. The surface is re-closed with the addition of sections of the immersed-boundary. This is used to split the surface used in the continuity equation of a cut-cell into implicit fluid contributions and explicit immersed-boundary contributions. The explicit immersed-boundary contributions are given by the dot-product of surface normal and immersed-boundary velocity.

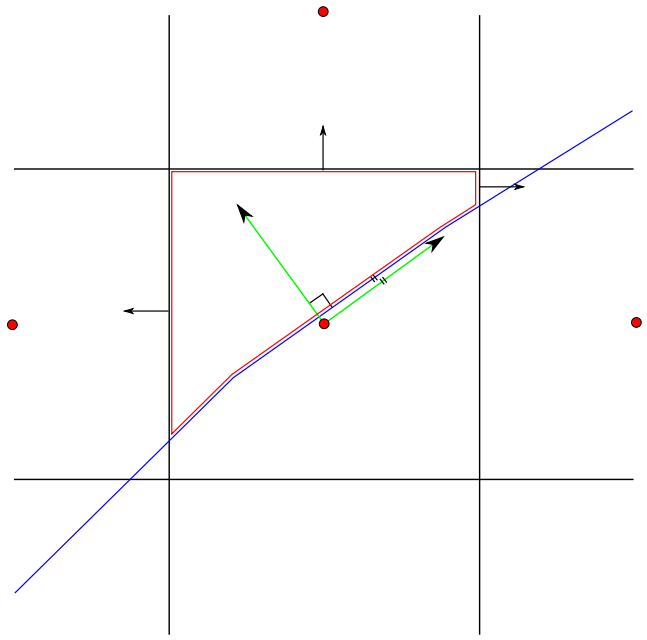


Figure 3.14.: Diagram of the equation system governing boundary-interior cell. The red boundary of the cell shows the boundary over which continuity is solved and the green lines show the splitting of the components of velocity used for mirroring.

Global continuity In some ghost-cell immersed-boundary methods there exists cell faces for which there is a continuity equation governing one side of the face but does not govern the other. This can be problematic as the flux across this face is represent a global contribution to the continuity equation. It is generally difficult to ensure these contributions sum to zero in a meaningful and physical way. Figure 3.15 shows an example of how the asymmetric cell-faces can give rise to mass loss error. The problem is worse for a moving immersed-boundary as the sum of continuity contributions, and their governing equations, is likely to fluctuate. The fluctuations in the continuity contributions from the immersed-boundary are to be expected for the same reason as they do not sum to zero, they are difficult to control directly. The fluctuations are additionally problematic as they cause large oscillations in pressure field near the immersed-boundary. This is discussed in more detail later in the rest of this section.

A benefit of the approach developed in this work, is that global continuity is consistent. This is because the closed surface bounding the fluid to form the immersed-boundary has a strict continuity equation that is consistent with the continuity equations of the rest of the fluid domain. There are three types of contribution to the continuity equation from the boundary of any cell, given by the type of the boundary as specified earlier (figure 3.11), namely:

- The immersed-boundary
- Cut faces
- Un-cut faces

In the method developed, none of either cut or un-cut faces separate a cell which is not governed by a continuity equation (*i.e.* complete-interior cells and boundary-interior cells that are completely inside the immersed-boundary) and a cell that is governed by a continuity equation (all remaining cells). This allows the method to ensure that the continuity equation contributions, from cell faces, have an equal an opposite contribution to the other side of the face. Hence global continuity is ensured. The requirement of being able ensure global continuity in this way is the primary motivation for enforcing a continuity equation for all cells that are partly outside the immersed-boundary.
There are contributions from the immersed-boundary itself, however these are analytic and therefore do not affect global continuity significantly. In particular, as the continuity contributions from the boundary are analytic, the error associated with mass-conservation should be no greater than if the immersed boundary where not present. Not only is the improved mass-loss error beneficial directly it also prevents the formation of pressure oscillations associated with the fluctuation in continuity equation. Figure 3.16 shows a schematic of the developed continuity equation and its effect on the conservation of mass. An additional benefit of this approach is that if the immersed-boundary were to deform in such a way that internal volume varied, the method would automatically accommodate the volume change appropriately. Other methods were global continuity is conserved, such as solving a continuity equation for all the cells including those inside an immersed-boundary, do not accommodate such changes as easily.

Motivation Without the a continuity equation governing boundary-interior cells a smooth transition between cell types is difficult. This is because: as an immersed-boundary moves, and a cell transitions from boundary-interior to boundary-exterior, it must instantaneously satisfy the continuity equation. Additionally there was no continuity equation for this cell at the previous

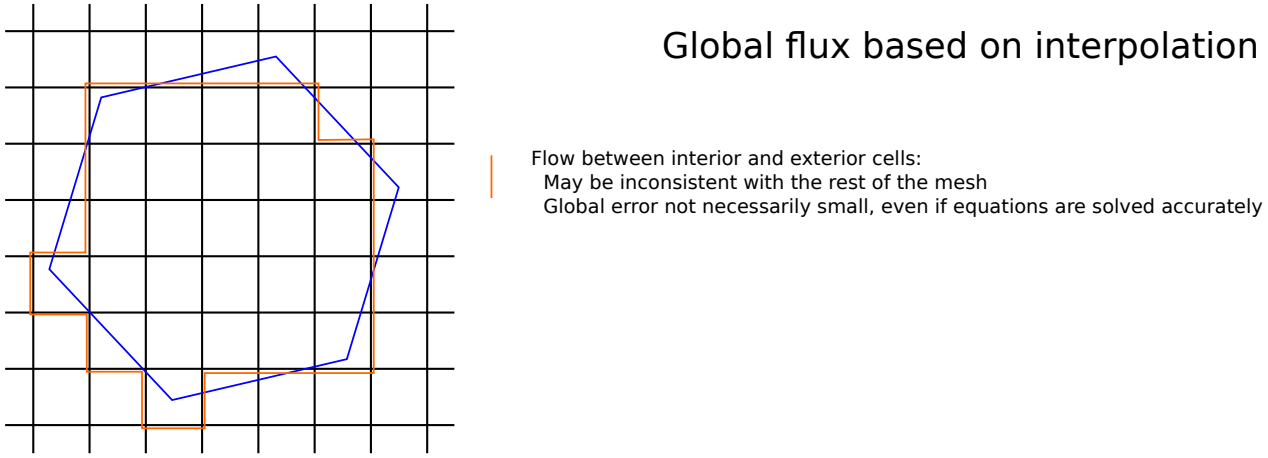


Figure 3.15.: Diagram showing the possible inconsistency of the continuity equation resulting form faces for which a continuity equation contribution is made on one side, but not accounted for on the other. The result of these cell faces is that the flux through the immersed-boundary is not guaranteed to sum to zero in a way that is consistent with the exterior continuity equations. Hence there may be mass loss errors locally and globally.

time instance, as it was boundary-interior. Hence it would be unfeasible to ensure that the old time-instance solution would converge to the current time instance solution as the time-step reduces to zero, as a different equation set was being solved. This type of discontinuity is particularly problematic for a number of reasons. The principle problem is that the correction to the velocity field, to satisfy the new continuity equation, occurs over a single time-step. As the correction is not determined by the size of the timestep (δt), if δt is small, the acceleration may be large. As a result the force necessary to correct the finite, fixed, velocity change is *inversely* proportional to the time-step. This can make the prediction of the motion of the body difficult, and result in oscillatory behaviour. Additional problems associated with the discontinuity of the continuity equations include:

- Un-physical velocity fluctuations around the immersed-boundary.
- Potential for global mass source/sink, varying over time, may have long range effects.

In general the discrepancy in the solutions is determined by the change in the geometrical description of the problem, hence it is governed by mesh spacing. It is therefore possible for the jump to be made small by refining the mesh, however, not only does this increase the computational cost of the method, it also puts constraints on the relative spacial and temporal length-scales. In particular there is a dependence on the CFL number of the method. This phenomenon has been noted by a number of papers.

When using the developed method, the transition of cells from different types no longer causesa discontinuous transition in the continuity equation of the cell. The difference in the coefficients of the equations tend to zero as immersed-boundary movement becomes small. This has the consequence that the old time-instance flow field is a close solution to current time-instance equations. Hence no finite change is required to be made over a potentially arbitrarily small time-step.

Advantages of the partial mirroring The advantage of the partial mirroring method is it allows the solution of a continuity equation for all cells that contribute to the overall solution.

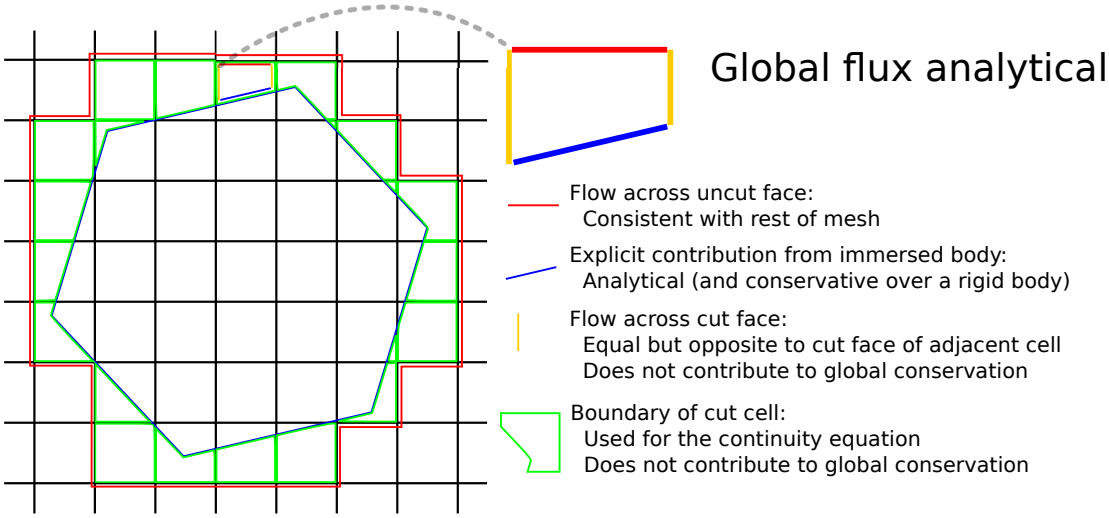


Figure 3.16.: Diagram showing the avoidance of the inconsistency of the continuity equation locally and over the entire domain seen in figure 3.15. This is achieved by restricting all contributions to the continuity equation from cell boundaries to be either analytical (*i.e.* from the immersed-boundary, blue) or matched exactly to another cell, as is the case for un-cut, red, of cut, yellow, faces. In this way mass loss errors are kept small locally and globally, in a way that is consistent with the continuity equation for the un-cut cells.

Advantages of the cut-cell variation of the continuity equation The advantages of the cut-cell variation of the continuity equation are:

- The method does not depend on empirical predictions of flow inside the body.
- It provides a physical way of removing any contribution of complete-interior cells to any continuity equation.
- It enables a smooth transition of the continuity equation as the cells change type, for all allowable transitions. Hence it removes the 'jump' in the continuity equation when a cell transitions between cell types. This is the principle advantage, and is discussed in more detail later in this section.

Implementation to ensure smoothness An number of possible transitions of cell and face types are possible. The types of transitions between cells are assumed to be limited to change by 'one type', *i.e* a cell can only become a type is is possible for it to be currently adjacent to. Namely:

- complete-interior to boundary-interior and vice-versa.
- boundary-interior to boundary-exterior and vice-versa.
- boundary-exterior to complete-exterior and vice-versa.

This assumption is valid if the CFL number is less than unit as the immersed-boundary cannot travel more than the separation of two adjacent cells. As this is normally the case for other reasons, this assumption is not particularly limiting however it is important that it be observed in order for the developed method to function correctly.

Each transition presents the possibility of causing a discontinuity to the discretisation of the Navier-Stokes equations. The method developed address these potential discontinuities.

Transition from boundary-interior to boundary-exterior The continuity equation for boundary-interior and boundary-exterior cells is the same for a fixed geometry. Hence there is no discontinuity as a direct result of this type of transition. This transition may result in a differing geometry or cause a transition of another cell from boundary-interior to complete-interior or vice versa. The smoothness of these transitions is considered separately.

Transition between partially, completely and un-cut faces In order to accurately estimate the flow through a cut cell-face, an interpolation of the flow to the 'live' portion of the face, the section that lies outside of any immersed-boundary, is necessary. This is as a result of the centre of the fluid part of the face no longer being at the original face centre. It is desirable to maintain the simplicity and accuracy of the un-modified Navier-Stokes equation discretisation while maintaining second order accuracy. An interpolation scheme is developed in this section and adopted to evaluate the flow-field at the fluid boundaries of a control volume, in order to achieve a smooth transition of face types. For numerical reasons the interpolation depends on whether the cell face is intersected, or cut, by the immersed-boundary as follows:

Un-cut face Face values are interpolated from the centres of the two cells the face divides.

Cut face A larger stencil is used capture first order variation. This allows second order accuracy of the interpolation to the live face centre to be maintained.

The motivation for the split in treatment is to maintain consistency with complete-exterior cells without a transitional discontinuity. It is necessary for complete-exterior cells to have face values interpolated from the centres of the two cells the face divides for numerical reasons.
There are a number of ways to interpolate from a large stencil to a cut face however care must be taken in choosing the interpolation method as there is the potential for a discontinuity to arise from various necessary transitions. In particular, the case that an un-cut cell-face becomes slightly intersected by the immersed-boundary, the face will transition between the method used for cut-faces and the method used for un-cut faces. The same can happen in reverse, as an immersed-boundary leaves a cut face.
When the large stencil interpolation scheme is used to calculate the interpolation coefficients to the a face, when previously the small stencil interpolation scheme had been used, or *vice-versa* if the interpolation schemes do not provide the same or similar interpolation coefficients in both cases, there is a discontinuity. If the large stencil interpolation scheme provides continuous interpolation coefficients as the immersed-boundary moves, continuity overall is guaranteed if, and only if, the transitional case gives rise to the same value for all interpolation points, for either interpolation scheme. This condition can be seen as requiring the larger stencil interpolation to reduce to the small stencil interpolation for an un-cut face. If this condition was not satisfied, the resultant discontinuity would have the same adverse consequences as the discontinuities in the complete mirroring method. There are a number of ways in which this problem can be addressed:

1. Interpolate to all faces with the method used for the cut face but keep the 'live' face centre as the centre of the entire face, for un-cut cells.

2. Interpolate to the cut face as though it had its 'live' face centre as the centre of the entire face.

3. Choosing an interpolation scheme that is guaranteed to match in the transitional case.

However these methods have the following disadvantages:

1. Requires a modification to all fluid cells. Additionally the increased stencil is likely to decrease the diagonal dominance of the solution matrix.
2. Does not maintain the second order accuracy of the interpolation method.
3. Despite looking promising, this is difficult to achieve for an arbitrary shaped mesh and limits the interpolation methods that can be used.

An alternative scheme in which the flow through the 'live' face and the remaining 'dead' face are both considered is developed. The benefit of considering these additional fluxes is that they can be used to ensure a smooth progression of used flux as the cell face in question transitions, from cut to un-cut, and, cut to eclipsed. This is done by ensuring:

- As a cell is cut, the flux is calculated as the original less the flux through the 'dead' section.
- As a cell is about to become eclipsed, the flux is calculated as the flux through the 'live' section.
- A smooth combination of the two is used for a cut cell.

Figure 3.17 depicts the scheme.

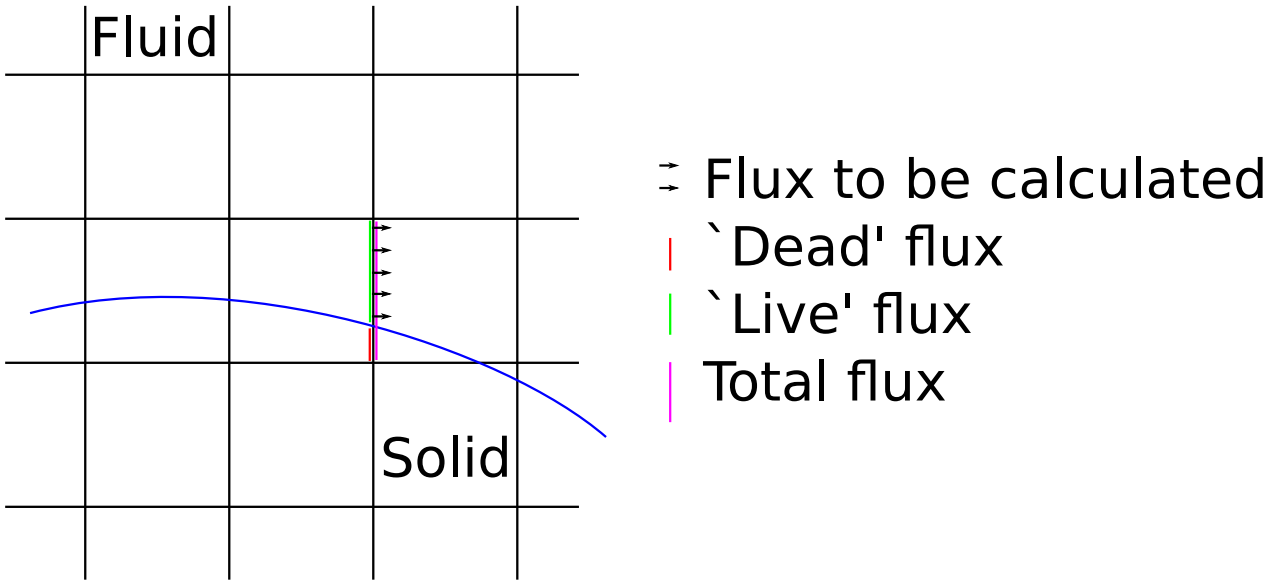


Figure 3.17.: Diagram of the calculation of a continuous cut-cell implicit fluid contribution to the continuity equation. The diagram shows how the flux through the 'live' portion of the face can be calculated directly or from the 'un-cut' flux minus the 'dead' flux. As a result, a weighted combination can be taken. The weights can be varied such that the more accurate value can be used while maintaining smoothness.

In this scheme the flow through the 'live' face is both second order and transitions continuously as the immersed-boundary intersects and leaves the face. To achieve this the interpolated flow is considered as a combination of:

- The a direct prediction (F_l).
- The difference between:

– An estimation of what the flow through the total (un-cut) face would be if the immersed-boundary where not present (F_t).

– An estimation of what the flow through the 'dead', or interior, section of the face would be (F_d).

The difference is a valid prediction for the flow through the 'live' face as the hypothetical flow through the entire face is the sum of the flow through the 'live' and 'dead' parts. *I.e.* $F_t = F_l + F_d$ hence $F_l = F_t - F_d$ and $\mathbf{u}_l = (\mathbf{u}_t\, a_t - \mathbf{u}_d\, a_d)/a_l$. Therefore the total flux (F) may be considered as:

$$F = \left[w^{\text{Live}}\,(F_l)\right] + \left[w^{\text{Dead}}\,(F_t - F_d)\right] \tag{3.12}$$

provided $w^{\text{Live}} + w^{\text{Dead}} = 1$. The method is second order accurate as each of the three interpolations are.

To ensure the continuity during the transitions the following criteria are sufficient:

F_t must be calculated in the same way as it is for un-cut cells.

$w^{\textsf{Live}}$ must:

- Vary continuously as the immersed-boundary moves.
- Tend to 0 as the area of the 'dead' face tends to zero.
  I.e. $0 = \lim_{A^{\text{Dead}} \to 0} \left(w^{\text{Live}}\right)$.

$w^{\textsf{Dead}}$ must:

- Vary continuously as the immersed-boundary moves.
- Tend to 0 as the area of the 'live' face tends to zero.
  I.e. $0 = \lim_{A^{\text{Live}} \to 0} \left(w^{\text{Dead}}\right)$.

To achieve this, in the developed method, F is taken as:

$$\begin{aligned} F = \quad & \left[w^{\text{Live}}\,(\mathbf{u}^{\text{Live}} \cdot \mathbf{a}^{\text{Live}})\right] \\ + \quad & \left[w^{\text{Dead}}\,((\mathbf{u}^{\text{Face}} \cdot \mathbf{a}^{\text{Face}}) - (\mathbf{u}^{\text{Dead}} \cdot \mathbf{a}^{\text{Dead}}))\right] \end{aligned}$$

Here the velocity $\mathbf{u}$ is the interpolated velocity from the surrounding cells, to the centre of area of the section of the face in question. *I.e.* $\mathbf{u}_i^{\text{Live}} = \sum \phi_j^{\text{Live}}\, \mathbf{u}_i^j$. This method ensures continuity at the limits in 'live' face contribution with 'live' area as:

Nearly un-cut $[\lim_{A^{\text{Dead}} \to 0}]$

In the case the 'dead' area tends to zero:
$F \to \left((\phi_i^{\text{Face}} \cdot \mathbf{a}^{\text{Face}}) - (\phi_i^{\text{Dead}} \cdot \mathbf{a}^{\text{Dead}})\right)$
$= (\phi_i^{\text{Face}} \cdot \mathbf{a}^{\text{Face}})$
This is same as the un-cut case.

Nearly fully cut $[\lim_{A^{\text{Live}} \to 0}]$

In the case the 'live' area tends to zero:
$F \to \phi_i^{\text{Live}} \cdot \mathbf{a}^{\text{Live}}$
$= 0$ (as $\mathbf{a} \to 0$, as $A \to 0$)
This is same as the fully cut case.

Smoothness as the interpolation stencil points change there type The larger interpolation stencil used for cut faces includes cells other the cells either side of the face. Further more it may be that one of the cells in the stencil is a complete-interior cell, hence cannot be used for the interpolation of the velocity, as these cells do not contain reliable velocity data. Hence complete-interior cells are removed from the stencil. As the immersed-boundary moves, one of these additional cells, who's velocity is used to determine the flux though a 'live' cut-face, may change type. This can change at any time, *i.e.* without the cells either side of the face changing their type. However if one of the cells were to transition to become a complete-interior cell and be removed suddenly, the removal from the stencil could cause a discontinuity, as it necessarily changes the stencil abruptly. This is rectified by using a weighted interpolation from the points in the stencil in such a way as interior cells are gradually faded out of the stencil. In this work, a weighted least-squares linear interpolation is used, however the principle of weighting the contributions should be equally applicable to other methods. The weight for the least squares fit is the product of both an inverse distance weight D and volume fraction weight F. *I.e.* the interpolation weight w is given by $w = FD$ where:

$$D = \frac{1}{d}$$

Where d is the distance from the stencil point in question to the interpolation point.

$$F = \frac{v}{v_{\text{original}}}$$

Where v is the volume of the stencil point cell that is fluid and v_{original} is the total cell volume.

Completely interior cells can be excluded explicitly however, unless cells close to the surface are removed prior to the transition, the removal may lead to a discontinuity. If the developed scheme is used, before this transition occurs the cell must have been removed from the stencil, as v must be 0, hence the transition is smooth.

It is possible a cell in the stencil transitions from boundary-exterior to boundary-interior. In this case the component of velocity normal to the current cell's immersed-boundary may change, and hence alter the calculated flux. However it is not possible to remove the boundary-interior cells from the interpolation as they may be required, as with out these cells the stencil points may not surround the interpolation point. This is one of the reasons the 'average' mirroring is important, as it ensures no discontinuities occur as a result of this type of transition.

Immersed-boundary motion within a cut face Continuity as the immersed-boundary moves within the cut-face is ensured if:

- ϕ_i^{Live} and ϕ_i^{Dead} vary continuously with immersed-boundary position.
- w^{Live} (and hence w^{Dead}) varies continuously with immersed-boundary position.

Implementation The overall contribution to the continuity equation from a cut face, in its implicit form is given by:

$$\sum_{i}^{\text{stencil}} \left[\sum_{j=1}^{3} \left(c_j^i \, \mathbf{u}_j^i \right) \right] \tag{3.13}$$

Where:

$$c_j^i = w^{\text{Live}} \left(\phi_i^{\text{Live}} \, \mathbf{a}^{\text{Live}_j}\right) + w^{\text{Dead}} \left((\phi_i^{\text{Face}} \, \mathbf{a}_j^{\text{Face}}) - (\phi_i^{\text{Dead}} \, \mathbf{a}_j^{\text{Dead}}) \right) \tag{3.14}$$

For the current work:

$$\begin{aligned} w^{\text{Live}} &= \frac{A^{\text{Dead}}}{A^{\text{Face}}} \\ w^{\text{Dead}} &= \frac{A^{\text{Live}}}{A^{\text{Face}}} \end{aligned}$$

Interpolation coefficients (ϕ_i^{Dead} and ϕ_i^{Live}) are the weighted least squares linear interpolation coefficients. Figure 3.18 shows the interpolation scheme described.

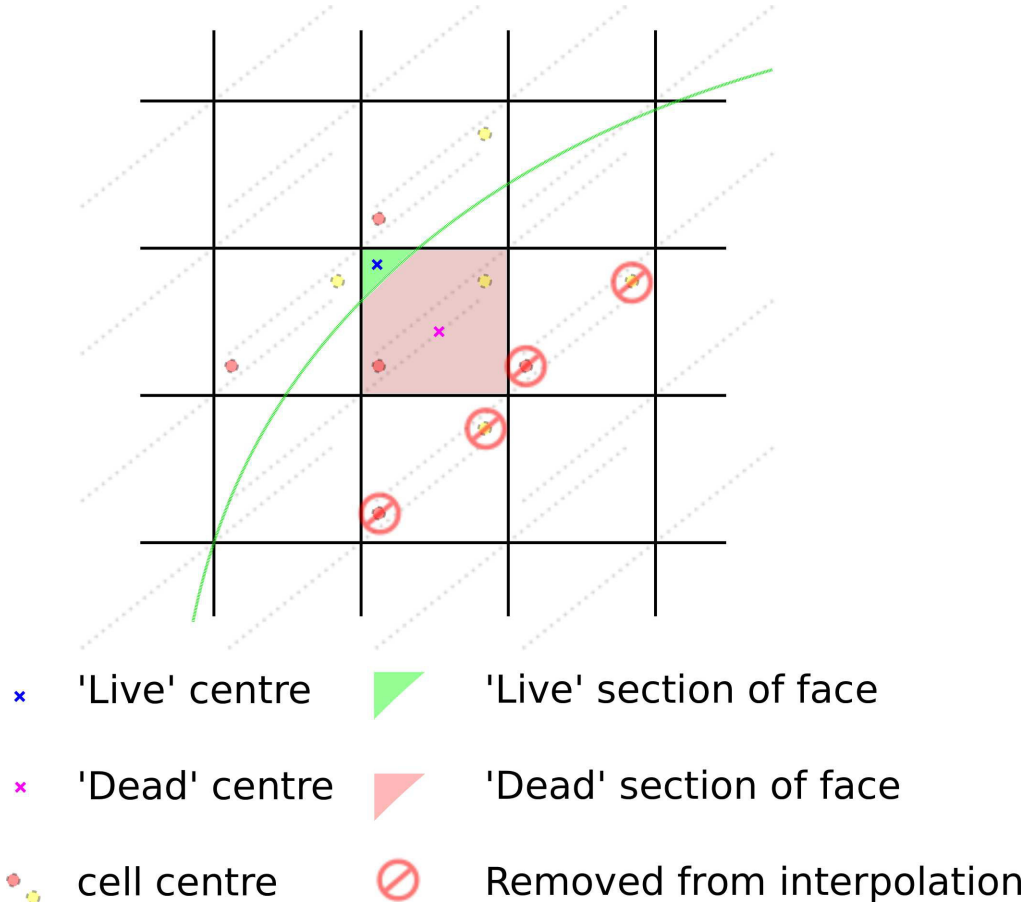


Figure 3.18.: Diagram showing a plane of faces in a three dimensional Eulerian mesh, cut by an immersed-boundary. The central cell face is split into 'live' and 'dead' sections, each with their own centre of area. The diagram shows the stencil used for the interpolation to the 'live' and 'dead' face centres, to calculate the cut-cell implicit fluid flux for the continuity equation.

Transition between partially and completely cut cells During the transition between the two cases the equations are consistent. This is because in the transitional case, the components of velocity in the non-normal direction have very little contribution to the continuity equation. Hence the out of plane component, as set by the continuity is equation, will converge to being the same as that of mirroring. This can be seen from the following observation of the transitional case:

- Both for implicit and explicit coefficients of the out-of-plane velocity in the continuity equation are proportional to the out-of-plane component of area.

- The contribution from the in-plane components of velocity are proportional to the area of the control volume, projected in their respective direction.
- However as the implicit and explicit contribution surfaces are closed in the in-plane direction, all of the first order contributions cancel.
- Hence the total contribution is proportional the projected area and depth, hence volume.
- The ratio of these contributions is therefore determined by the ratio of cut-cell volume to immersed-boundary area. This ratio is at most the maximum distance separating the cell faces and the boundary, in the boundary normal direction. This must be small, relative to the cell length-scale, by assumption as this distance is the minimum distance the boundary must travel to remove all exterior faces.

However in a particular implementation the discretisation of the two quantities may differ, *i.e.* the mirroring equation stencil and the continuity equation stencil may differ. In this case, the transition may be smoothed by solving a weighted combination of out-of-plane mirroring and the continuity equation. Such that:

- The weight of the classic mirroring equation is small or zero when the dead volume is small.
- The weight of the cut-cell mirroring equation is small when the live volume is small.
- The weight varies smoothly with live (or dead) volume.

It is advisable to limit the effect of the cut-cell mirroring to the last small section of the cell, as this recreates the cut-cell method better. It is also worth noting the eclipsed boundary interior cells do not effect the overall continuity equations directly so the effect is not as pronounced. However smoothness is still a concern, as the momentum equations are still effected.

Transition between complete-interior and completely cut boundary-interior cells No continuity equation is solved for either of these cells. Only cells that have portion outside of an immersed-boundary contribute, hence directly this transition will not cause a discontinuity. Depending on the interpolation scheme used, it is possible that such a cell is used as an interpolation point to determine the flux through a cut face nearby. If this is the case then such a transition may result in a discontinuity. Indeed it is certainly possible that the cell in question has a velocity fluctuation during the transition. Such a problem is avoided in this implementation by using only a compact stencil for interpolation to cut faces. In particular it is not possible a complete-interior cell used in the interpolation to a cut face.
However if for some reason an interpolation scheme was used such that complete interior cells where used, the problem would be limited as global continuity is still preserved. This phenomenon is the same as described at the beginning of section 3.2.3.2. Some overall discontinuity is still to be expected. Avoiding this is the main reason for choosing a compact stencil for interpolating to a cut face despite the potential benefits of higher order interpolation, possible with a larger stencil.

3.2.3.3. Adaptive mirroring

The mirroring scheme used for the partial mirroring method can be modified by the developed improvements, given in section 3.2.2.1. However the advantages of this modification are reduced as

the main source of reduced oscillation in the complete mirroring method was related to continuity. In particular the increased accuracy of the solution to the continuity equation of the previous time instance of a transitional cell leads to a smaller necessary jump. In the presented method, there is no longer a discontinuity in the continuity equation to improve. However there is still some advantage in relation to the smoothness of the momentum equations. The smoothness of the momentum equations is less important as the change can be accounted for locally and does not require a finite step-change in velocity over a time-step.

3.2.3.4. Parity based incrementation

If a given particle's CFL number is small, the temporal resolution of the position of the particle is excessive. In such a case, the enforcement of the immersed-boundary is not sensitive to the exact location of the interface as the position error is much smaller than the interpolation stencil length scale. Hence updating the particles position only every other time-step, or third etc, retains accurate enforcement of the immersed-boundary condition. The an example, demonstrating this justification, is shown in figure 3.19.

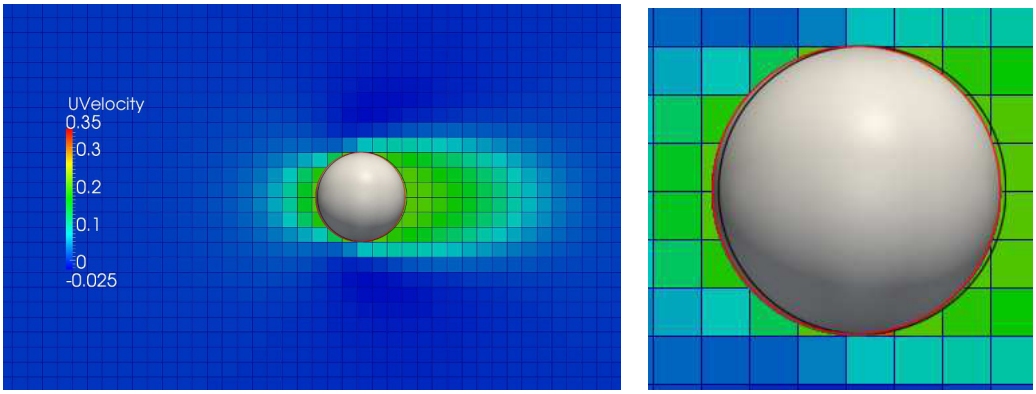
The benefit of this approach is it yields time-steps in which the stencil used for the boundary enforcement is unchanged. This gives time-steps in which the flow field is unaffected by these changes. Thes time-steps can be output and used to obtain a more accurate fluid particle interaction force. At very low CFL number this is of particular benefit as the in the unmodified case the spurious pressure field fluctuations that arise from 'fresh' and 'dead' cells grows approximately with $\mathrm{CFL}^{-1.5}$, for a fixed mesh resolution.

It is possible use this method only when a particle's CFL number obeys certain conditions. For example if the maximum distance travelled by any of the surface nodes in a time step, relative to the local grid refinement, is less than certain fixed number *e.g.* 0.05. While not currently implemented, it would be possible to adapt the movement further such that the particles had successive time steps without movement, to allow any residual effects of pressure oscillations to dissipate.

3.2.4. Cut cell momentum equation

The form of the momentum equation for boundary-exterior cells is not obvious. In the cut-cut partial mirroring method, the momentum equation can be discretised in the same way as for the complete-exterior cells. This is reasonable as the neighbouring cells have reliable values, hence the calculated momentum flux, shear and pressure forces are considered accurate. By construction, it is not possible for there to be a complete-interior cell adjacent to any boundary-exterior cell. However, this requires some dependency on boundary-interior cells, which can have discontinuities. To avoid this dependency, which could cause potential discontinuities a cut-cell version of the momentum equation system was implemented to complement the cut-cell version of the continuity equation.

An interpolation scheme, similar to that used for the continuity equation is employed. This is used to to interpolate velocity gradients to the cut face and in the calculation of the momentum convection. Explicit contributions are made from the immersed-boundary to the momentum flux. Semi-implicit contributions from the immersed-boundary are made to the viscous forces. The pressure contributions to the momentum equation are done in the same way as for an un-cut cell except they are weighted to account for the cut volume. The scheme can be summarised as follows:



(a) Velocity of flow in direction of motion of particle

(b) Close up to the vicinity of the immersed boundary

Figure 3.19.: Simulation of flow around a moving sphere using the mirroring immersed-boundary method method. The red circle represents the correct position of the sphere, the black circle an old time-instance used in lieu. Given the small difference relative to the mesh resolution, the difference in the discretisation is small. Using this method, the immersed-boundary is discreteised to have the position shown in black for two time-instances in a row. During the separating time-step the position of the particle did not change hence there were no 'fresh' or 'dead' cells, with the associated problems. This provides a time-instance with very reliable flow, which can be used for results and, in particular, for force calculation.

Transient term The transient term is based on the 'live' volume of the cell. *I.e.* the volume of the cell lying outside of the immersed-boundary.

Shear terms

Immersed-boundary contribution The shear stress is discretised for each triangle and summed. The individual triangle contribution is based on the intersecting triangle area and the distance in the direction normal to the triangle to the fluid cell centre.

Fluid phase contribution The fluid phase contribution to the shear terms is based on a linear fit of velocity through the stencil used in the cut-cell continuity equation (3.2.3.2). This is used in conjunction with the 'live' area of the cut face.

Convection

Immersed-boundary contribution An explicit contribution is made by each intersecting immersed-boundary triangle, given by the mass flux times the immersed-boundary velocity.

Fluid phase contribution Convection terms are based on an interpolated velocity field. The old time step value and the coefficients are interpolated using the same method used in the cut-cell continuity equation.

Sources These are added in a manner analogous to un-cut cells. However, instead of linearity in the cell volume, the terms are linear in the volume of the cell outside of the immersed-boundary.

Pressure terms In order to maintain consistency with the Rhie and Chow equation the pressure contributions to the momentum equations are scaled relative to the cut cell volume, as are the Rhie and Chow pressure coefficients in the continuity equation. This is equivalent to scaling the viscous and convective contributions relative to the ratio of original to cut cell volume, while

retaining the original volume elsewhere. Hence the pressure contributions are not discretised separately in terms of the immersed-boundary and the fluid phase separately. Instead the original (un-cut) pressure force contributions are used however, they are adjusted to account for the diminished control volume. In particular all of the contributions are multiplied by the ratio of 'live' volume to original (un-cut) cell volume.

Hence the momentum equation becomes:

$$S^i = \frac{V_{\text{cut}}}{V_{\text{org}}} \left[\rho V_{\text{org}} \frac{\partial u^i}{\partial t} + \sum_{f}^{f \,\epsilon\, \text{Original faces}} p_f n_f A_f \right] + \sum_{f}^{f \,\epsilon\, \text{Fluid faces}} (\rho_f u_f^i \text{Conv}_f + \mu A_f^j \tau^{ij}) + \sum_{F}^{F \,\epsilon\, \text{IB faces}} (\rho u_F^i \text{Conv}_F + \mu A_F^j \tau^{ij}) \tag{3.15}$$

The advantage of using the same interpolation scheme as used for the continuity equation is to ensure the same smoothness during transition of cell type.
The reason a cut-cell version of the pressure contributions is not used is it would require using extrapolated coefficients from a distance away from the boundary. As the coefficients are implicit this would create an non-physical linking of the pressures in the vicinity of the boundary. Additionally the extrapolation would require a larger stencil than is readily available.

3.2.4.1. Increased frequency of solid phase to resolve stiff collisions

An advantage of the form of the immersed-boundary contribution to the shear terms is that, as the immersed-boundary approaches the cell centre, strong shear terms correct the fluid velocity to be close to that of the immersed-boundary. This increases the smoothness of the transition. However, there is the possibility the calculation of force involves the division by a very small number (the normal distance to the immersed-boundary). To prevent the possibility of a singularity, if this distance is less than a small fixed fraction of the cell spacing the distance is replaced with the small fixed fraction. Although it is very unlikely to occur often naturally, by artificial construction (*i.e.* coincidental construction of immersed-boundary and mesh) or chance, it is possible. In addition to the potential benefit of incrementing the particle at a reduced frequency, it is also possible that there is an advantage of incrementing the particle at an increased frequency. In particular in the case of stiff colliding particles, it may be necessary to have a very small particle timestep. If it became necessary to increment the fluid at high frequency, the cost may become prohibitive. Additionally it may force the CFL number to be small, which is associated with oscillations in the pressure field. These problems can be overcome, by introducing a sub-time-step iteration scheme. Where the particle moves many times per time-step. It is possible to do both sub-step-incrementation and parity-based-incrementation. This is done by alternating between not-incrementing and repeatedly-incrementing the particle, during fluid time-steps. The advantage of this is that the particles relative motion is limited/reduced, at the same time as maintaining the existence of time-instances after a non-changed stencil.

3.3. Force calculation

As introduced in section 1.3.7 and discussed in section 2.3.2, the mirroring immersed-boundary method requires an additional force calculation based on integrating the fluid stresses over the boundary. See equation 1.3. This is complicated as a result of the fluid variables being known at points on the fluid mesh, not the particle surface. Hence an interpolation scheme is necessary in order to compute this integral. The choice of interpolation scheme is important in this case, for a number of reasons:

- As the particle moves, the stencil of points used for interpolation must change. This can lead to discontinuities in the forcing of particles.
- Some of the cells corresponding to the stencil may be intersected by the immersed-boundary. Any oscillation or error in the flow as a result of the immersed-boundary will be picked up by these cells and contribute to the resultant force. The changes is which cells are intersected by the immersed-boundary happens in addition to the change of stencil itself, with similar consequences.
- In the region surrounding an immersed-boundary the 'true' fluid properties, namely velocity and pressure, vary rapidly, complexly and anisotropically. Consequently the interpolation scheme is extremely important as it significantly effects the values of the fluids properties, used at boundary, determined from these important regions. This is particularly the case when two particles are in close proximity, namely in the case of a lubrication force, see section 1.3.4 and chapter 5. In this case the problem is exacerbated by the lack of exterior fluid points to use, when two particles are in close proximity. Hence particular attention is paid to situations in which two particles are in close proximity.
- Accurate prediction of the force is a critical component of the system.

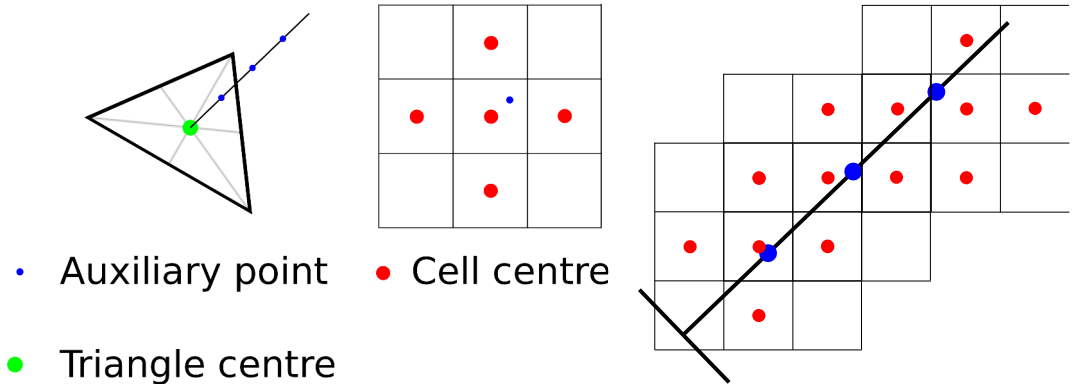
3.3.1. Stencil used for interpolation

The stencil used to interpolate the fluid properties to each triangle is based on the construction of three 'auxiliary' points. The use of these points is derived form a method presented in [33]. Each auxiliary point has a location on the line normal to its triangle, passing through the triangle centre. The points are uniformly separated, starting one separation away from the triangle centre. The separation is taken as the maximum cell diagonal distance, for the cell the triangle centre lies in. If the mesh spacing is uniform this ensures each auxiliary point lies in a different cell. The stencil is comprised of the cells in which the auxiliary points are found, and their face neighbours. Figure 3.20 shows this configuration and an example in two dimensions.

3.3.2. Velocity gradient

The velocity gradient is determined by fitting a polynomial through each of the components of the velocity field over the stencil given in section 3.3.1. The polynomial is comprised of all of the zeroth, first and second order coefficients of each direction.

The zeroth order contributions are weighted by the square the maximum cell diagonal, and the first order contributions weighted by the maximum cell diagonal. This ensures each contribution has



(a) Auxiliary points (b) 2D example of the resultant stencil

Figure 3.20.: Schematic of the process used to determine the points used to interpolate fluid properties to an immersed-boundary element.
Figure 3.20a: Three auxiliary points are created along each triangle of the immersed-boundaries. For each of these auxiliary points, the nearest cell centre and its face neighbours are added to the stencil.
Figure 3.20b: A two dimensional example of the process, for all three auxilliary points.

the same dimension and hence is not unduly weighted depending on the lengthscale of the problem. Adjustments could be made on top of this weight to effect the importance of zeroth order, linear or quadratic fit differently. These adjustments were tried but did not significantly effect the results and hence were not used.

The fit is made using a least squares linear regression fit of the coefficients of the polynomial, evaluated at each of the stencil points, to fit the known values there. This is done using the LAPACK routine `dgels` [85]. Details are given in the appendix, A.1.2.1. In addition to the weighting used in the appendix, a weight based on the fluid fraction of the stencil points was used. In particular each points weight was multiplied by its 'live' fraction. In this way non-physical interior points are removed from the stencil smoothly. The velocity gradients are found from the gradients, in each, direction of the polynomial fit, evaluated at the triangle centre.

3.3.3. Pressure

Similarly to the velocity gradient `dgels` is used to determine a polynomial fit through the data, except here the value directly, and not the gradient, is required. The interpolated value of pressure is set as the polynomial evaluated at the triangle centre point. The same weighting scheme is used as for the velocity gradient calculation. Sufficient points exist in the stencil to fit a third order component in the out-of-plane direction. However to do this requires re-orienting the relative positions into in, and out-of-plane components. This was tested, and did not make a significant difference to the calculation of pressure in the free stream case. Additionally, as the method is not practical when particles are in close proximity, the third order component requires removal when not enough fluid cell points are available. There was a small discontinuity as this transition occurred and hence the method was not used.

3.3.4. Special cases

The presence of an additional immersed-boundary complicates the interpolation of fluid properties to a given surface, as the interpolation scheme uses a relatively large stencil. This stencil may protrude past the additional boundary. In this case accurate fluid data is not obtainable.
In order to maintain an accurate estimation of the fluid properties, the interpolation scheme must be adapted. This is generally done my lowering the order of the interpolation scheme.

Pressure For pressure this is relatively simple as the value is interpolated directly from known pressure field. In the case that the some of the cells that would be used are unavailable, the order of the interpolation is reduced. Depending on the number of removed points this is either:

- Weighted linear interpolation.
- Shepard's interpolation.

Velocity gradient Estimating the velocity gradient at the surface is more problematic. In the case that the is sufficiently few remaining reliable fluid data points; it particularly difficult to resolve the velocity gradient and hence shear force. This is because a first order fit, at least, is required ascertain the gradient. Hence the span of the avaliable reliable cloud of points is calculated first. The velocity is fitted linearly in the spanned directions. In these directions the components are taken directly as the gradients of the fit. In the remaining direction, if there are any, the velocity gradient is taken as zero.

In the case the data points protrude past an immersed-boundary, when two particles are in close proximity, the resultant force can be of particular importance. However in this case it may be possible to reconstruct reliable data to represent the interstitial fluid. This is discussed in chapter 5.

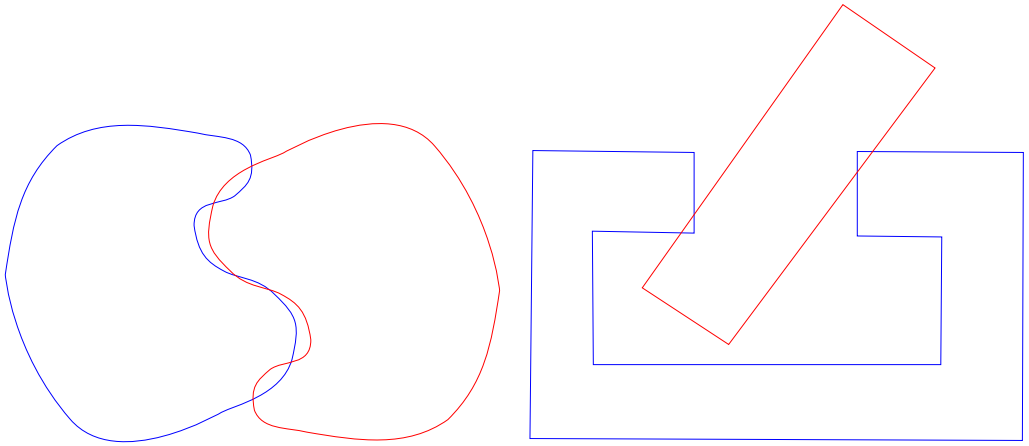
3.4. Contact model

3.4.1. Problems associated with the use of multiple particles

The principle difficulty assosiated with the use of multiple, arbitrarily shaped, particles is the numerical framework in which to encapsulate the collision data. In particular, there is no simple parametrisation.
For example:

- No canonical centre or contact direction.
- May have multiple contact patches.
- May be that a single contact point does not describe the contact, for example can have no net force, but a torque.
- After contact, the particle no longer has a closed surface with the fluid.

Some of these complexities are illustrated in figure 3.21.



(a) No obvious contact position or direction (b) Multiple contact patches cannot be seen as one

Figure 3.21.: Two dimensional representations of contacts that do not have a simple parametrisations.

3.4.2. Discretisation Scheme

- Breakdown the overlap of a pair of particles into simple pieces, where each piece is bounded by only two triangles (one from each body).
- Each piece simply bounded piece is parametrised in terms of:
  - Total volume (V)
  - Centre of volume (c_v)
  - Contact area A
  - Contact direction n_c
  - Effective curvature $C_{\text{effective}}$
- These pieces are recombined, and used in conjunction with the material properties, to produce complex contact behaviour.
- Closed surface for pressure re-obtained by using an additional internal pressure, calculated by smoothing the external pressure.

3.4.2.1. Breakdown of overlap

Basic geometrical properties The contact model requires:

- Volume of overlap section
- Centre of volume of overlap section
- Contact area of overlap section

However the calculations of these quantities from an arbitrary contact piece is non-trivial. This is because of the complex geometries that can arise from the overlapping triangular pieces that make up a contact. An example is given in figure 3.23.

In order to determine the necessary properties following steps are made:

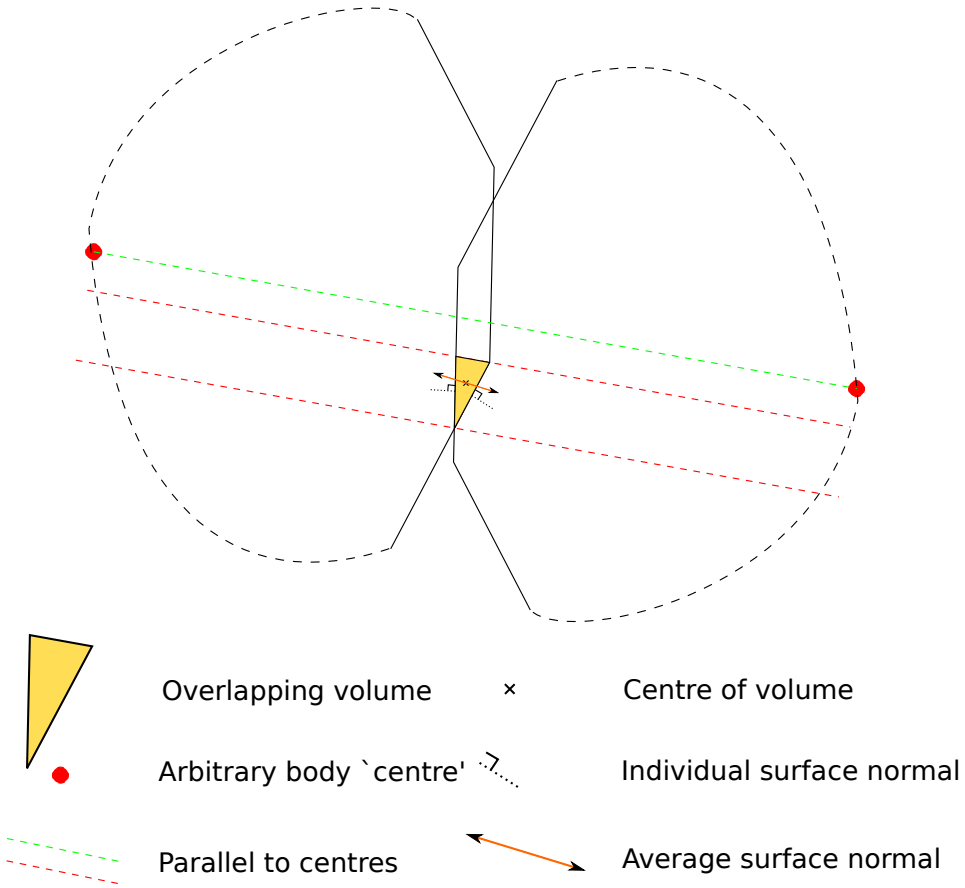


Figure 3.22.: Decomposing a contact overlap into sections bounded by only two triangles. The triangles are projected down the centre-line, joining the centre of each of the colliding bodies. Triangle-triangle intersection/overlap in two dimensions is used to produce a polygon of the projection of the intersection. This polygon can then be broken into triangles which can, in turn, be used to calculate the actual three dimensional overlap.

- The intersection of the triangle edges with the other triangle are made, in order to determine the projected boundries of the volume.
- The resultant volume is sub-divided into individual triangular pieces (see figure 3.24).

For each of these simple triangular pieces, the individual properties can easily be found, from which the properties of the whole piece can be determined.

The contact direction of overlapping section is taken as the normalised average of the normals of the two triangles forming the collision piece. This is given by the normalised difference of the two triangles' normals, as the orientations of the two normals are reversed.

Curvature estimation The curvature for each triangle is calculated as shown in figure 3.25. Curvature is fitted by comparing the in-triangle plane and out-of-plane components of the relative position of adjacent triangles. The in-triangle plane can be considered as is or fitted to account for skewness in the surrounding triangles.

Overall 'effective' curvature ($C_{\text{effective}}$) for the contact volume is considered by combining the curvatures of the two triangles that make it (C_A and C_B). This is done in the usual way, harmonic averaging, *i.e.* equation 3.16. The motivation for the use of harmonic averaging is that it captures the rate of seperation of the two surfaces, in relation to the distance from the epicentre of the contact.

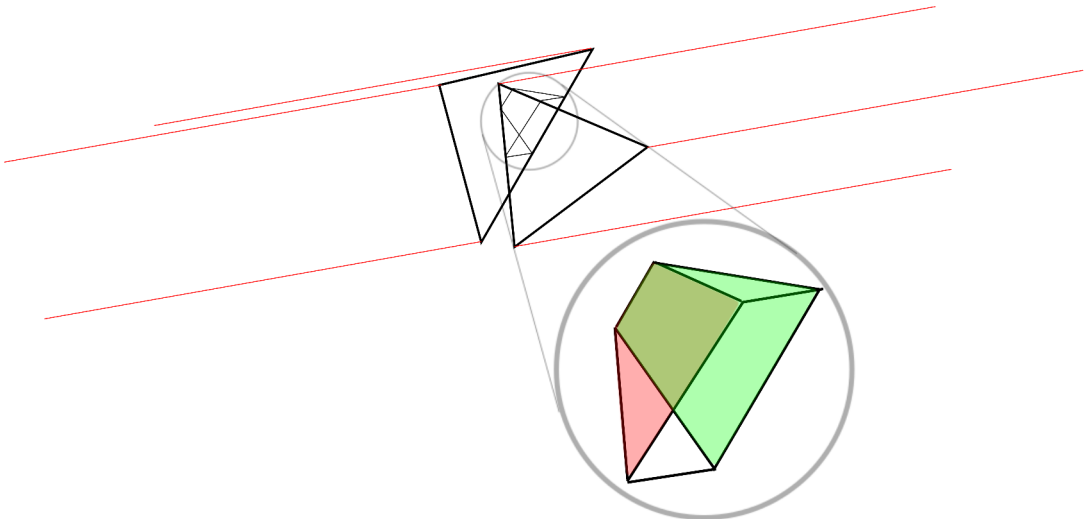


Figure 3.23.: Example of a contact piece used in the discretisation of a collision, with boundaries at previously unknown points.

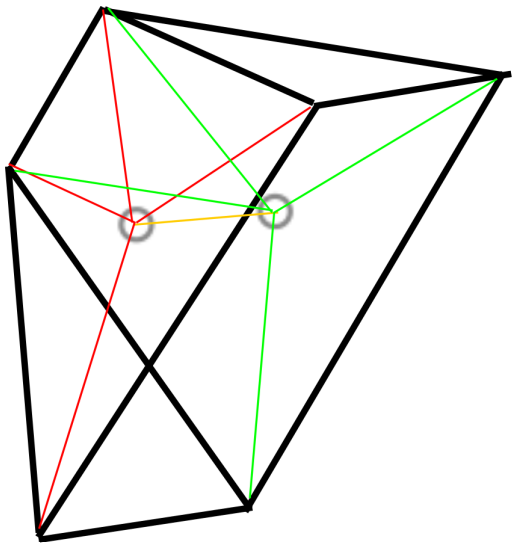


Figure 3.24.: Example of a contact piece used in the discretisation of a collision, as found in figure 3.23, expanded into triangular prisms.

$$C_{\text{effective}} = \frac{1}{\frac{1}{C_A} + \frac{1}{C_B}} \tag{3.16}$$

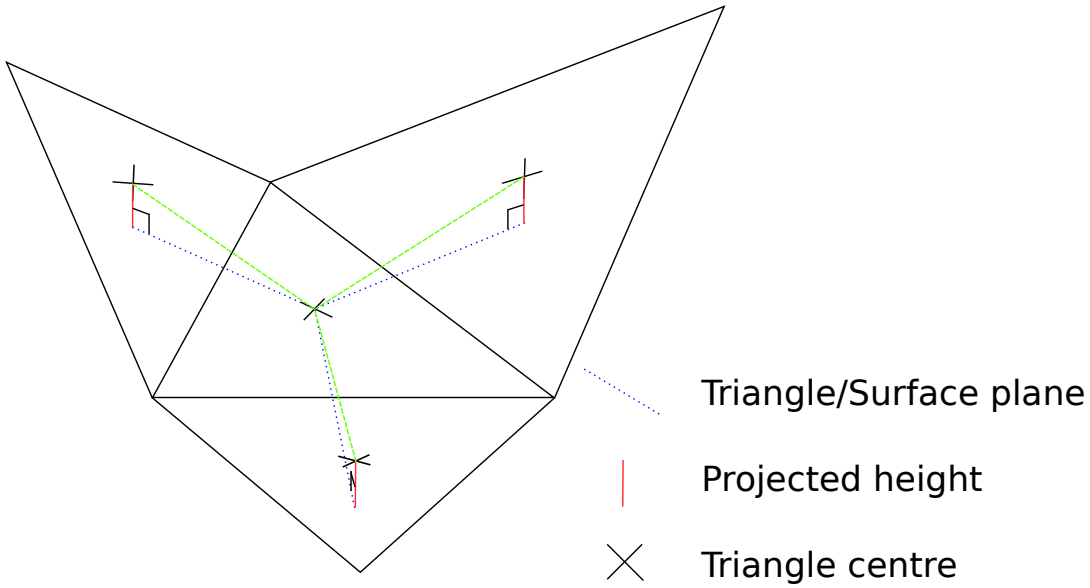


Figure 3.25.: Decomposing of neighbouring triangle centres, into in and out-of plane components, in order to estimate curvature. Height is expected to vary with the square of tangential distance from the central triangle's centre. This relationship can be used in reverse to compute curvature.

3.4.3. Force calculation from the discretised pieces

For each discretised piece, using the material properties of the two bodies, a force is calculated. The force is calculated based on Herztian contact dynamics.

3.4.3.1. Elastic component

For a constant radius of curvature *i.e.* locally an ellipsoid, the repulsive force between two bodies is given by:

$$F = a^3 \frac{4\ E^*}{3\ R} \tag{3.17}$$

over an area of πa^2, where E^*, the effective Young's modulus, is given by the harmonic average of the Young's modulii of the two contacting bodies, *i.e.*:

$$E^* = \frac{1}{\frac{1}{E_A} + \frac{1}{E_B}} \tag{3.18}$$

Hence an average pressure of:

$$P = a \frac{4\ E^*}{3\ \pi\ R} \tag{3.19}$$

Here a is the entire contact radius, which must be estimated for each contact from the pieces. It is advantagious for numerical reasons for this to be calculated only from the individual pieces. There are two reasons for this. Firstly the contact is not necessary circular, hence the total area may not be indicative of the local contact. Secondly, due to parallelisation, there are notable tractability advantages related to not requireing all of the contact pieces between particles to be known, in order to calculate the force from an individual piece.
The safest estimation of a, via an individual contact, is the contact depth, as this is not subdivided by the pieces that discretise the contact. Over an entire contact, the area and contact depth are

related by:

$$\delta = \frac{a^2}{R} \tag{3.20}$$

Hence $a = \sqrt{\delta\ R}$. Using this gives:

$$P = \sqrt{\delta\ R}\frac{4\ E^*}{3\ \pi\ R} \tag{3.21}$$

Although it is not possible to know $\sqrt{\delta}$ directly from a single piece of the contact, the correct area average can be obtained from the current overlap distance. Indeed the average square root of local overlap ($\overline{\sqrt{delta_l}}$) is directly related to the maximum overlap (δ) via equation 3.22:

$$\frac{\overline{\sqrt{\delta_l}}}{\sqrt{\delta}} = \int_0^{2\pi}\int_0^a r\ \delta\ \sqrt{1-r^2}\,\mathrm{d}r\mathrm{d}\theta/\delta\ \pi a^2 = 2/3 \tag{3.22}$$

Hence each piece of the discretisation is taken to contribute a force, F, given by equation 3.23 to the overall force.

$$F = A\ \sqrt{\delta_l}\frac{2\ E^*}{\pi\ \sqrt{R}} \tag{3.23}$$

It can be verified using equation 3.22 that, if the contact is sufficiently refined, this forces sum to the correct Hertzian limit. Indeed this force is also a good estimation locally as the pressure distribution is supposed to be given by:

$$p = p_0\ \left[1-(r/a)^2\right]^{1/2} \tag{3.24}$$

Where $\left[1-(r/a)^2\right]$ is proportional to the local displacement at r. Hence the pressure from each piece should indeed be proportional to $\sqrt{\delta_l}$.

3.4.3.2. Viscous component

In order to add damping to the collisions a viscous component is added. The viscous component of force is added according to the model of Tsuji *et al.* [7]. When discretised for each individual contact piece this becomes:

$$P = a\frac{4\ E^*}{3\ \pi\ R} \tag{3.25}$$

3.4.3.3. Friction

Hertzian contact, and a viscous correction, provides the normal force for a contact. The normal direction of the contact for each piece is taken as the average of the surface normals of the two triangles that bound it. The normal component is taken as zero if the net contact force is attractive. A cohesive model could be implemented, perhaps based on contact area, however cohesion is not applicable to this study.

The in plane component of force is determined using a coefficient of friction, the normal force and the relative motion between the two particles at the contact. In order to prevent unphysical oscillations, if the relative velocity is sufficiently small, the force is capped to prevent an unphysical change to the relative motion. This is done by calculating the effective mass m_E of the connecting bodies, related to the mass, moment of inertia and distance to the centre of mass of each body.

$$m_E = 1/\left(\frac{1}{m_1}+\frac{1}{m_2}+\frac{r_1^2}{i_1}+\frac{r_2^2}{i_2}\right) \tag{3.26}$$

Where i is the rotation axis given by $i = \frac{\mathbf{I}\cdot(f\wedge d)}{|f\wedge d|}$. With $\mathbf{I}$ the inertia tensor, f the direction of the tangential component of the force (relative to the centre of mass) and d the vector between the centre of mass and the collision.

As the overall collision is not known until the contribution from all the pieces is added, a second loop over the collision pieces is required. If it can be assumed that two bodies only have a single contact, as is the case for a spherical body, this can be avoided. Instead the overall collision could be calculated in a single loop and be treated in the same way as an individual piece with regards to a single tangential force calculation. This is only possible with a single contact. Additionally as the force calculation is not particularly expensive, the benefits are slight. However if the particles where sufficiently densely packed and resolved at a fine timestep relative to the fluid, *e.g.* in order to resolve a stiff particles the cost saving may be significant.

It would be possible to store an in plane deformation history of each particle surface and use this to determine a more complex tangential force model. This would allow a more detailed model and could be used to model a wider range of solids. Indeed it is possible to link the immersed-boundary method with any force based contact model, with or without deformation, as the fluid particle coupling is independent of force. However for the purpose of embankment structures, the elastic deformation in the collision plane is negligible and is suitably modelled by opposing relative motion directly. There are a number of reasons the elastic contributions to frictional forces are not significant. Principally, provided stability is maintained, the precise nature of the frictional forces is not critical as:

- In embankment structures a stable packing array is reached independently of frictional forces.
- The velocity of the particles in the system is usually very small, hence there is little sliding motion to oppose.
- As the particles are lubricated, the coefficient of friction between them is low.
- The evolution of the particles is prevented by forces in the collision interface normal direction as the particles are non-spherical.

Additionally, as the particles are stiff and the coefficients of friction typically low, the tangantial deformations are small.

3.4.3.4. Re-combination to calculate force and torque

Torque Calculation of the centre of the force for Herztian contact for each collision piece is complex. The determination of the centre of the force for each collision piece is significant as it affects the torque calculation. The main dependency of the torque calculation on the centre of force is its in-plane component. This is for two reasons:

- The force usually acts predominantly in-line with the collision normal. This is significant as the component of position in the direction of the force does not effect torque.
- The contact is invariably wider than it is deep. As a result the normal component of the centre of force is tightly bounded.

The determination of the in-plane component of the centre of force is affected by the collision depth dependence of the force. As the collision piece is already limited to second order accuracy by the assumption of piecewise linearity of the body, the centre of volume of the overlapping pieces is taken.

While this deviates a small amount from the Hertzian limit, in an unrefined case, this may not be unphysical. Indeed the majority of particles likely to be modelled are not perfectly smooth at or near the body's lengthscale. Local heterogeneities are likely to add more noise than this approximation. If required, the Hertzian limit can be obtained by appropriate resolution of the body. Alternatively, if the body is simply parametrised, as is the case for a sphere, the analytical force can be used instead.

3.4.4. Time incrementation of the solid phases

As the solid phase may be stiff, it is possible a finer timestep than that of the fluid may be necessary to accurately capture the collisions. Hence it is sometimes necessary to increment the body a number of sub-steps per fluid steps. The number of sub-steps (n) is chosen such that the subdivided time step is less than the maximum permissible timestep (δt).

δt is chosen such that:

- No significant change in overlap can occur per time step.
- The time period (corresponding to the natural frequency) of each contact is subdivided a few times.

In both cases a numerical parameter is necessary. In particular a maximum permissible change in overlap as a fraction of particle radius and maximum permissible fraction of a period of osculation that can occur per time step. These are chosen to balance their influence on the particle's trajectory and the computational cost.

For cases in which particles are in constant contact the natural frequency determines the minimum step however in energetic collisions the maximum overlap is usually limiting. As these parameters are adjustable individually, it is possible to use a single set for a wide range of cases. It was found that permitting an overlap change of 0.005 of a particle radius and requiring 20 timesteps per period was sufficient to remove dependency. As the particle phase was still not the major determinant of tractability this setup was used throughout except for validation studies.

Depending on the significance of the collisions to the results of a particular setup, it may be beneficial to alter these. An overlap fraction of 0.02 and 5 timesteps per period where found to give sensible results though some effect on restitution was found.

It is noteworthy that the refinement of the bodies in question is of significance as more finely resolved particles have a smoother repulsive force. Hence it may be necessary to resolve more coarsely refined bodies with a higher frequency. For this reason it is recommended that bodies be discretised such that the largest surface element has lengthscale less than 5-10% of the bodies radius.

3.5. Equation system solution

As introduced in section 3.1.1 the discretisation scheme used is presented in [60]. A small number of modifications are needed to account for the developments to the mirroring immersed-boundary method.

3.5.1. Momentum weighted interpolation

Momentum weighted interpolation is used throughout the method. However in the vicinity of an immersed-boundary this can be problematic. This is due to the absence of a consistent range of

pressure values required by the stencil for some exterior faces. The following modifications are made to overcome this difficulty:

- Rhie and Chow continuity equation coefficients are determined via extrapolation from away from the body for faces of boundary-interior cells, and removed for complete-interior cells.
- Given the wide range of possible immersed-boundary geometries, this does not exclude a non-zero coefficient from a boundary-interior cells. Hence a Laplace equation is solved inside for the complete-interior cells. This gives rise to a smooth pressure field over the entire domain. Hence there are sensible values even if any complete-interior cells are taken in the stencil. The impact of the exact value of the pressure for complete-interior cells is negligible, they do not appear in the momentum equations. An alternative, for complete-interior cells, would be to use some sort of reference pressure or to guarantee their exclusion from the Rhie and Chow terms. The use of a reference value is less than ideal, as it requires one to be present and determined. Further more it does not guarantee smoothness in the pressure field. Without any correction a large gradient could occur and skew the continuity equation. Guaranteeing their removal from the Rhie and Chow equations is complex and potentially causes problems for continuinty.

Figure 3.26 shows the reconstructed interior pressure. There are two additional benefits of this method with regards to the availability of a semi-reliable internal pressure field:

- Firstly in calculating pressure force contributions during collisions. The pressure field is of consequence for reconstructing the pressure force acting on the surface of the immersed-boundary. In the unlikely event of two bodies overlapping by more than a single cell spacing the pressure field is still available over the surface of the body. It also aids in reconstructing the pressure at the surface for smaller overlaps. It is unlikely that this will make a significant difference to the overall collision force provided a remotely consistent pressure field is present inside the body, as the collision force itself will likely dominate. In a similar manner, prior to collision, the interstitial pressure may not be known. Even prior to collision, it may be that there are no external cell centres separating the bodies. Hence the use of internal points is required and a smoothed field is beneficial, even if it is avoided when possible.
- Secondly extrapolating the pressure field to the boundary-interior points may require use of complete-interior cells, when two immersed-boundaries are in close proximity to each other. While it would be possible to revert to a different method in this case, there are advantages for smoothness and simplicity to having reliable values for pressure at each point.

3.5.2. First order time scheme

The immersed-boundary method involves transitional periods, as immersed-boundaries cross cells and cells change their type and hence discretisation scheme. While effort is made to prevent discontinuities in time in the velocity field, discontinuities in *gradients* over time are not always avoided. Hence extrapolating from two time instances ago is not recommended. A further problem regarding the use of old-old time-instance values is the possibility of including the velocities of cells which, at the time, where complete-interior. Indeed while the CFL constraint prevents the possibility of a boundary-exterior cell having been complete-interior the previous time instance, no such guarantee

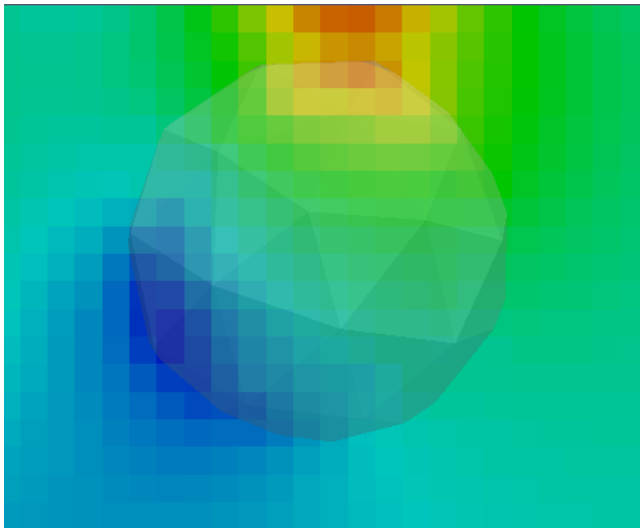


Figure 3.26.: A translucent immersed-boundary showing a pressure field inside the body, compared with the external flow field pressure.

is made for two time instances ago. As a result there are a number of cases where the 'old-old' values of the velocity field may not be accurately used to extrapolate the velocity field. Hence a first order forward Euler time integration scheme is used. It would be possible to reconstruct a smooth interior velocity field in order to have a less problematic extrapolation however this is not necessary particularly physical. Indeed any benefit would be subject to the non-physical 'flow' deep inside the body. An alternative would be to use a second order method only for complete-exterior cells. However given the limited number of such cells in a typical granular material case, and the lack of evidence to support substantial improvement, this technique was not adopted.

3.5.3. Matrix formulation for cut cells

For boundary-interior cells, to which partial mirroring applies, there is no canonical ordering of the equations in the Navier-Stokes equation matrix. The ordering used below was found to maximise diagonal dominance, improving solution time.

- The equation for pressure, described in section 3.2.2.2 it assigned to the 'pressure' row in the matrix for that cell.
- The two in-plane vectors are dot producted with the basis vectors.
- The largest dot product magnitude determines the first assignment. *I.e.* if the largest dot product of any in-plane vector, with any basis vector product came from:
  - in-plane vector $\mathbf{a_i}$
  - basis vector $\mathbf{e_j}$

  Then mirroring equation governing $\mathbf{u}_j$ would become placed in the row of the matrix corresponding to the velocity in the direction $\mathbf{e_i}$.
- The largest dot product of the remaining in-plane vector and basis vectors determines the second assignment in the same way.
- The continuity equation is set as the remaining velocity equation.

4. Particles settling under gravity

4.1. Aims and Motivation

The primary concern of this research is the behaviour of densely packed arrays of particles. However, due to the complexity of such cases, direct validation is near impossible. Conversely the individual particles these arrays are comprised of are easily parametrised in certain situations. This offers way to test the overall validity by testing the component parts. Some experimental data exists, relating to individual particles under varying situations and the resultant behaviours. In particular, ten Cate *et al.* [86] describes a number of experiments in which a particle settles under gravity in a viscous fluid. In this chapter the developed method is validated by scrutinising simulations involving individual particles. To this aim a number of validation cases were run in order to ascertain the correct functioning of the method across the range of cases likely to be encountered. These cases can be subdivided into physical variations, determined by experimental setup, and numerical variations relating to the method of simulation. For the physical cases the subdivisions are as follows:

- Steady vs unsteady flow.
- Particles in free stream and those in close proximity.
- Reynolds number of the flow, *i.e.* low Reynolds number and high Reynolds number.

The numerical conditions investigated focus on:

- Resolution of the particle.
- CFL number.
- Motion of an immersed-boundary relative to the Eulerian mesh.

4.2. Drag forces on a particle travelling through a viscous fluid

Experimental data on the spherical particles moving through viscous fluids is provided by ten Cate *et al.* [86]. The experiments consist of dropping a Nylon sphere in a tank filled with various grades of silicon oil, and measuring the velocity of the particle.

The results show the establishment of the terminal velocity of a heavy sphere falling through a range of Newtonian fluids, for an otherwise fixed experimental setup. As the particle reaches a steady terminal velocity, and the gravitational force and buoyancy forces are known, the experiments establish relations between particle velocity and drag forces though fluid at a range of Reynolds numbers. These values, along the with the knowledge that terminal velocity has been reached, and is independent of the influence of the far (end) walls, allows the verification of the developed method for steady flow. In particular the relation between the hydro-dynamic interaction force acting on a

Case	ρ_f [kg/m^3]	μ_f [Ns/m^2]	Re [−]	St [−]
E1	970	373	1.5	0.19
E2	965	212	4.1	0.53
E3	962	113	11.6	1.5
E4	960	58	31.9	4.13

Table 4.1.: Fluid properties used for determining the flow around a settling sphere and the resulting Stokes number and Reynolds number estimations, as predicted by Abraham [87].

particle and its velocity, in the geometry used for the experiments can be investigated. This relation is used to validate the developed models. In addition, the accelerating and decelerating phases of the flow gives a benchmark for the unsteady case and how the particle interacts with the wall. The developed methods are tested under a range of circumstances including the effects of:

- Different Eulerian mesh resolutions.
- Different time increments.
- Galilean transformations of the case (to control the relative velocity of the Eulerian mesh and immersed-boundary).
- The resolution at which the immersed-boundary is tessellated.
- Alternate conditions:
  - Fixed velocity, measuring force.
  - Fixed force, measuring velocity.

This allows the validation of the correct relation between force and velocity, the sensibility of the flow field and the non-dependence on numerical results.

4.2.1. Experimental setup

The experiments performed by ten Cate *et al.* [86] consist of releasing a Nylon sphere suspended in a tank containing a viscous fluid. The tank is of dimensions: $100, 100, 160$ mm for width, breath and height respectively. The sphere is of diameter $d_p = 15$ mm , density $\rho_p = 1120$ kg m^{-3} . The sphere is released from within the fluid at a height of 120 mm from the base of the container, as measured from the bottom of the sphere, centrally in the horizontal plane. The experimental setup is shown in figure 4.1. Table 4.1 shows the grades of silicon oil and their properties, used to perform the experiments. Also shown in table 4.1 are estimations of Stokes number and Reynolds number for each case. Reynolds numbers and Stokes numbers are based on drag coefficient predictions of Abraham [87]. Table 4.2 shows velocity predictions from the drag coefficient predictions of Abraham, which assume zero Reynolds number and infinitely removed domain boundaries. Table 4.2 also shows how the experimental results differed from the predictions.

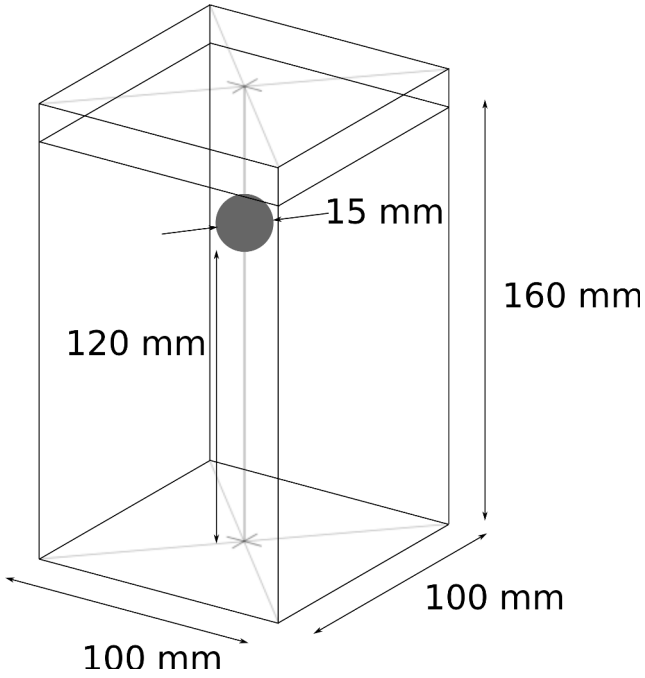


Figure 4.1.: Setup of a tank containing viscous fluid and a solid sphere used by ten Cate *et al.* to perform experiments.

Case	u_{inf} [ms^{-1}]	u_{max} [ms^{-1}]
E1	0.038	0.034
E2	0.060	0.057
E3	0.091	0.087
E4	0.128	0.121

Table 4.2.: Experimental results showing the terminal velocities [u_{inf}] of settling spheres through each of the grades of silicon oil. Additionally, for each case the predictions of Abraham [87] for the terminal velocity if the boundary effects and inertial effects are ignored are given [u_{inf}].

4.2.2. Numerical setup

4.2.2.1. Boundary conditions

Boundary conditions were set as follows: At the free surface, pressure and velocity gradient in the normal direction set to zero. Walls where set to no slip and zero pressure gradient, in the normal direction.

For the cases were the particle is stationary relative to the mesh, the free surface condition is as above while walls are set to have the characteristic velocity, as well as zero pressure gradient. Particle-wall collisions require a complex numerical framework. Hence, as the particle is expected to hit the bottom of the tank, the fluid domain is extended downwards and a cubic immersed-boundary placed over the end section to replicate the bottom wall in the simulation. This provides the necessary data and framework required to extrapolate the flow field to the falling particle's surface as well as to resolve the collision. In this case the pressure boundary condition is slightly different as the pressure gradient is not fixed but implicitly extrapolated from slightly away from the surface. This is only

	Cells per radius	Total number of cells
Low	3	102400
Medium	4	200000
High	5	480323

Table 4.3.: Resolutions of meshes used for settling sphere simulations.

significant when the two immersed-boundaries are in close proximity.

4.2.2.2. Gravity

Gravity, where applicable, is accounted for by a buoyancy force acting directly on the body. This is preferable to adding a momentum source to the Navier-Stokes equations, linear in density and volume, as it reduces the interpolation error of the pressure field over the body. In particular, calculation of both the volume and pressure interpolation is a potential source of error. This is particularly important for the current cases as the difference in density between the particle and fluid phases is small relative to either densities. Hence a small error in either the buoyancy force or the gravitational force could be large relative to overall force imbalance acting on the particle.

4.2.2.3. Refinement

An equi-distant Cartesian mesh is used to discretise the Eulerian phase for all of the simulations. The refinement is varied for each of the simulations, to test mesh dependency. The refinements are given in table 4.3.

4.2.3. Stationary particle

Stationary particle simulations are simple test cases of a forced-stationary spherical particle, over which the fluid is forced at the experimentally determined terminal velocity. This setup can be seen as a Galilean transform of the original experimental case at the terminal velocity stage. The advantage of these cases is it removes the problems associated with moving the particle relative to the Eulerian mesh.

The aim of these simulations is to determine a baseline for the moving particle cases. In particular it is important that the fluid force, acting on the particle, converges to the expected force required to obtain the given relative velocity, namely the force imbalance due to gravity. This is clear as the fluid force acting on the particle and the velocity field relative to the particle should be independent of the cause of motion.

In order to take advantage of the known terminal velocity of the falling sphere at certain locations, the sphere is placed at a location at which the falling sphere is at its terminal velocity. The location of 80 mm from the base of the tank is taken. This setup is illustrated in figure 4.2.

4.2.3.1. Equivalence of methods for the stationary case

In the developed method, the stationary case should not vary significantly from other validated methods. That is, for a given mesh resolution and timestep, the flow around a stationary immersed-boundary should be independent of the immersed-boundary method used. Figure 4.3 shows the

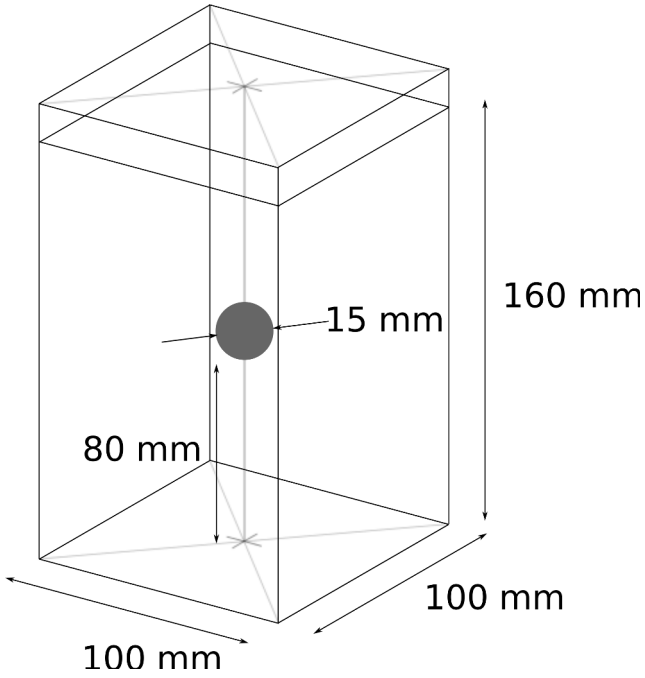


Figure 4.2.: Simulation setup of a solid sphere falling through a viscous liquid equivalent to that used by ten Cate *et al.* [86]. The setup differs form the original by lowering the particle such that a known the relation between force and velocity is known, *i.e.* such that the particle is removed from the boundary conditions. While this is not the starting position of the sphere, the sphere passes this point during the experiment, and the simulation is comparable to this point.

Case	E1
Refinement	4 cells per radius
Max CFL number	~ 0.03

Table 4.4.: Details of simulation setup for drag force comparison (fig 4.21).

force acting on the immersed-boundary for an instantly accelerated sphere for both the developed method and that of Mark and van Wachem [58]. The setup, including geometry and velocity is that of case E1 see table 4.4. Both forces are those calculated using the least-squares polynomial fit method developed, see section 3.3. Both forces are nearly identical to each other at each time step in the simulation. Figure 4.4 shows the same result decomposed into pressure and shear stress contributions. Figure 4.5 shows the same experiment but with additional simulations using different aspects of the developed method. Both the fully cut-cell partial-mirroring method *without* averaged mirroring and the method of Mark and van Wachem [58], with the addition of averaged mirroring are included.

It is clear that in all cases the resultant force is very similar at all stages of the flow's development. This is to be expected as the variation in the discretisation schemes are largely to reduce the effect of motion of the immersed-boundary relative to the Eulerian mesh. In the static case the difference is less significant. In the case the immersed-boundary does not move relative to the Eulerian mesh the immersed-boundary is recreated solely in terms of the geometric interpolation used to represent the set points lying on the immersed-boundary. Indeed the representation of the immersed-boundary is solely by means of the fixed interpolation coefficients from the Eulerian mesh to this set of points.

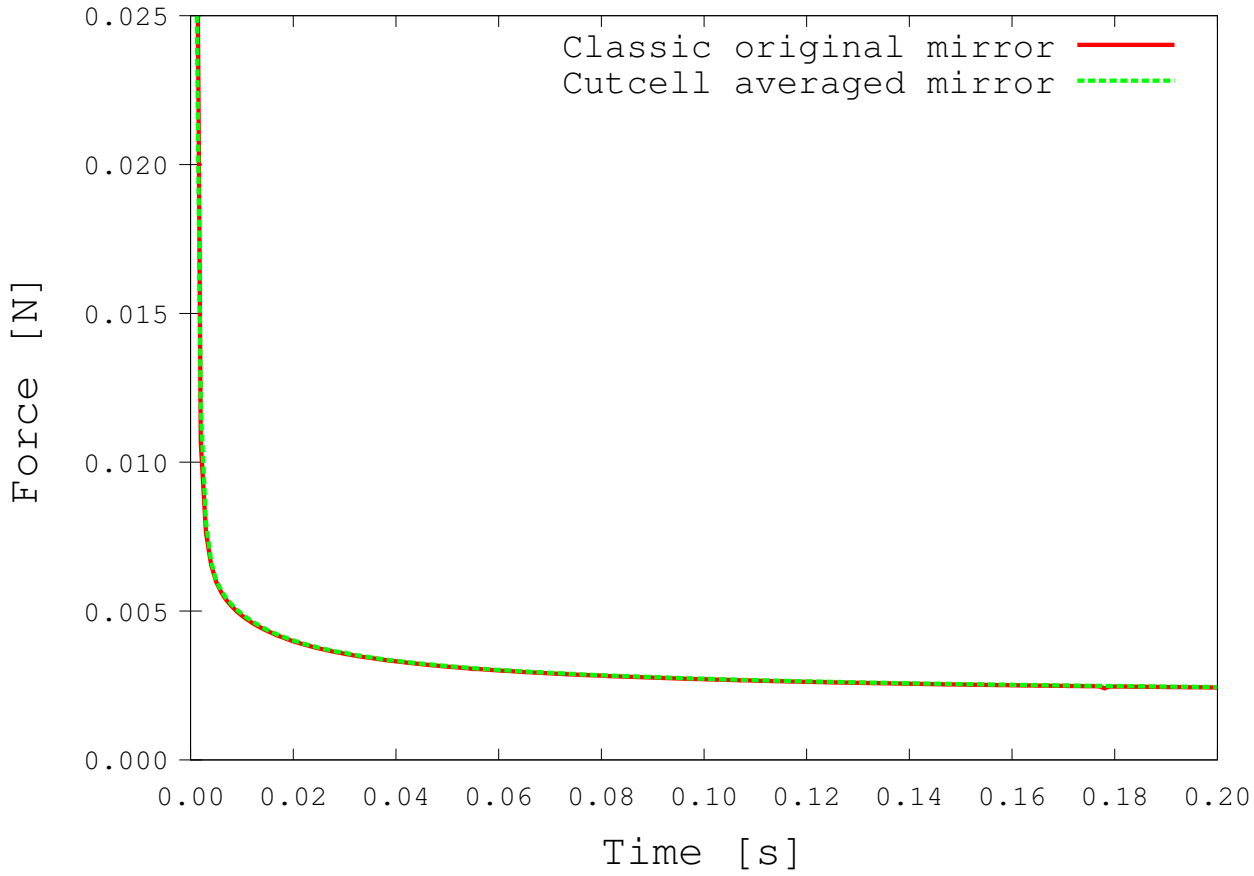


Figure 4.3.: A plot of the drag force on a particle in a suddenly accelerated fluid phase. Two cases differ by which discretisation scheme is used, the green curve uses the developed partial-mirroring, cut-cell momentum and averaged-mirroring modifications to the method. The red curve shows the predictions of the mirroring immersed-boundary method without these modifications. In both cases the same force calculation method is used. The force calculation method used is described in section 3.3. A good agreement is found, the lines are practically super-imposed, which corroborates the sensibility of the developments to the method.

However differences in these coefficients could also be seen as representing small changes in the geometry of the immersed-boundary. Hence the changes in the choice of discretisation scheme are equivalent to small (relative to cell-spacing) perturbations to the geometry of the immersed-boundary. Hence, for a suitable choice of interpolation function, the discretisation is determined by refinement alone.

There are slight differences for the initial, rapidly evolving, period of the simulation between the cut-cell partial mirroring method and the complete mirroring method. This is consistent with the differences between the cut-cell and the un-cut version of the momentum equations. Indeed while the momentum exchange over the exterior surface of the boundary-interior cells should be the same, how this is is distributed over the fluid and solid phases is dependent on the momentum equations. In the un-cut method, the transient term in the momentum equation of boundary-interior assumes the fluid in that cell has the mass of an *entire* cells worth of fluid. In the cut-cell momentum equation only the correct mass of fluid is accounted for. Hence in the un-cut case the acceleration of the fluid in boundary-interior cells balances a slightly different (greater) momentum exchange, and hence less

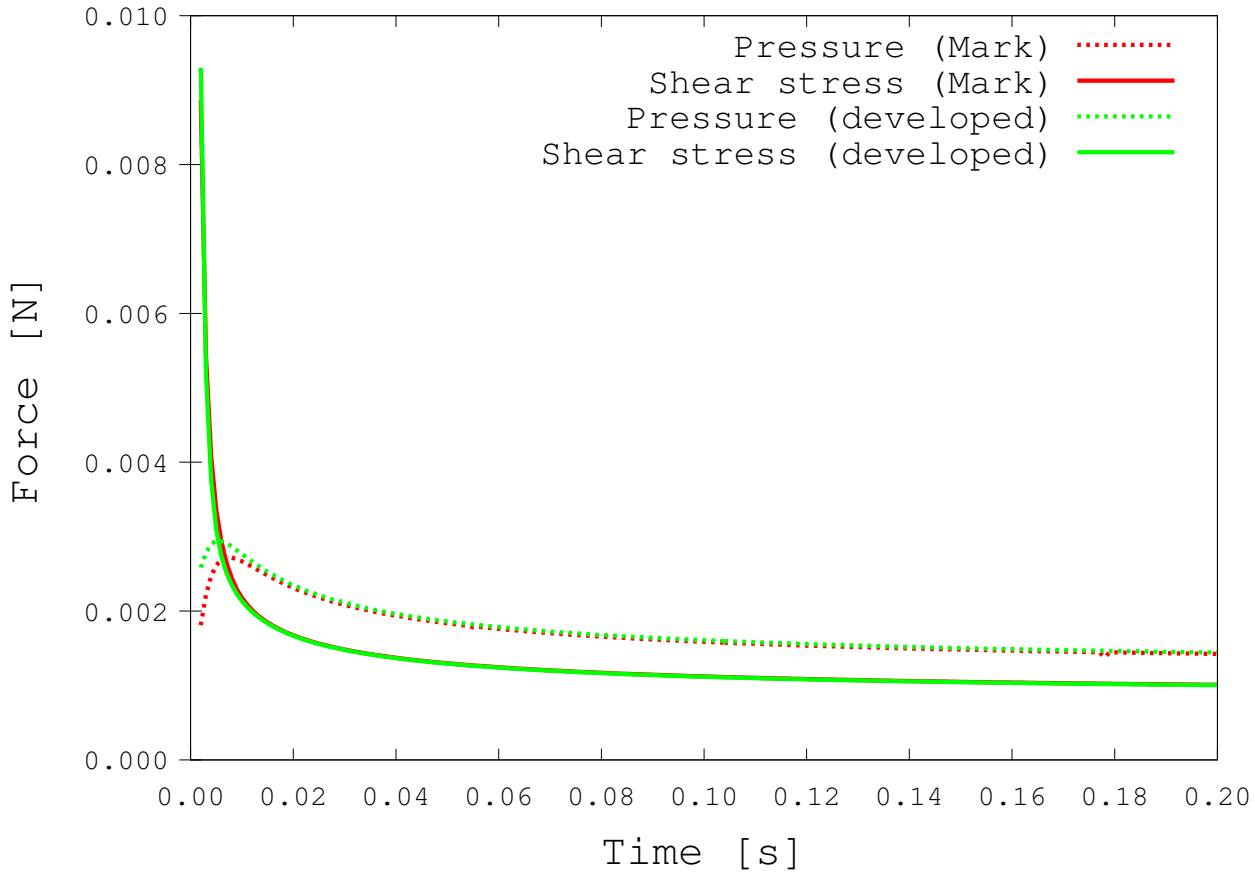


Figure 4.4.: A plot of the separation on hydro-dynamic interaction force (drag force) into pressure and shear stress contributions for the cases shown in figure 4.3. A very good agreement is found once the flow has developed however, there is a slight difference in the separation of the forces in the initial 2 hundreds of a second were inertial effects are substantial. This is not unexpected given the use of a cut-cell momentum equation.

remaining to be balanced by the solid phase.

4.2.3.2. Dependence on refinement

Refinement dependence is of obvious importance to the immersed-boundary method. The usually limiting dependence is for the moving particle case as the continuity equation variation due to motion of the immersed-boundary is strongly linked to refinement. Section 4.2.4 discusses resolution dependency of the moving case. A refinement study of the stationary case give an interesting result. Figures 4.6, 4.7 and 4.8 show the fluid force acting on a body, seperated into contributions from shear stress and pressure and the combined total force. In each figure the force acting on the same physical case, with differing Eulerian mesh resolutions (2, 3 and 4 cells per Eulerian grid cells per radius), and for a different velocity gradient and pressure calculation method is shown. The different methods of calculating the gradients are: the methods described in section 3.3.2 and 3.3.3, and similar but constantly-weighted polynomial interpolation. The results show dependence on refinement, an increased refinement increases the viscous shear force leaving the pressure force unchanged. The results do not show a dependence on the interpolation method used, with the force contributions, and overall force, indistinguishable for each given refinement. The lack of dependence on interpolation

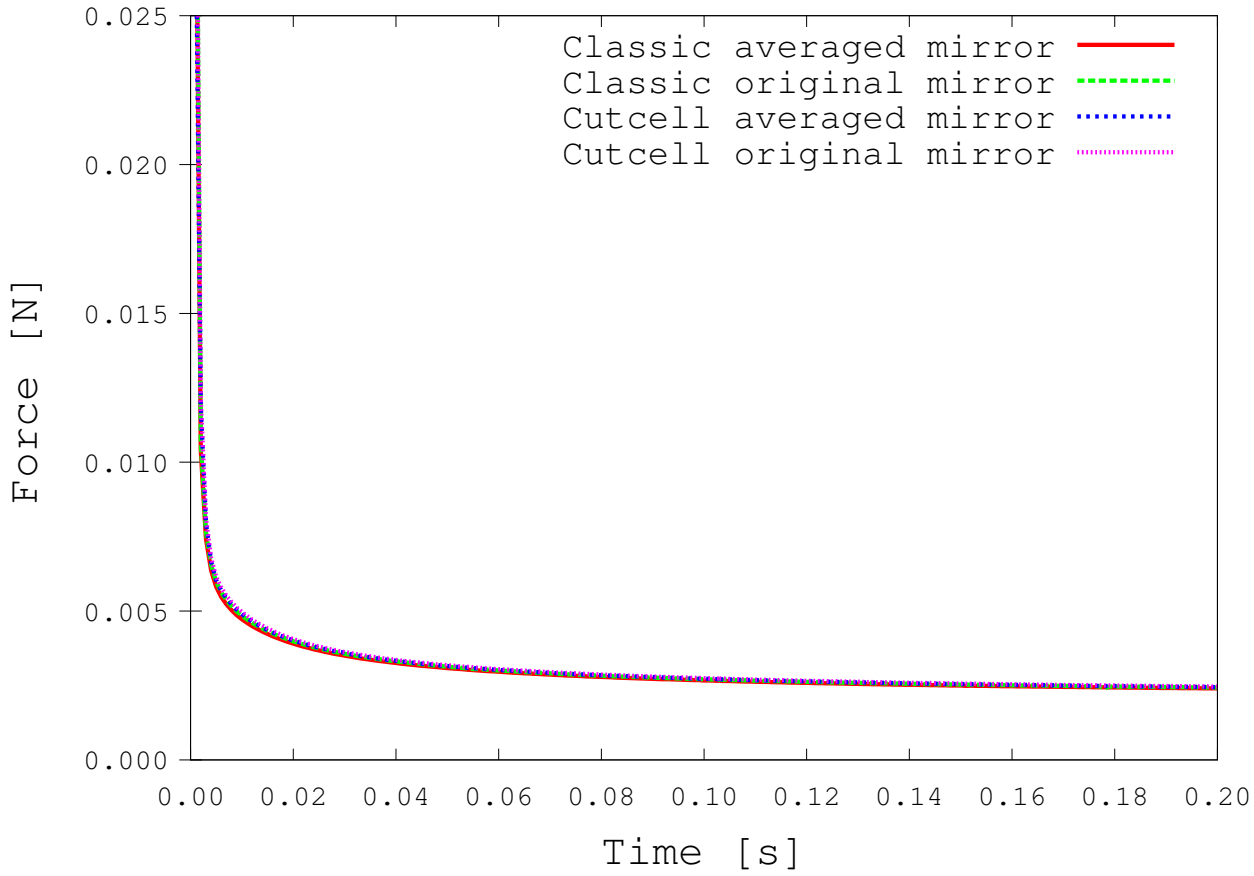


Figure 4.5.: A plot of the total hydro-dynamic interaction drag force on a particle in a suddenly accelerated fluid phase for four numerical setups. The numerical setups differ by which modifications are included. The setups used are: the un-modified mirroring immersed-boundary method as shown in figure 4.3 (green), with the addition of only the averaged mirroring modification (red), with the addition of the cut-cell partial mirroring method (blue) and with the addition of both the cut-cell partial mirroring method and with the averaged mirroring modification (magenta).

scheme of velocity gradient and pressure is un-surprising given the smoothness of the fields. Given the non-dependence on resolution of the flow field, even in the moving case see section 4.2.4, and the lack of dependence on the weighting used, it is surprising that the resolution impacts the force so strongly. The variation in force is as a result of the differing interpolation of velocity gradient at the immersed-boundary surface. There are a number of explanations for this, even given the general good agreement between the velocity field at different resolutions. One possibility is that while the low resolution is able to capture the smooth variation in the far field, the near field variation is under-refined. However if this were the major cause of the discrepancy it would be unlikely that such good agreement between the interpolation methods would be found. A more plausible explanation is the increased distance over which the interpolation takes place for lower refinement, as a result of the fixed number of points being more spread out, as described in section 3.3. This could affect force calculation as the non-quadratic components of the field skew the interpolation more heavily if a larger domain is used. The potential to reduce this dependence is of interest, and a number of approaches to this may be possible. One possibility would be to increase the order of the interpolation, however, this would require the use of more interpolation points, further extending the domain over which

the interpolation takes place. This approach may be beneficial at increasing the accuracy of refined simulations however is unlikely to be able to reduce the refinement required to give reliable results. For the free stream case the functional form of the expected variation of the flow field around the immersed-boundary is predictable, particularly at low Reynolds number. In such a case it may be possible to exploit this form to be able to capture the complex flow with fewer parameters hence requiring a smaller domain. However, this method is not applicable to modelling particles in close proximity, as the assumptions regarding the variation of the flow field around the immersed-boundary are invalid.

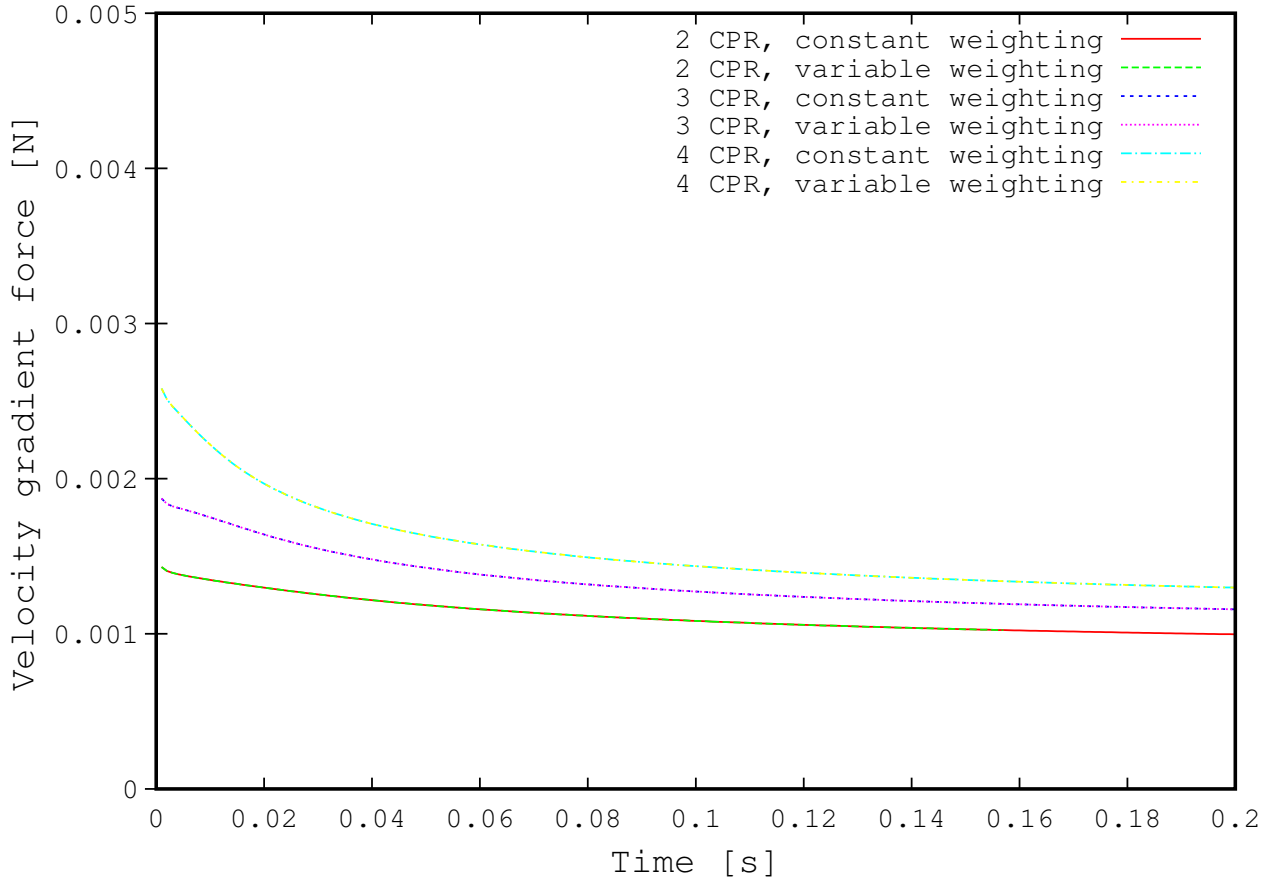


Figure 4.6.: A plot of the shear stress contribution to drag force on a particle in a suddenly accelerated fluid phase. The cases differ by which scheme is used to interpolate fluid properties to the immersed-boundary surface, to determine the hydro-dynamic interaction force; either a constantly weighted polynomial interpolation (constant weighting) or either the variably weighted polynomial interpolation as described in section 3.3 (variable weighting). For each of the interpolation methods, three Eulerian phase resolutions are given; 2, 3 and 4 Cells Per Radius (CPR). The three distinct line correspond to these resolutions, for each of which the interpolation method has little effect.

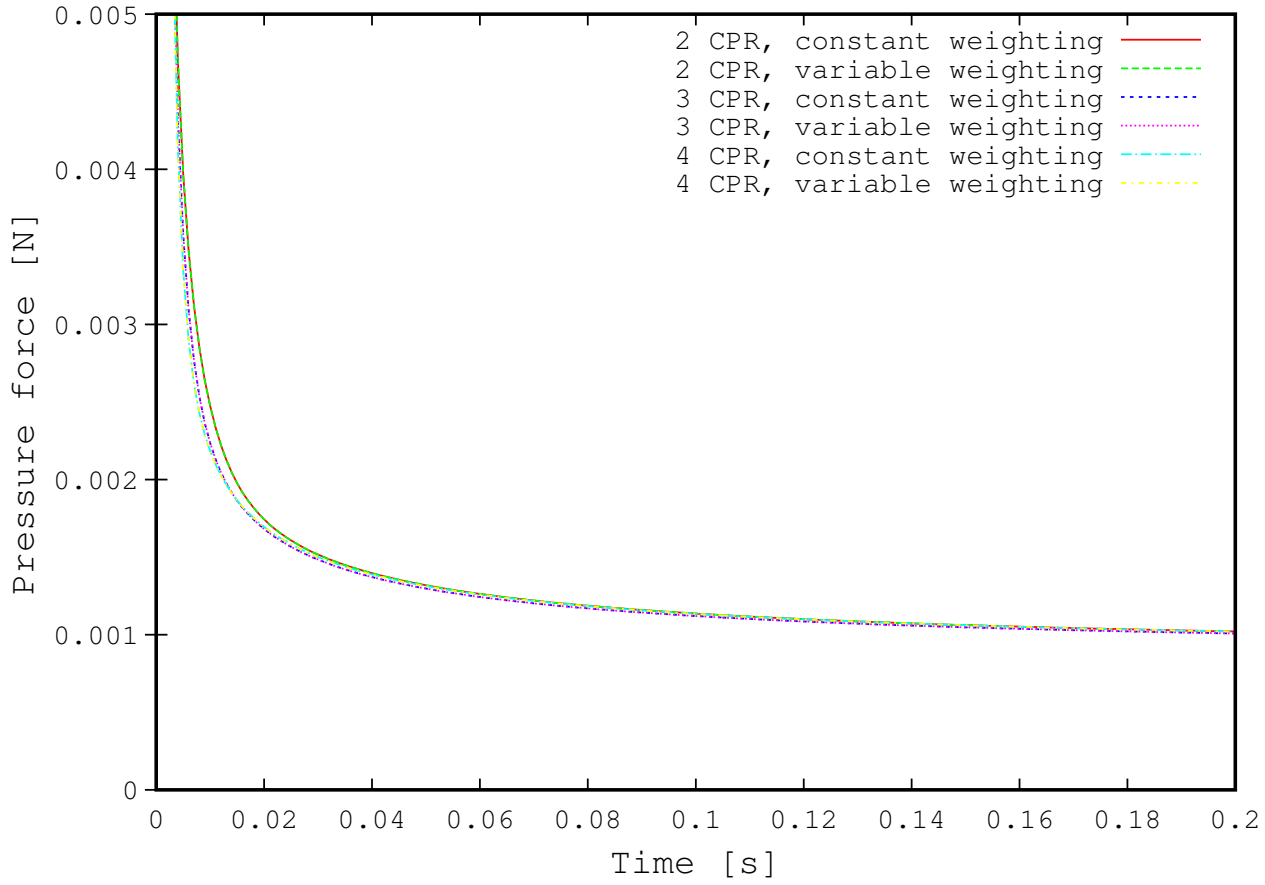


Figure 4.7.: A plot of the pressure contribution to drag force on a particle in a suddenly accelerated fluid phase. The cases differ by which scheme is used to interpolate fluid properties to the immersed-boundary surface, to determine the hydro-dynamic interaction force; either a constantly weighted polynomial interpolation (constant weighting) or either the variably weighted polynomial interpolation as described in section 3.3 (variable weighting). For each of the interpolation methods, three Eulerian phase resolutions are given; 2, 3 and 4 Cells Per Radius (CPR). All of the lines on the graph are effectively indistinguishable.

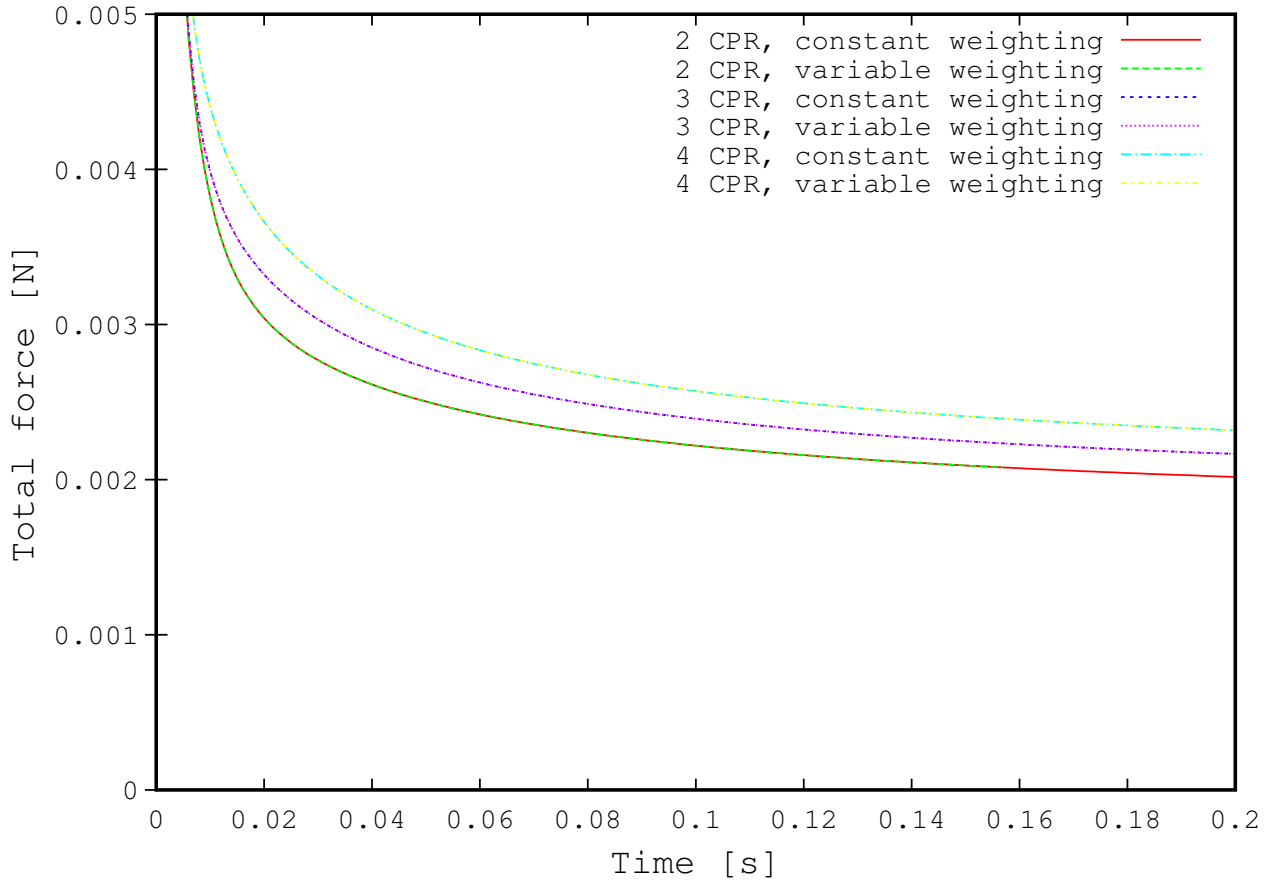


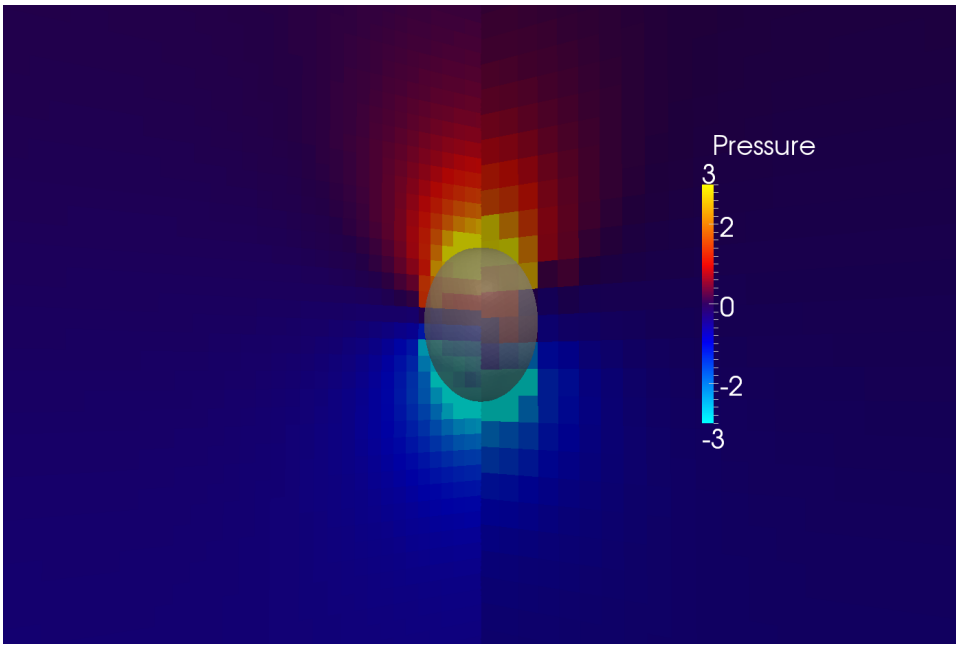
Figure 4.8.: A plot of the total drag force on a particle in a suddenly accelerated fluid phase. The cases differ by which scheme is used to interpolate fluid properties to the immersed-boundary surface, to determine the hydro-dynamic interaction force; either a constantly weighted polynomial interpolation (constant weighting) or either the variably weighted polynomial interpolation as described in section 3.3 (variable weighting). For each of the interpolation methods, three Eulerian phase resolutions are given; 2, 3 and 4 Cells Per Radius (CPR). As expected given figures 4.6 and 4.7, refinement is the only significant determinant of hydro-dynamic interaction force.

4.2.4. Moving particle

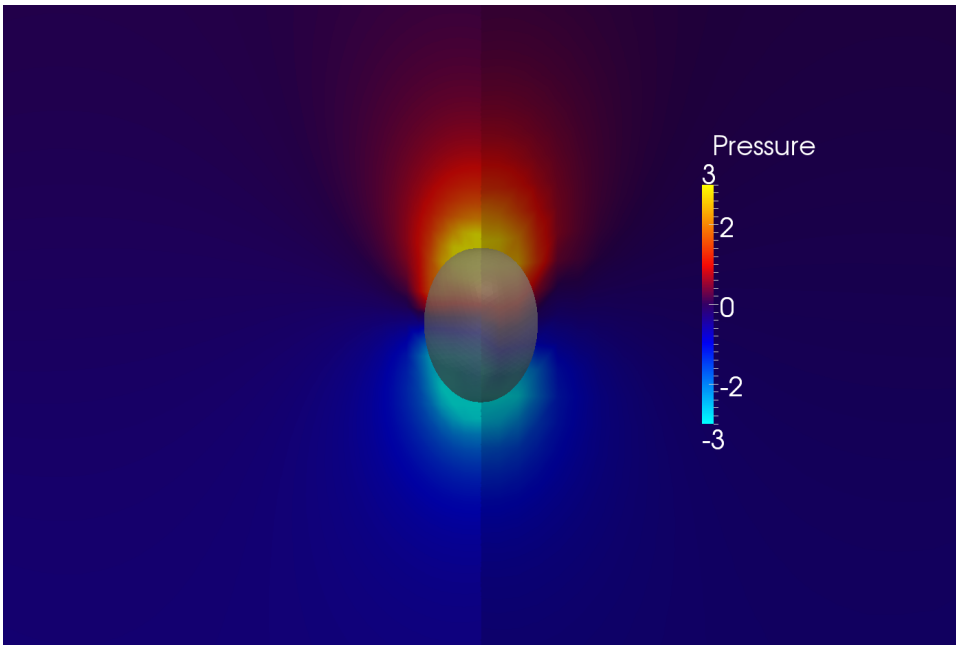
In this section moving particle simulations, the direct analog of the stationary particle case (section 4.2.3), via Galilean transformation are performed and analysed. The simulation setup is that given in section 4.2.3, with the exception that instead of the particle being held in place the particle moves relative to the fixed geometry. The flow over the particle is maintained by holding the fluid, as in the experiment. The difference between the cases can be seen as the effects of the motion of the immersed-boundary relative to the Eulerian mesh, an entirely numerical difference. The influence of the end wall boundary conditions is negligible which can be seen from the experimental results and the dynamic simulation, shown in section 4.2.6

4.2.4.1. Mesh dependence

The pressure field is often sensitive to the numerical implementation of the immersed-boundary method, as noted by a number of authors *e.g.* Schneiders *et al.* [43], especially in the case the immersed-boundary moves relative to the Eulerian mesh. Here, for two different resolutions (3 and 5 cells per radius), the same moving sphere case is run. In both cases the method used is that with the cut-cell partial mirroring method and the averaged mirroring modification. Parity based stepping is not used. Figure 4.9 shows the pressure field near the particle for both cases. Two images are show, both containing data from the high and low resolution simulations. The left image shows raw (cell-centre) data for both simulations, right the shows the interpolated field to enable a more direct comparison. The interpolation and visualisation is performed using ParaView [88, 89]. Each image is split into two, with the right half showing the lower resolution, $n/r = 3$, and the left showing the higher resolution, $n/r = 5$. Despite significant variation between individual cells due to the under-refinement of the body a very good agreement is found between the fields.



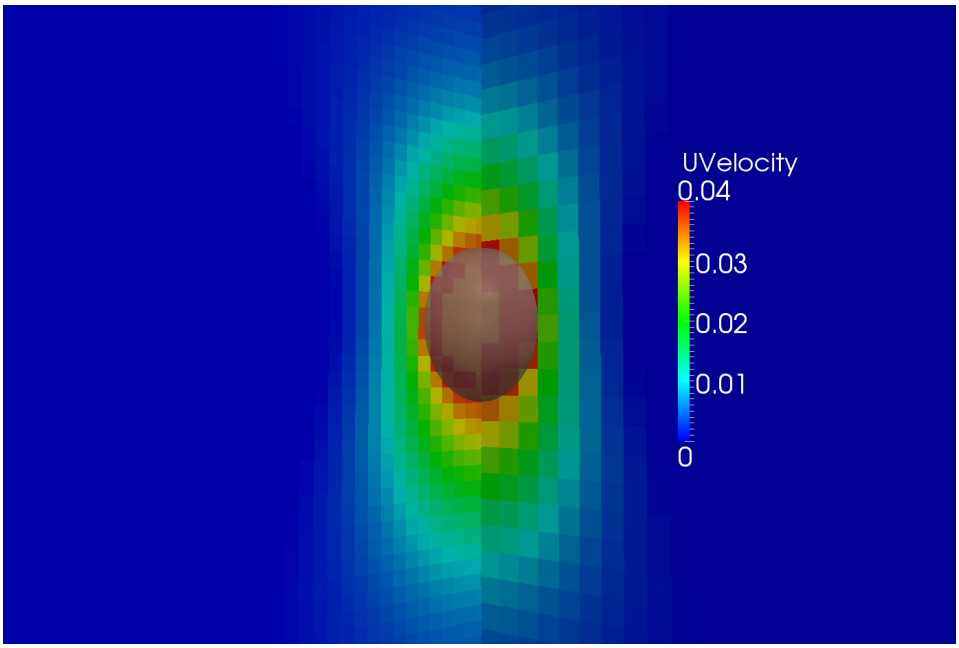
(a) Raw (cell centre) data



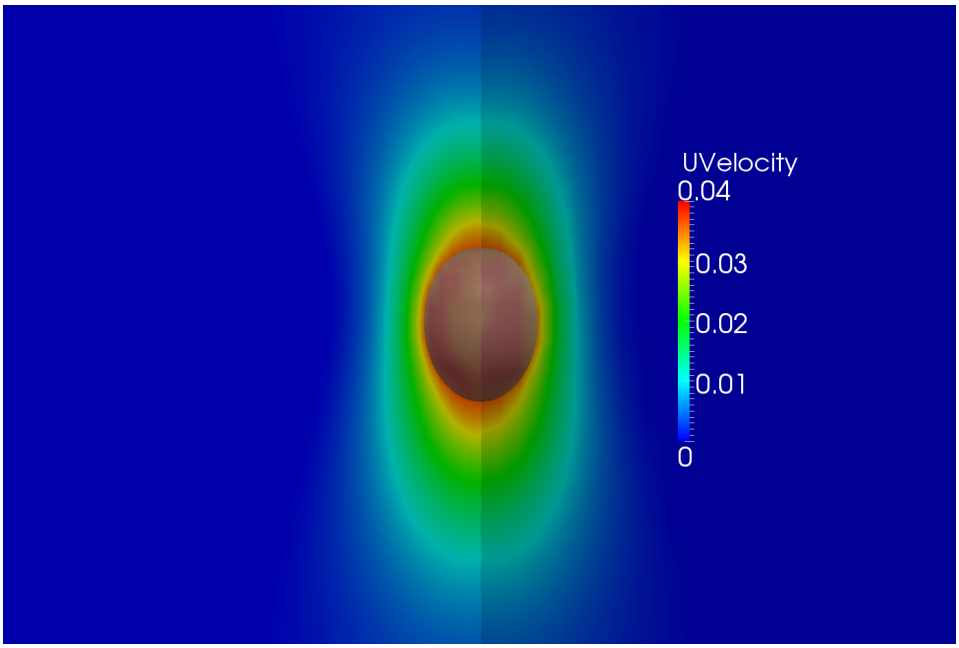
(b) Interpolated field

Figure 4.9.: The pressure field around a moving particle simulated using the developed cut-cell partial mirroring method with the averaged-mirroring modification. Results are shown for two Eulerian mesh resolutions, 3 and 5 cells per immersed-boundary radius. Both cell data (fig 4.9a) and interpolated data (4.9b) are shown. In each figure the lower resolution, 3 cells per radius, is shown on the right of the image and the higher resolution, 5 cells per radius, is shown on the left. In both cases the pressure is taken along a slice with normal orthogonal to the motion of the particle and passing through the centre of the immersed-boundary. The motion of the immersed-boundary is vertical in the images.

The velocity field is compared for the same case as pressure field. Figure 4.10 shows the velocity component parallel to the particle's motion. Figure 4.11 shows the velocity component normal to the particle's motion. Velocity components are in the plane on which they are displayed. Streamlines of the fluid phase velocity, relative to the velocity of the particle, are also shown in figure 4.11. For each of the components of velocity there is excellent symmetry between the differing resolutions, similar to the pressure field, as shown in figure 4.9.



(a) Cell data



(b) interpolated field

Figure 4.10.: The effects of refinement on velocity on a moving sphere, simulated using the developed cut-cell partial mirroring method with the averaged-mirroring modification. The figures show the component of the velocity field around a moving particle in the direction the particle is moving relative to the fluid. Results are shown for two Eulerian mesh resolutions, 3 and 5 cells per immersed-boundary radius. Both cell data (fig 4.10a) and interpolated data (4.10b) are shown. In each figure the lower resolution, 3 cells per radius, is shown on the right of the image and the higher resolution, 5 cells per radius, is shown on the left. In both cases the velocity is taken along a slice with normal orthogonal to the motion of the particle and passing through the centre of the immersed-boundary. The motion of the immersed-boundary is vertical in the images.

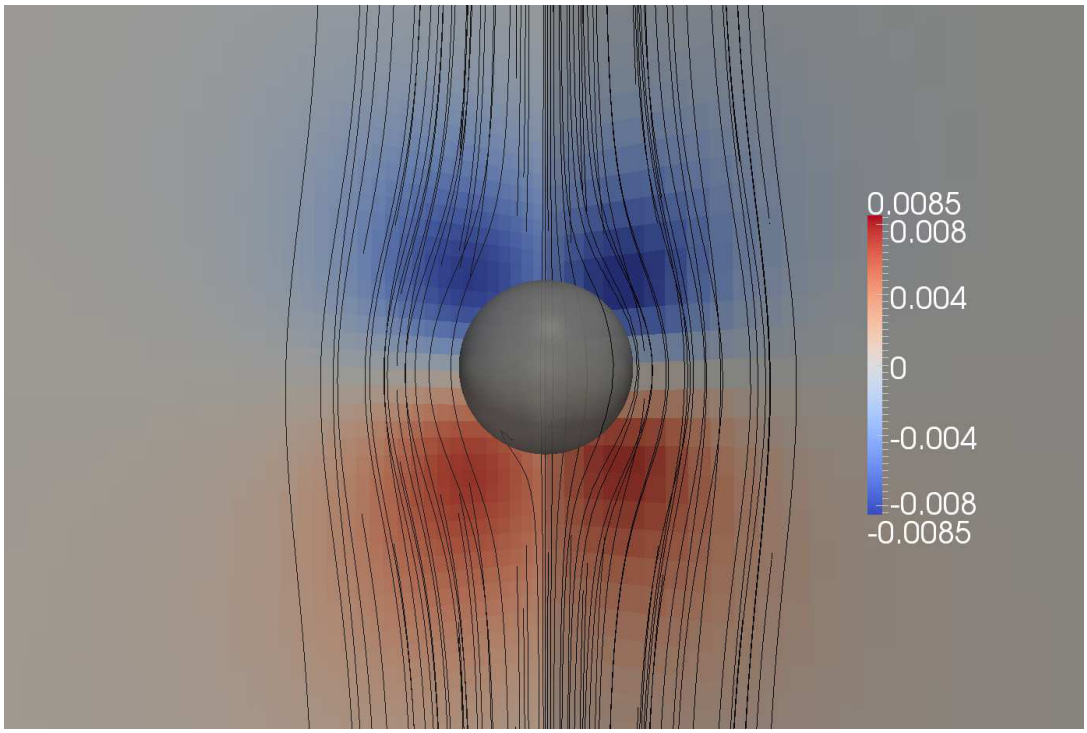


Figure 4.11.: The effects of refinement on velocity on a moving sphere, simulated using the developed cut-cell partial mirroring method with the averaged-mirroring modification. The figures show the cell centred data for two simulations, differing only by refinement of the Eulerian mesh used, 3 and 5 cells per immersed-boundary radius. The image is split in two, with the lower resolution, 3 cells per radius, shown on the right of the image and the higher resolution, 5 cells per radius, is shown on the left. Streamlines are added for each half, calculated from the relative velocity of the fluid phase and the immersed-boundary. The velocity is taken along a slice with normal orthogonal to the motion of the particle. The slices do not pass through the centre of the immersed-boundary, but are separated by 1.5 particle radii. The motion of the immersed-boundary is vertical in the image.

4.2.4.2. Force calculation

In order to determine the accuracy of the force prediction for moving particles a series of simulations are performed using a setup similar to that given if figure 4.2. Instead of fixing the particle and moving the flow, the particle is forced with a constant velocity through a stationary fluid. The premise of the test is that by forcing the particle to maintain terminal velocity, the resistive force must meet the gravitational force required to reach terminal velocity *i.e.* $f = mg$. In the current method, gravity is not applied to the fluid phase and is accounted by a correction to the gravity force on the bodies as detailed in section 4.2.2.2. Hence the force acting on the particle is only expected to balance the gravity force less the buoyancy force i.e.: $f = gV(\rho_p - \rho_f)$. In the E1 case $f = 9.81\frac{4}{3}\pi 0.0075^3\ (1120-150) \approx 0.00260$ N . Figure 4.12 shows a the development of force acting on the immersed-boundary in the described case at high refinement (*i.e.* 5 cells per radius). The method used is the cut-cell partial mirroring method, with the averaged mirroring modification. Figure 4.13 shows the relative error from each of the refinement levels run, on a log-log scale. Convergence is approximately second order, the gradient of the fitted line is -2.30 *i.e.* relative error, e, relates to refinement, n, by $e \propto n^{-2.30}$.

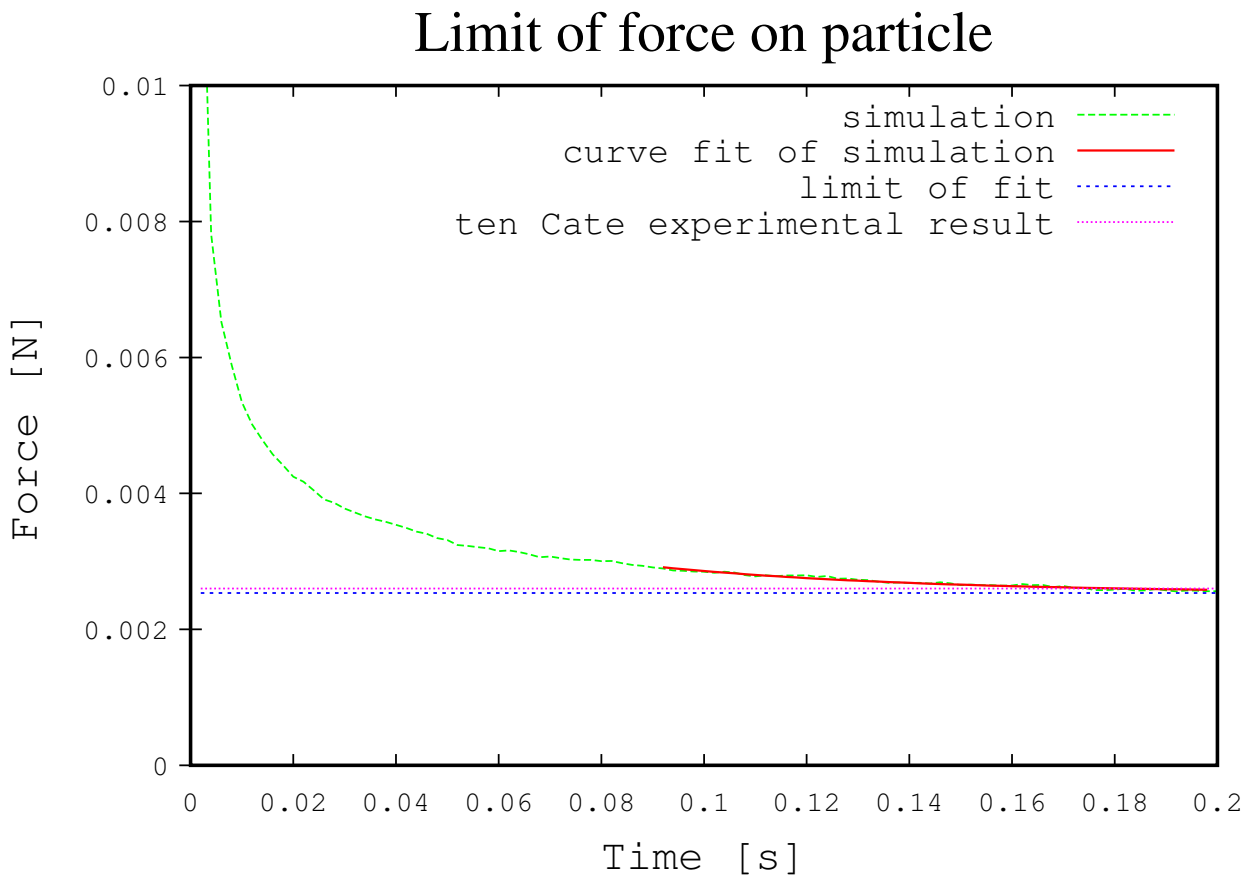


Figure 4.12.: A plot of the components of forces between the particle and fluid, compared with experimental results from ten Cate *et al.* [86]. The drag force [simulation] is fitted to an exponential decay curve [curve fit of simulation], using data over the region plotted. The limit of this exponential decay [limit of fit] is shown along with experimental predictions of ten Cate *et al.* [ten Cate experimental result].

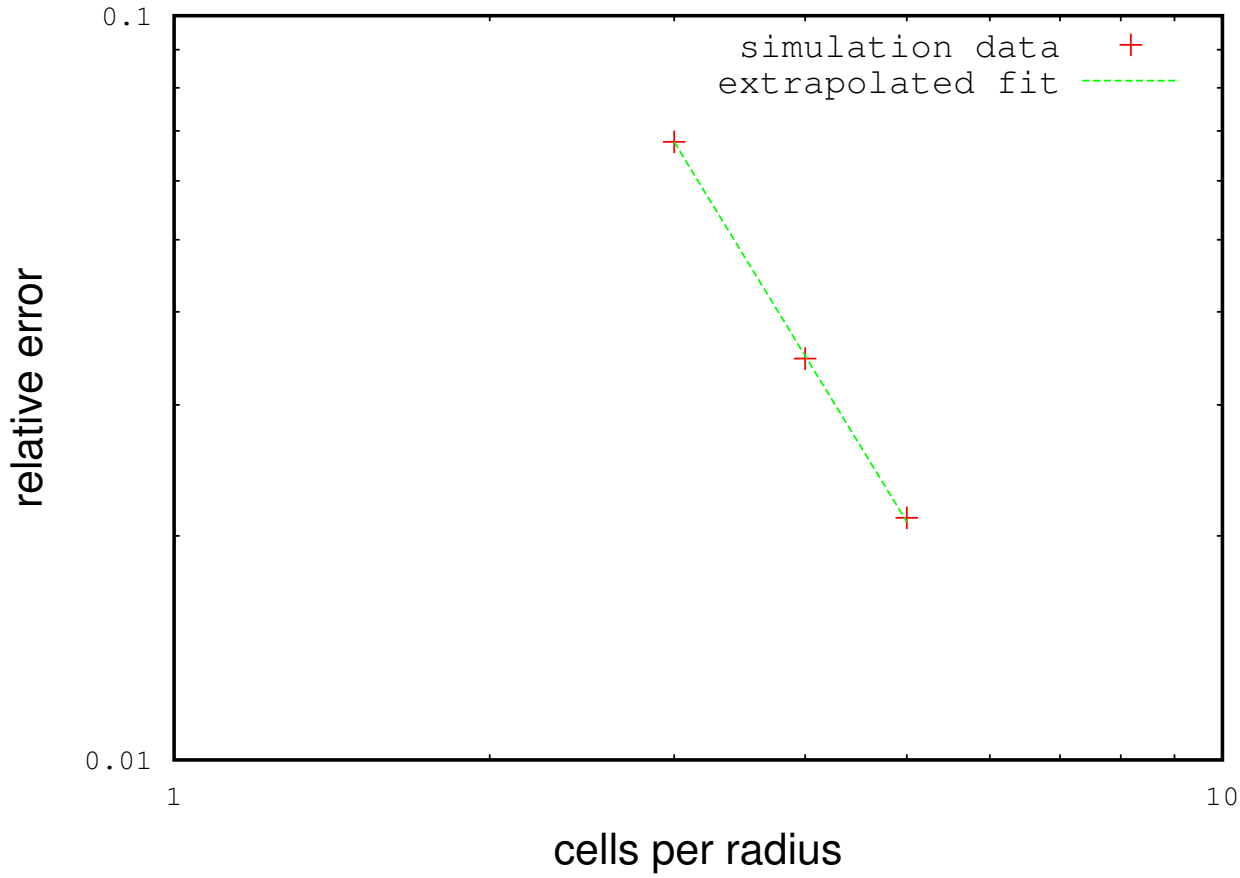


Figure 4.13.: A plot of the convergence of forces between the particle and fluid with refinement, based on the method described in section 4.2.4.2. Using the fit shown, the error is given by: $0.846n^{-2.30}$.

4.2.4.3. Smoothness

Unlike when the particle is held stationary relative to the mesh, in the case of the a moving particle the discretisation scheme is the cause of significant influence on the results. In particular the smoothness of the pressure field in time is problematic, *i.e.* it un-physically oscillates over time.

Figures 4.14 and 4.15 show the forces acting on a the particle moving through the fluid simulated with variations of the partial mirroring method. The variations are:

- The partial mirroring method without additional modifications, [Partial mirroring].
- The partial mirroring method with the addition of the averaged mirroring modification, [Averaged mirroring].
- The partial mirroring method with the immersed-boundary incrementing every other timestep (as described in section 3.2.3.4), [Parity based stepping].

The stationary particle equivalent, [Stationary equivalent], is simulated using the partial mirroring method without additional modifications. Figure 4.14 shows the simulation run with 4 cells per radius and figure 4.15 with 5 cells per radius. The CFL number for the high resolution case is approximately 3.52e−2 and 2.93e−2 for the medium resolution case. In each case the results are shown with the average mirroring modification and with parity based incrementation modification as

described in section 3.2.3.4. The force, calculated on an equivalent stationary body where the flow is accelerated, is also shown on these figures. Section 4.2.5 explains the significance of this comparison. Good agreement between the methods is found and the oscillations in the force are small. The good agreement is positive given the low CFL number as low CFL numbers tend to be associated with larger oscillations [64]. A slight trend is seen where by the position of the particle relative to the mesh determines the force, however these variations are small ($\sim 1\%$), even in the lower-refined case.

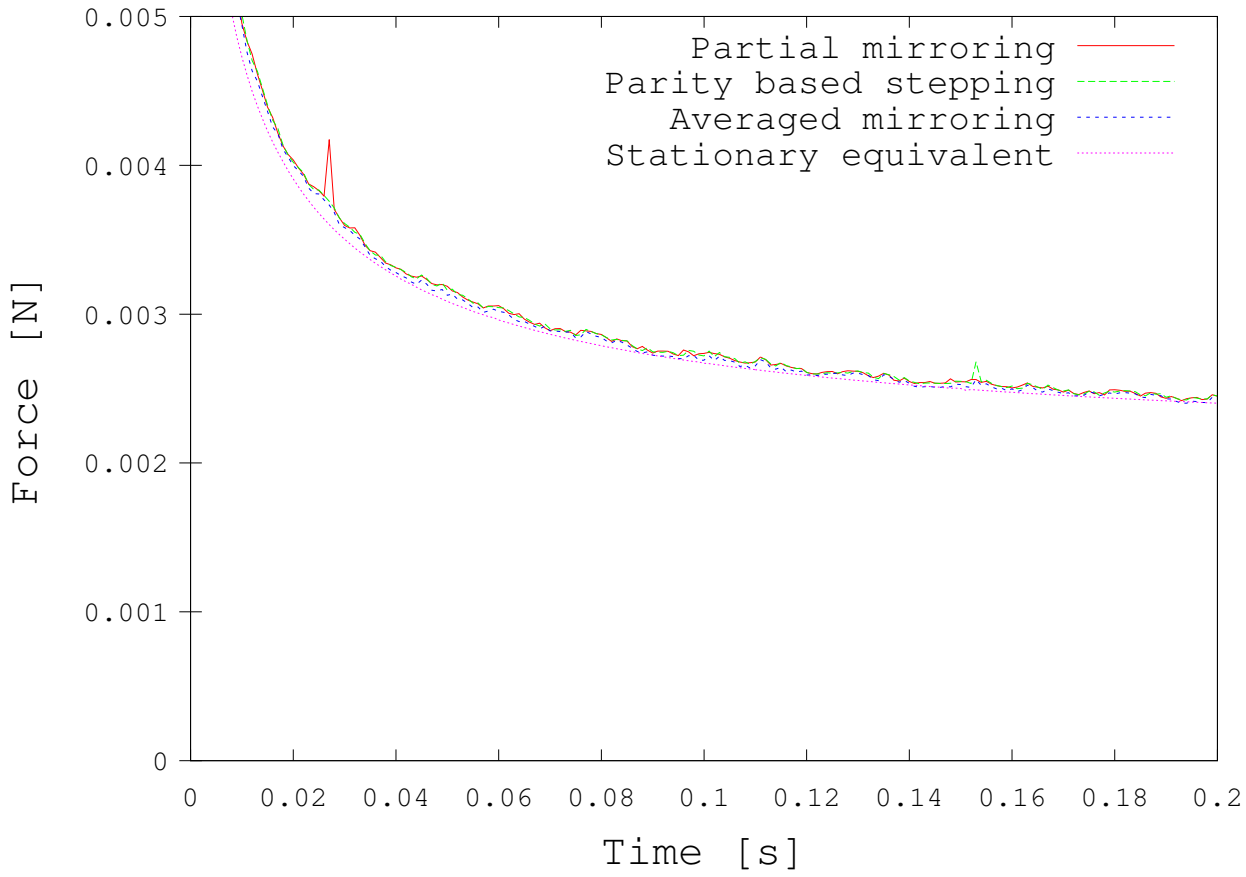


Figure 4.14.: Force acting on a moving particle, calculated with the partial mirroring method with 4 cells per radius refinement

Figures 4.16 and 4.17 show the forces acting on a the particle moving through the fluid simulated with variations of the complete mirroring method. The variations are:

- The complete mirroring method without additional modifications, [Partial mirroring].
- The complete mirroring method with the addition of the averaged mirroring modification, [Averaged mirroring].
- The complete mirroring method with the immersed-boundary incrementing every other timestep (as described in section 3.2.3.4), [Parity based stepping].

The stationary particle equivalent, [Stationary equivalent], is simulated using the partial mirroring method without additional modifications. Figure 4.16 shows the simulation run with 4 cells per radius and figure 4.17 with 5 cells per radius respectively. For these methods, significant oscillations in the force acting on the particles are seen. Significant improvement is made with the addition of

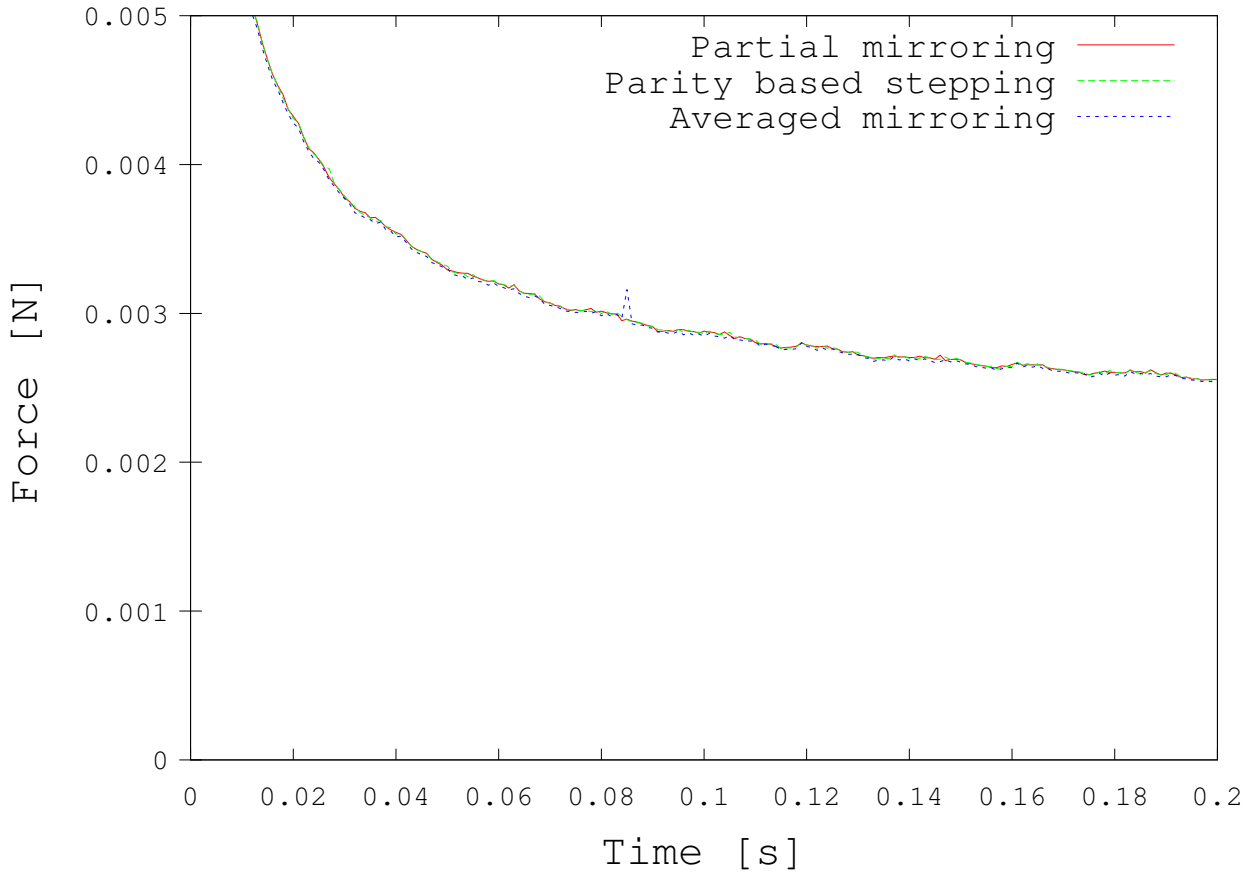


Figure 4.15.: Force acting on a moving particle, calculated with the partial mirroring method with 5 cells per radius refinement

either the averaged mirroring method, or the parity based incrementation method.
Oscillations in the force are significant, deviating by ±15% from the stationary equivalent, for the unmodified medium refinement case. Despite the significant peak differences between the force acting on the stationary and moving particles, the overall agreement is quite good.

Figure 4.18 shows the force predicted by sphere moving through a viscous fluid, as simulated using variations of partial mirroring method developed. The variations use the cut-cell momentum equation and differ by their inclusion of the average mirroring modification and the parity based incrementation technique, for which each combination is presented. The convergence to the experimental data is shown for the variation of average mirroring and parity based incrementation, as discussed in section 4.2.4.2. Figure 4.19 shows an enlarged section of the same data.
The convergence to the experimentally predicted force of ten Cate *et al.* [86] of the variation using average mirroring and parity based incrementation is very similar to the equivalent without the cut-cell momentum equation, both being around 2% under predicted at 5 cells per radius. Similar analysis for the other variations also yield similar force predictions. A slight improvement is seen from the use of averaged mirroring, however this is not significant. The results show little difference is obtained using the parity based incrementation technique on a case already using partial mirroring and averaged mirroring. Hence, for moving cases, this is not used as it presented an additional constraint on CFL number, inparticular that it must be less than 0.5 else a cell could transition two 'types' in a single timestep. Even if the constraint of requiring the maximum CFL number

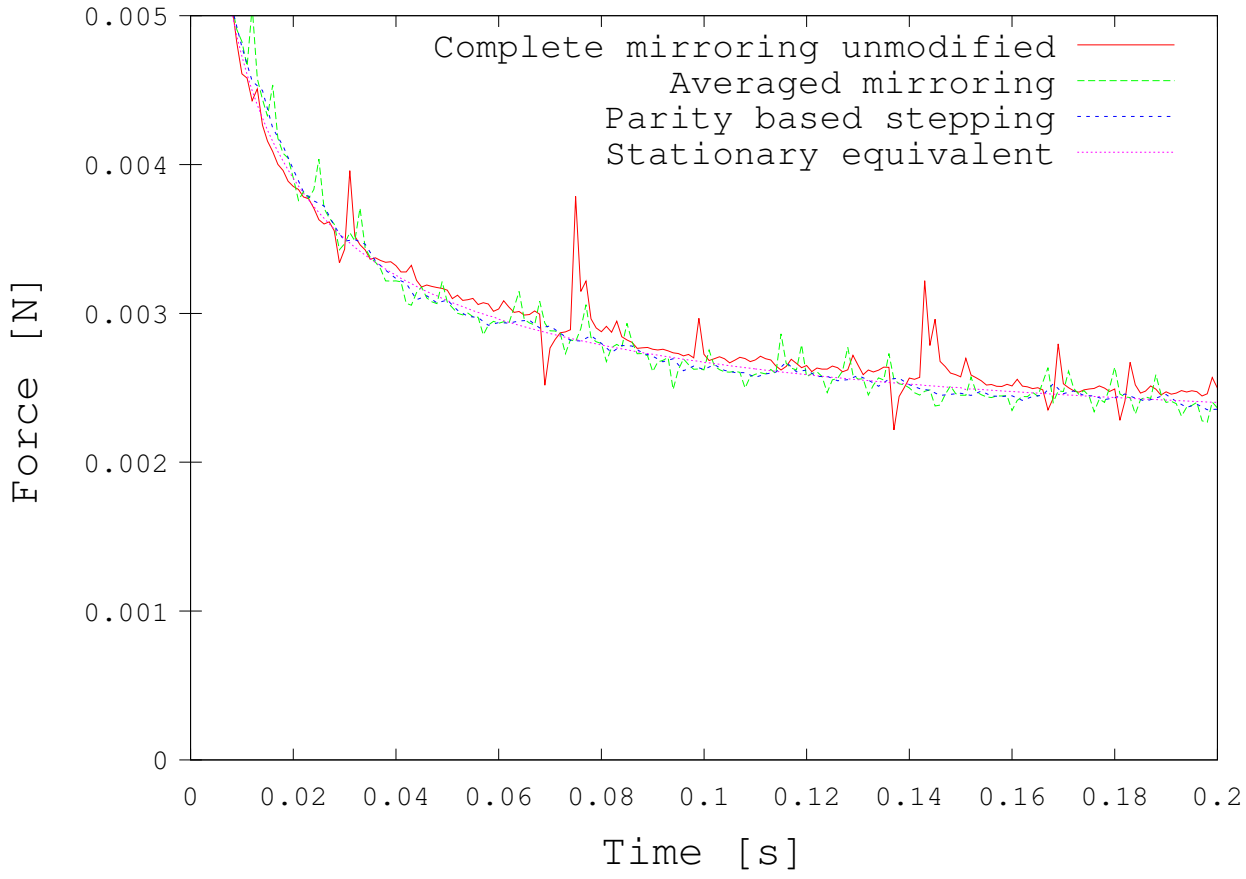


Figure 4.16.: Force acting on a moving particle, calculated with the complete mirroring method with 4 cells per radius refinement

to be less than 0.5 is satisfied, the use of parity based incrementation is not recommended for high CFL numbers. The reason for this is that parity based incrementation effects the accuracy of the interpolation of the fluid properties to the immersed-boundarydependent of the CFL number. Having an interpolation point move by a small fraction of the stencil (cell-spacing) lengthscale is substantially less significant than being out by a larger fraction of the cell-spacing. The particle CFL number directly measures this. The major effect of the inaccuracy of the position of the particle is to change the interpolation coefficients associated with the mirroring equation. This gives some insight into the impact of moving the position of the particle, namely that it is equivalent to changing the interpolation coefficients. Changes in weighting for the interpolation, which do not significantly impact the method, give rise to changes to interpolation coefficients of roughly the same magnitude as moving the interpolation point a few percent of the separation of the data points. On this basis, it is recommended that parity based incrementation is advantageous for particle CFL numbers less than 0.05 but not if the particle CFL number is greater than 0.2. It is noteworthy that the use of parity based incrementation does not have to be applied globally but can be applied selectively to individual particles. However if parity based incrementation is applied selectively, the 'missed' incrementations should occur in tandem.

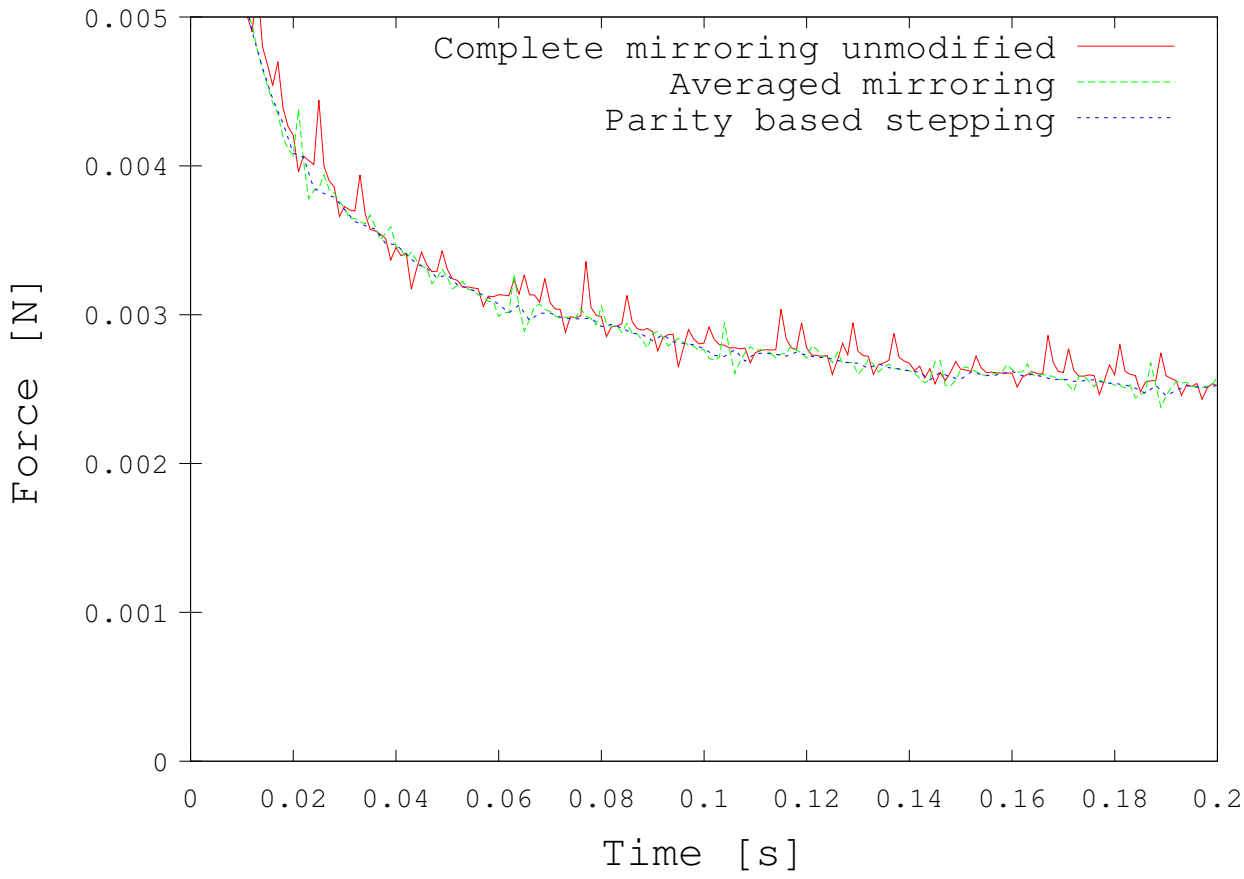


Figure 4.17.: Force acting on a moving particle, calculated with the complete mirroring method with 5 cells per radius refinement.

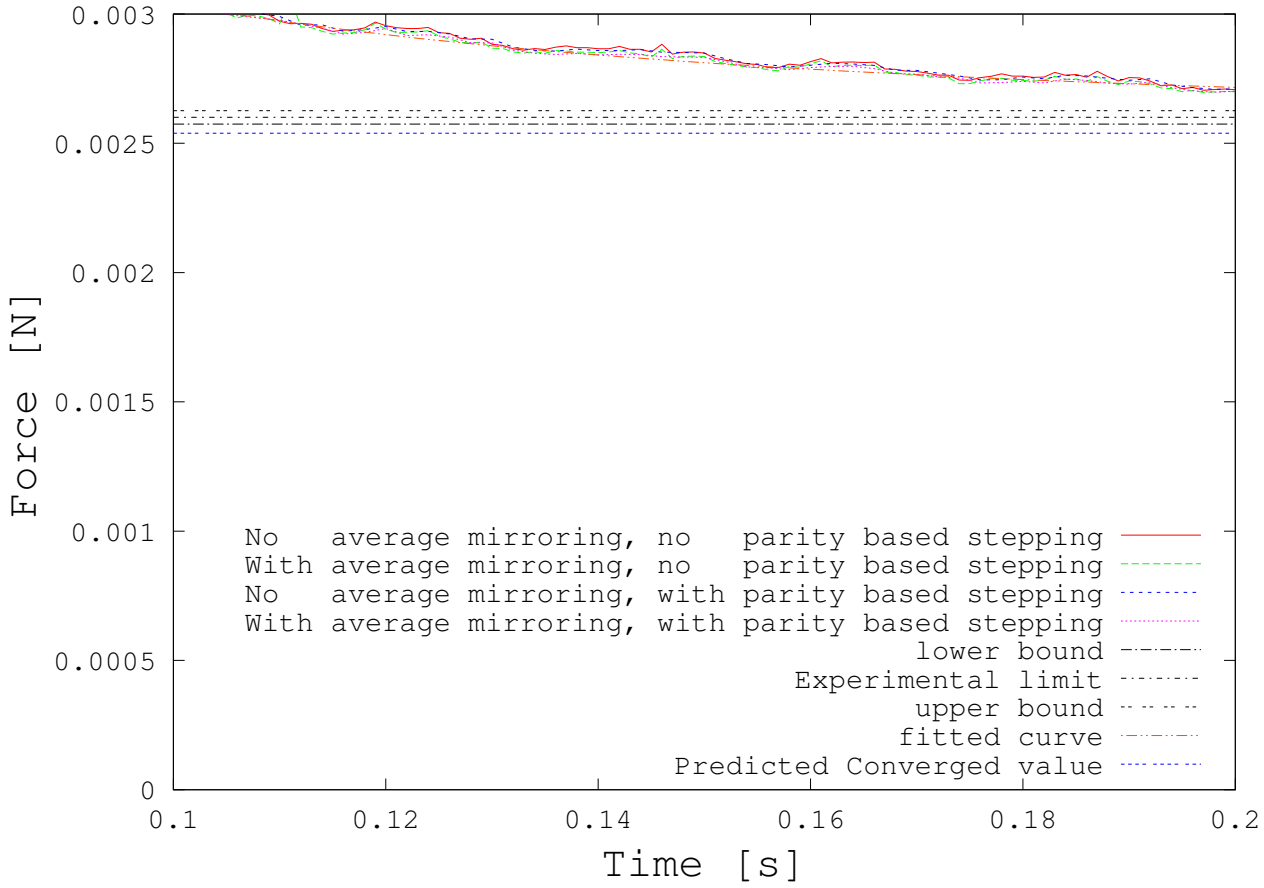


Figure 4.18.: The partial mirroring method and cut-cell momentum equation, with the different combinations of averaged mirroring and parity based incrementation are used to simulate a sphere moving at fixed velocity. For each case, the resultant force, calculated as described in section 3.3 is plotted. Similar behaviour is found for each of the methods tested. A curve is fitted through the plot of force resulting from the use of both average mirroring and parity based incrementation. The limit of this fit is plotted [Predicted Converged value] along with the predicted result from experimental data (ten Cate *et al.* [86]) [Experimental limit]. An estimation of the experimental error is given [lower bound] and [upper bound]. An enlarged section of this plot is shown in figure 4.19 which shows the small variations between the methods.

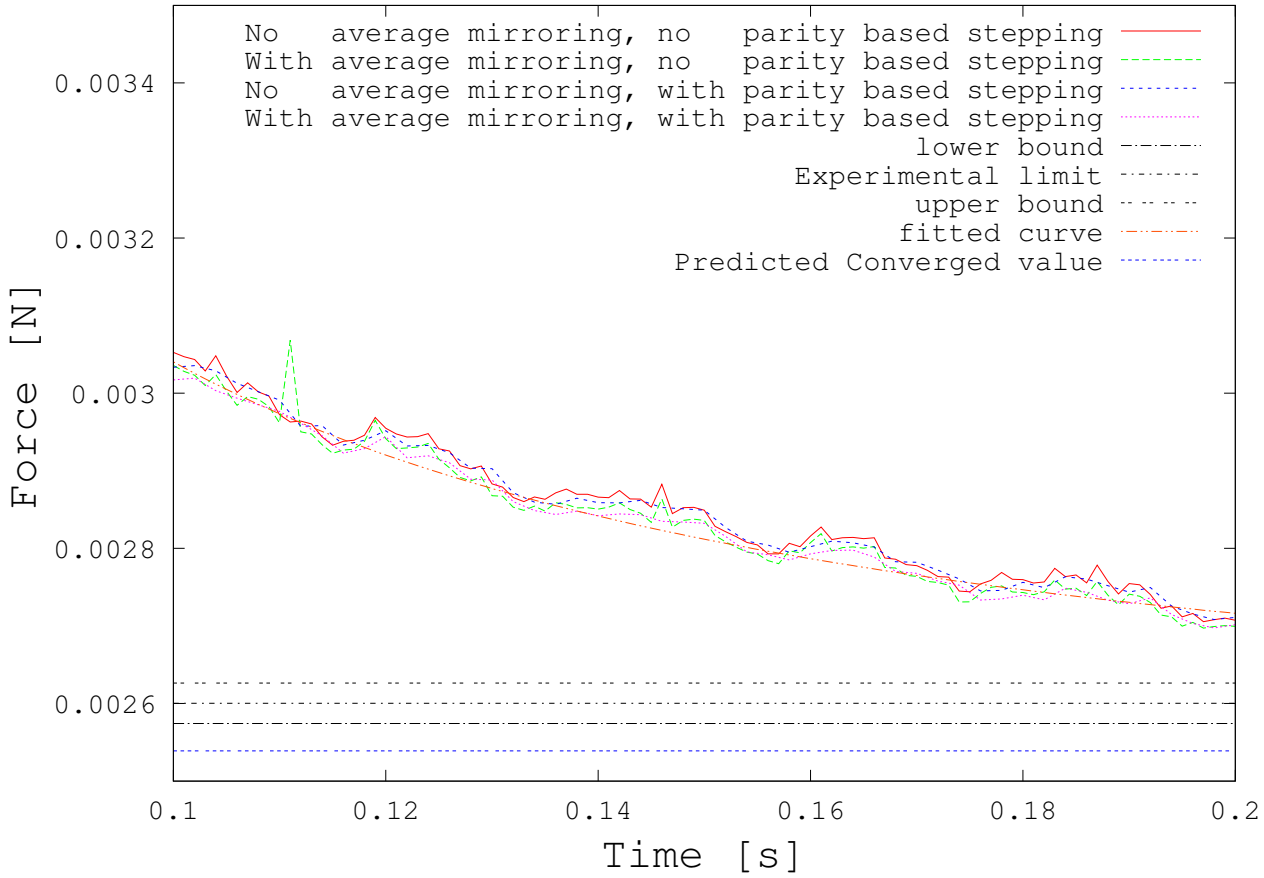


Figure 4.19.: An enlarged section of figure 4.18; the plot of force acting on a moving immersed-boundary for a variation of different modifications to the partial mirroring method.

Figure 4.20 shows a comparison of the developed partial mirroring method with cut-cell momentum equation and the averaged mirroring modification but without parity based incrementation and the complete mirroring method with both the parity based incrementation technique and averaged mirroring. The CFL number for this case is approximately 2e−2. Despite the improvement to the method there are still significantly greater fluctuations in the pressure field. However, given that the oscillations are less than 5% most of the time, the method is not untenable and makes a noticeable computational expense saving by not having to calculate the intersection of the triangles of the immersed-boundary with the Eulerian cell faces. The more problematic issue at this refinement is the overall inaccuracy of the force, as discussed in section 4.2.3.2, for which little dependence on the method is found.

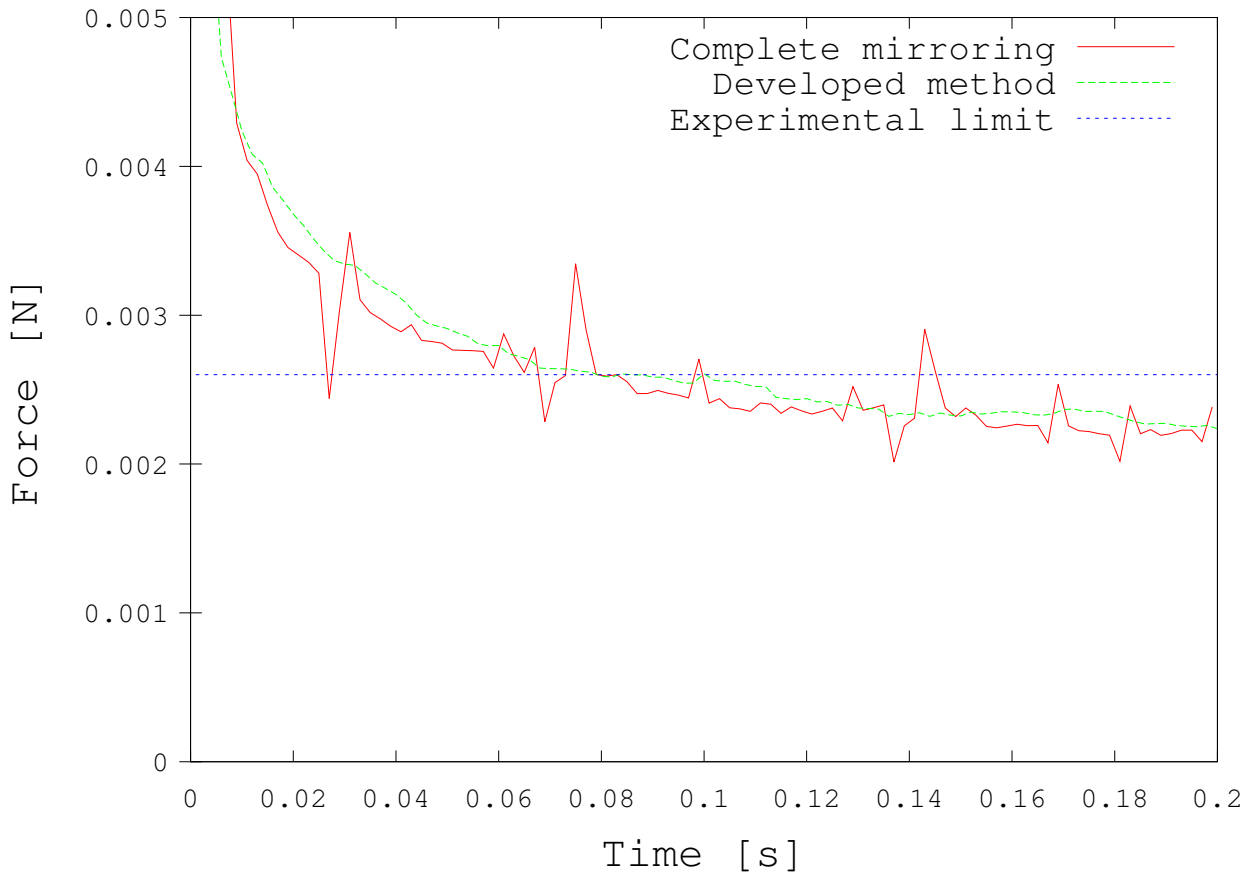


Figure 4.20.: Comparison of methods at at low resolution. The force on a moving particle is shown as simulated with two methods; the complete mirroring method with both the parity based incrementation technique and averaged mirroring [Complete mirroring] and the developed partial mirroring method with cut-cell momentum equation and the averaged mirroring modification but without parity based incrementation [Developed method]. In both cases the Eulerian grid refinement is such that there are 3 cells per particle radius with CFL $\approx$ 2e−2.

Case	E1
Refinement	4 cells per radius
Max CFL number	∼0.03

Table 4.5.: Details of simulation setup for drag force comparison, as shown in figure 4.21.

4.2.5. Comparing stationary and moving cases

One of the problems often encountered in the immersed-boundary framework is the dependency on the movement of the immersed-boundary relative to the Eulerian mesh. In particular that the flow field in the near vicinity of a moving body is subject to spurious oscillations. Figure 4.21 shows the force resulting from simulating case E1, as detailed in table 4.5, using the developed cut-cell partial mirroring method with the averaged mirroring modification. Results are shown for a moving sphere, compared to that of a the same method for the stationary particle. An excellent agreement is found between the stationary and moving cases. This agreement is thought to be one of the principle advantages of the developed method.

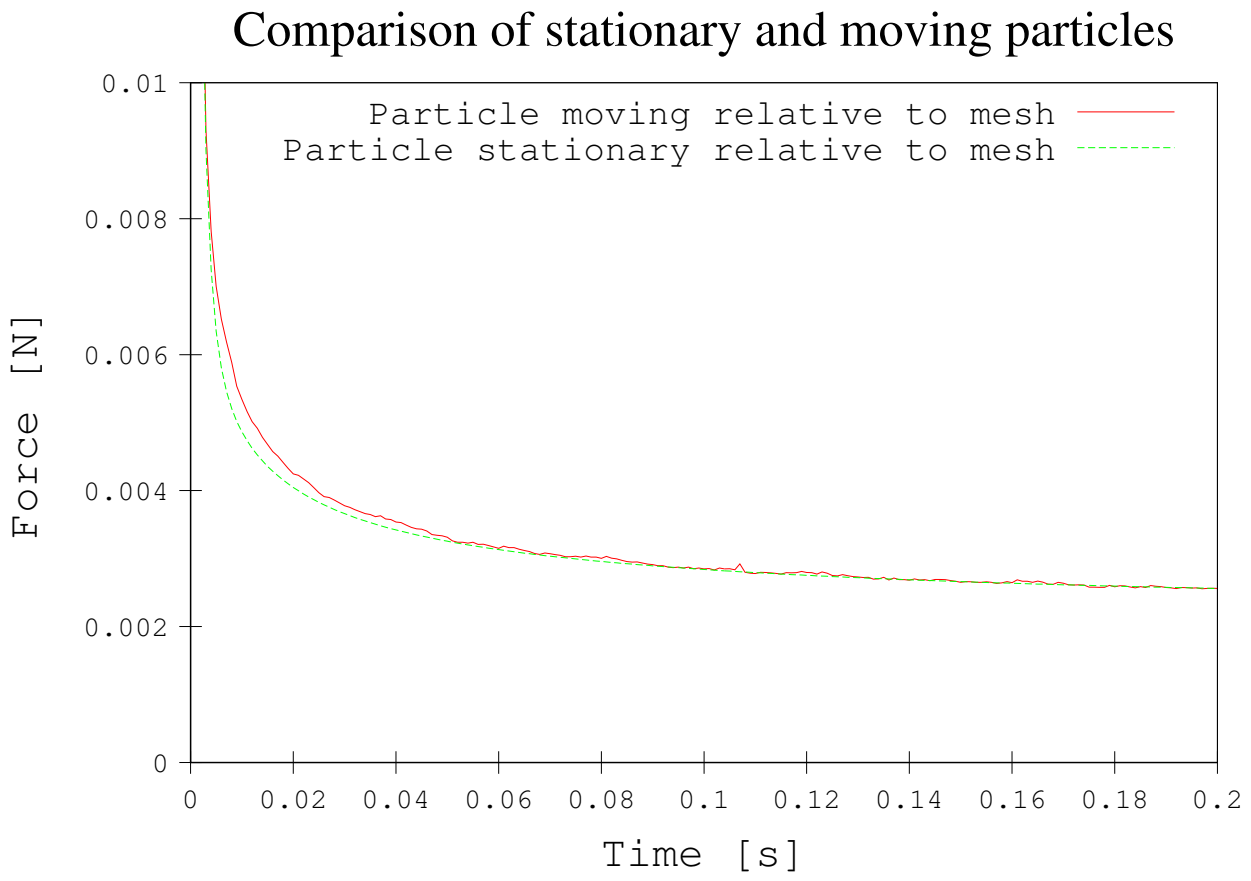


Figure 4.21.: A plot of the components of forces between the particle and fluid phases. The two cases differ by whether the immersed-boundary moves relative to the mesh.

4.2.6. Dynamic case

Figure 4.22 shows the falling velocities over time of simulations using setups E1, E2, E3 and E4 of the experiments performed by ten Cate *et al.* [86], as detailed in table 4.1. The experimental data is

plotted on the same graph. The method of simulation in each case is the cut-cell partial mirroring method with the averaged mirroring modification. In each case, medium refinement is used (*i.e.* 4 cells per radius) and the CFL number is approximately 0.1. The correction to the lubrication force, developed in chapter 5, is used. Data is plotted until the particle has come to rest for cases E1, E2 and E3 as the particle in these cases do no incur a significant collision prior to coming to rest. Case E4 is plotted up to the point at which collision with the bottom of the tank occurs. Data after the collision with the bottom wall is omitted as it depends on the compliance of the collision of the sphere and the tank, for which details are not known.

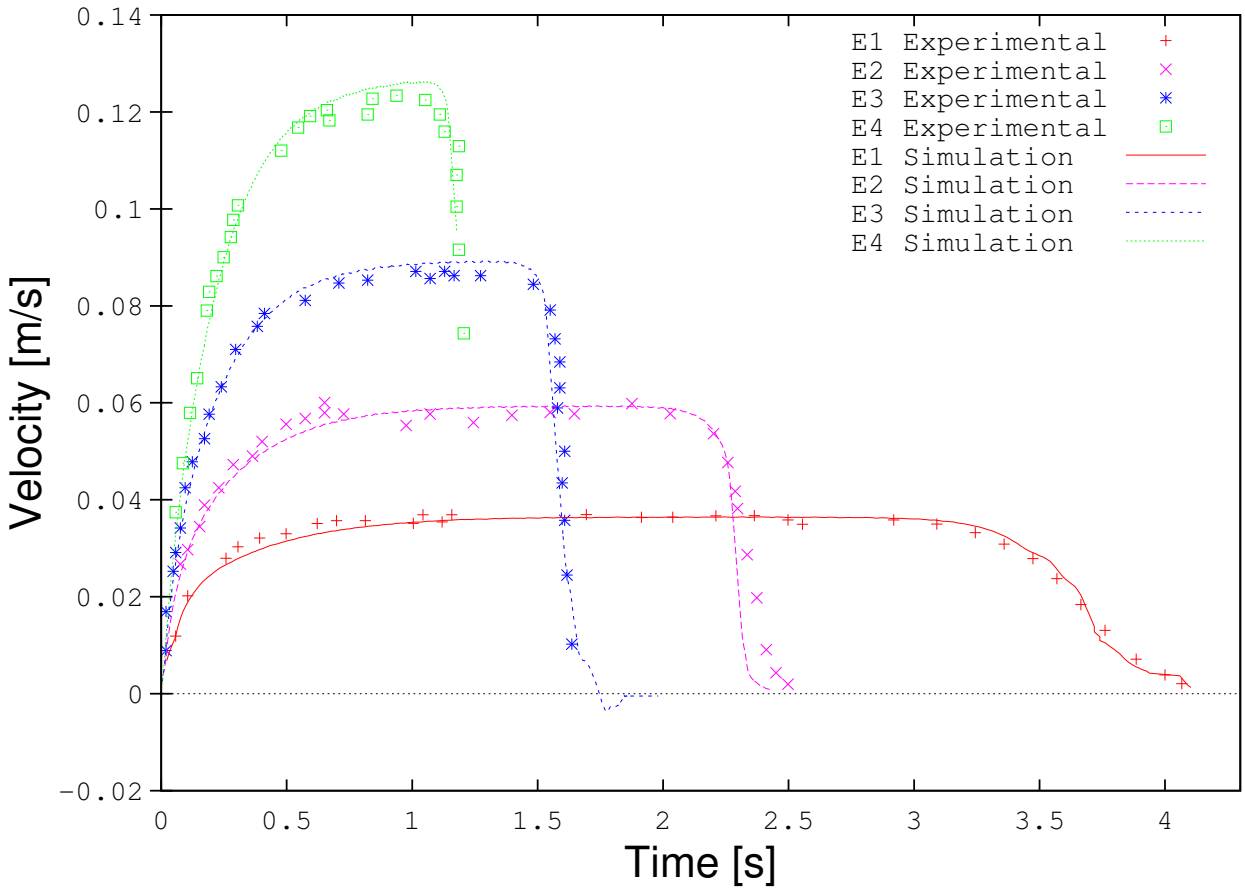


Figure 4.22.: Plot of the velocities of each of four cases (E1, E2, E3 and E4) of particles settling though a viscous fluid. Results are plotted for experimental results of ten Cate *et al.* [86] and simulation using the developed cut-cell partial mirroring method with the averaged mirroring modification. Data is plotted until the particle has come to rest for each case except case E4, which is plotted up to the point at which collision with the bottom of the tank occurs.

Figure 4.23 shows the force acting on the particle for the E3 case. The force acting on the immersed-boundary during the final deceleration phase is composed mostly of the fluid interaction however a very slight impact occurs. This differed from cases E1 and E2 where the velocity is negligible during the final approach. Case E4 gave rise to a significant impact in the absence of a lubrication correction and a less significant impact with the lubrication force correction method, depending on the implementation factors such as which radius is taken for the discretised sphere.

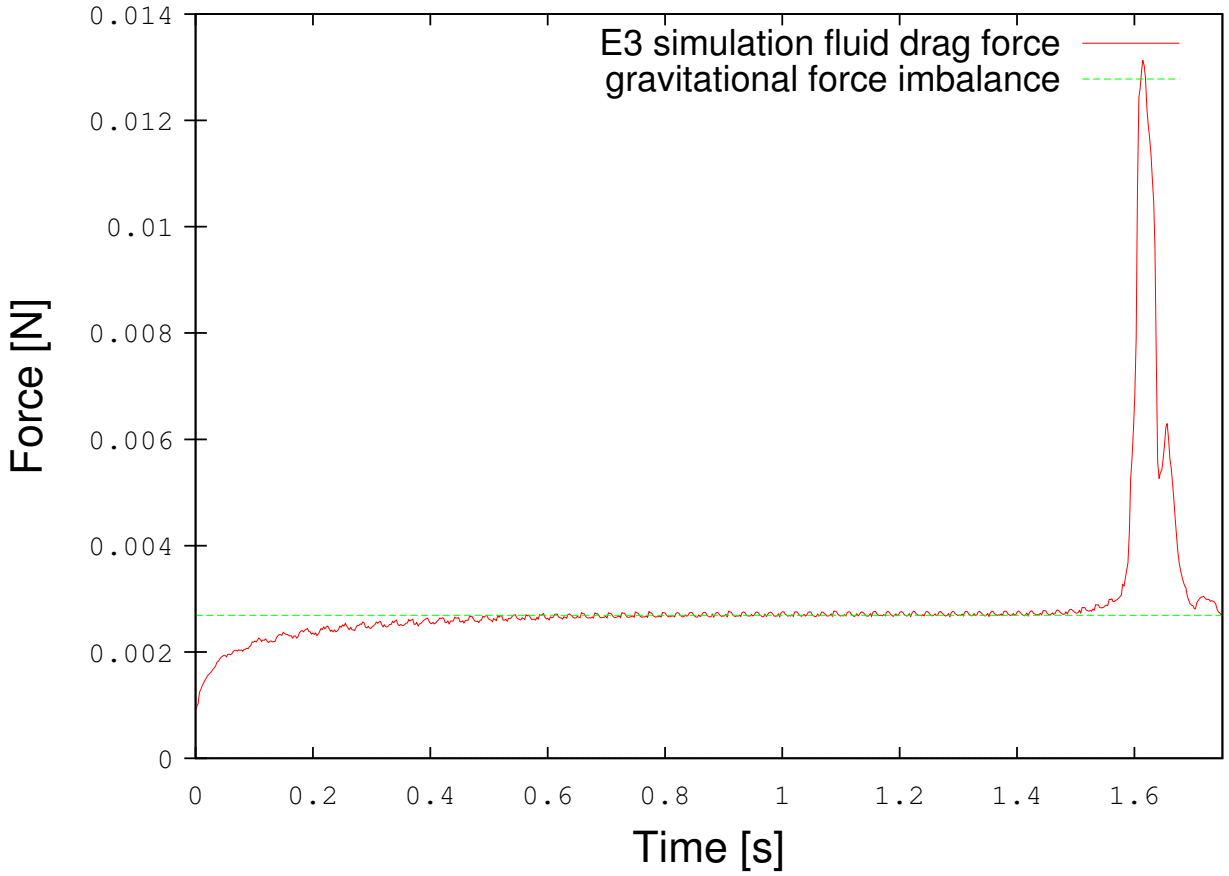


Figure 4.23.: Plot of the fluid drag force during the E3 simulation using the cut-cell partial mirroring method with the averaged mirroring modification. On the same plot is the gravitational force imbalance used to drive the particle's motion.

4.3. Oscillating sphere

A further investigation into the effects of the cut-cell partial mirroring method on the flow around a particle is determined by simulating a sinusoidally oscillating sphere. Figure 4.24 illustrates the setup in which a single particle is given a fixed sinusoidal velocity. For each method, two Eulerian meshes are used. One is refined to three cells per radius [Low res], the other is refined to four cells per radius [Med res]. The timestep varies between the meshes, with the coarser having a coarser timestep, such that the maximum particle CFL number is approximately the same in all cases: 7e-3. The particle moves by a quarter of its radius either side of its central position. The sphere oscillates inside a square prism, with all walls at least four particle-diameters removed from the particle. Mei [90] details theoretical work relating to the expected force acting on the moving particle and concludes a sinusoidal drag force is expected given the low Reynolds number (≤ 200) of the simulation. Figure 4.25 shows the response force of the developed method and a fully mirrored equivalent, also using the averaged mirroring technique. For each of the methods a sinusoidal curve is fitted through the force acting on the particle.

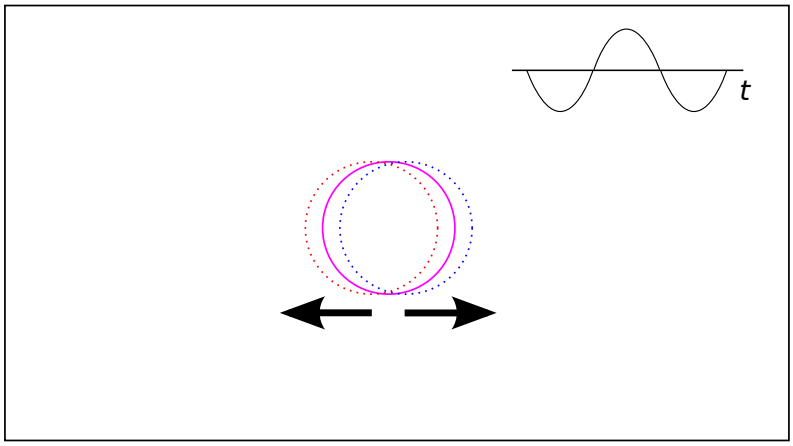


Figure 4.24.: Illustration of the setup used to to calculate the force on an oscillating sphere. The position and velocity are prescribed to be that a sinusoidally moving particle.

While generally good agreement is found in both methods at both refinements, with very similarly parametrised fitted sinusoidal curves, significantly less deviation from the sinusoidal fit is observed in the developed method. The formation of 'fresh' and 'dead' cells causes some deviation from the predicted force in the partial mirroring method however these deviations are small. A few notable exceptions occur, at a similar point oscillatory phase, corresponding to the formation of numerous 'fresh' and 'dead' cells in rapid succession. The deviation is much larger when predicted using the complete mirroring method. The deviations are frequently around 50% of the expected peak force and occasionally significantly greater. At one time instance, around 0.1 seconds, the error is around 200% of the peak force. Additionally greater dependency on the resolution is found on the overall force. While the partial mirroring method predicts nearly identical force magnitudes at both resolutions, some dependence on resolution is found for the complete mirroring method.

Very good agreement between the forces at the two refinements, despite the refinement being a significant influence on the force for the steady case, is notable. This likely due to the fact that pressure terms are comprise a large fraction of the hydro-dynamic interaction force acting on the body. Figure 4.26 shows the split of the force prediction for the partial mirroring method at both

refinements into pressure force and viscous shear force. The pressure force has been found not to vary significantly with refinement in the developed method, see figures 4.8, 4.6 and 4.7. This corroborates that the calculation of the velocity gradient, hence viscous shear force, at the immersed-boundary particle surface is the major source of forcing error of the partial mirroring method.

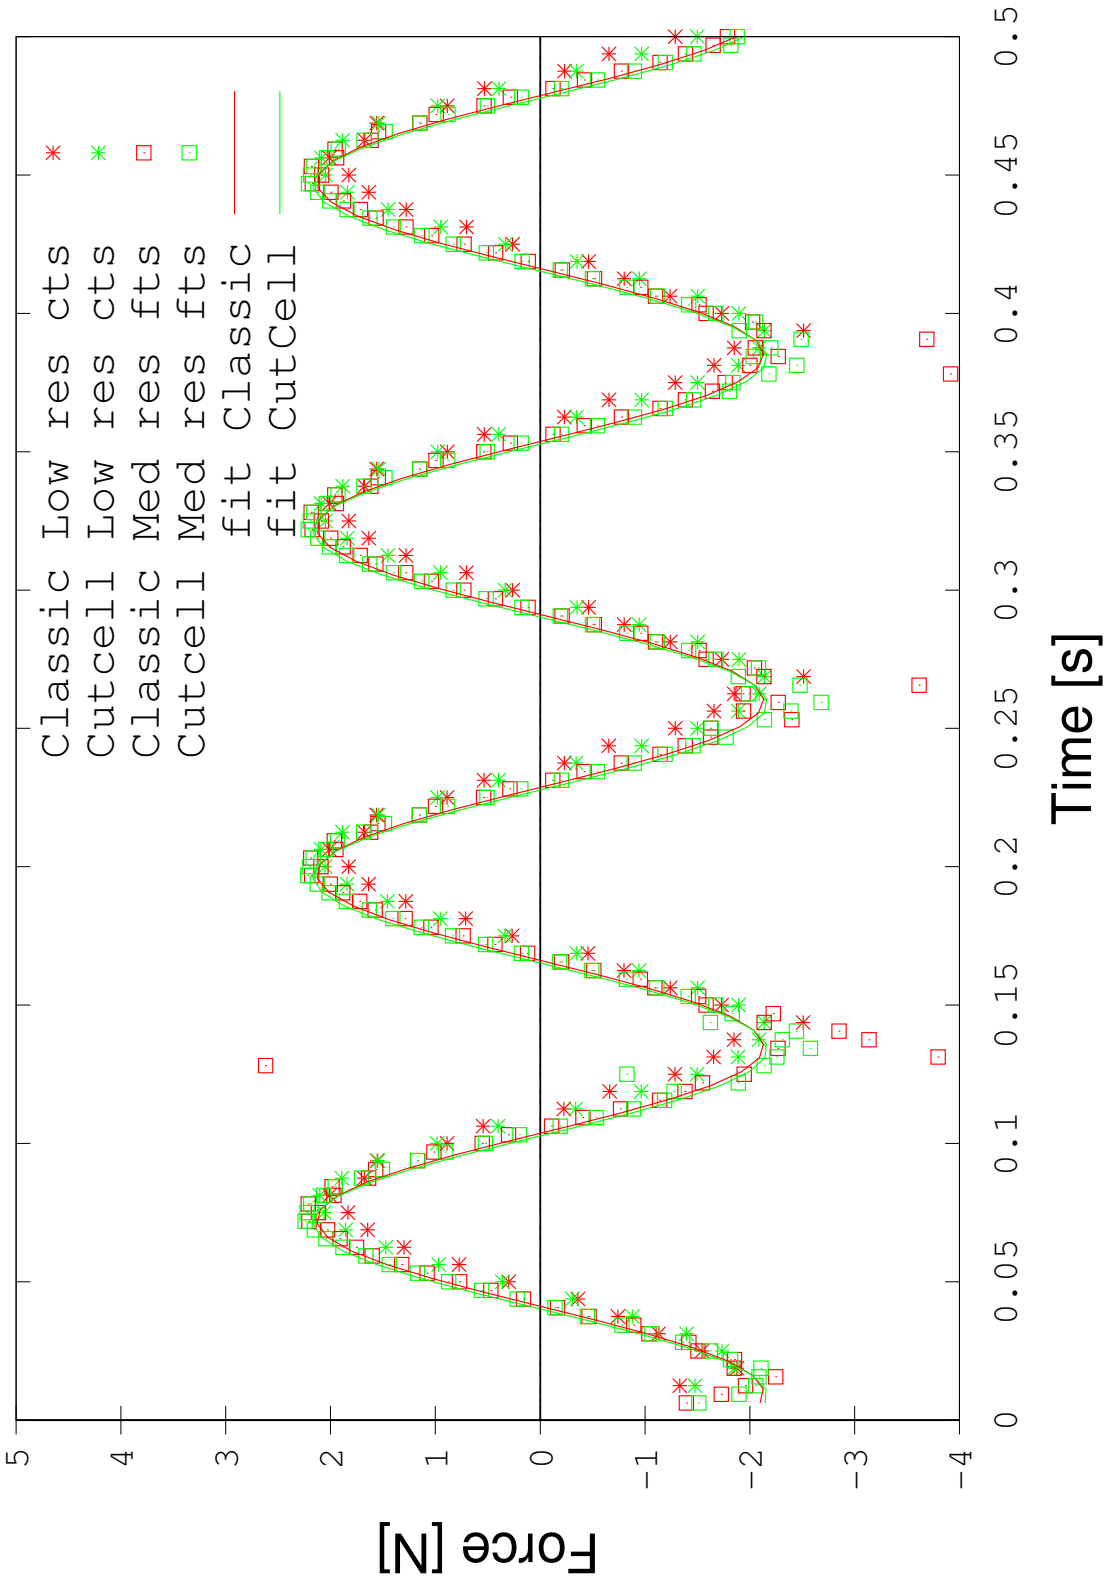


Figure 4.25.: Graph showing the force acting on a sphere forced at a sinusoidal velocity. Two simulation methods are used: complete mirroring and cut-cell partial mirroring, both with the averaged mirroring modification. For each method, results are presented for two different resolutions: a low resolution case with 3 cells per radius and a higher resolution case at 4 cells per radius. The timestep for the coarser mesh is also coarse to maintain the same CFL number. A sinusoidal fit though each of the two higher resolution cases is also shown.

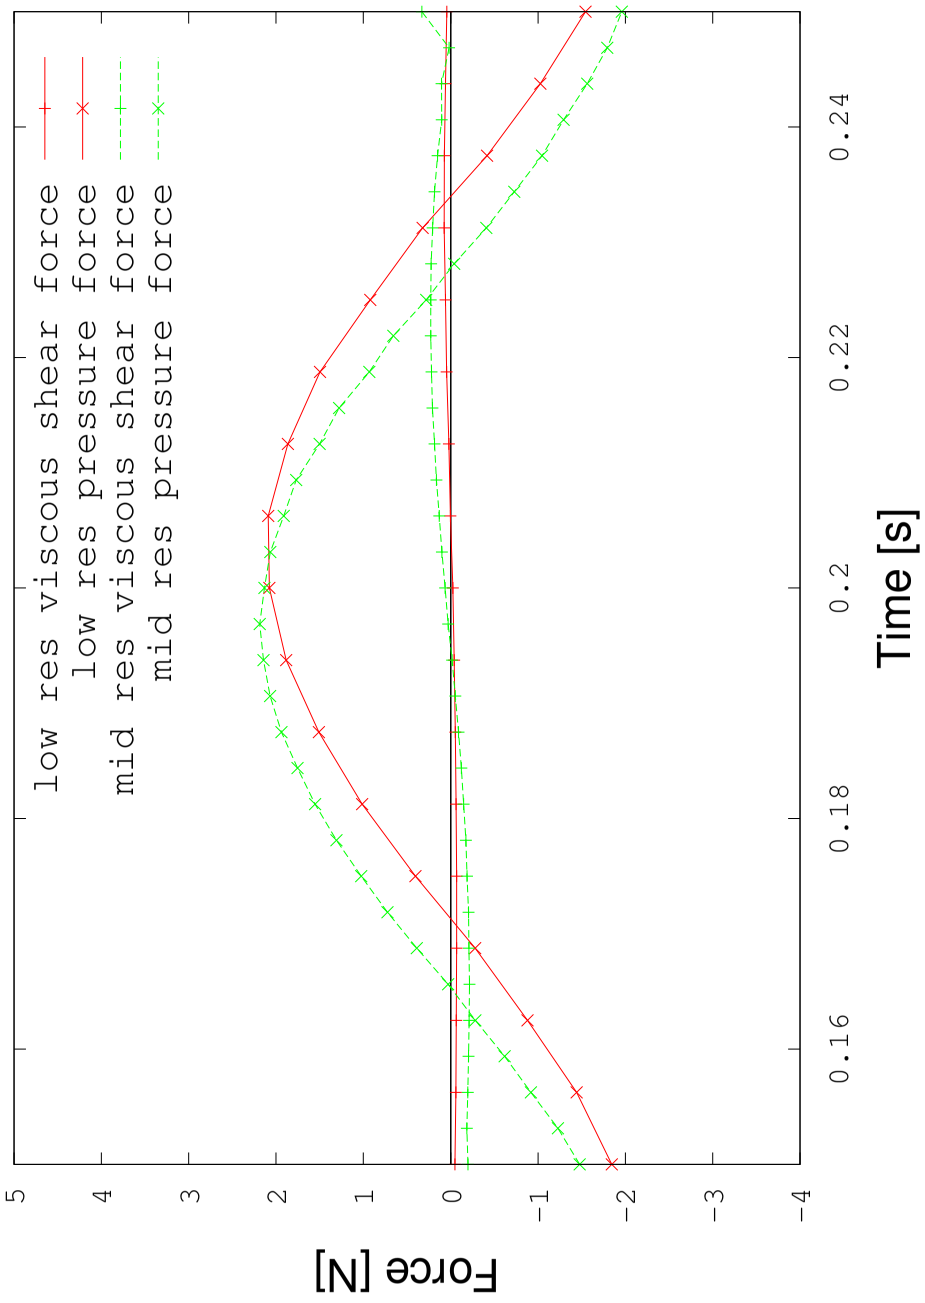


Figure 4.26.: Graph showing the contributions to the force from pressure and viscous shear force on an oscillating sphere. Relatively large changes to the viscous shear stress has little effect on the overall force as it is not a significant contributor. Pressure force and total force are largely independent of refinement.

4.4. Comparison with other methods

Previous attempts to use a cut-cell version of the continuity equation are tested. In particular, the method proposed by Zastawny [33], developed form the ideas of Seo and Mittal [64] is investigated due to the overlap of some elements of the methodology used in this thesis. The principle of the methods of Zastawny and Seo and Mittal is to drop one of the continuity equation constraints, rather than a mirroring constraint as is the case in this work. Overall continuity is restored by 'merging' the continuity equation contribution of the cell with one of its neighbours.

It is found that the continuity equation is satisfied over an area bounding the immersed-boundary, hence not yielding significant pressure oscillations. However it is also found to be un-suitable to the current requirements due to a lack of consistency of the flow field in the near vicinity of the immersed-boundary. This is significant as a large fraction of the flow in a granular material is close to an immersed-boundary. The following simulation is of a sphere advecting through a viscous fluid at $u_p = 0.2$ m s^{-1} , $Re \approx 200$. Figure 4.27 shows the component of velocity in the direction of the relative motion of the particle a two adjacent time instances. Three things are apparent from this figure:

1. There is a a sharp gradient in the velocity of the flow around the body.

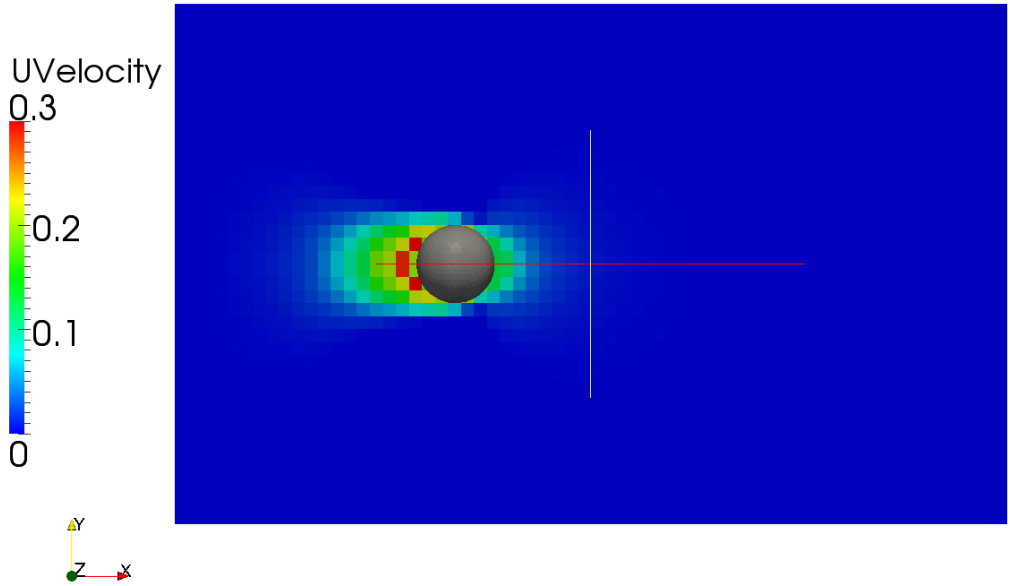
2. The velocity behind the immersed-boundary is significantly greater than it should be at some points. Given the low Reynolds number the velocity should not exceed the advecting velocity of the particle, 0.2 m s^{-1} , yet the maximum velocity is greater than 0.3 m s^{-1} . This is an overestimation of $\sim 50\%$.

3. There is significant difference between the velocity field relative to the body between time instances. Again this is unphysical given the low Reynolds number of the simulation and that the flow has run long enough for approximately steady flow to be expected.

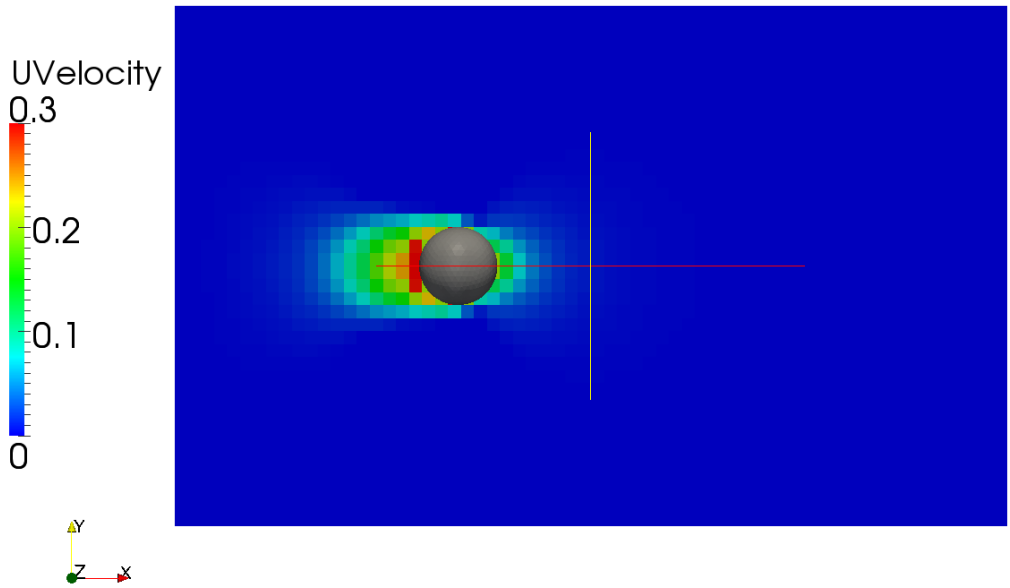
Figure 4.28 shows the component of velocity normal to the direction of the relative motion of the particle. Form this figure it is also clear that the nature of the flow is not ideal. In particular the velocity field near the immersed-boundary appears to be in the wrong direction. However, when a larger radius around the particle is considered (highlighted in white), the flow is sensible. Specifically outside of the 'merged' cells, a good agreement with the expected value is found. This is reasonable as it represents the volume over which continuity is evaluated. Figure 4.29 shows two sets of stream lines highlighting the issue. In the first panel, streamlines originating from upstream of the body are shown. The second panel shows streamlines from the same case, originating from the near wake of the body. While the first set appear correct, un-physical noise is observed behind the particle in the second set.

Given the high numerical cost associated with the large domains typical of granular materials and the inherently high cost of fully resolved methods, required to accurately model the complex flow through granular materials, minimising computational expense is paramount. This is relevant to the choice of method as, given that a large proportion of the fluid is in close proximity to one or more particles, it is expected that a large fraction of the fluid cells will be subject to merging. It is important to discretise the region of flow close to a particle compactly and accurately, even at high refinement, as the flow is expected to vary rapidly in this region and is significant for determining the coupling with the immersed-boundary. As merging cells increases the number of cells required to

remove the non-physical influence of the immersed-boundary, an increase in the overall refinement is necessary to accurately predict the majority of the flow. It can be concluded that techniques that require multiple cells to discretise the flow around the immersed-boundary should be avoided if possible when granular materials are to be modelled. Indeed all possible attempts should be made to minimise the non-physical range of influence of immersed-boundary in any method, however it is particularly important for granular materials. Hence, while the merged cut-cell method may be practical for simulating free stream particles, the method is not suitable for modelling flow through granular materials.

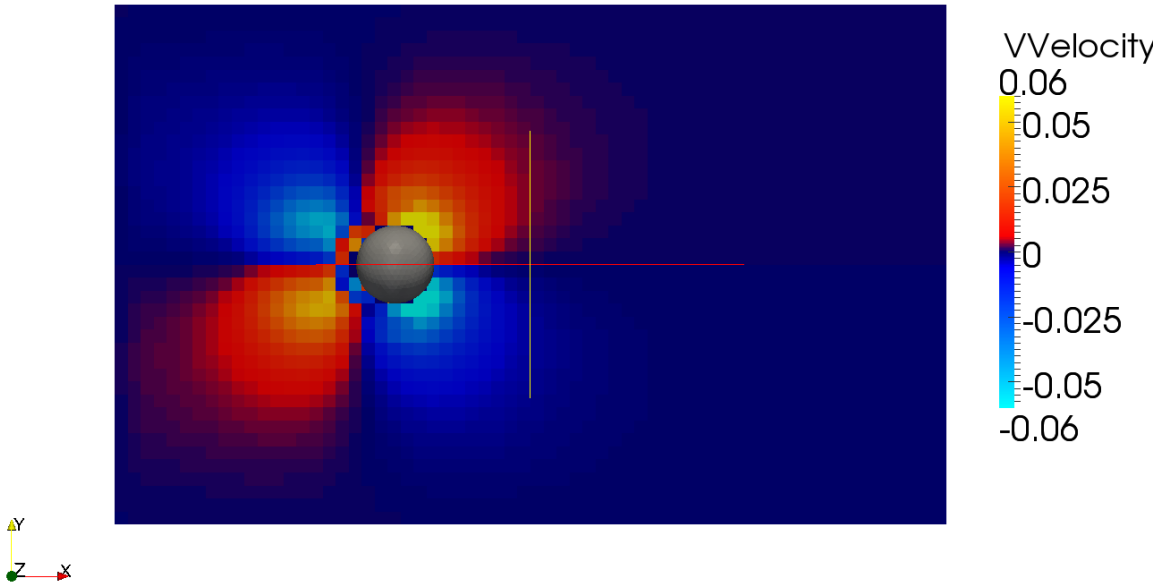


(a) Velocity normal to relative motion of an immersed-boundary at t_0.

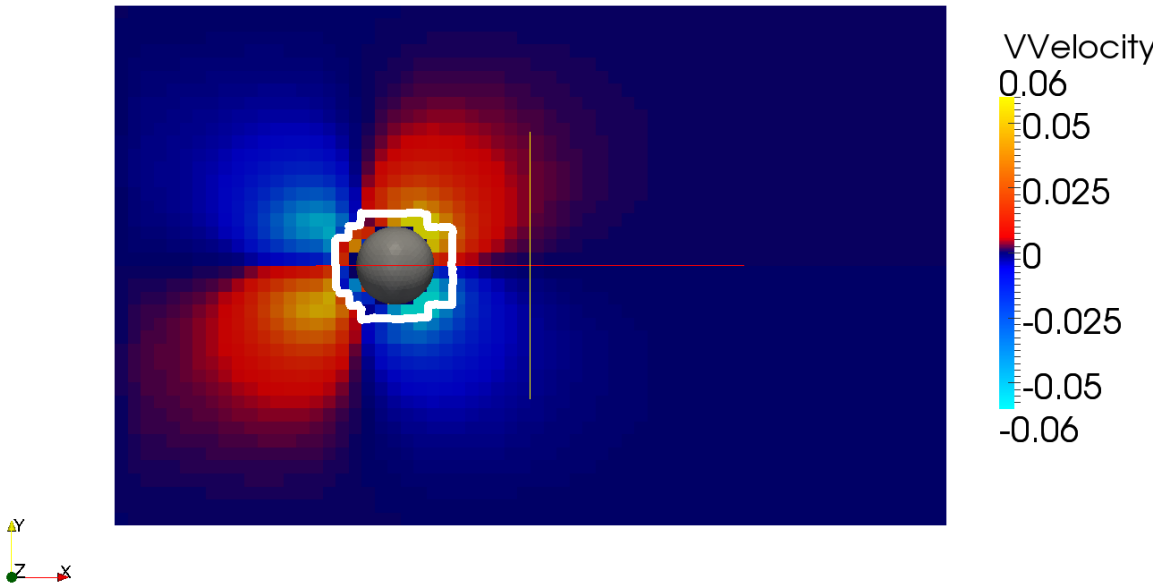


(b) Velocity normal to relative motion of an immersed-boundary at t_1.

Figure 4.27.: Normal velocity around a particle moving at 0.2 m s^{-1} in the cut/merging method, at two adjacent time instances. The velocity field around the immersed-boundary is non-physical, in particular the variation in time, between t_0 and t_1 and the maximum velocity are erroneous.

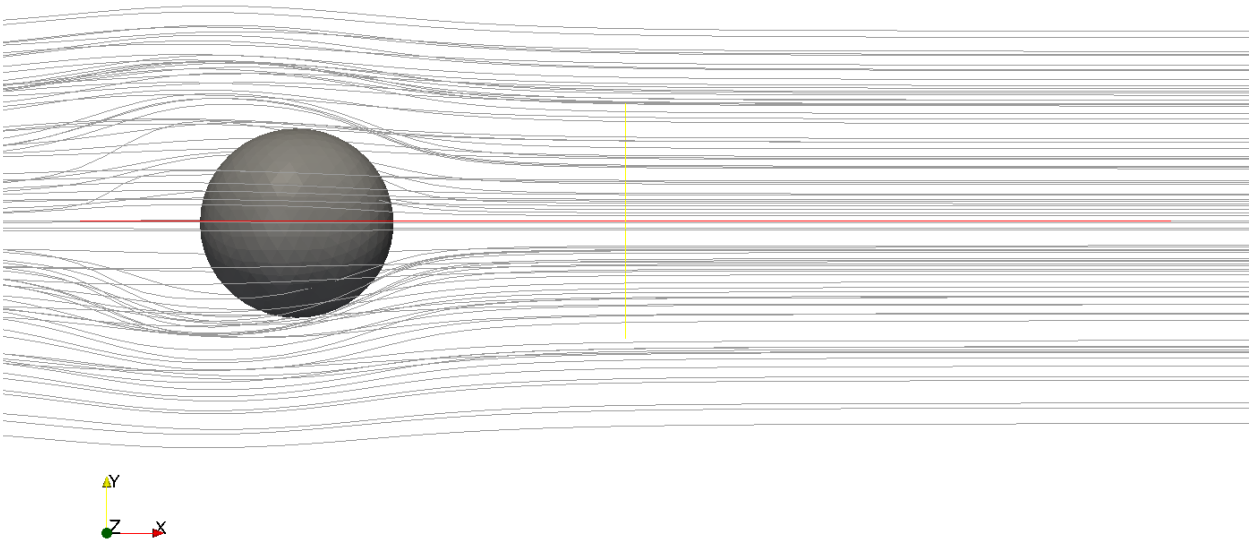


(a) Velocity tangential to relative motion of the immersed-boundary.

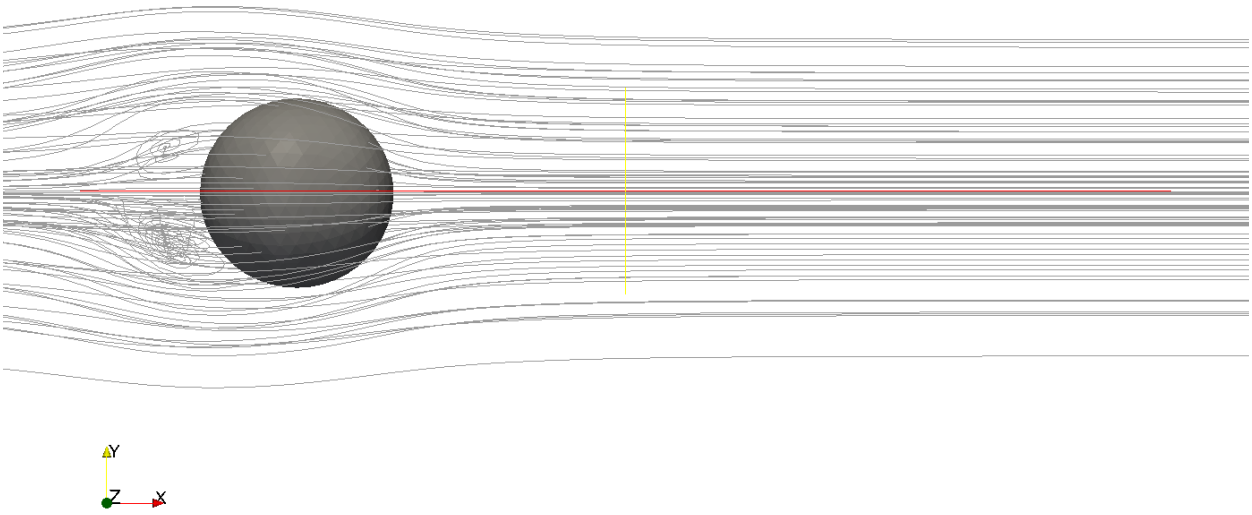


(b) Highlighted outside of the range of merging.

Figure 4.28.: Normal velocity around a particle moving at 0.2 m s^{-1} in the cut/merging method. Figure 4.28a shows the cell-centre data for all cells outside the body. Figure 4.28b shows the separation of cells over which merging is permitted to occur. This figure shows the influence of merging on the velocity field, in particular that while sensible results are present outside of the range of merging, at low resolution, inside of the merged cells erroneous trends are found.



(a) Originating from ahead of the body.



(b) Originating from immediately behind the body.

Figure 4.29.: Streamlines of the flow around a particle moving at 0.2 m s^{-1} in the cut/merging method, relative to the velocity of the body. Figure 4.29a shows the stream lines of fluid particles ahead of the body, out of the range of cell merging. Figure 4.29b shows the stream lines of fluid particles near to the surface of the body, within the range of cell merging. Disparity between flow far removed and close to the cell illustrated by adding streamlines originating close to the body. The far removed case shows sensible results which is consistent with the premise that continuity is solved over a two cell range.

4.5. Discussion

In this chapter the effects of the developed methods are investigated by comparison to experimental data. In general good agreement is found with the experimental test case presented in the literature, in particular to the work of ten Cate *et al.* [86].

The hydro-dynamic interaction force is found to agree well with the experimental data however there is a dependency on refinement of calculation of the velocity gradient at the particle surface and hence force. The convergence is good but this is the major limiter of refinement required for an accurate simulation. For the current method of force calculation, 3-4 cells per radius are necessary to calculate the force accurately.

The spurious pressure oscillations, caused by the changing numerics as the immersed-boundary mesh moves relative to the Eulerian mesh [64] are improved by the addition of modifications to the mirroring equations and the continuity equation. The cut-cell partial mirroring method, described in section 3.2.3, was found to significantly diminish the spurious pressure oscillations for moving particles. The averaged mirroring modification, described in section 3.2.2.1, is designed to insure that the velocity of boundary-exterior points is similar to that of the immersed-boundary when the immersed-boundary is close to the points in question. This modification to the mirroring equation improves the smoothness of the force acting on a particle without modifying the trend of the force. In particular the was virtually no effect on the flow around, or force acting on, a stationary particle.

For all of the developed methods, a good agreement is found for the flow around a stationary sphere. This gives confidence that the flow in this case is correct and that the flow serves as a good test case for general motion of an immersed-boundary. Importantly, good agreement is found between simulations, using the developed cut-cell partial mirroring method and averaged mirroring modification, differing only by the relative motion between the immersed-boundary and the Eulerian mesh. This was not the case in the unmodified method for which the motion of the immersed-boundary relative to the Eulerian mesh causes significant influence. Along with the good agreement with the experimental data of ten Cate *et al.* [86] the combination of the cut-cell partial mirroring method and averaged mirroring is shown to be robust at simulating individual particles at relatively low refinement over a range of Reynolds numbers.

Low dependency on the resolution of the Eulerian mesh is found in the developed method. While the force calculation is somewhat dependent on the resolution of the Eulerian mesh, the flow field is accurately captured with 3 cell per particle radius.

The parity based incrementation technique, described in section 3.2.3.4, was found to be advantageous in some cases. However the cases in which it was advantageous were those in which a significant oscillation in the pressure field is present (without the parity based incrementation modification). In these cases the use of the cut-cell partial mirroring method and the averaged mirroring modification has significantly greater effect at reducing the oscillation in the pressure field. The further use of the parity based incrementation technique does not significantly further improve the results.

The cut-cell momentum equation, presented in section 3.2.4, gives very similar results to the unmodified momentum equation, in the case where the cut-cell partial mirroring method and averaged mirroring modification is used. This is reassuring as it suggests the velocity and pressure field are sufficiently smooth that small differences in the discretisation scheme do not play a significant role. A slight benefit to the cut-cell momentum equation is it accurately accounts for the volume outside of an immersed-boundary, therefore is better suited to unsteady/transient flows. There is a cost

associated with finding the intersection of the immersed-boundary mesh and the Eulerian mesh, necessary in order to use the cut-cell momentum equation. However if this intersection is already required, *e.g.* for the cut-cell continuity equation, there is little disadvantage associated with the use of the cut-cell momentum equation.

5. Lubrication forces

The necessity of lubrication force and its resolution in various numerical frameworks is discussed. In particular the significance of the phenomena and the limitations of the current modelling techniques are addressed. Proposals for improvement in modelling the lubrication force in the immersed-boundary framework are presented.

5.1. Determination of lubrication force

5.1.1. Prediction using dimensional analysis

If the assumptions of Buckingham π theorem hold, this could be exploited to determine a relation between the non-dimensional groups describing lubrication force.

5.1.1.1. Low Reynolds number collisions

The lubrication flow is often characterised by low Reynolds number flow. If this is the case the lubrication force can be considered a function of the reduced set of variables: μ, u, b and h. Here u is the rate of approach of the two particles in the direction normal to the interface normal. h is the minimum separation between the two bodies *i.e.* the distance separating the closest points of the two bodies. b, the harmonic average of the radius of curvature given by equation 5.1 is taken as a way to compare the collision of various types of surfaces. This is used to compare two different curvatures, as it accurately describes the separation of the two surfaces as a function of radial distance from the closest point.

$$b = \frac{2}{\frac{1}{r_1} + \frac{1}{r_2}} \tag{5.1}$$

The assumptions of Buckingham π theorem, and the assumption that these parameters establish the lubrication force, hold then the form of the lubrication force is heavily restricted. In particular; the lubrication force between two spheres can be considered a function only of the kernel of A:

$$A = \begin{array}{c} \\ m \\ s \\ kg \end{array} \begin{array}{c} \begin{array}{cccc} \mu & u & b & h \end{array} \\ \begin{pmatrix} -1 & 1 & 1 & 1 \\ -1 & -1 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{pmatrix} \end{array}$$

A has kernel spanned by:

$$\begin{array}{c} \mu \\ u \\ b \\ h \end{array} \begin{pmatrix} 0 \\ 0 \\ 1 \\ -1 \end{pmatrix}$$

Corresponding to a single dimensionless parameter:

$$\left[\frac{b}{h}\right]$$

Based on this restriction, the corresponding relation is proposed:

$$\Phi\left(\left[\frac{b}{h}\right]\right) = \frac{f}{\mu\ u\ b} \tag{5.2}$$

Assuming a power-law relation Φ can be parametrised as $\Phi(x) = Cx^B$ or some constants B and C.

$$C \times \left(\frac{b}{h}\right)^B = \frac{f}{\mu\ u\ b} \tag{5.3}$$

Estimates of the numerical values of B and C are determined using high resolution simulation of two spheres, colliding at fixed speed. The simulation has free stream particle Reynolds number, Re_p, sufficiently low to justify Stokesian dynamics ($Re_p = 0.06$). Table 5.1 shows the properties of the simulation. Figure 5.1 shows the setup and resulting flow-field during a near collision.

Particle radius	0.03 m
Dynamic viscosity	1.0 N s m^{-2}
Fluid density	1.0 Kg m^{-3}
Particle velocity	1.0 m s^{-1}
Refinement	15 cells per radius
Max CFL number	0.1

Table 5.1.: Details of simulation setup for computation of lubrication force for low Reynolds number.

The non-dimensionalised lubrication force as a function of the non-dimensional parameter h/b predicted by this simulation is given in figures 5.2 and 5.3. Assuming equation 5.3, it remains to find constants A and C. These constants can be found using a least squares regression line through a log-log plot of the results of a simulation (see figure 5.2). Doing so gives:

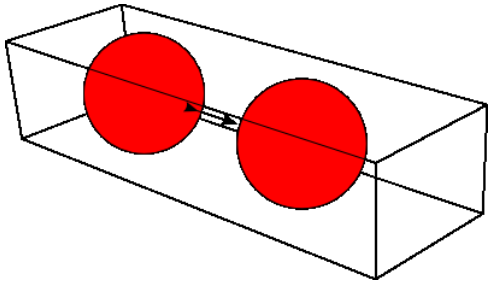
$$B \approx 1$$
$$C \approx 4.7$$

Hence:

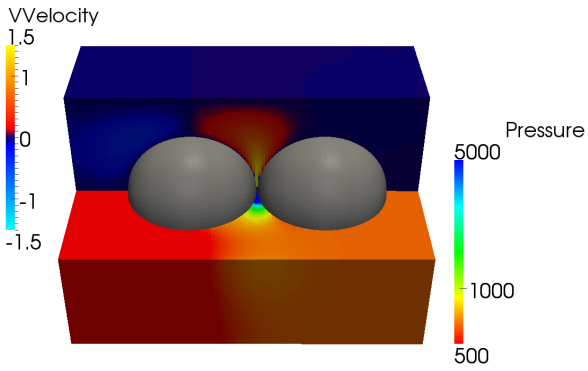
$$\frac{f}{\mu\ v\ b} = 4.7 \times \left(\frac{b}{h}\right) \tag{5.4a}$$
$$f = 4.7\ \mu\ v\ b \left(\frac{b}{l}\right) \tag{5.4b}$$

It is relevant that the simulation used to model the lubrication flow is expected to be inaccurate when h is small relative to cell spacing (hence the need for the correction discussed later in this section). Contrarily, the model has the correct behaviour at large h/b. Indeed this model predicts a particle stationary in a flow is free of force if it is sufficiently far removed other moving objects. However, while the asymptotic limit is correct, it is clearly possible that the convergence is not. This is taken into account when computing the least squares regression line, only results with

($h \geq 1.5 \times$ grid spacing and $h \leq 2.0b$) are used.



(a) Setup



(b) Flow around particles

Figure 5.1.: The setup used to determine the lubrication force between two particles and resultant flow field. 5.1a (left) shows a single fixed particle being approached by a second moving particle inside a square prism. 5.1b (right) shows the resultant fluid properties, including the pressure field and the component of velocity in the vertical direction.

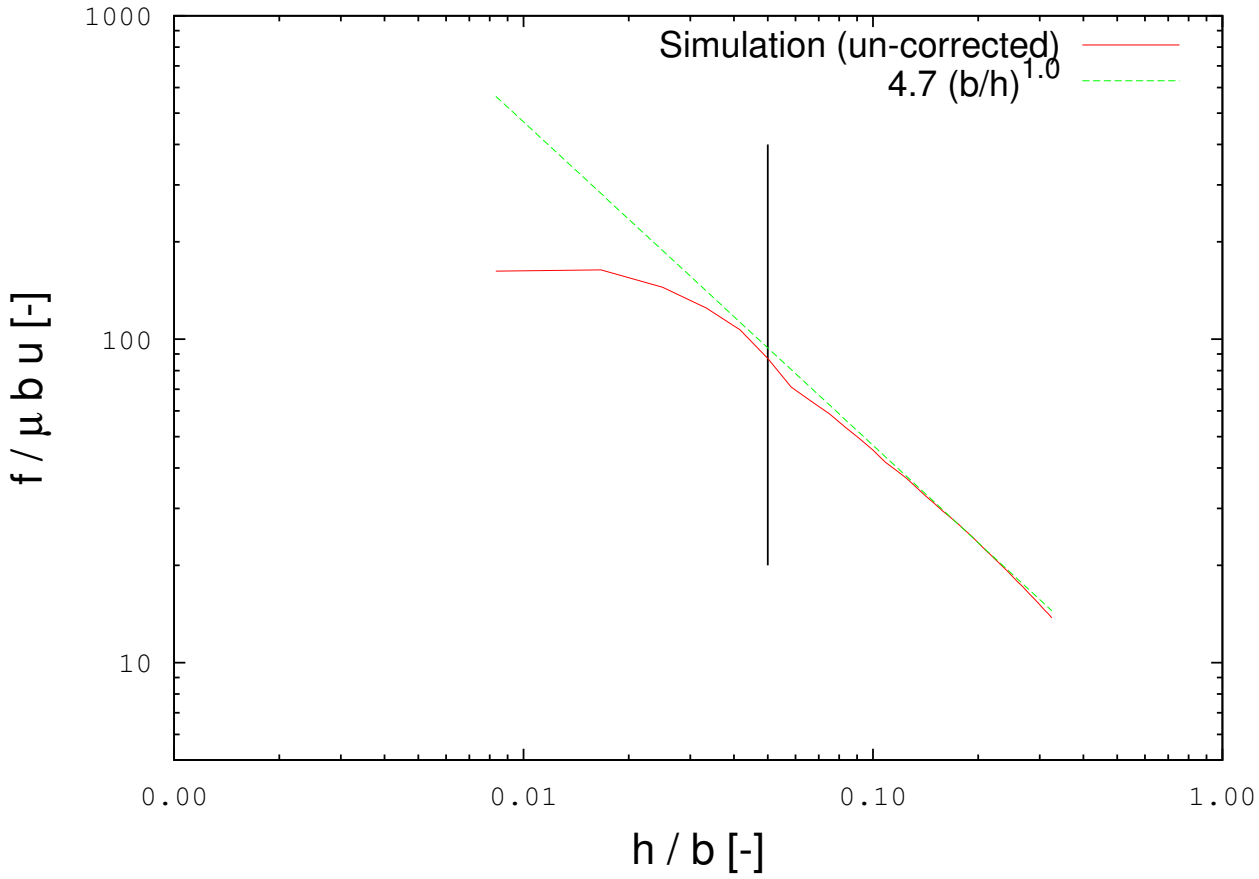


Figure 5.2.: Graph showing a log-log plot non-dimensionalised lubrication force $\frac{f}{\mu\, u\, d}$ against $\frac{h}{b}$, for colliding spheres with $Re = 0.06$. The vertical line indicates a single cell spacing The plot also shows and the fitted line through the curve at the reliable points.

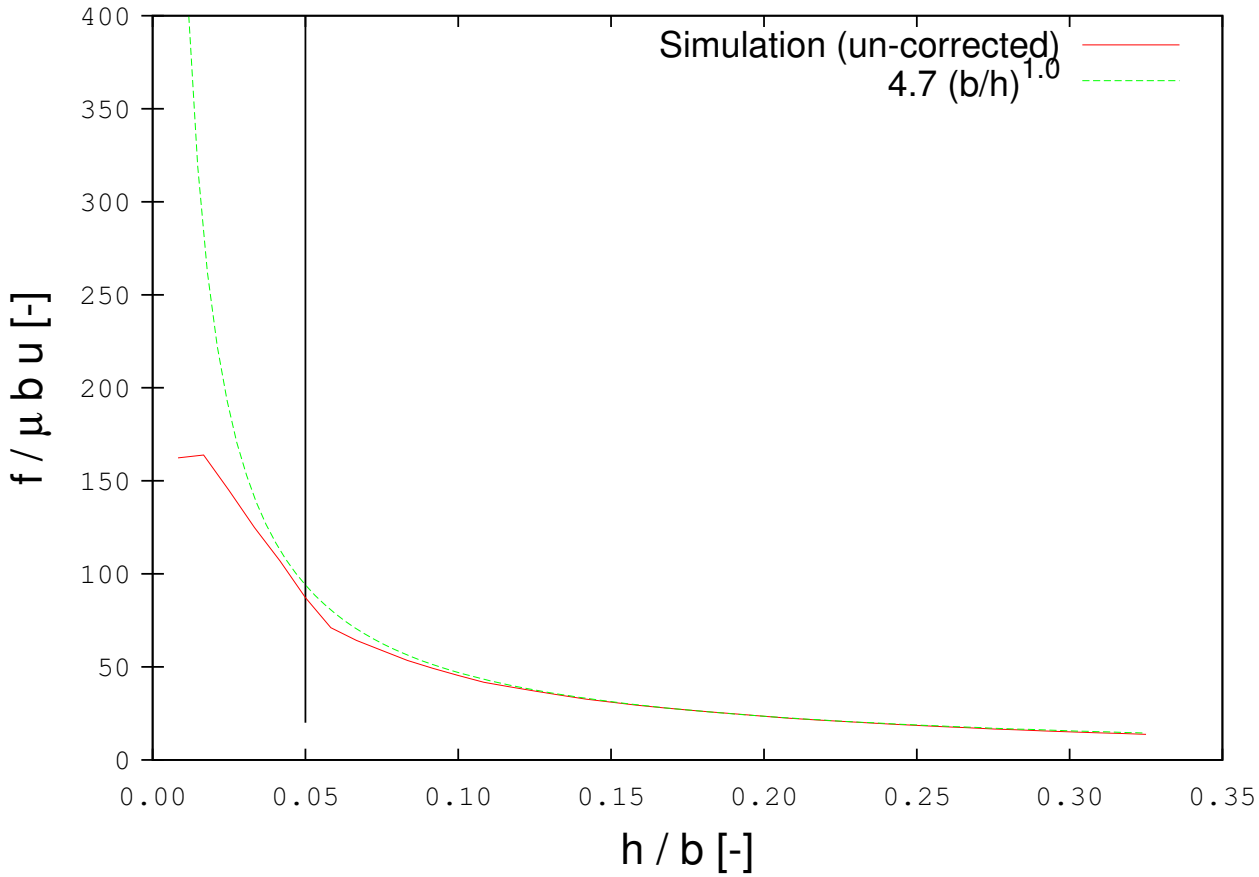


Figure 5.3.: Linear scaled plot of the non-dimensionalised lubrication force $\frac{f}{\mu\, u\, d}$ against $\frac{h}{b}$ For colliding spheres with $Re = 0.06$.

The theoretical work of Brenner [65] is not directly applicable, as the work concerns the collision of a particle and a wall. However as discussed the results are comparable using the harmonic average of the radii of curvature of the two surfaces. Indeed in this case we have the same form, and the coefficients become:

$$B = 1$$
$$C = 1.5\pi \approx 4.7124$$

The values of B and C obtained are remarkably close to the fitted parameters of the simulation. Ideed, both are comfortably within half a percent of the simulation-predicted values.

Despite very good agreement with non-dimensional predictions and theory the results break down as the spheres come within a single cell spacing. While this is to be expected, given the inevitable breakdown of the method when significant variation is expected to occur within the a single cell spacing, this is problematic. The breakdown of the method at small separations occurs as the flow is expected to vary significantly non-linearly, or indeed quadratically, over a range smaller than the cell spacing. Further more the flow in the region of significant non-linearity is expected to contribute significantly to the method. The situation is illustrated in figure 5.4. The reason this is potentially problematic is that lubrication forces between particle in close proximity can strongly affect the dynamics. Indeed in an idealised case, a perfect sphere in contact with a smooth wall immersed in an incompressible viscous fluid cannot be separated using a finite force [66]. While this is clearly not true of an actual sphere and wall, due to roughness etc, the affect is clearly not negligible. In order to account for this, a correction is made to the developed method where the predicted inaccuracy of of lubrication force is corrected. The correction assumes a linear variation between one cell spacing and contact, where at a single cell spacing the value is correct and at contact the repulsive force is twice that at a single cell spacing. Hence the corrective force f_c is given by:

$$f_c(h) = \begin{cases} 0 & h \geq s \\ l(h) - l(s)(2 - \frac{h}{s}) & h \leq s \end{cases} \tag{5.5}$$

where l is the force predicted by Brenner and s is the cell spacing. The smoothness of the corrective force is guaranteed as $l(h) - l(s)(2 - \frac{h}{s}) = 0$ when $h = s$. Figure 5.5 shows the implemented correction.

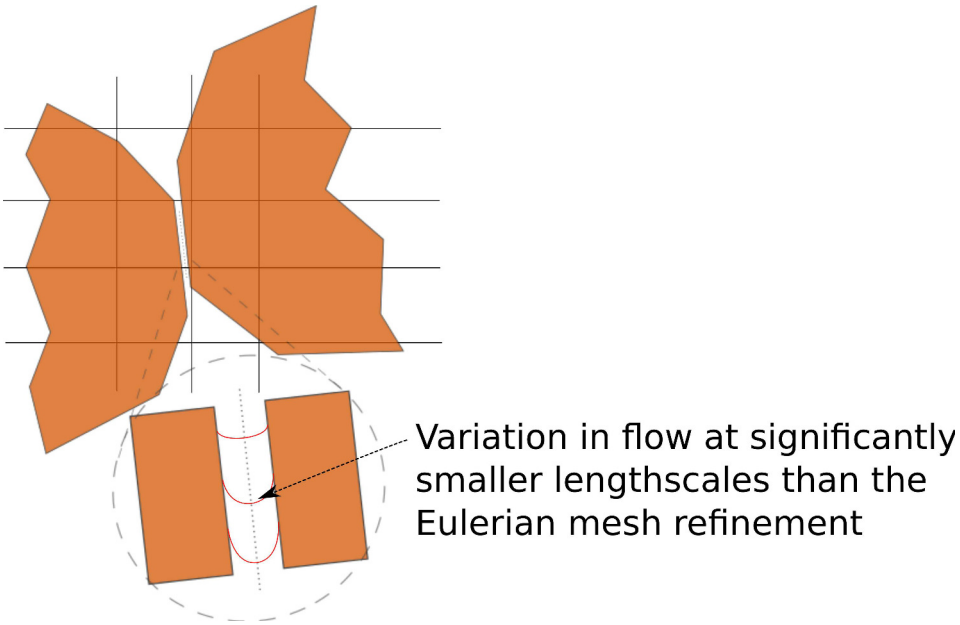


Figure 5.4.: Diagram showing two particles in close proximity in a two dimensional immersed-boundary framework with a representative flow field in the interstice, as a result of the relative motion of the two particles. The thin interstitial fluid layer has significant variation throughout the gap separating the two particles. This is unlikely to be accurately resolved by the Eulerian mesh resolution shown. Even if the resolution where to be increased such that the interstitial flow is well resolved for this specific example, the problem would re-occur once the particles had got closer. Additionally such refinements would not be feasible.

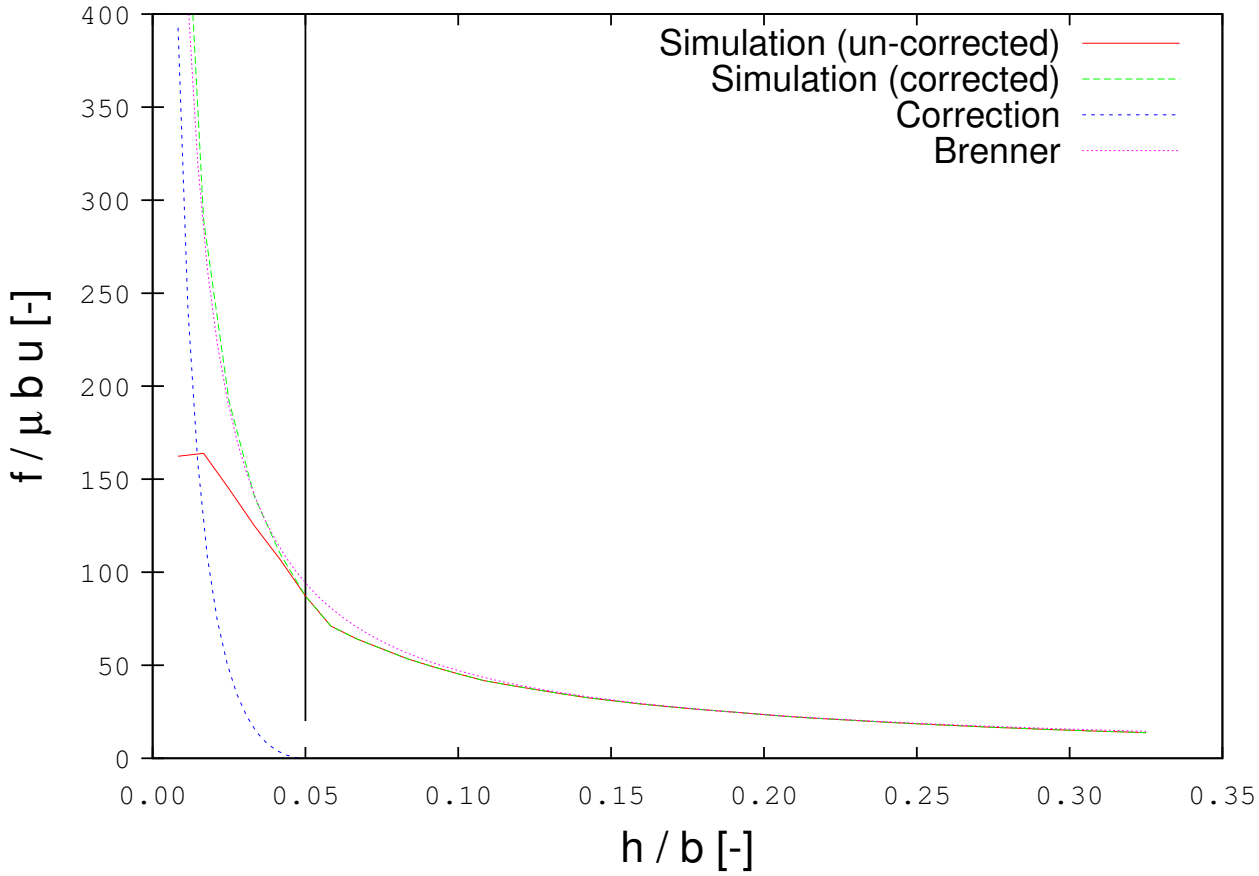


Figure 5.5.: Linear scaled plot of the non-dimensionalised lubrication force with and without the developed correction as a function of $\frac{h}{s}$, the non-dimensionalised separation between the two spheres. The correctional force used is also plotted.

5.1.1.2. General case

The lubrication force, between two spheres at any Reynolds number, can be restricted using dimensional analysis in a similar way to that described in section 5.1.1.1. In the case of non-negligible Reynolds number the lubrication force can also depend on the density of the interstitial fluid ρ. Hence the overall inputs are: μ, ρ, u, b and h. As a result the predictions of Buckingham π theory for the lubrication force between two spheres can be considered a function only of the kernel of A':

$$A' = \begin{array}{c} \\ m \\ s \\ kg \end{array} \begin{array}{c} \begin{array}{ccccc} \mu & u & b & h & \rho \end{array} \\ \begin{pmatrix} 2 & 1 & 1 & 1 & -3 \\ -1 & -1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix} \end{array}$$

A' has kernel spanned by:

$$\begin{array}{c} \\ \mu \\ u \\ b \\ h \\ \rho \end{array} \begin{array}{c} \begin{array}{cc} \end{array} \\ \begin{pmatrix} -1 \\ 1 \\ 1 \\ 0 \\ 1 \end{pmatrix}, \begin{pmatrix} 0 \\ 0 \\ 1 \\ -1 \\ 0 \end{pmatrix} \end{array}$$

Corresponding to dimensional groups of spanned by two non-dimensional parameters:

$$\left[\frac{\rho\ u\ b}{\mu}\right] \text{ and } \left[\frac{h}{b}\right]$$

These are recognisable as the particle Reynolds number and the non-dimensional separation from the low Reynolds number case in section 5.1.1.1. These two parameters give a restriction corresponding to the proposed relation:

$$\begin{aligned} f \quad &= (\mu\ u\ b) \times p(\tfrac{h}{b}, \tfrac{\rho\ u\ b}{\mu}) \\ &+(\rho\ u^2\ b^2) \times q(\tfrac{h}{b}, \tfrac{\rho\ u\ b}{\mu}) \end{aligned} \tag{5.6}$$

For some functions p and q.

Fitting the functions p and q is problematic given the large number of possible forms. However, based on the success of the harmonic averaged curvature modification to the results of Brenner [65], a similar proposal is made based on the same type of modification to the theoretical result concerning sphere-wall collisions given by Cox and Brenner [66]. Cox and Brenner [66] calculate a correction to Stokesian dynamics based on a small inertial correction, which is used to ascertain the functional form of p and q. A small Reynolds number approximation is necessary for the results of Cox and Brenner. Hence their predictions are only valid for small Reynolds numbers. The extent to which the low Reynolds number limit is valid is investigated. However the low Reynolds number condition is expected during the cases of interest as in order to have high Reynolds number during a near collision, the initial velocities must be very high.

The prediction of Cox and Brenner [66], is that:

$$p(x, y) = 1.5\pi \left[\frac{1}{x} + 0.2 \ln\left(\frac{1}{x}\right)\right] \tag{5.7}$$

with an asymptotic limit:

$$p(x,y) = \frac{1.5\pi}{x} \tag{5.8}$$

$$q(x,y) = \frac{3\pi}{20} \ln\left(\frac{1}{x}\right) \tag{5.9}$$

Using a setup identical to that of the low Reynolds number simulation (*i.e.* as described in figure 5.1 with parameters given in table 5.1) except varying the density of the fluid similar results where found for a range of particle Reynolds numbers. Table 5.2 gives the simulations run.

Case	ρ_f [Kg m^{-3}]	Re
L1	10	0.6
L2	20	1.2
L3	30	1.8
L4	40	2.4
L5	50	3.0
L6	60	3.6
L7	70	4.2
L8	80	4.8

Table 5.2.: Details of simulation setup for computation of lubrication force for a range of Reynolds numbers.

Figure 5.6 shows the resultant repulsive forces, non-dimensionalised assuming low Reynolds number, relative to the Stokes flow prediction. The figure validates the assumption that the low Reynolds number approximation gives a reasonable approximation, even without the correction proposed by Cox and Brenner, in the Reynolds number range 0 to 4.8. However, it also shows that the correction is applicable as the difference between the simulations and the low Reynolds number simulation appears linear in Reynolds number or ρ. Figure 5.7 shows the difference in lubrication force predicted by the simulations in table 5.2 and the lubrication force predicted by the simulations in table 5.1. For each of the simulations in tables 5.1 and 5.2, $\mu\ u\ d$ is the same. Hence the differences in forces, f, are directly comparable.

Figure 5.8 shows the difference between the lubrication force of simulation of higher Reynolds number collisions, table 5.2, and lower Reynolds number collisions, table 5.1. In this figure, the difference is normalised by the difference in the density of the fluid used in the simulations and subsequently non-dimensionalised by dividing by u^2b^2. The results are plotted against the non-dimensional separation, h/b. The prediction of Cox and Brenner, modified using the harmonic averaged curvature method, for what the differences between the lubrication forces should be is given by:

$$\frac{f_a - f_b}{(\rho_a - \rho_b)u^2b^2} = 0.6\pi \ln(\frac{b}{h}) \tag{5.10}$$

for collisions a and b each with the same value of u, b, and μ. This prediction is also included in figure 5.8. The fit between the simulations and the prediction of Cox and Brenner is erratic but the trend of the fit is good and the error is not large when the particles are separated by more than two Eulerian mesh cells.

The erratic nature is perhaps predictable as the differences fitted are a small fraction of the overall force. Indeed relative to the overall force, the perturbations from the predictions of Cox and Brenner

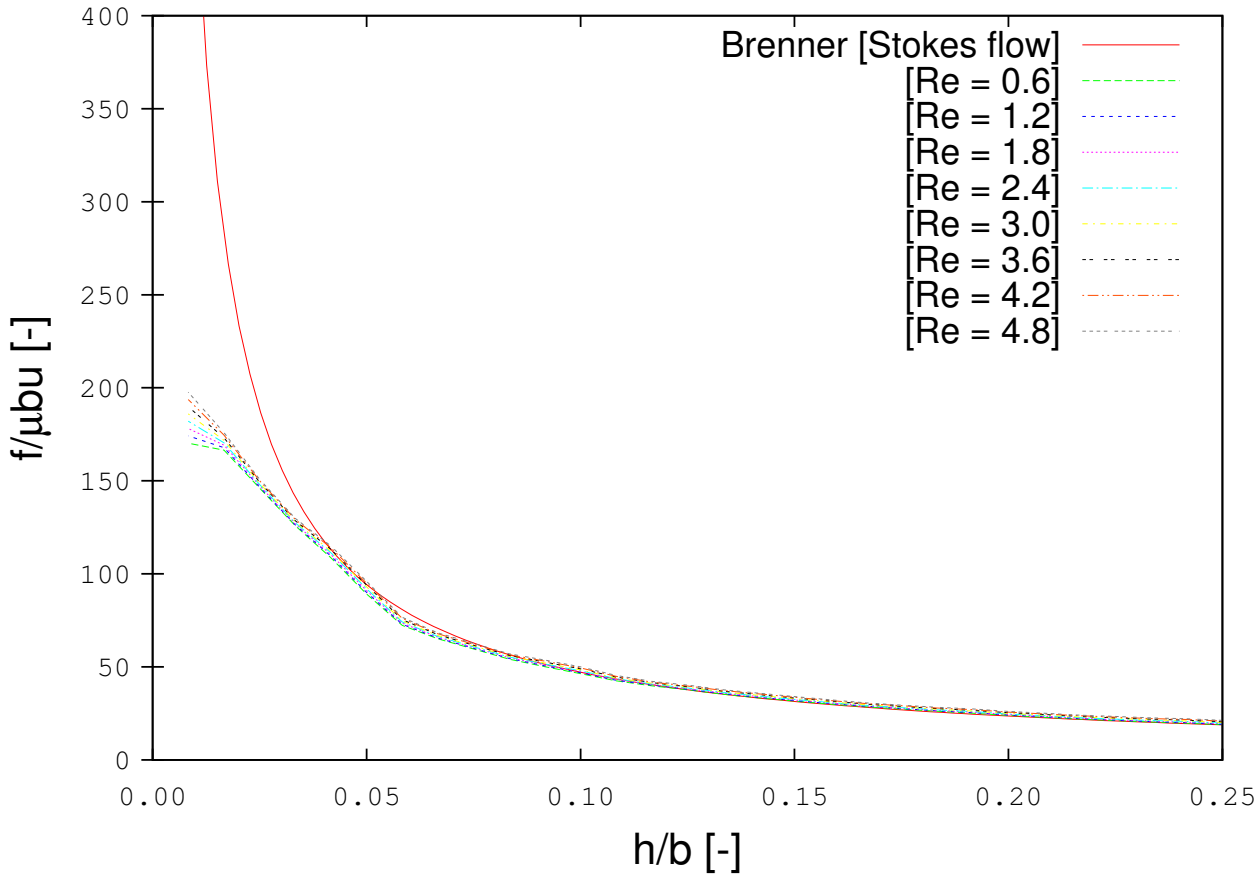


Figure 5.6.: Results of the simulations of colliding spheres of with particle Reynolds number in the range 0.6 to 4.8, given in table 5.2. Graph shows the non-dimensionalised lubrication force $\frac{f}{\mu\, u\, d}$ against non-dimensionalised separation, $\frac{h}{b}$. Also plotted on the same axis for comparison is the prediction given by substituting the harmonic average surface curvature into the prediction of Brenner [65].

are small. A correction based on this is not implemented in this thesis as, although the results are not perfect, the error is unpredictable and hence cannot be accounted for reliably. The prediction that the difference is linear in density is excellent and suggests the prediction can be used across a range of Reynolds numbers. Additionally this difference is small for the range of Reynolds numbers of interest. It can be concluded that the accuracy of the low Reynolds number prediction over this range of numbers is sufficient to justify its applicability.

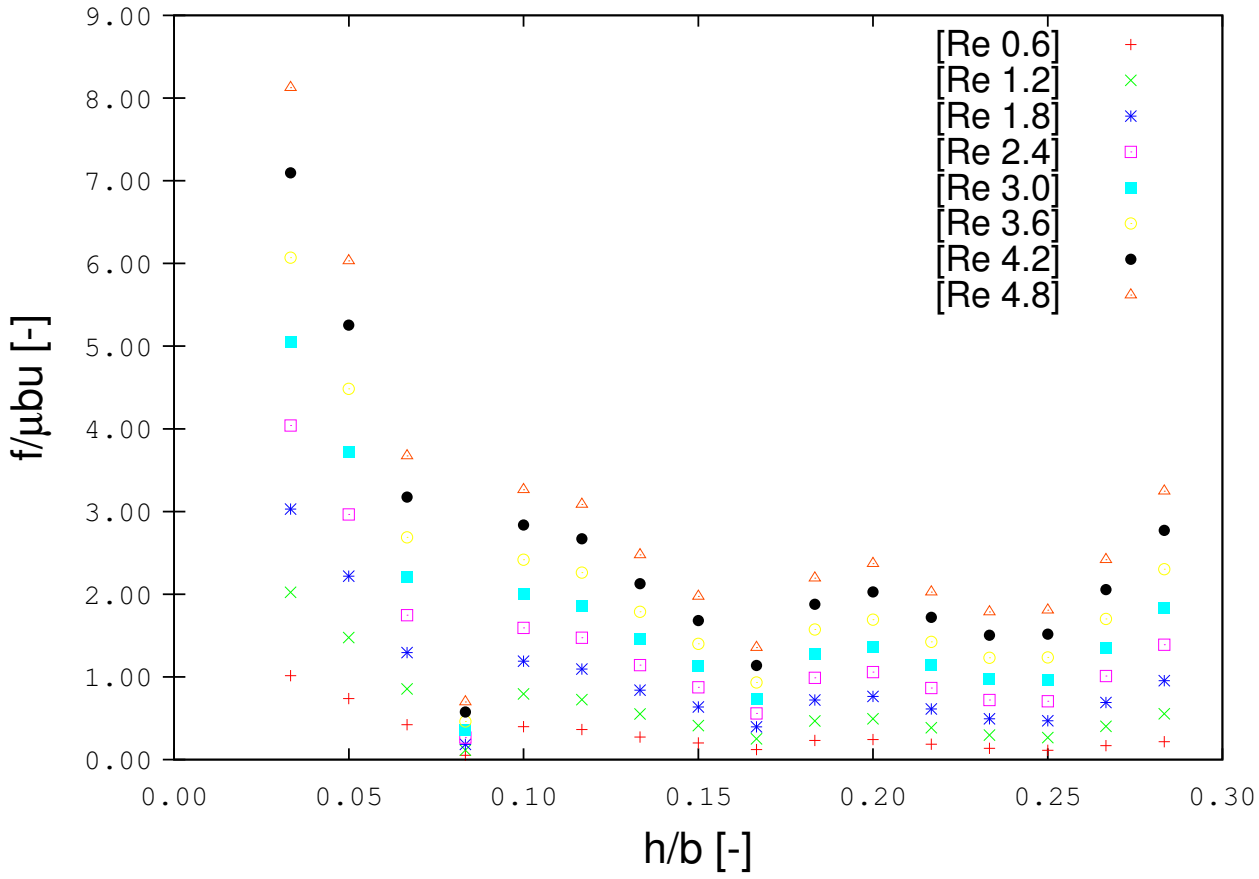


Figure 5.7.: Graph showing the difference in the simulated lubrication force between colliding spheres with Reynolds number up to 4.8 and the simulated lubrication force between two colliding spheres with very low Reynolds number ($Re = 0.06$, *i.e.* Stokes-flow). The lubrication force is non-dimensionalised as $\frac{f}{\mu\, u\, d}$ and plotted against non-dimensionalised separation, $\frac{h}{b}$.

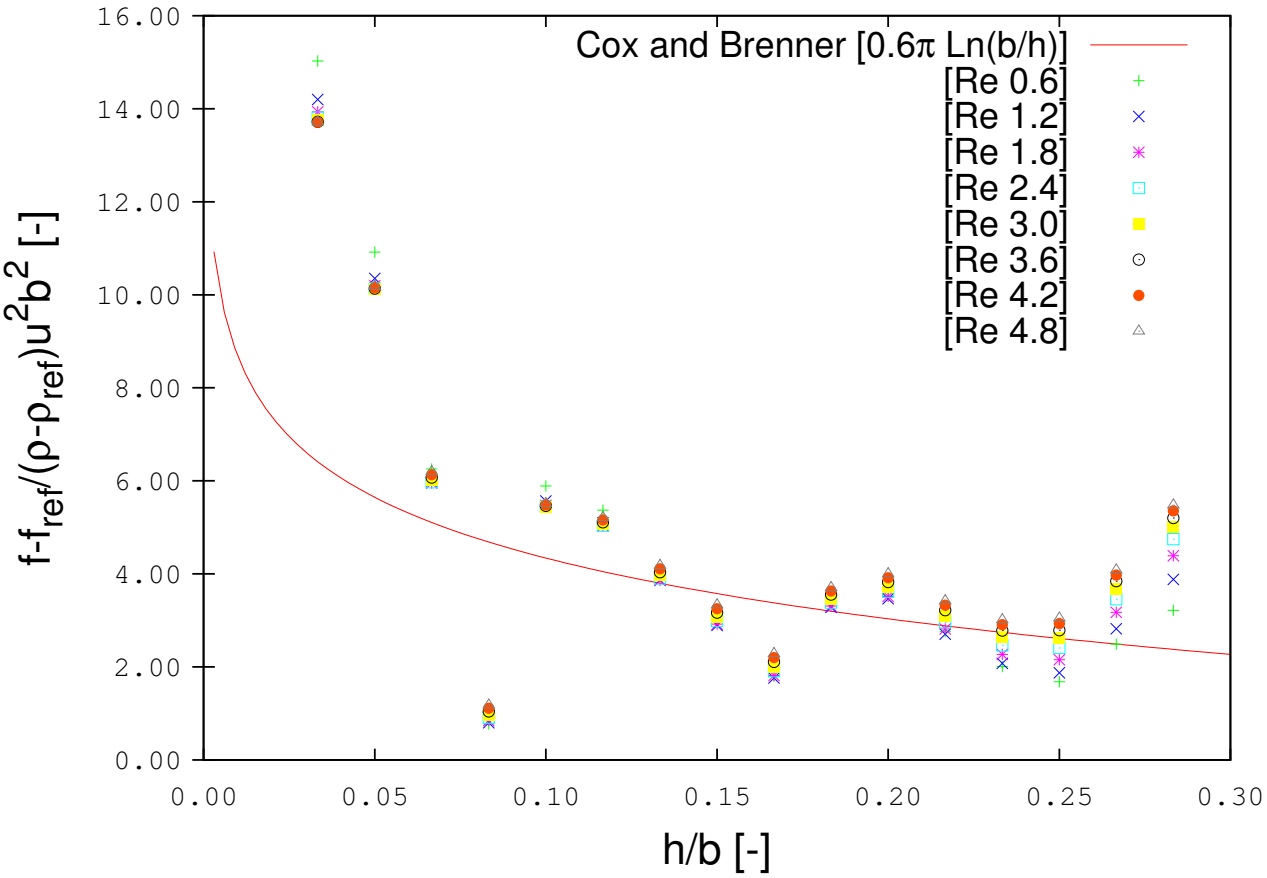


Figure 5.8.: Graph showing the effect of the variation of Reynolds number via fluid density, ρ, on the lubrication force of two colliding spheres relative to the predictions of Cox and Brenner [66]. In particular the difference between the lubrication force of a low Reynolds number simulation (table 5.1) and a range of higher Reynolds numbers (table 5.2) are shown. The difference between the lubrication forces are normalised relative to the fluid density difference and non-dimensionalised to allow an easier comparison to the theoretical predictions of Cox and Brenner [66].

The relative errors in lubrication force of the original and corrected methods are given in figures 5.9 through 5.12. The errors are given relative to Stokes dynamics predictions (Brenner [65]) and to the inertially corrected Stokes dynamics predictions (Cox and Brenner [66]). In each case the relative error is given by:

$$\frac{\text{calculated} - \text{expected}}{\text{expected}}$$

As expected, the un-corrected lubrication forces show substantial disagreement with the theoretical limit (with and without inertial correction), when the particles are in close proximity. However, very good agreement is found outside of a single cell spacing. With the implemented correction factor the error relative to the inertially corrected predictions of Cox and Brenner is small. The maximum relative error in forcing is 10% at a single peak just prior to when the cells are separated by a single cell (at $h/b = 0.05$). When considering that this is a channel flow being resolved at a single cell width, this is not unexpected. Outside of this peak the error is less than 5%.

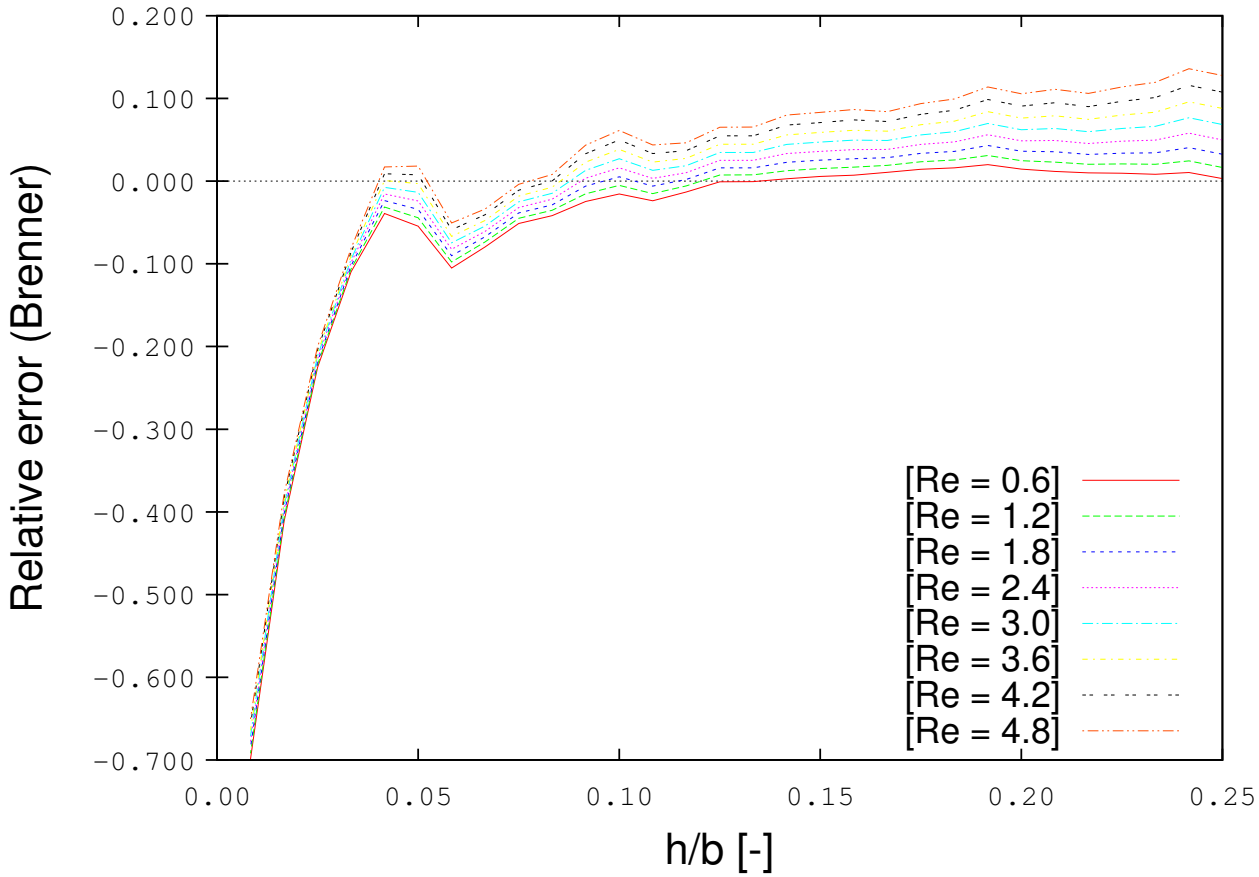


Figure 5.9.: Graph showing the error of lubrication force for the un-corrected method relative to the zero Reynolds number prediction (Brenner), for a range of particle Reynolds numbers.

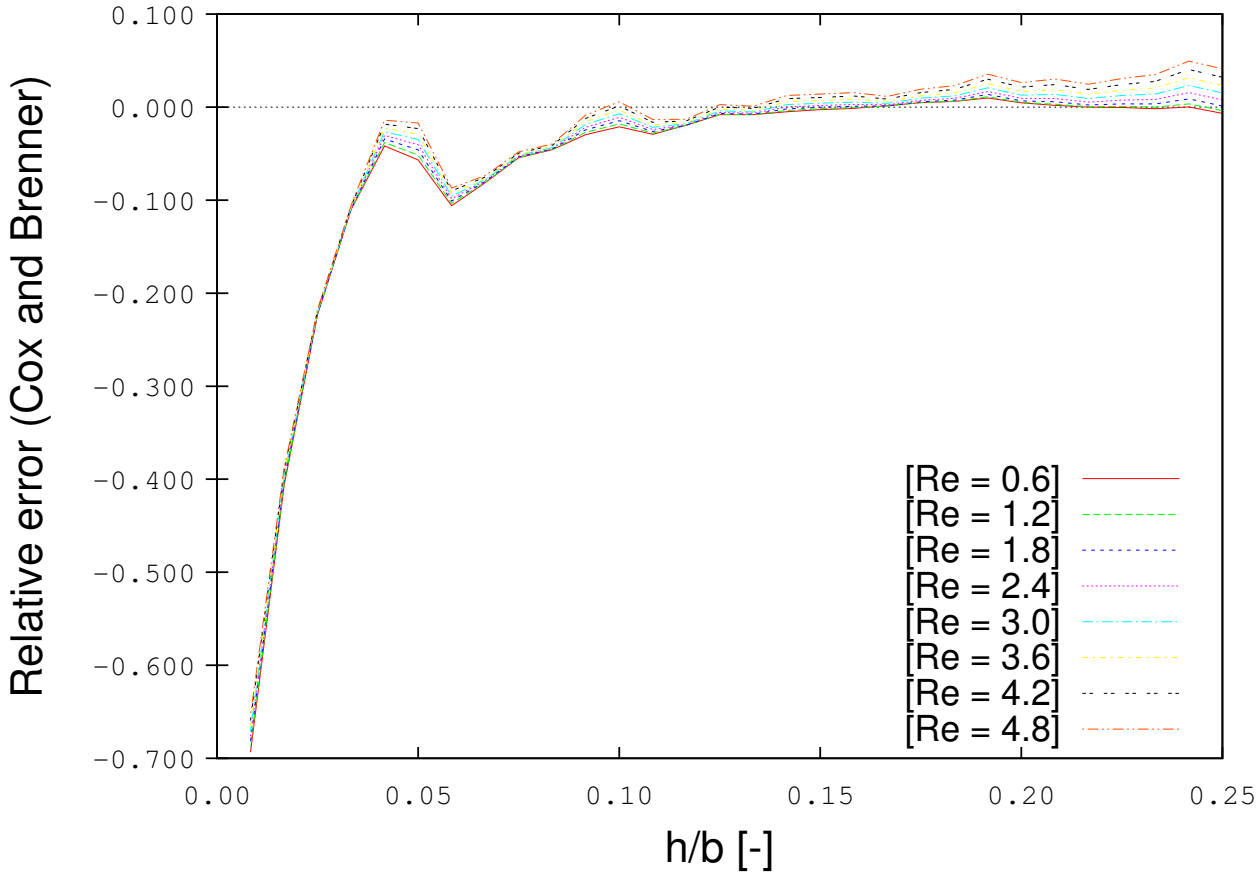


Figure 5.10.: Graph showing the error of lubrication force for the un-corrected method relative to the low Reynolds number prediction of Cox and Brenner, for a range of particle Reynolds numbers.

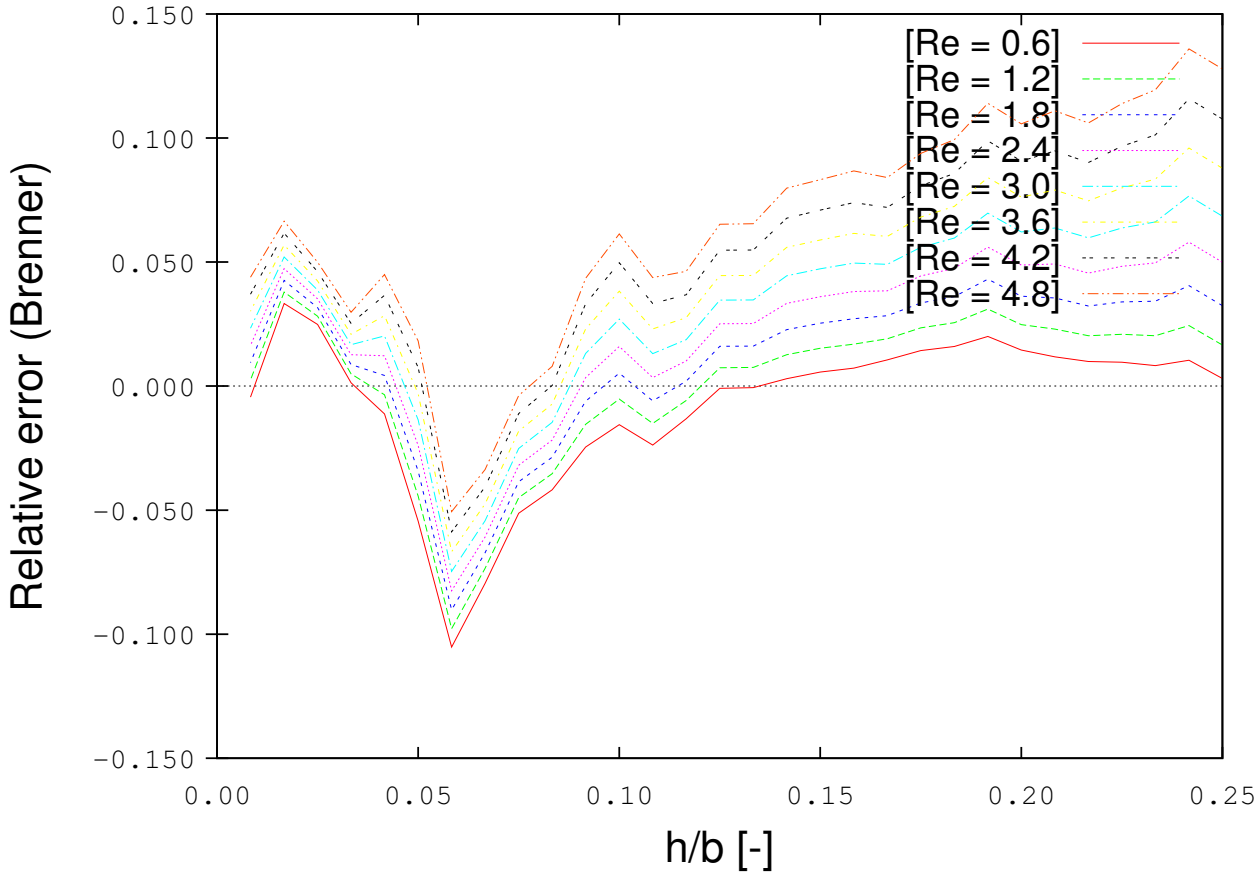


Figure 5.11.: Graph showing the error of lubrication force for the corrected method relative to the zero Reynolds number prediction for a range of particle Reynolds numbers.

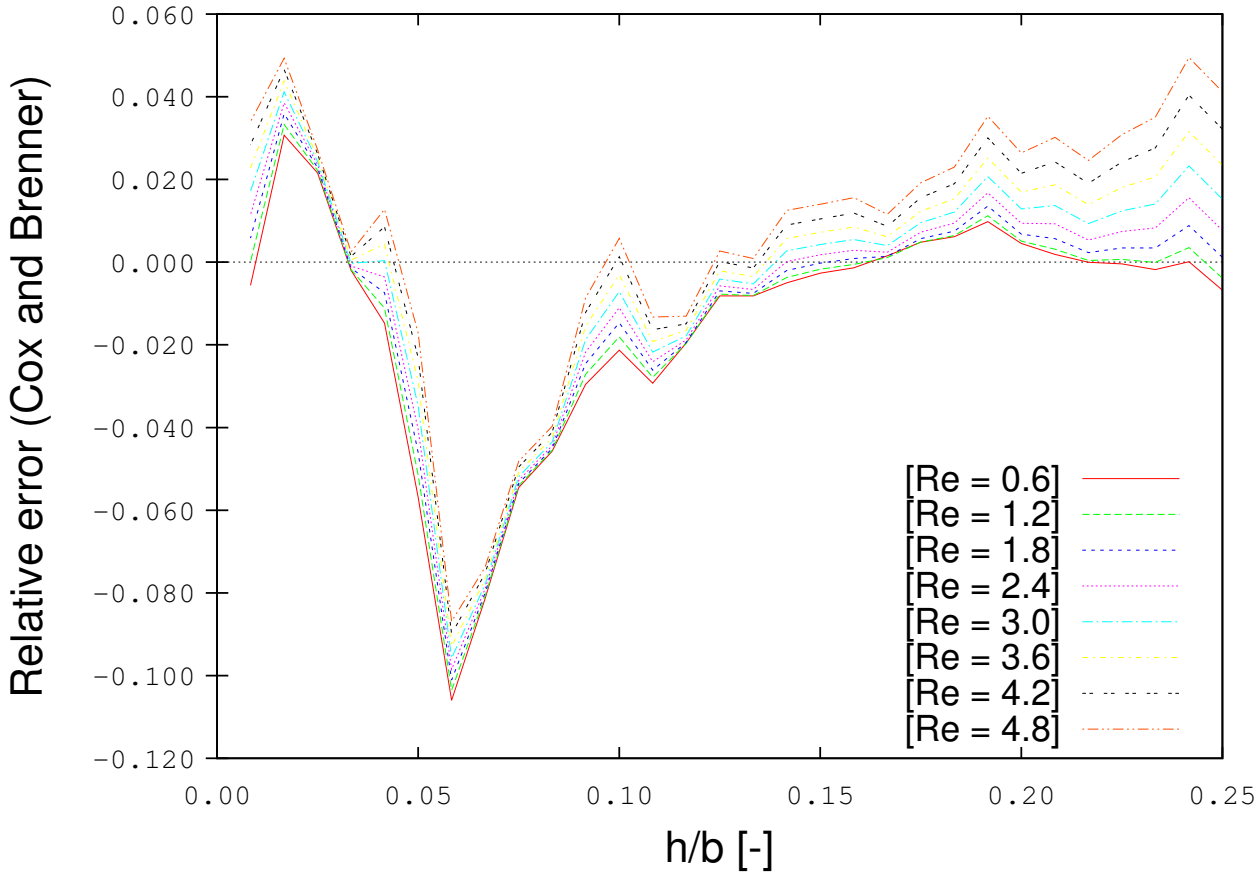


Figure 5.12.: Graph showing the error of lubrication force for the corrected method relative to the low Reynolds number prediction of Cox and Brenner, for a range of particle Reynolds numbers.

5.1.2. Dynamic case

Previous simulations, section 5.1.1, consisted of particle approaching each other with constant velocity. In this section the results of simulations with a free moving sphere colliding with a stationary sphere are presented. The aims of these simulations are:

- Validate the force velocity relation over a range of velocities
- Determine the influence of the short range lubrication force on the overall dynamics.
- Determine the correct time integration of the force.

Two simulations of a moving particle initialised to be on a collision course with a fixed particles are run. In both cases the moving particles start a at a substantial distance form the fixed particle and held at a fixed velocity, moving towards the stationary particle. The fixed velocity is maintained such that the flow has developed well before the particles come into contact. Once the particles are within a third of their radii from each other (*i.e.* $h = b/3$) the fixed velocity constraint is dropped and the particles are governed by the forces acting on them. The properties used are given in table 5.3. The difference between the simulations is the initial rate of approach and the properties of the particles:

Case: Slow, low density has initial approach velocity of $0.1\ \mathrm{m\,s^{-1}}$ and a particle density $1\ \mathrm{Kg\,m^{-3}}$.

Case: Slow, high density has initial approach velocity of $0.1\ \mathrm{m\,s^{-1}}$ and a particle density $10\ \mathrm{Kg\,m^{-3}}$.

Case: Fast, low density has initial approach velocity of $1.0\ \mathrm{m\,s^{-1}}$ and a particle density $1\ \mathrm{Kg\,m^{-3}}$.

Case: Fast, high density has initial approach velocity of $1.0\ \mathrm{m\,s^{-1}}$ and a particle density $10\ \mathrm{Kg\,m^{-3}}$.

Cases **Case: Slow, high density** and **Case: Fast, low density** are expected to behave similarly. This expectation is as a result of the fact that they have a comparable momentum. Additionally, **Case: Slow, high density** is expected to have a lower force, but act for a longer time than **Case: Fast, low density** with a similar overall momentum exchange per unit distance travelled.

Figures 5.13 and 5.14 show the velocities of the moving particles. Figures 5.15 and 5.16 show the repulsive force between the moving particles and the surrounding fluid. In each of the cases the force on the approaching body and on the stationary body are given.

The cases show three different regimes. **Case: Slow, low density** shows a rapidly decelerating particle. In this case, while the relationship between force and velocity is dependent on the proximity

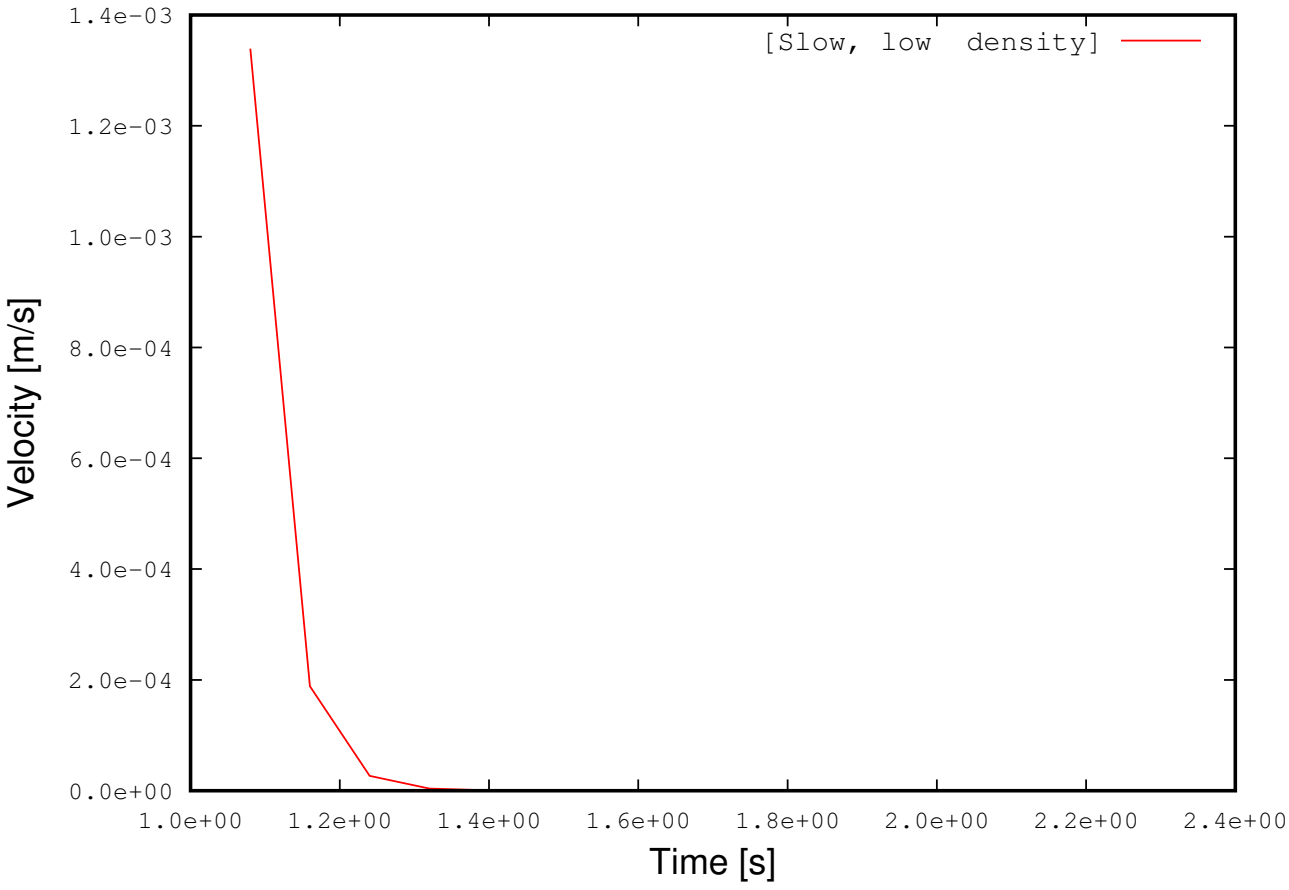
Particle radius	0.015 m
Dynamic viscosity	$1.0\ \mathrm{N\,s\,m^{-2}}$
Fluid density	$1.0\ \mathrm{Kg\,m^{-3}}$
Refinement	7.5 cells per radius

Table 5.3.: Fluid and particle properties for simulation of low Reynolds number colliding spheres.

of the particle, the exact form of the lubrication force is not as significant as for other cases. This is because the expected behaviour of the particle, *i.e.* coming to a halt near to, but not colliding with, the sphere would occur with despite significant under or over-estimation in lubrication force when the two particles are in close proximity.

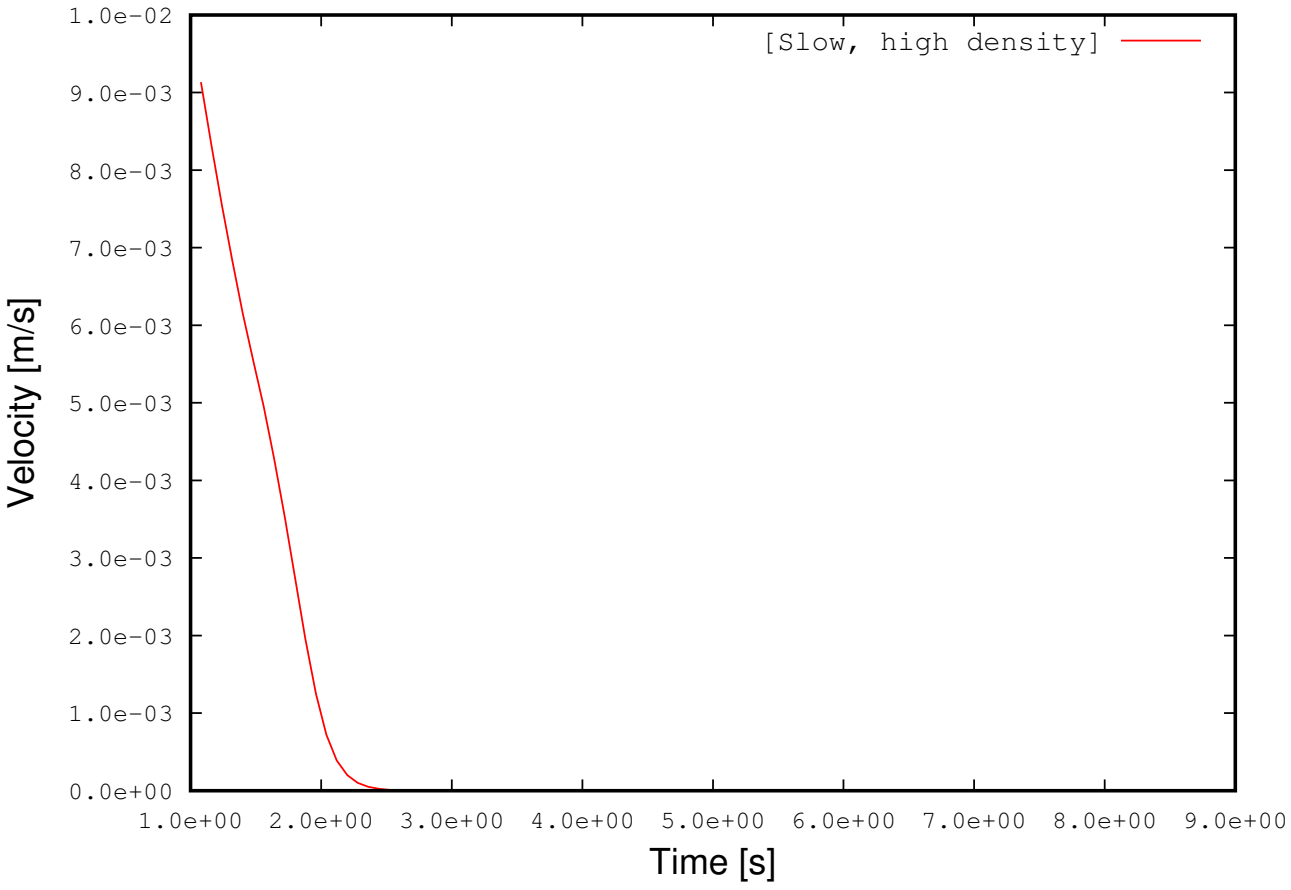
Case: Fast, high density depicts a substantially different case where the exact form of the repulsive force is significant. The majority of the deceleration of the particle occurs within a very small fraction of the particle's radius, as seen in figure 5.19. Indeed in this case the outer radius of the two discretised surfaces overlap. In this particular case, the lubrication force correction is critical. Figure 5.19 shows the difference in velocity and force over time, in case **Case: Fast, high density**, with and without the developed lubrication force correction. In this case, without the correction, a significant particle-particle collision takes place and the particle rebounds with a significant velocity. Hence in this case the correction developed in section 5.1.1 is necessary to accurately predict the behaviour of the system.

Case: Slow, high density and **Case: Fast, low density**, depicted in figures 5.14 and 5.16 exhibit transitional behaviour. In particular, the effect of the close range lubrication force, and hence the correction, is significant but does not completely dominate the behaviour of the system. The transition of hydro-dynamic interaction forces as the particles approach each other gives rise to the complex form of the repulsive force observed *i.e.* decreasing followed by a sharp increase followed by a decay to zero. This can be explained by the decreasing velocity of the particle, coupled with the close range interactions of the bodies. The difference in the forces acting on the stationary and moving particles shrinks rapidly during, and flowing, this peak in force. This is as expected, as this represents the point at which the majority of the force acting on the particles is as a result of the thin film of fluid between the particles.



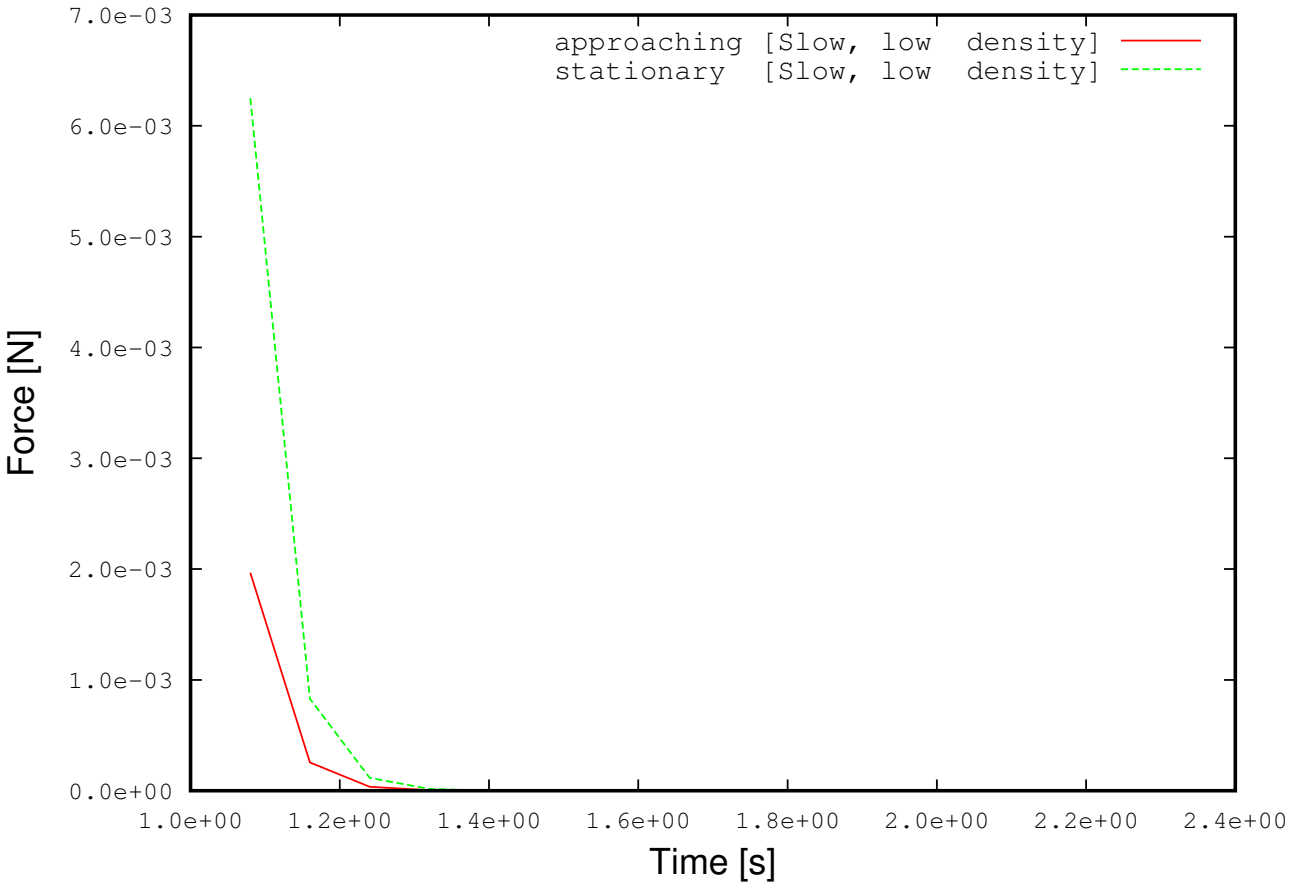
(a) Slow, low density

Figure 5.13.: Graphs showing the velocities of colliding spheres, decelerating under lubrication force for an impacting [5.13b] and a non-impacting [5.13a] case.



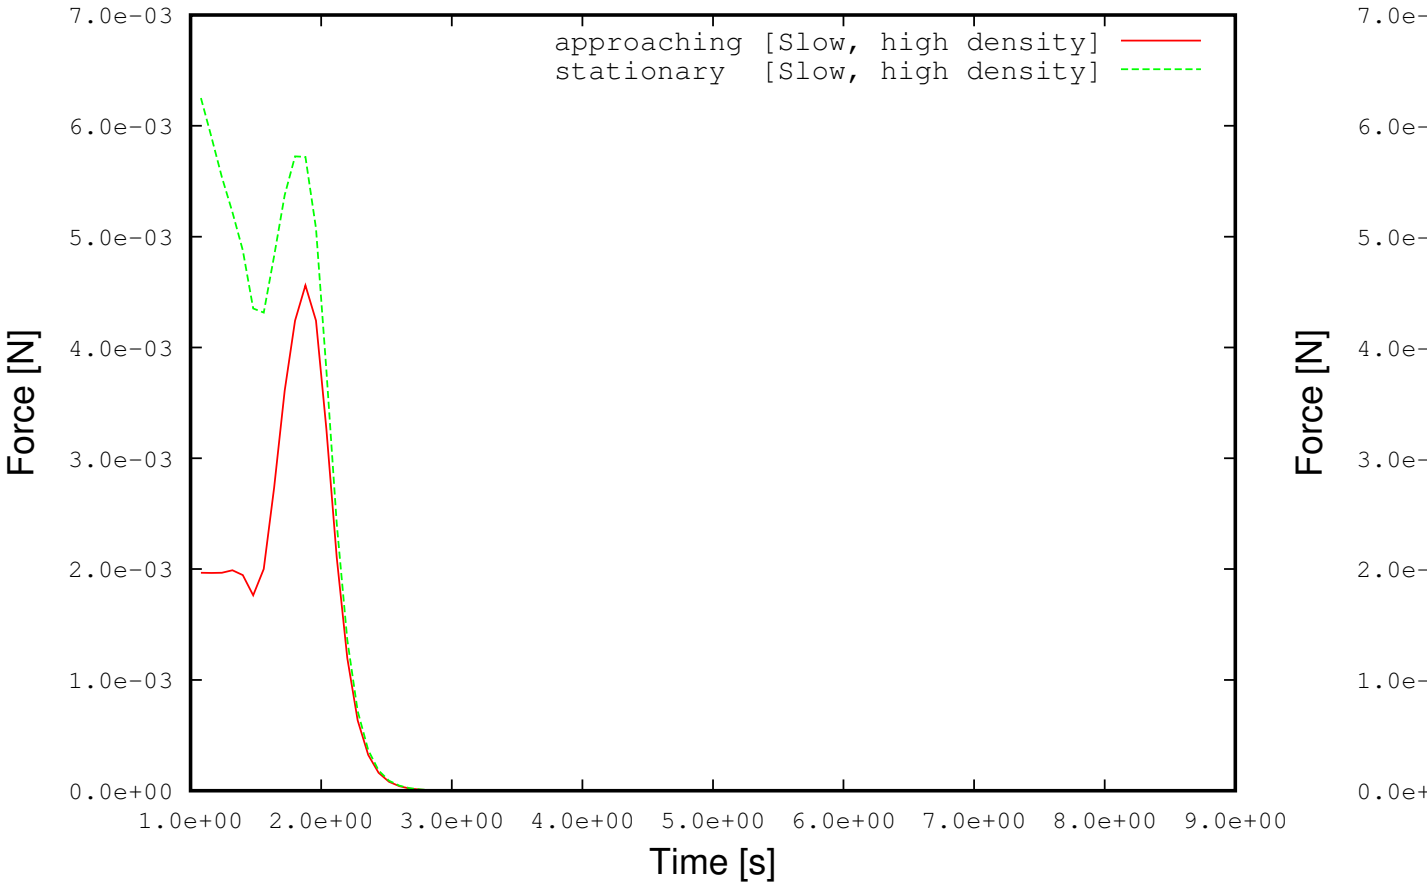
(a) Slow, high density

Figure 5.14.: Graphs showing the velocities of colliding spheres, decelerating under lubrication force for two transitional cases.



(a) Slow, low density

Figure 5.15.: Graphs showing the repulsive forces between colliding spheres, decelerating under lubrication force for an impacting and none impacting case.



(a) Slow, high density

Figure 5.16.: Graphs showing the repulsive forces between colliding spheres, decelerating under lubrication force for two transitional cases.

Figures 5.17 and 5.18 give the force as a function of non-dimensionalised separation (h/b) for **Case: Slow, high density** and **Case: Fast, low density** respectively. In each case the force on the stationary and moving particles are given, along with the contribution from the correction and the theoretical result for close contact. The form of this graph is of interest as it determines where the energy dissipation takes place. In both cases the forces acting on the colliding spheres is negligible within 2% of the radius of each other, as the particle as slowed to a stop. Hence the particles are not substantially affected by the lubrication forces at this or smaller distances.

Due to the small size and slow seepage velocity, this is thought to be more energetic than the most energetic type of collision expected inside a granular material between a free flowing particle and the fixed matrix. However if a clump of particles were to collide, the effective mass of a large group might make this type of high-momentum-impact case significant.

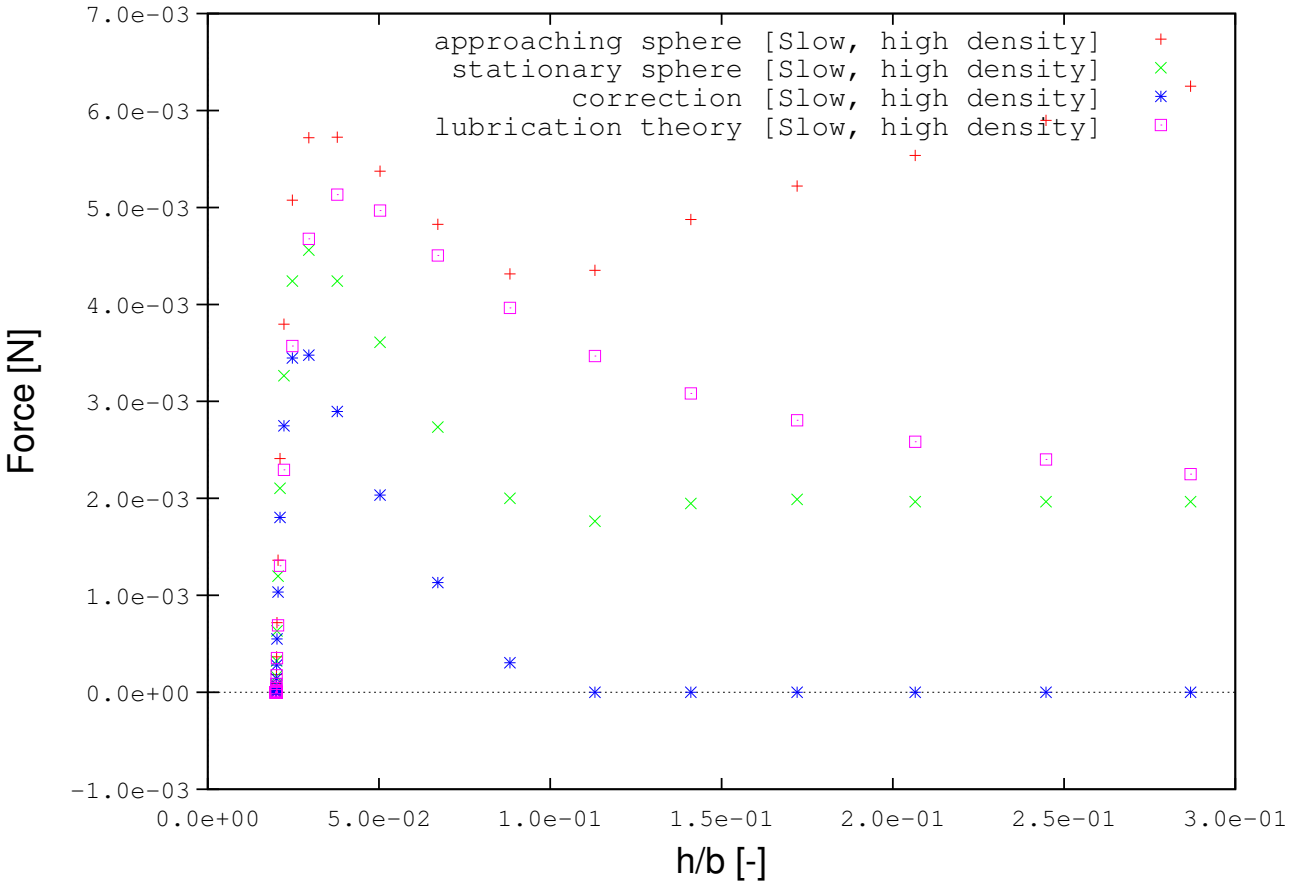


Figure 5.17.: Graph showing the dependence of the lubrication force on separation for a slow collision. The force on the stationary particle, and the moving particle, is shown. For comparison, the corrective sub grid scale contribution, and the singular term from theory (Cox and Brenner) are shown.

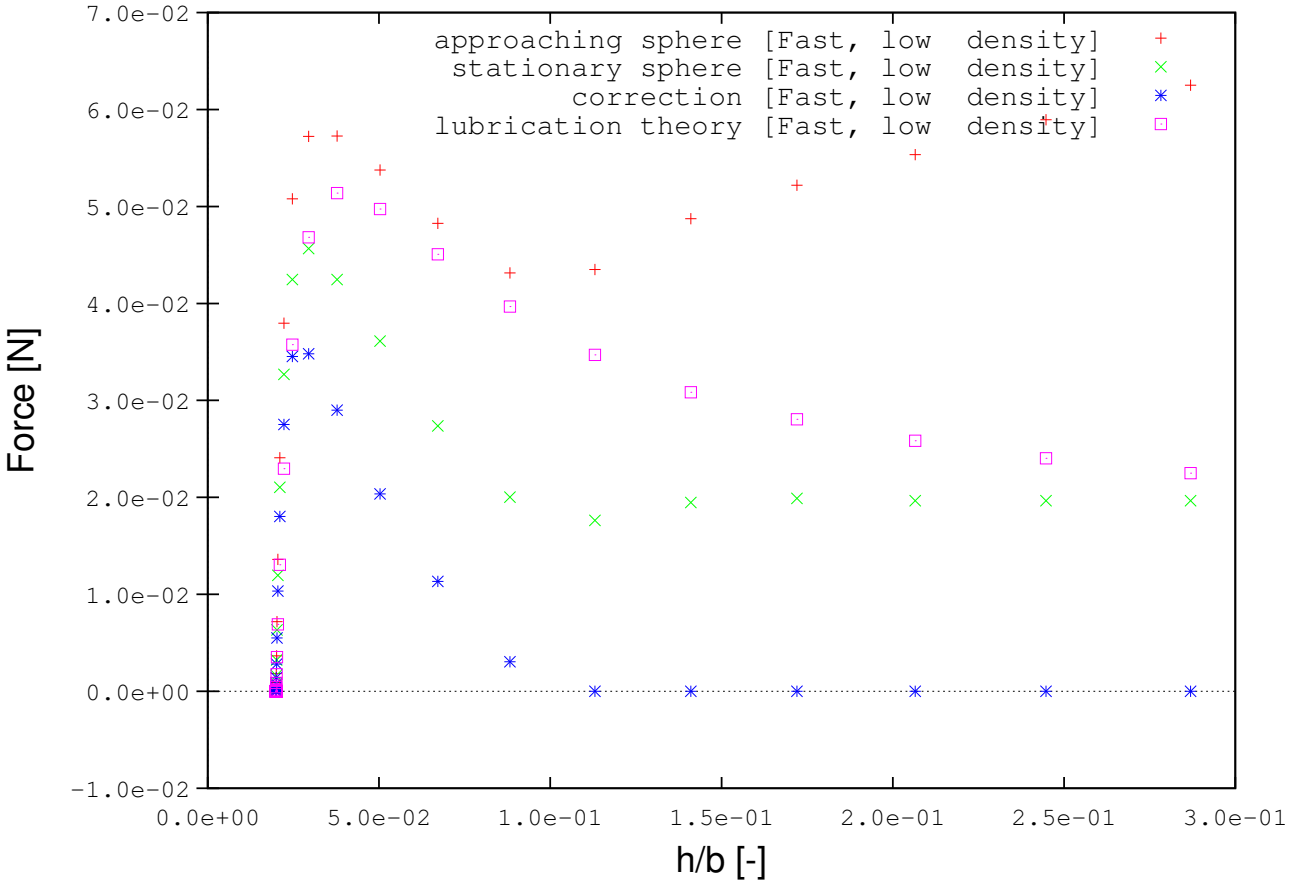
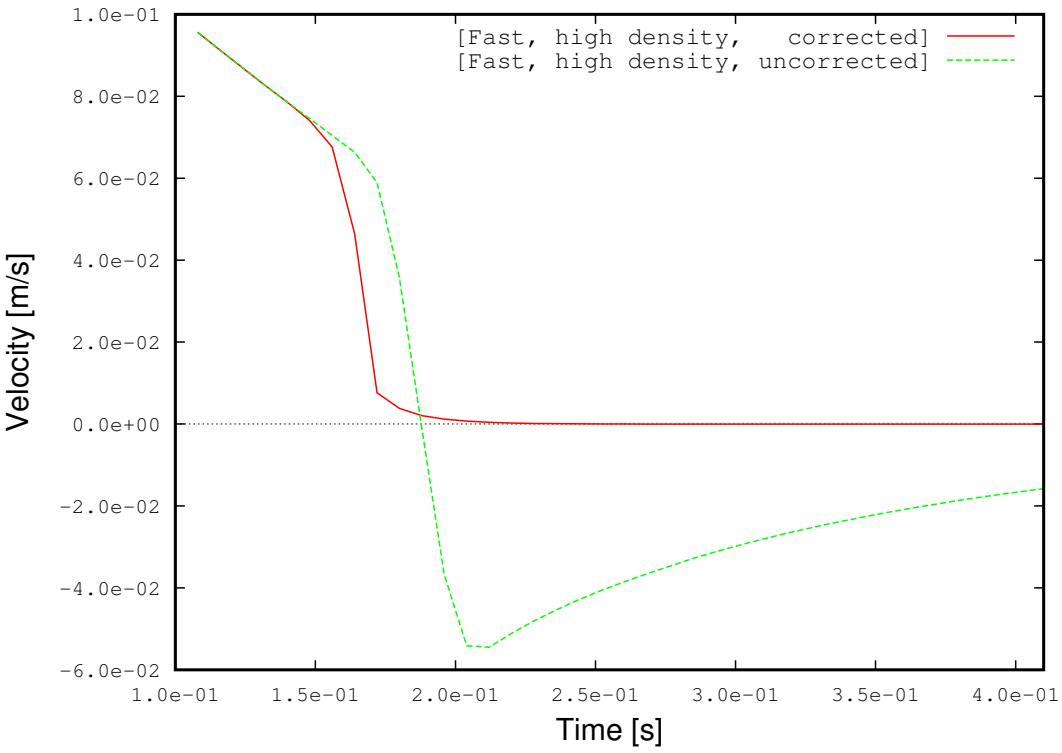


Figure 5.18.: Graph showing the dependence of the lubrication force on separation for a fast collision. The force on the stationary particle, and the moving particle, is shown. For comparison, the corrective sub grid scale contribution, and the singular term from theory (Cox and Brenner) are shown.



(a) Velocity

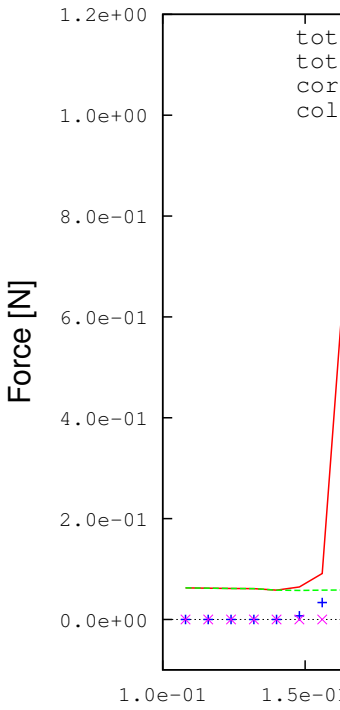


Figure 5.19.: Graphs showing the repulsive forces and relative velocity between colliding spheres (**Case: Fast, high density**) with and without the correction to lubrication. The uncorrected case gives rise to a significant particle-particle contact and a rebound occurs. When the corrective force is present it halts the particle before significant collision hence no rebound is made, which agrees with theory.

5.1.3. Impact of force calculation method

The force calculation method, used to estimate the repulsive forces acting on the particles, is a potential source of error. Indeed while the momentum equations are solved using a cut-cell method, the resultant pressure and velocity gradients are interpolated back to the triangulated body over a larger stencil. This interpolation introduces the potential for error in the force calculation as the assumption that the fluid properties vary smoothly over these regions is likely to breakdown. While attempts have been made to account for this (*i.e.* by using a reduced stencil if the points no longer lie in the fluid domain see sections 3.3 and 3.3.4) the potential for inexact results due to interpolation remains. Using a more finely tessellated mesh to describe the particle's surface improves the accuracy of the force calculation slightly as the larger number of triangles, of smaller area, allow the rapidly changing properties to be integrated more precisely.

5.2. Collision dynamics

While lubrication force is the major determinant of force acting on particles in close proximity, but not touching, granular materials are invariably dominated by particle-particle contact. Using the method detailed in section 3.4 simulations of colliding pseudo-spherical particles of varying refinements are used to determine the accuracy. The results are compared to the expected force predicted by Hertzian dynamics for two equivalent spherically idealised particles. This equivalent for the triangularly tessellated particle is not obvious, as, for lower refinements, the surface's distance from its centre is not constant. In this section, the equivalent sphere is taken to be one whose radius is half way between the minimum and maximum distance between the tessellated surface and its centre. The correction is significant as, for the expected small overlaps, this difference may be a significant fraction of the overlap. Figure 5.20 shows the force between two particles, during a collision, for a range of particle refinements.

During what could reasonably considered as 'contact' between two immersed particles there may be some interstitial fluid separating the particles. However, in this case a solid-solid contact model is still applicable as the resitutive force is determined by the physical properties of the granular phase. The dominance of the particle phase properties is as a result of the major source of compliance between the particle pair being the solid phase deformation. Figure 5.21 illustrates an example of such a wet-contact. In a case such as illustrated in figure 5.21 no lubrication correction is required as the particle-particle contact force is an accurate representation of the combined 'actual' contact and contact through a very thin interstitial fluid. The transitional period is more complex. Ideally a separate equation system, governing the interstitial gap size, δ_l, would be solved to predict the combined repulsive force of contact and lubrication force by balancing the deformation force, f_d, with lubrication force, f_l.

$$f_t(\delta_t) = f_d(\delta_d) = f_l(\delta_l) \tag{5.11}$$

$$\delta_t = \delta_d + \delta_l \tag{5.12}$$

However the deformation for is not known in closed-form hence solving this relation is difficult. Furthermore, the curvature of the new particle contact surface curvature is not explicitly known and will alter the lubrication force for a given interstitial gap size. Instead a the two forces are mollified

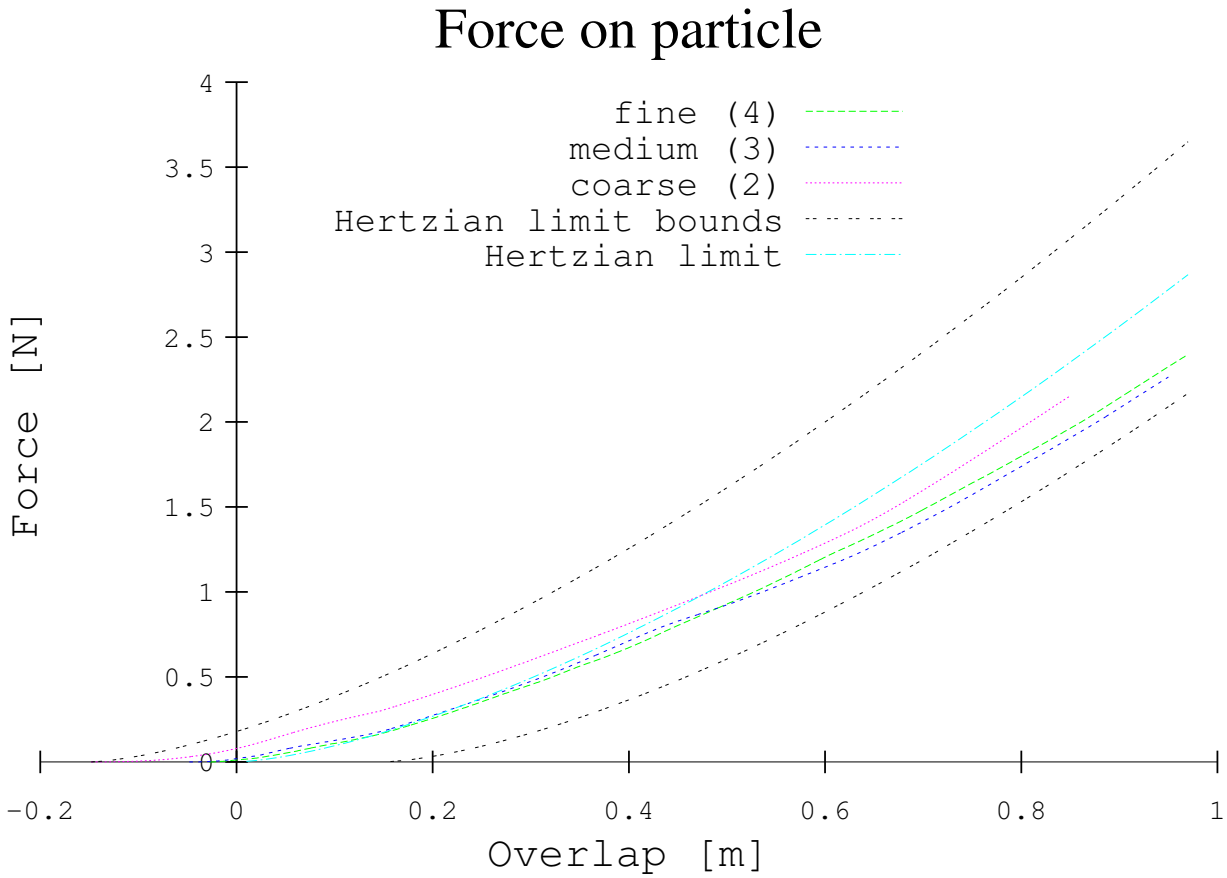


Figure 5.20.: A plot of the collision forces between two particles of radius 3m. Also included is the result as predicted by Hertzian contact. Bounds are given as, especially for the unrefined case, there is no clear analog of radius. Refinement is indicated by the number of subdivisions from a base sphere (n: 2, 3 and 4). The total number of triangles for each sphere is 20×4^n.

together such that before the un-deformed particles are in close proximity the lubrication force is dominant, after the un-deformed contact, the collision force is dominant. An example of this is given in figure 5.22 where the lubrication force of the un-deformed particle (blue line) and the collision force ignoring lubrication (red line) are combined into an overall force (black line). In this work the combined force (f_t) is given by:

$$f_t = f_d + \phi_L f_L \tag{5.13}$$

Where f_d is the dry collision force and f_L the lubrication force of the non-deformable particles. ϕ_L is a weighting coefficient for the lubrication correction given by:

$$\phi_L = \begin{cases} 1 & \text{if } h < 0 \\ \frac{kb+h}{kb} & \text{if } h > -kb \\ 0 & \text{if } h < -kb \end{cases} \tag{5.14}$$

For some constant, k, the fraction of the averaged particle radius, b, over which the transition takes place. The choice of k is fairly arbitrary and depends on the expected relative deformations of the bodies and interstitial fluid separations. In this work k is taken to be 0.005.

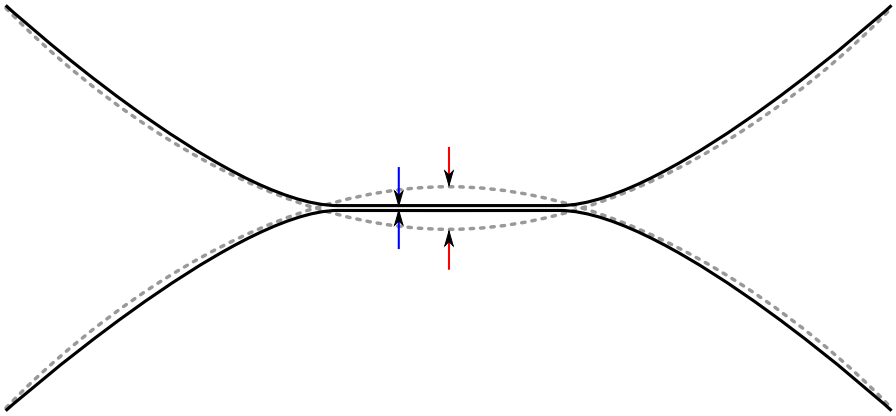


Figure 5.21.: An example in two dimensions of wet contact between two particles showing a fluid interstice and the deformation relative to the original particles. The small amount of interstitial fluid separation (blue) does not make a significant difference to the overall deformation (red). Hence calculating the restitutive force based on the overall deformation is valid.

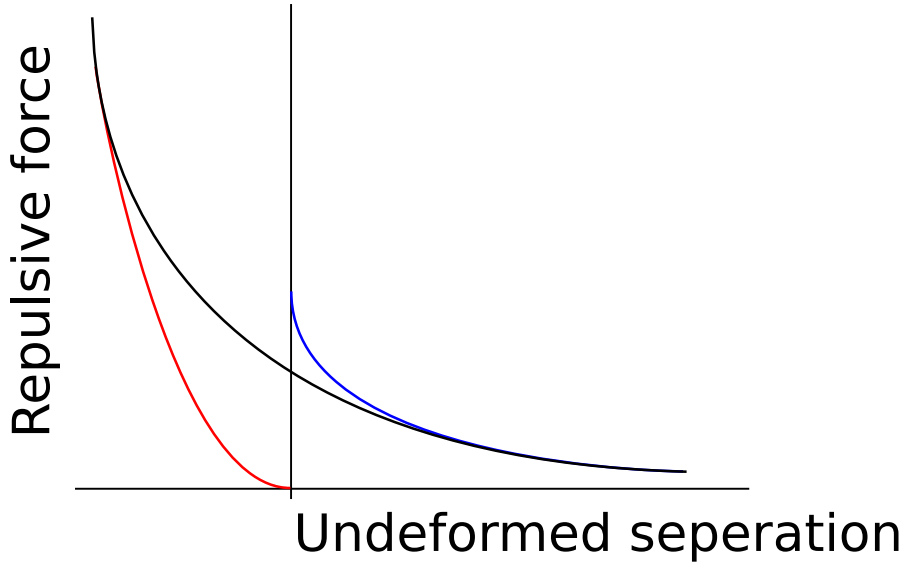


Figure 5.22.: Example of how the lubrication force of the un-deformed particle (blue line) and the collision force ignoring lubrication (red line) can be combined into an overall force (black line).

5.3. Discussion and conclusions

The lubrication forces acting between close particles is potentially significant in determining the dynamics of a granular system. The effects of lubrication force are important as, depending on the momentum of the particles and their separation, potentially very large forces may occur. Such large forces *may* occur during collision of the particles if the lubrication force were not present. However, as the force due to hydro-dynamic interaction is dissipative and not elastic, the effect of lubrication force is still significant.

The developed cut-cell partial mirroring method shows an excellent ability to model the repulsive force between two approaching spheres when they are removed from each other by a few Eulerian cells. In particular the predictions of Cox and Brenner [66] are well matched. However the lubrication force is under predicted for particles in close proximity relative to cell-spacing. The inaccuracy in

the prediction of lubrication force, when particles are in close proximity, is to be expected as its determination is based on flow that occurs at sub-grid scales. To account for this, a sub-grid scale correction is applied which serves to correct the overall prediction of hydro-dynamic interaction forces. With this sub-grid scale correction, good agreement is found with theoretical results. In particular both the predictions of Cox and Brenner [66] and the predictions of dimensional analysis are found to hold extremely well at all separations.

A collision model is developed that does not rely on the sphericity of particles, only the local curvature. Hence arbitrary shaped bodies can be used both for fluid-particle and particle-particle interaction. Collision forces between tessellated surfaces in dry contact are found to have excellent agreement with Hertzian theory, upon which they are based. This agreement is demonstrated by the prediction of force between two approximately spherical particles, which is found to be in good agreement with an idealised case. Indeed the force is found to be within the range predicted by the inscribed and circumscribed spheres, for a range of tessellation refinements.

Combining the lubrication force with a particle-particle contact model is somewhat complex, as the theoretical predictions for lubrication force assume non-deformable surfaces. Furthermore the repulsive forces between particles in an actual collision are dependent on subtle variations, such as surface roughness. Lubrication force and collision forces are combined in the developed model by assuming a deformation of both the interstitial fluid and the particles. While this separation is not exact, as the dynamics in the case of particle-particle contact are somewhat independent of which force contributes as both forces act to stop relative motion, this is not prohibitive. However some differences may occur with regards to tangential forces and collision damping hence further work in this regard would be welcomed.

6. Flow through granular materials

6.1. Constrained particles

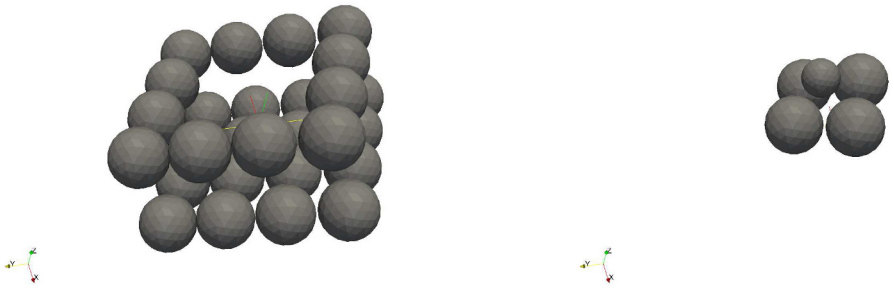
6.1.1. Simulation and objectives

Two simulations are used in order to show the effects of the fidelity of the immersed-boundary method on the behaviour of a granular material. Each case simulates five freely moving particles, restrained by a 'cup' of 28 fixed stationary particles. The system is driven by a fluid flow, with constant volumetric flow rate between the inlet and the outlet. Figure 6.1 shows the arrays of particles used. The case is repeated using both spherical particles and non-spherical particles, constructed by the randomised addition of asperities to initially spherical particles. The method for constructing these non-spherical particles is given in appendix A.2. The parameters used are given in table 6.1. Figure 6.2 shows an example of a non-spherical particle constructed with the method in appendix A.2, with the parameters used for the simulations (table 6.1). Also shown alongside the non-spherical particle, in figure 6.2, is the 'base' sphere from which it was constructed. For the purpose of comparison, the constructed non-spherical particles are assigned a effective radius, r^*, equivalent to the radius of the base sphere. The intention of these simulations is to provide insight into the effects of the sensitivity of the response of a granular system to small-scale adjustments. In particular the effects of sub-particle scale alterations are of interest as these are often neglected by geo-technical models. This is achieved by comparing the results to an idealised case and also by comparing the results of the two similar simulations, differing by the roughness of the particles used. The simulations also serve to verify the sensibility of the simulation method.

The motivation for the use of a combination of a fixed array of particles and some freely moving ones are two fold. Firstly it is tractable, as only a few particles are required. Secondly it replicates a case of interest namely suffusion, introduced in section 1.2.1.1, which is widely regarded a major threat to internal stability, Garner and Sobkowicz [3]. Suffusion is governed by the interaction of particles that are free to move relative to a matrix of particles that are constrained by the stresses acting on the material. Hence the modelling of this type of case is of interest.

In both of the simulations, the free particles are initialised with a velocity equal to the average flow rate. The particles start with their centres slightly staggered in the direction of the flow (x). Hence the cases are not exactly symmetric in the x–y or x–z planes. Table 6.2 gives some of the properties of the simulation.

Figure A.1 in appendix A.3.1 shows a similar simulation for a larger number of freely moving particles. Interpreting the results meaningfully in this case is difficult as the complexity of the case makes phenomena difficult to isolate.



(a) Fixed constraining particles

(b) Freely moving particles

Figure 6.1.: Setup of particles used to simulate the motion of constrained particles. Figure 6.1b shows the five free moving particles of interest. Figure 6.1a shows 28 fixed particles used to prevent the through flow of the unconstrained particles. These fixed particles represent a constrained matrix of particles characteristic of a granular system. The five free moving particles are placed upstream, negative x direction, of the 'cup' of static particles. The flow is allowed develop around all of the particles before any significant particle-particle interaction occurs. Figures 6.4a and 6.7a show the combined array of particle near their initial arrangement, for the spherical case.

Number of asperities	32
Parameter a	0.3 r^*
Parameter b	0.1 r^*
Parameter c	0.6
Parameter d	0.0

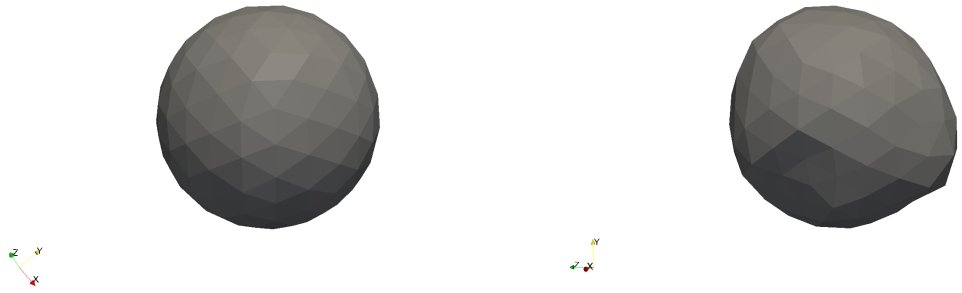
Table 6.1.: Parameters used to determine the roughness of the non-spherical particles, as detailed in appendix A.2, with r^* the 'base' radius of the particle.

6.1.2. Model used

The model used for the simulations of spherical and non-spherical particles is the developed mirroring immersed-boundary method using both the cut-cell partial mirroring method and the average mirroring modification. A lubrication correction force was also employed in each case, in particular the correction used in conjunction with collisions, as described in section 5.2, was used. The correction did not account for the roughness of the non-spherical particles. While this is not ideal, as the particles are approximately spherical, the approximation is unlikely to be a significant source of error. There are two cases of concern with regards to the prediction of lubrication force for non-

Reynolds numbers	~ 45
Stokes number	~ 60
Density ratio $\left(\frac{\rho_p}{\rho_f}\right)$	7.3
Refinement large spheres	3.2 cells per radius
Refinement small sphere	2.4 cells per radius
Max CFL number	~ 0.2

Table 6.2.: Properties for the flow over constrained particles. Reynolds number is calculated for the larger of the two types of particle, using the average flow rate.



(a) Spherical particle (b) Non-spherical particle

Figure 6.2.: The two types of particle for which the simulation is run. Figure 6.2b shows the tessellation of a non-spherical particle as described in appendix A.2. Figure 6.2a shows the tessellation of an almost spherical particle with radius, r, equal to the effective radius of the non-spherical particle, r^*.

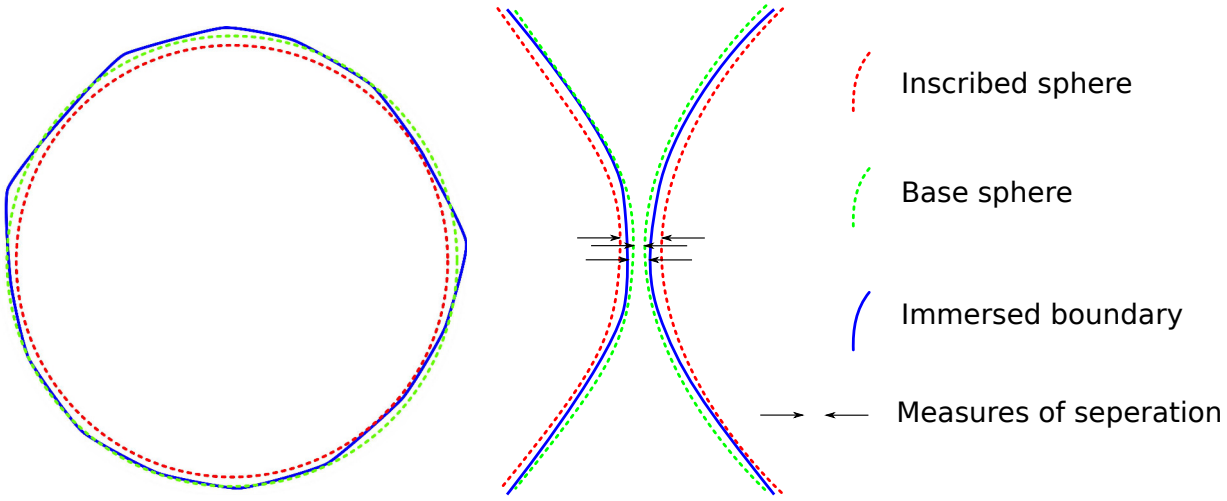
spherical particles. For each case, for small asperities, it is reasonable to compare the changes in lubrication force as a result of asphericity directly with the sub-grid correction as the un-corrected hydro-dynamic interaction force does not vary significantly as a result of the geometry in the vicinity of the collision. These cases are:

- In the case that the separation distance is smaller than expected due to the asperities, *i.e.* the particles are close to contacting as a result of their protrusions, the correction will be lower than if the correct spacing is used. However in this case the curvature of the surfaces is likely to be much higher than for the base spheres, as the protrusions have more curved surfaces. Hence the actual lubrication force should be lower than the prediction using the base sphere's curvature and the rough sphere's separation.
  The overall effect of the combination is an under-prediction of the magnitude of the force hence is likely to lead to an under-prediction of the necessary correction. Despite the inaccuracy, the use of this under-correction improves the accuracy of the lubrication force prediction relative to the use of no corrective sub-grid contribution at all, a greater under-prediction.

- It is possible the minimum separation between two non-spherical particles is greater than that between two equivalent spherical particles. This is not a likely occurrence as it requires the absence of any protrusions in the vicinity of the contact. In this case the sub-grid scale contribution may be greater than would be correct for the given case. It would be possible to prevent this over prediction by applying the sub-grid scale contribution to the lubrication force based on a reduced particle radius. The reduced particle radius, r^m, being the minimum separation from the particle centre to the surface. An example of the scheme using r^m is shown in figure 6.3, alongside r^* and the actual separation.

However, in the particular situation of the non-spherical simulation in this section, the sub-grid contribution acts over a distance greater than the asperities' height. Indeed the sub-grid contribution is applied over a single grid cell – approximately a third of a particle radius – which is greater than the size of the asperities' height. Hence the overall effect of the asperities is likely to be small relative to the overall sub-grid contribution. Additionally, while the sub-grid scale contribution may be larger

than appropriate, the contribution is bounded by the contribution from a slightly deformed particle. Hence any over-prediction is unlikely to cause unphysical behaviour.



(a) Non-spherical particle (b) Non-spherical particle collision

Figure 6.3.: Diagram of an immersed-boundary representing an non-spherical particle (blue), showing the base radius of the particle r^* (green) and the reduced radius r^m (red). Figure 6.3a shows the individual particle and figure 6.3b shows a two such particles in close proximity. Figure 6.3b also shows some of the possible measures of separation used for calculating the sub-grid contribution to the lubrication force.

One possibility to predict the lubrication force sub-grid scale contribution would be to use a similar principle to that used in the collision force (section 3.4) of either an expanded or translated copy of the immersed-boundaries to discretise the interstice between two arbitrary bodies. However the appropriate form of the lubrication force, given the resultant information, may not be obvious in all cases. For example if the curvature is varies significantly relative to the minimum separation distance or the particles have significant concavities. In this case is difficult to know the correct formulation of lubrication force as the collision is not dominated by a mutually convex approaching surface with constant curvature over a short radius. The cases shown in chapter 5 and the cases commonly dealt with in the literature are both restricted to this case. While the resolution of lubrication forces in cases without this restriction is certainly possible, perhaps the most efficient method would be highly resolved CFD, in which case the sub-grid correction term is redundant.

6.1.3. Particle motion

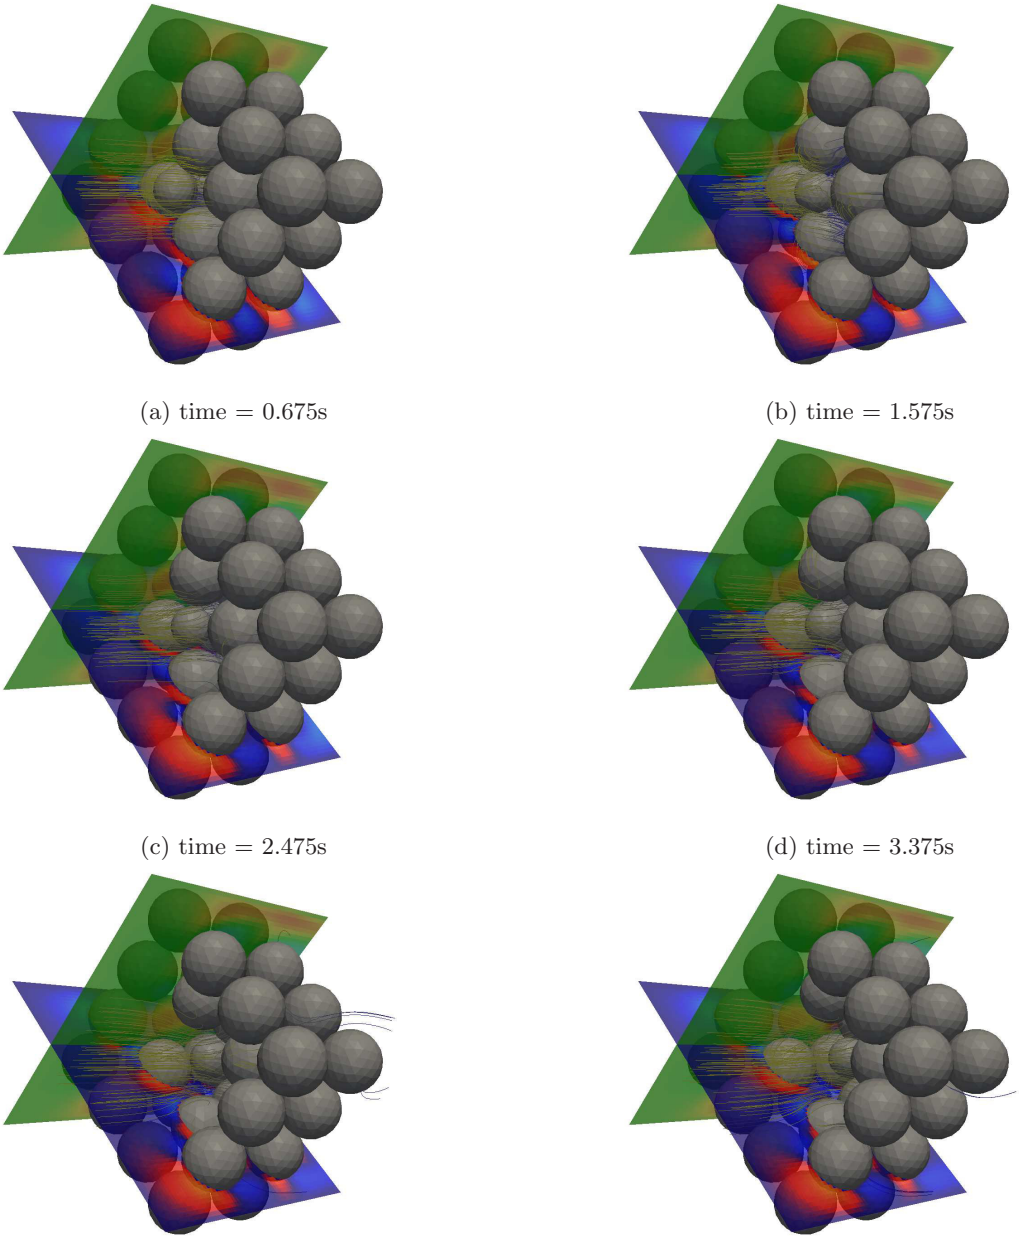
Figures 6.4 and 6.5 show the evolution of the five free particles in both of the simulations, spherical and non-spherical particles respectively, over time. Initially the five particles are advected with the flow, without substantial motion relative to each other. After the four larger particles are held up by the stationary particles the smaller central particle continues, colliding with the larger particles and driving them radially away from the centre.

In both cases the smaller of the five moving particles lodges in the central pore of the static restraining particles. The four larger particles spread out radially, due to the central high pressure flow caused by blockage of the central pore by the smaller particle.

Figure 6.7 shows a side view of the simulation using spherical particles and figure 6.8 the non-spherical particle case. These figures show differences between the two cases. A re-orientation of the

non-spherical particles allows some of them to move further outwards than the equivalent spherical particles. This can be seen by comparing the lateral displacement of the near particles in figures 6.7 and 6.8. This difference in resting position is significant as the further displaced particles more effectively restrict the flow through some the corner pores and side pores. This not only effects the overall flow through the pores but also the resulting force acting on the particles as the further constricted pore requires a greater pressure gradient to drive the flow.
The flow through the bottom layer of fixed particles, shown in figure 6.9, also shows this restriction. The non-spherical case shows a significant difference between the four corners, corresponding to the position of the larger particles. The two particles on the $y-$ side restrict the flow more significantly than the two on the $y+$ side, which is seen by the reduced flow through the corresponding pores. Despite the slight initial asymmetry of the original setup, the spherical particle case has roughly symmetric flow through the pores.

Figures 6.9 and 6.10 show the components of flow through the bottom layer of static particles. Figure 6.9 shows the component of the flow in the direction of the overall mass flow which is 'into the page' on the figure. Figure 6.10 shows a component of the flow tangential to the direction of the overall mass flow. The figures (6.9 and 6.10) show that the flow field varies both over the particle lengthscale and the domain lengthscale. In particular not only does the flow vary around individual particles but there is significant variation in the flow from particle to particle. It is particularly note worthy that in the current cases, the flow through four of the pores between the layer is significantly greater than the remaining five.



(a) time = 0.675s

(b) time = 1.575s

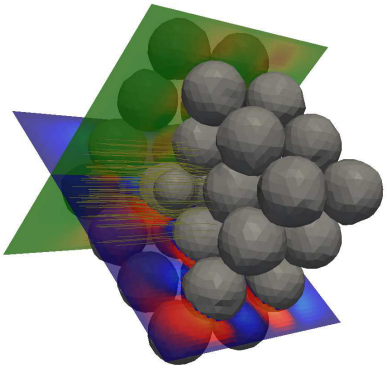
(c) time = 2.475s

(d) time = 3.375s

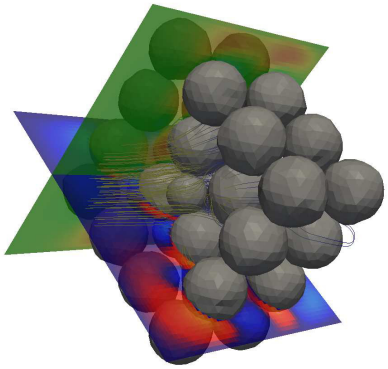
(e) time = 4.275s

(f) time = 5.175s

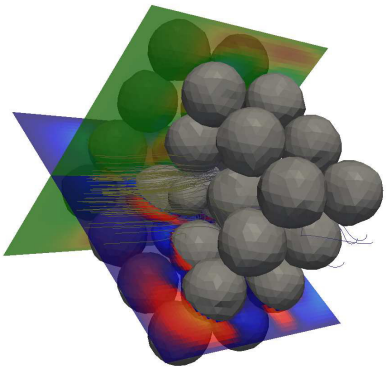
Figure 6.4.: Time evolution of the five spherical particles constrained by an array of fixed spheres at 0.9 second intervals. Streamlines show the flow relative to the fixed array of spheres, coloured by pressure. Semi-opaque slices show the normal and tangential components of flow through some of the channels between the particles. Figure 6.6 shows the scale for the fluid properties.
A contact between the central particle and the remaining larger particles starts at around $t = 0.9$ s , and ends around $t = 2.1$ s . Final contact between the central particle the constraining array of particle at the outlet starts at around $t = 5.0$ s .



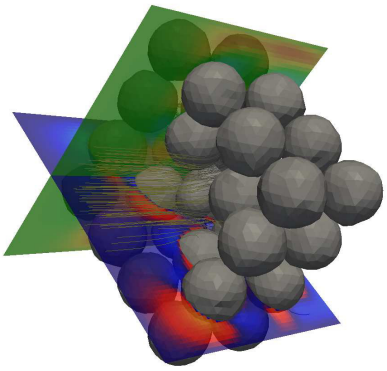
(a) time = 0.675s



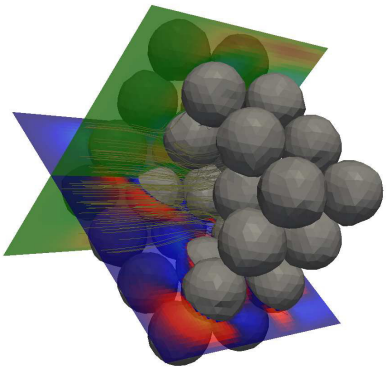
(b) time = 1.575s



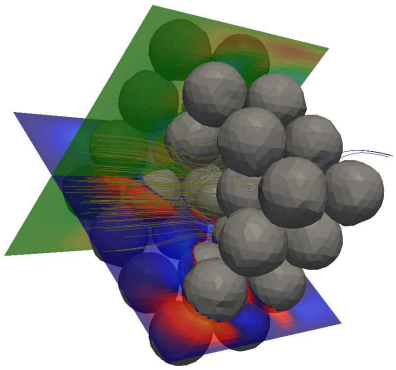
(c) time = 2.475s



(d) time = 3.375s



(e) time = 4.275s



(f) time = 5.175s

Figure 6.5.: Time evolution of the five non-spherical particles constrained by an array of fixed non-spherical particles at 0.9 second intervals. Streamlines show the flow relative to the fixed array of non-spherical particles, coloured by pressure. Semi-opaque slices show the normal and tangential components of flow through some of the channels between the particles. Figure 6.6 shows the scale for the fluid properties.
A contact between the central particle and the remaining larger particles starts at around $t = 1.0$ s , and ends around $t = 2.0$ s . Final contact between the central particle the constraining array of particle at the outlet starts at around $t = 4.6$ s .

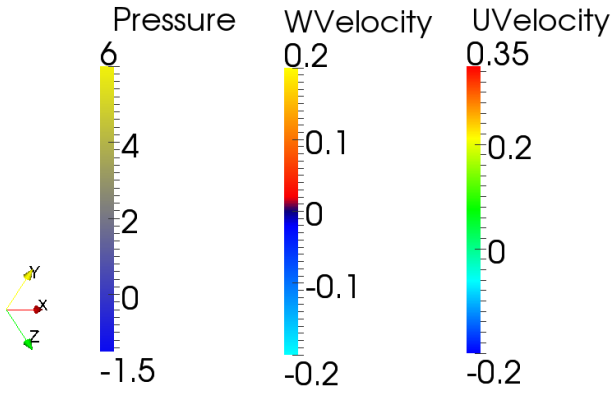
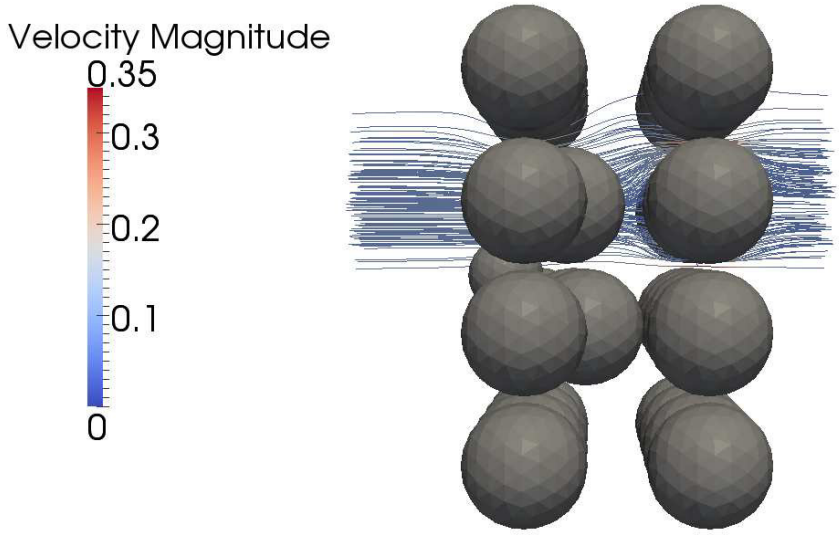
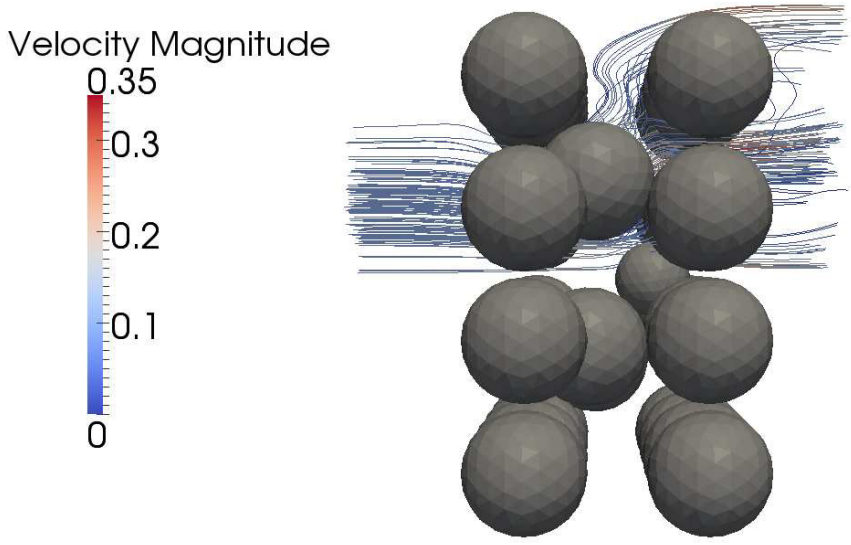


Figure 6.6.: Key for evolution of constrained particle figures: 6.4 and 6.5. The orientation axes and scales for pressure and the components of velocity are shown.

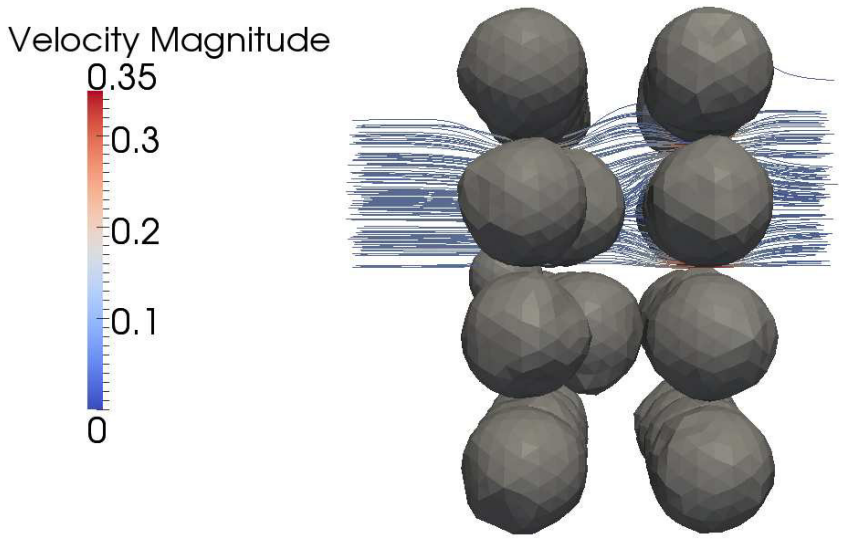


(a) time = 0.225s

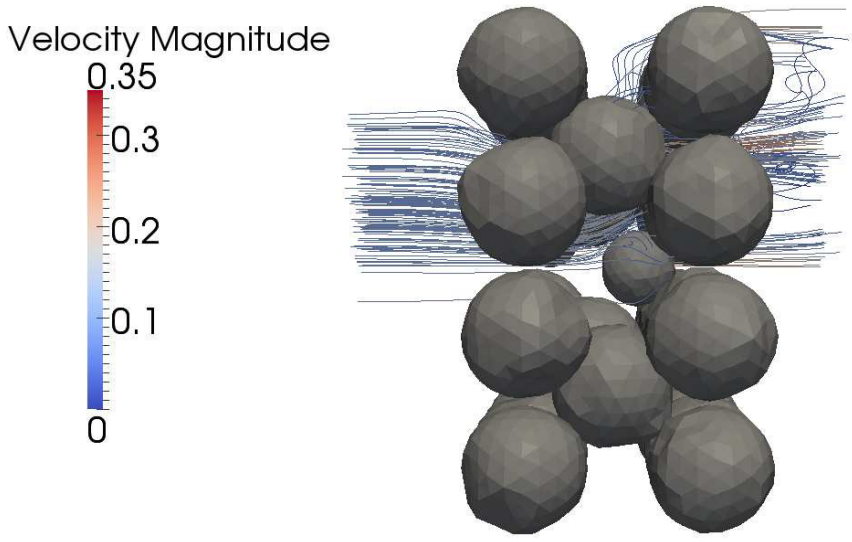


(b) time = 5.4s

Figure 6.7.: Side view of the progression of spherical particles with streamlines coloured by velocity magnitude. Figure 6.7a shows the beginning of the simulation and figure 6.7b shows a period after the particles have come to rest.

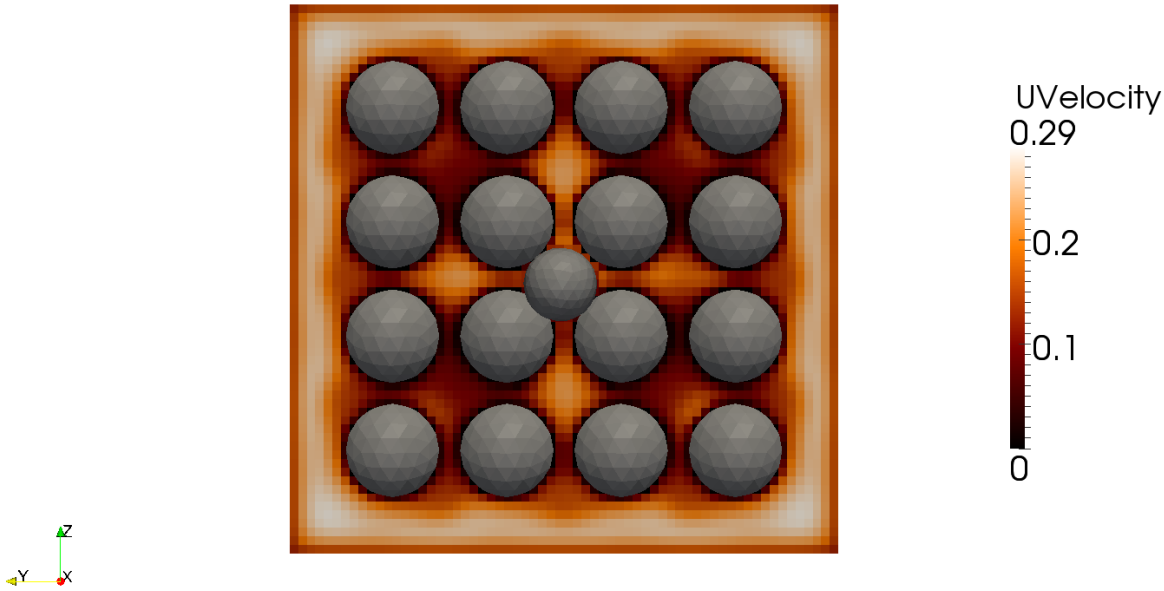


(a) time = 0.225s

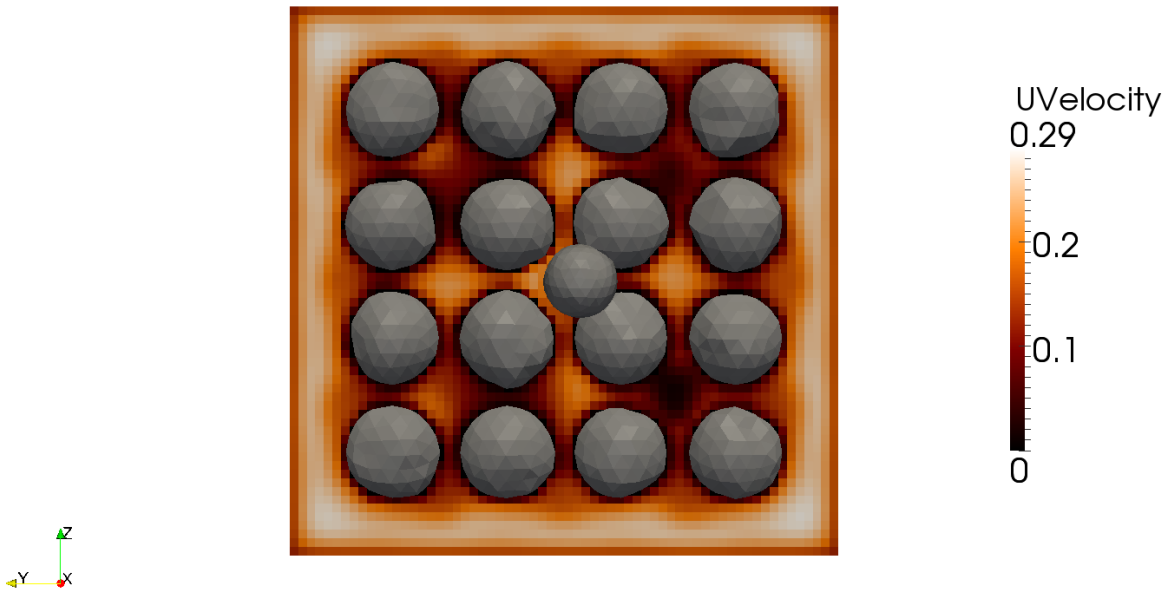


(b) time = 5.4s

Figure 6.8.: Side view of the progression of non-spherical particles with streamlines coloured by velocity magnitude. Figure 6.8a shows the beginning of the simulation and figure 6.8b shows a period after the particles have come to rest.

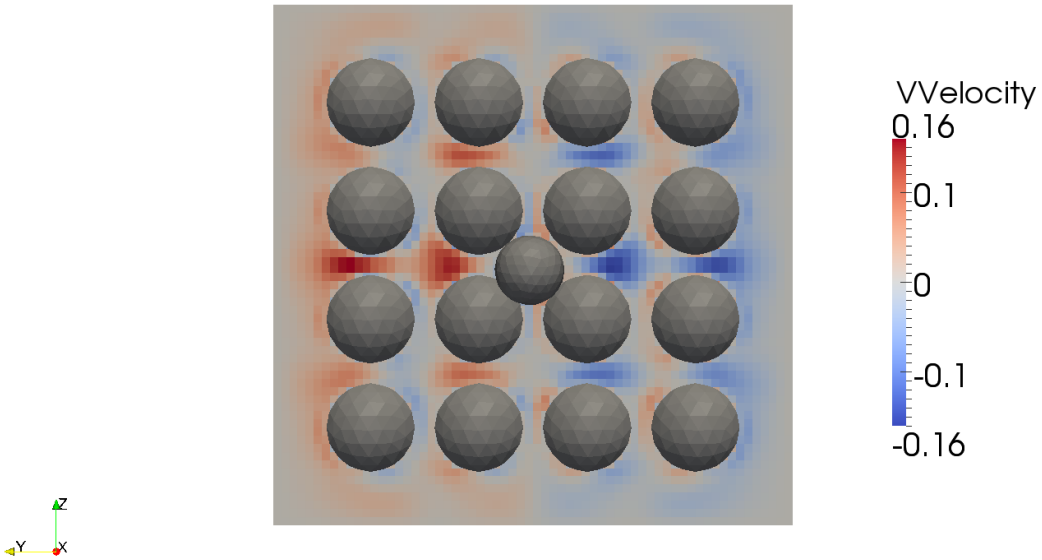


(a) Spherical case

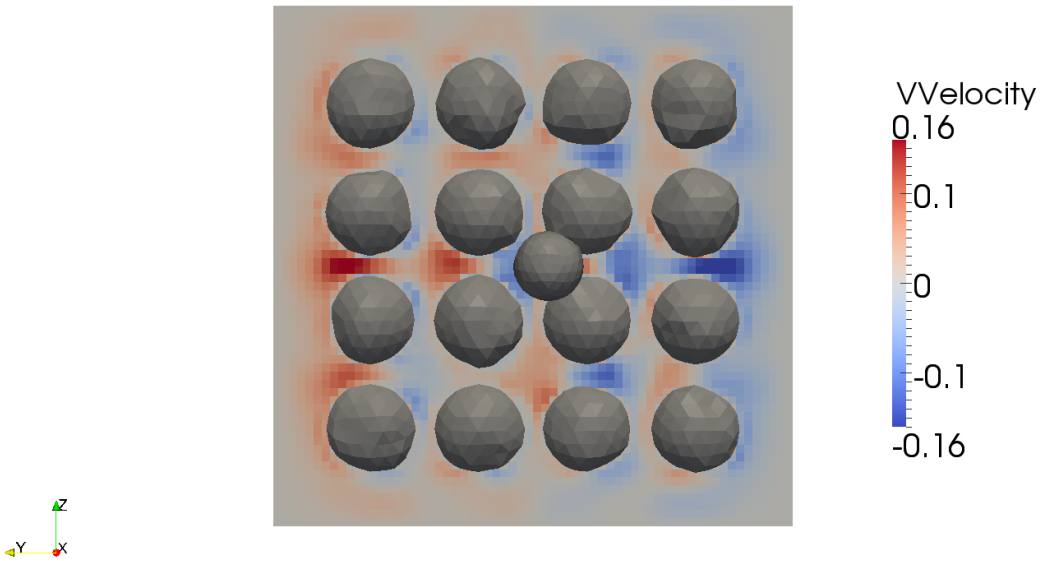


(b) Non-spherical case

Figure 6.9.: Component of velocity in the direction of the bulk fluid flow, through the final constraining layer of particles, shown for the spherical and non-spherical case



(a) Spherical case



(b) Non-spherical case

Figure 6.10.: Tangential component of velocity through the final constraining layer of particles, shown for the spherical and non-spherical case.

6.1.4. Hydro-dynamic interaction forces

Figures 6.11 through 6.16 show the hydro-dynamic interaction forces acting on the five free moving particles for the spherical and non-spherical cases. Figures 6.11 through 6.13 show the forces for the spherical case with figures 6.14 through 6.16 showing the forces for the non-spherical case. Figures 6.11 and 6.14 show the component of force in the direction of overall mass flow (denoted the x component). Figures 6.12 and 6.15 show one of the remaining components (y) and figures 6.13 and 6.16 show the other (z). The five particles are denoted by the relative starting positions:

- [Central]: The central, smaller particle
- $[y\pm, z\pm]$: One of the outer, larger particles. The sign of y and z determined by the starting location of the particle *i.e.* $y+, x+$ would be the particle on the top left in figures 6.9 and 6.10.

In each case the force is non-dimensionalised relative to the fluid viscosity (μ), the average flow rate (u) and the radius, or the effective radius, of the particle in question (r, r^*). Here the average flow rate, u, is the volumetric flux per cross sectional area of the entire domain. The effective particle radius, r^*, is taken as the equivalent 'base' radius, as discussed earlier in this section.

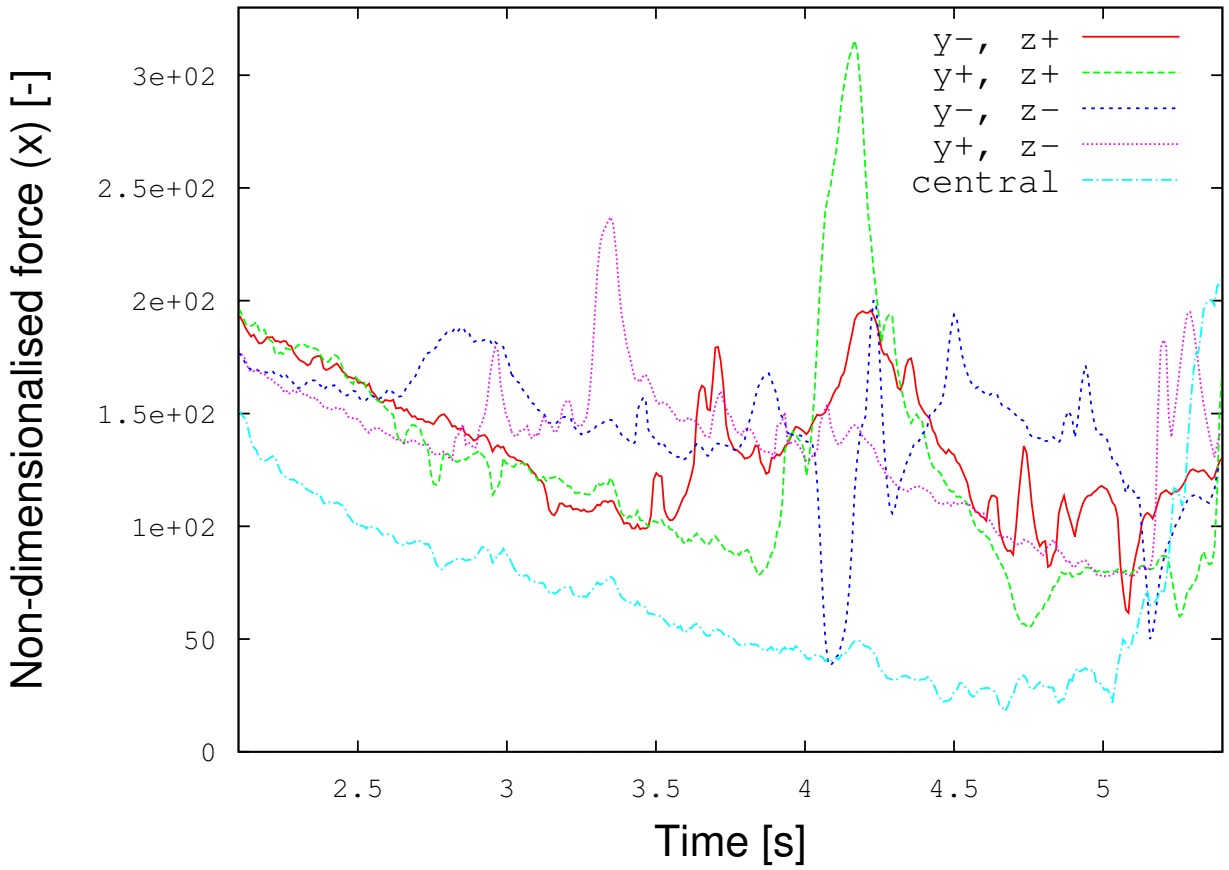


Figure 6.11.: Hydro-dynamic interaction force in the direction of mass flow, x, acting on the five free particles in the spherical case, over time. The force is non-dimensionalised relative to ($\mu r u$), which is proportional the Stokes drag of the equivalent free-stream case.

In both cases the hydro-dynamic interaction force varies erratically over time. However this is to be expected as when the particles move relative to their surroundings the relative geometry changes rapidly. The rapidly changing geometry results in significant changes in the force acting

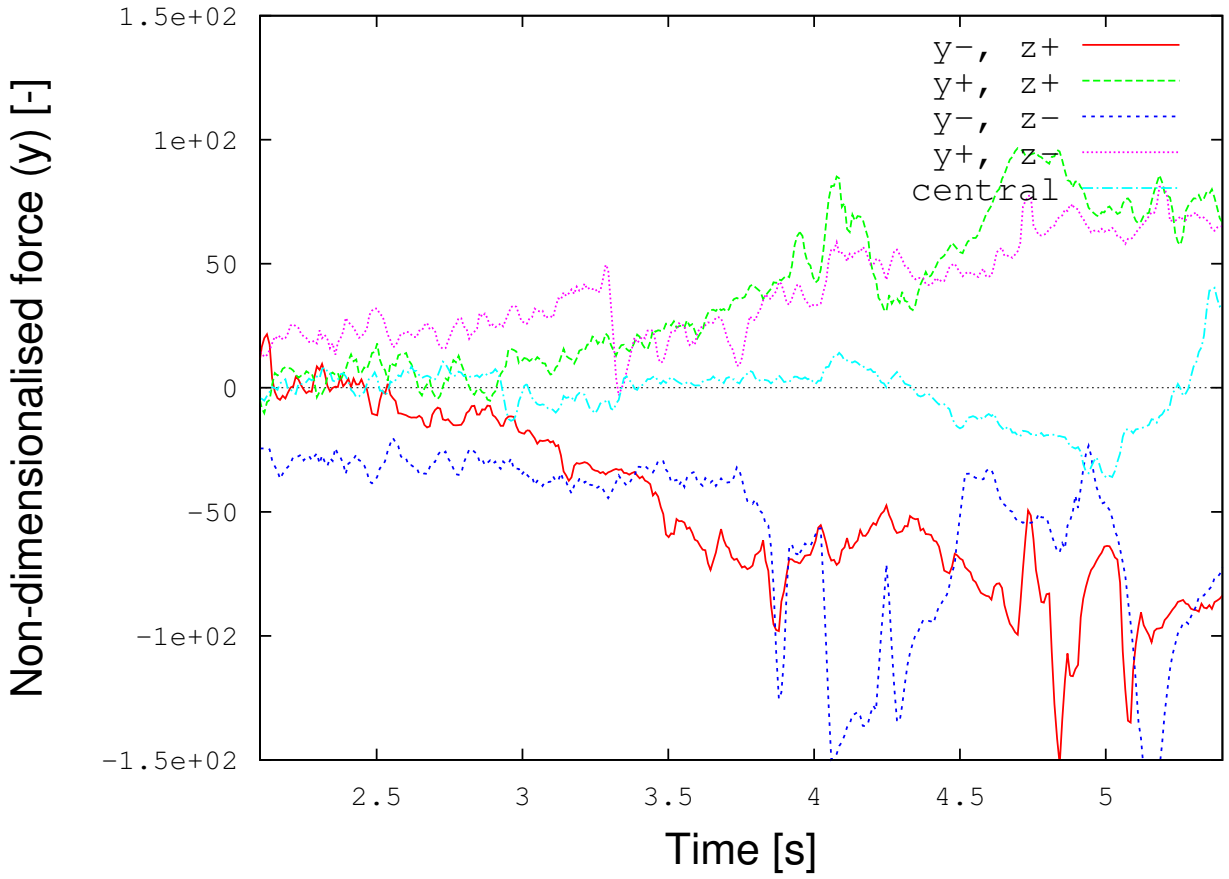


Figure 6.12.: Hydro-dynamic interaction force in a direction tangential to mass flow, y, acting on the five free particles in the spherical case, over time. The force is non-dimensionalised relative to $(\mu r u)$, which is proportional the Stokes drag of the equivalent free-stream case.

on the particles as the different geometries have substantially different porosities. Furthermore, collisions between the particles causes the transfer on momentum between the particles which further compounds the fluctuations. In both the spherical and non-spherical cases the component of force on the particles in the x direction is similar, both in terms of magnitude and trend. A notable feature in both cases is the peak force acting on the $y+, z+$ particle. The $y+, z+$ particle starts behind *i.e.* upstream of the other particles and the peak in force in each case occurs when the other free particles are in close proximity to the static array of particles. The increased force coincides with the time the remaining three particles 'block' the pores in the static mesh, increasing the flow rate in the remaining unblocked corner. Also common to both cases is the trend of the force acting on the central particle. In both cases, the force on the particle decreases over time as the four larger particles separate out, reducing the 'nozzle' effect they created. When the central particle is forced into the central pore between the static array of particles the hydro-dynamic interaction force acting on the central particle increases substantially, approximately by an order of magnitude. This large increase in force is consistent with the highly impermeable barrier this arrangement forms. This type of behaviour, where significant differences in force occur between particles in similar situations, highlights the importance of modelling the small scales of a granular system.

The forces in the tangential directions, y and z, are shown in figures 6.12, 6.15, 6.13 and 6.16. These forces all show comparable behaviour. In each case the hydro-dynamic interaction force pushes

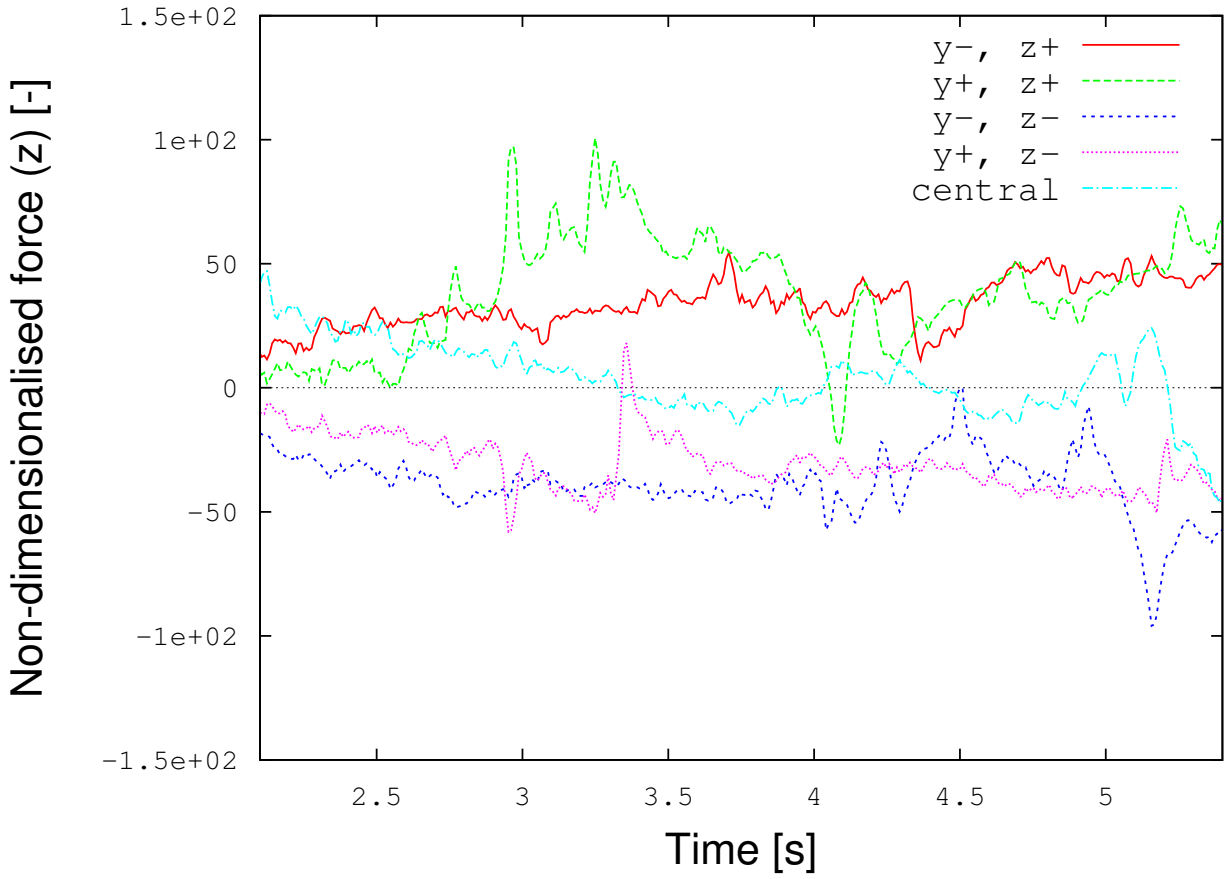


Figure 6.13.: Hydro-dynamic interaction force in a direction tangential to mass flow, z, acting on the five free particles in the spherical case, over time. The force is non-dimensionalised relative to $(\mu r u)$, which is proportional the Stokes drag of the equivalent free-stream case.

the particles radially outward, with the particles on the $y+$ side being forced further to the $y+$ side and similarly for the other directions. The magnitudes of these forces are quite large. Indeed, for the outer particles the tangential hydro-dynamic interaction forces are almost as great as the hydro-dynamic interaction force acting on the particles in the direction of the flow. This also emphasises the significance of modelling the flow at a sub-particle scale as without this the hydro-dynamic interaction force on each particle would be in the direction of relative motion, between the particle and the bulk fluid velocity. The significance of the shape of the particle is also seen. In the non-spherical case, the magnitude of the components of force in the y and z directions are initially different but converge to similar values to each other. However in the spherical case, despite the initial symmetry, the forces in the y direction become larger in magnitude than those in the z direction, which leads the particles to separate further in the y direction. This asymmetry is unstable and grows, as the further the particle deviates, the greater the asymmetry of the force becomes as the particles gets nearer the constriction in the wall of the static 'cup' of particles. Had the simulation been run for sufficiently long the same may behaviour may have occurred for the non-spherical case however, the increased ability of the spherical particles to roll over each other makes the effect more prominent and rapid, in the spherical case.

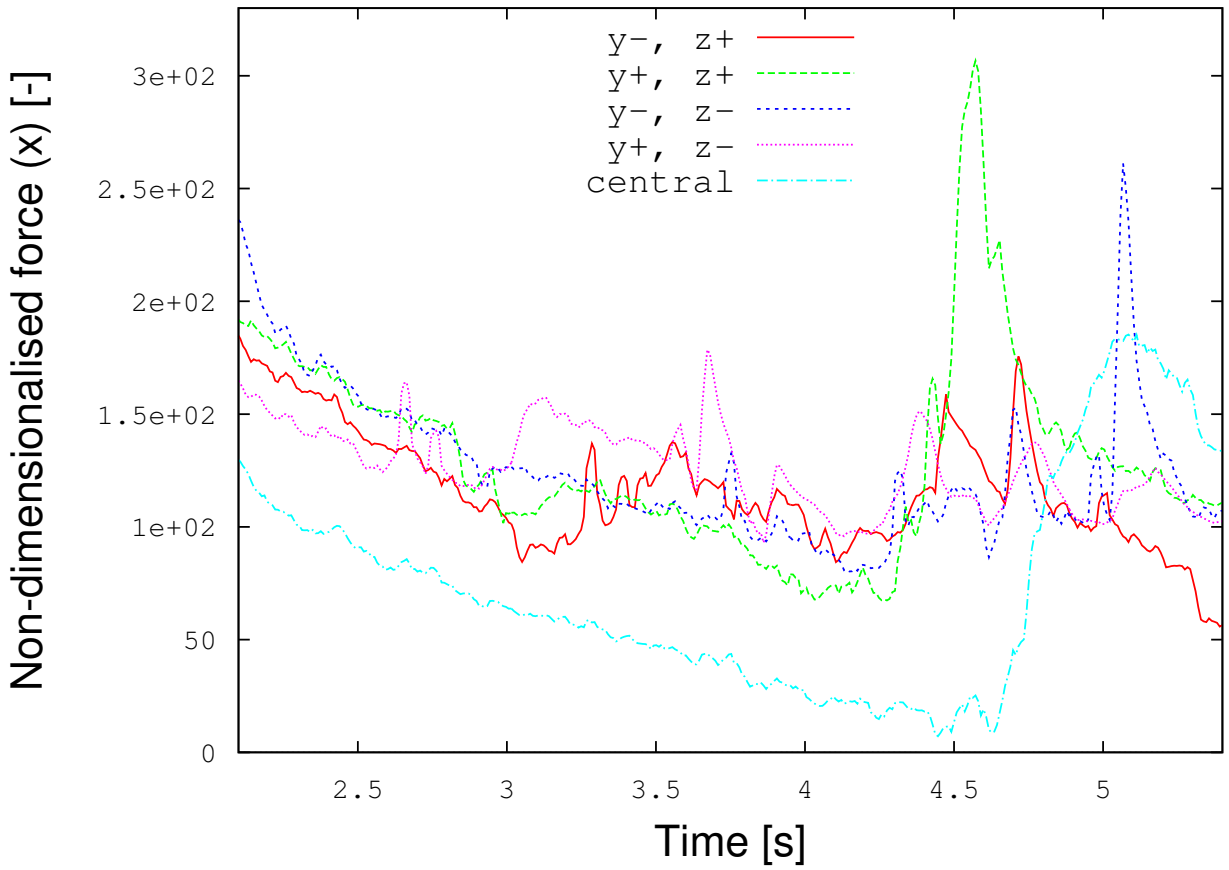


Figure 6.14.: Hydro-dynamic interaction force in the direction of mass flow, x, acting on the five free particles in the non-spherical case, over time. The force is non-dimensionalised relative to $(\mu r^* u)$, which is proportional the Stokes drag of the equivalent free-stream case. Here r^* is taken as the base radius of the particle, as detailed in appendix A.2.

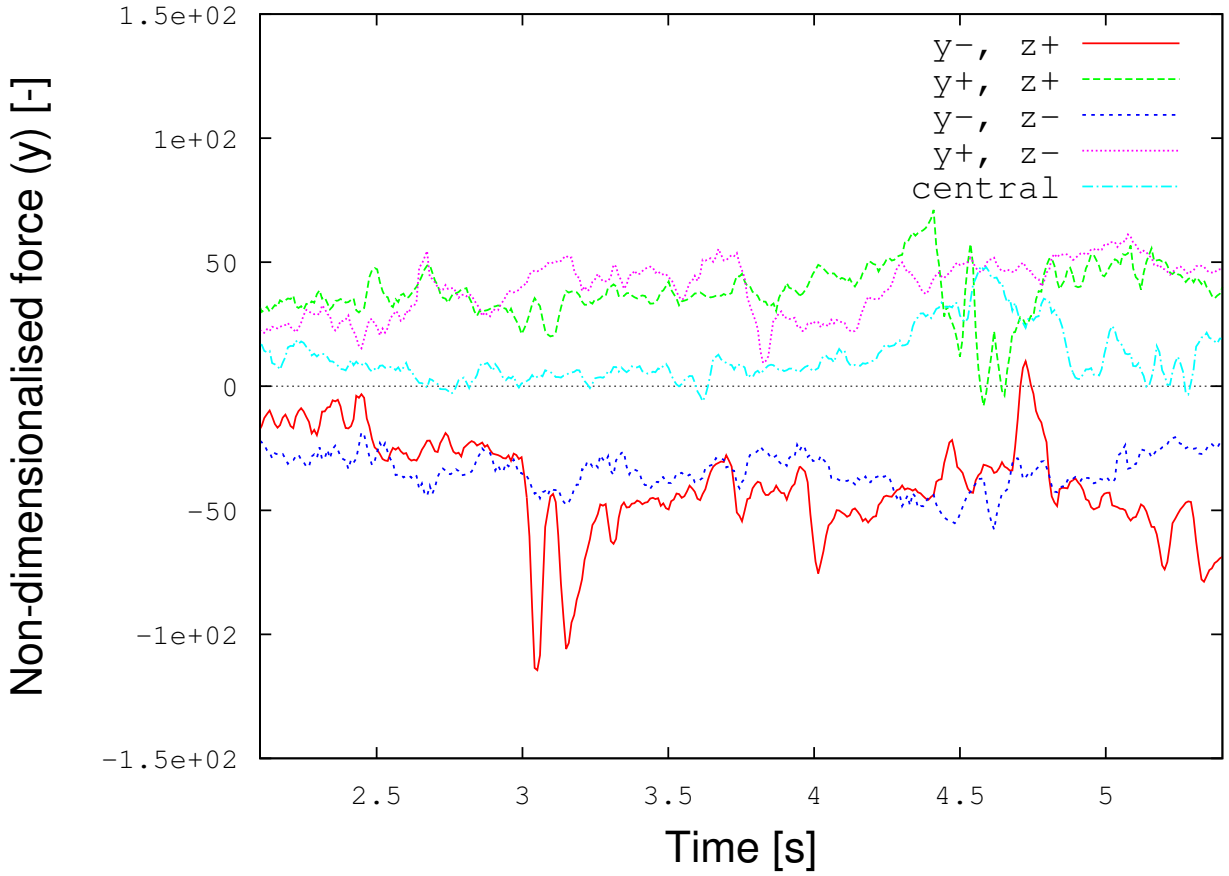


Figure 6.15.: Hydro-dynamic interaction force in a direction tangential to mass flow, y, acting on the five free particles in the non-spherical case, over time. The force is non-dimensionalised relative to $(\mu r^* u)$, which is proportional the Stokes drag of the equivalent free-stream case. Here r^* is taken as the base radius of the particle, as detailed in appendix A.2.

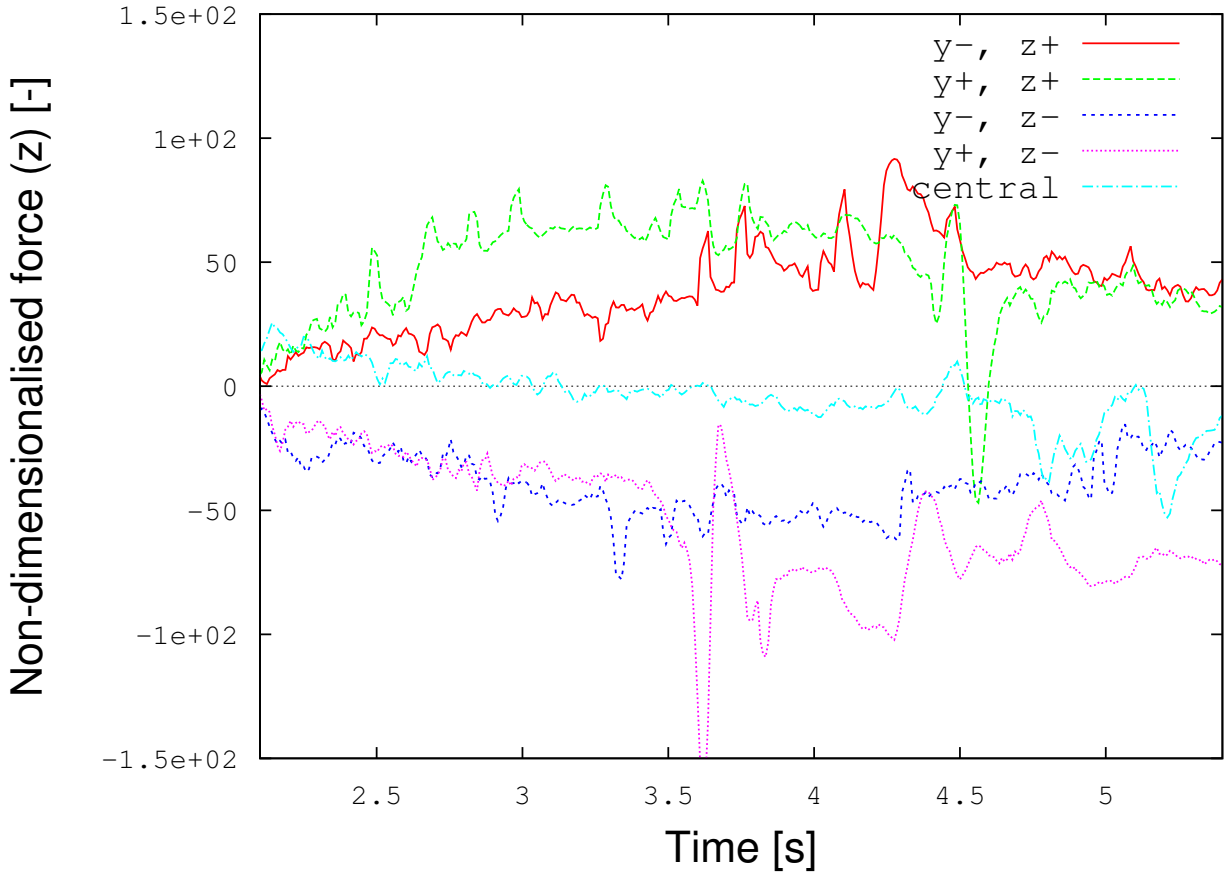


Figure 6.16.: Hydro-dynamic interaction force in a direction tangential to mass flow, z, acting on the five free particles in the non-spherical case, over time. The force is non-dimensionalised relative to $(\mu r^* u)$, which is proportional the Stokes drag of the equivalent free-stream case. Here r^* is taken as the base radius of the particle, as detailed in appendix A.2.

6.1.5. Collision forces

Figures 6.17 and 6.18 show the x component of the collision force acting on the central particle for the spherical and non-spherical case respectively. Both show a similar pattern; the central particle first collides with the other free particles, and subsequently into the static mesh of particles. In each case the collision with the free particles imparts a similar impulse to the central particle. However for the non-spherical particle the force is greater but acts for a shorter period of time. The distinct peaks of the spherical case correspond to collisions with the larger particles, which occur at different times. The particles in the non-spherical case happen to align such that the collisions occur at similar times hence the single large peak. These findings are concurrent with the conclusion that while asperities of the particles of a granular material make some difference to the collision forces, the spherical approximation is not a prohibitive limitation in modelling particles once they have become dislodged from an array and have started moving. However it may make a significant difference to the stability of a packed bed. Unfortunately a simulation of movement in a packed bed of particles is not tractable as, not only does the system require a large number of particles, but the instabilities are sporadic. Hence a long simulation period is required.

Further small differences are shown by the collision forces. In particular, the length of time elapsing between collisions. For the non-spherical case the collision of the final particle with the wall occurred earlier than in the spherical case, despite the earlier collisions finishing later. This seems sensible given the increased radius over which collisions can occur for a rough particle.

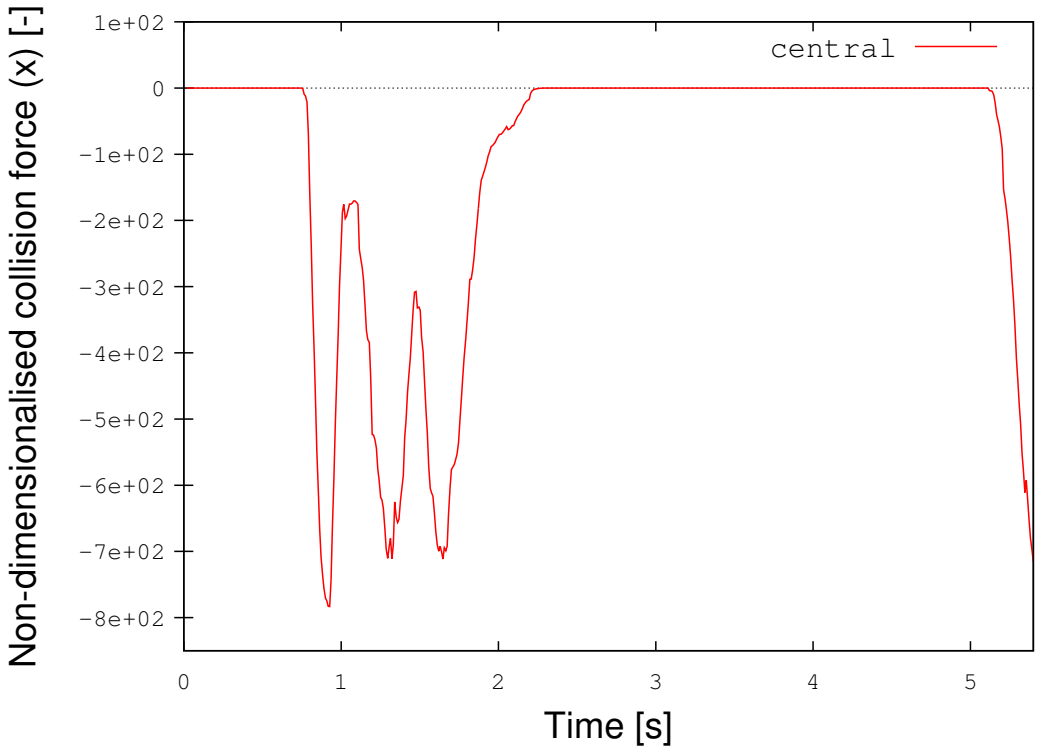


Figure 6.17.: Graph showing the x component of collision force on the central particle for the spherical case. The force is non-dimensionalised relative to (μru) for comparison to the hydrodynamic interaction forces given in section 6.1.4. The initial force is due to collision with the other moving particles $(y\pm, z\pm)$. The second collision is with the static mesh of particles.

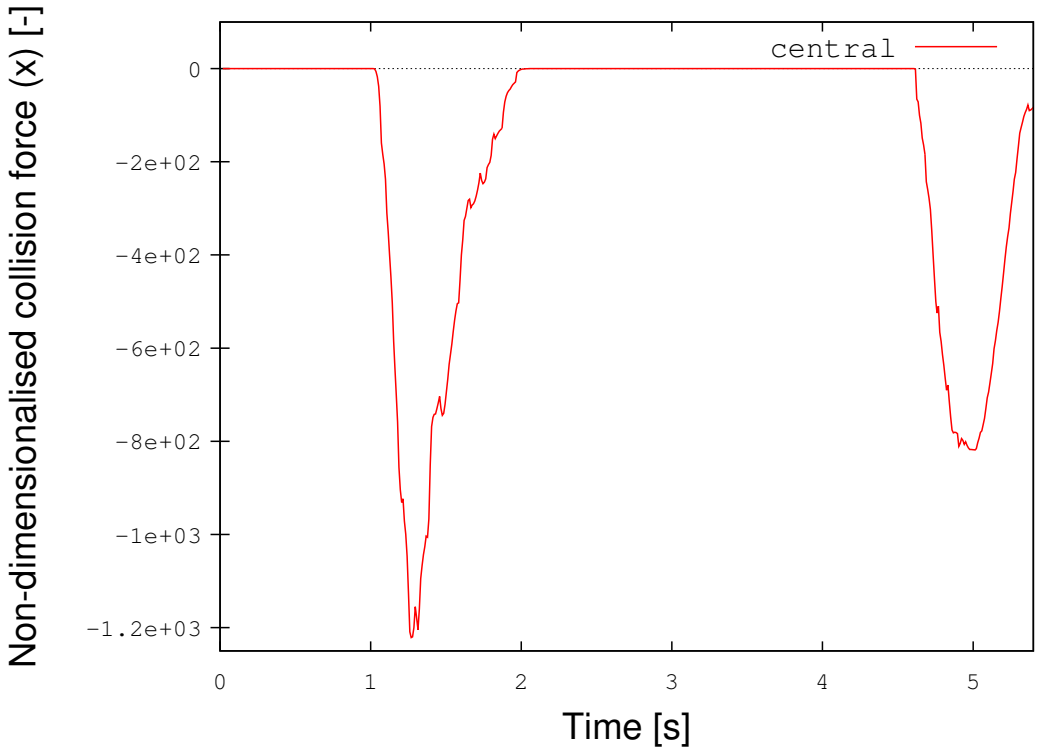


Figure 6.18.: Graph showing the x component of collision force on the central particle for the non-spherical case. The force is non-dimensionalised relative to $(\mu r^* u)$ for comparison to the hydro-dynamic interaction forces given in section 6.1.4. The initial force is due to collision with the other moving particles $(y\pm, z\pm)$. The second collision is with the static mesh of particles.

6.1.6. DEM model

An analogous simulation to the previous two, run using a DEM (Discrete Element Model) code implemented in MultiFlow shows the significance of resolving all of the scales. The major differences between the results of this simulation, and that of the two mirroring immersed-boundary method simulations are:

- The central particle does not impact the remaining four particles until after the collision of the four larger particles had collided with the restraining array of static particles.
- All of the forces acting on the particles are in the x direction until around 6 s into the simulation. Furthermore negligible deviation in the y–z plane occurs for any particle until later still in the simulation.
- The hydro-dynamic interaction force on the central particle is under-predicted through out the simulation.
- The hydro-dynamic interaction forces acting on the particles after they come to rest act in the x direction.

Figure 6.19 shows the simulation after 5.0 seconds for comparison with the other two simulations, in particular to figure 6.7. The hydro-dynamic interaction force is shown for each particle with an arrow. The differences are not surprising considering the significance of the unresolved scales in the

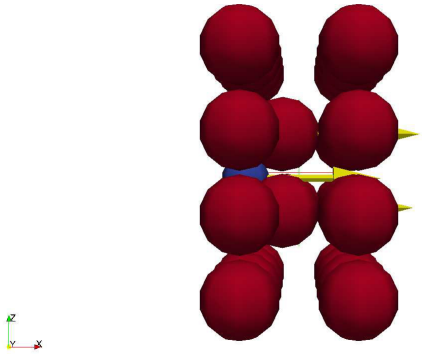


Figure 6.19.: Figure showing the positions and hydro-dynamic interaction forces (yellow arrows) acting on spherical particles, simulated with the DEM method, taken at time 5.0s. The setup and properties of of the simulation are identical to those given in section 6.1. Similar results using the developed method are shown in figure 6.7.

DEM model. Indeed, there are many cases in the resolved simulations where small variations in the geometry have significant effect of the force acting on a particle and, as a result, the motion of the particle. *E.g.* the tangential force acting the central particle, shown in figure 6.12, which experiences significant and sustained tangential load, as a result of small perturbations to what is effectively a symmetrical case. This is not, and cannot, be seen in the DEM model, which is not capable of accommodating such subtle changes accurately. Hence in the DEM model the complex flow, resulting from the complex and changing geometry, is not seen. In particular, the complex hydro-dynamic interaction force seen in figures 6.11 through 6.13 cannot be captured. This is significant in its own right however, the inaccuracy of the prediction is compounded by the lack of feedback. In the DEM model, the incorrect force leads to incorrect motion of the particles, a process which continues yielding substantially inaccurate results.
This phenomenon is particularly notable with regards to the complexity and heterogeneity of the flow. In a physical case, the spontaneous heterogeneities of the flow and particles feedback into the system yielding an intricate and complex flow through the particle array. This is not seen in the DEM model until particle-particle interactions are the dominant forces acting on the particles.
The differences observed, between the developed method and the DEM model, are largely accounted for by this phenomenon. While there are other differences, in particular the collision model, this does not have anywhere near as significant an effect.

In addition to a lack of accuracy, there is also a lack of information regarding the flow in the DEM simulation. In the developed immersed-boundary method the flow between the pores is refined, and the local direction of the flow depicted accurately through the individual pores. This information is missing in the DEM method. Even if the aggregated flow predictions of the method could be improved the limitation of the lack of information about how the flow varies is difficult to overcome.

6.2. Periodicity

The motivation for the use of periodicity is order to simulate the effect of further removed boundaries, *i.e.* a larger effective domain, than would be tractable without the use of periodicity. This is possible because, with the use of periodic boundary conditions, only the particles of interest, and not those needed to recreate realistic boundary conditions, are required. Hence in order to run similar

simulations with a greater number of interesting particles, the ability to use periodic boundary conditions is beneficial.

Periodic boundary conditions within the immersed-boundary framework require modification of the numerical framework. Two things are required to implement periodicity in the developed mirroring immersed-boundary method. Firstly; the fluid boundary conditions, and secondly; the particle boundary conditions.

For the particle boundary conditions, the situation is complicated by the use of non-spherical particles. However the developed solution is very similar to existing techniques – periodicity is achieved by using multiple copies of each of the immersed-boundaries. In particular in addition to the original immersed-boundary particles, for each particle, for each direction that is periodic, two additional copies of the original are added. These are placed a period in front of and a period behind the original immersed-boundary. In this way, as the body moves across a periodic boundary, the additional body moves into the domain at the opposite periodic boundary. When an arbitrary point in the main body crosses the boundary it, and its copies, are instantly moved back by the period distance, such that the original immersed-boundary remains inside the domain. This scheme maintains the presence of all of the necessary copies of the immersed-boundaries, where they may be needed.

The implementation of the fluid boundary conditions, necessary to achieve periodicity, is more complex. The current immersed-boundary method requires large, irregular and variable stencils. These stencils are used for interpolation of fluid properties near an immersed-boundary, to the immersed-boundary itself. In order to replicate the necessary interpolation, when the stencil crosses domain boundaries, a framework in which the 'looped' cells (the other side of a periodic domain-boundary) must be available. Additionally, they must be updated such that their position is in accordance with the period they represent. The provision of such information is complicated by the necessity for parallelisation. Indeed the necessary positions and fluid variables may be required by a processor that does not directly have access to such data.

Hence periodic fluid boundary conditions are not implemented in the current work. However this implementation is feasible and should not require further modification to the method provided only either one or two directions are required to be periodic. The use of three periodic directions is potentially problematic. This is because having all three directions as periodic means there can be no pressure inlet. The usual solution, the use of a pressure point, is complicated by the presence of immersed-boundaries. The complication arises because a pressure point is required to be outside any immersed-boundaries, *i.e.* be a complete-exterior cell centre however these are not constant.

A possible approach, to overcome the problem, is to use a pressure point which is capable of moving to another cell when necessary. This was tried with a simpler implementation of the immersed-boundary method, namely without partial or average mirroring. *N.B.* this simpler implementation did not require such complicated interpolation stencils, and could be run without parallelisation. In this case the pressure field jumps as the pressure point moves, and pressure velocity decoupling occurs near the pressure point. The mass-loss errors in cases without partial mirroring may have been the cause of the decoupling.

This was corroborated by the use of a test case in which the partial mirroring method was used to simulate a moving particle in a domain bounded by no-slip walls in all directions. The test case included a fixed pressure cell, away from the path of the particle. In this case, the pressure decoupling

was substantially reduced.

It is believed a periodic implementation of the developed method is feasible. Such an implementation would be welcome further work and may allow more complex granular systems to be simulated.

6.3. Discussion and conclusions

The developed immersed-boundary method with cut-cell partial mirroring and average mirroring is used to calculate the fully resolved flow for two test cases. Each test case consists of five freely moving particles impacting a ridged array of particles. The two variations of the simulation, one with spherical particles and one with non-spherical particles, showed potentially significant differences. These differences show the strong dependency of granular systems on small-scale variations. While many of the observed differences resulted in similar overall behaviour some aspects were significantly different. One of the principle differences between the two simulations is the final stress state. The stress on the non-spherical particles was significantly more symmetric than the stress of the spherical particles once the particles come to rest, despite the initial asymmetry of the non-spherical case. Further more, the final overall stress was higher in the spherical case. This difference in stress state is particularly significant to the stability of granular materials as the resultant deviatoric stress, which is significantly higher in the spherical case, is important in determining stability. Further more the particle shape itself is of interest in determining the potential for deformation of a granular material, given the stress state.

The results presented in this chapter also demonstrate the significance of modelling the fluid at a resolution that captures the flow through the pores of the system. A quantification the effect of high-resolution modelling is complex, as it depends on the coarse grid model to which it is compared. However certain conclusions can be drawn. A significant consequence of the use of a pore-scale resolving method is the prediction of the hydro-dynamic interaction forces acting on the particles. Not only does the magnitude of hydro-dynamic interaction forces get effected by the fabric of the granular material – the arrangement of the particles within the array – but also the direction of the forces changes. In both-fully resolved simulations, for the majority of the particles, the radial component of hydro-dynamic interaction force is almost as great as the axial component. This is likely to have a significant impact on the modelling of internal instability as the direction of the hydro-dynamic interaction forces acting on the particles is a major factor in whether a particle will move relative to its constrained neighbours. It is unlikely that this behaviour can be modelled accurately without the use of a pore-scale resolving method as the hydro-dynamic interaction force depends on subtle variations between cases. This subtle variation is shown in part by the effects of particle-asperities on the behaviour of the system.

7. Conclusions

7.1. Reduction of spurious pressure oscillations

The presence of spurious pressure oscillations has been a persistent problem for sharp-interface resolving immersed-boundary methods. The cause of these oscillations has been identified as the changing numerics as the immersed-boundary mesh moves relative to the Eulerian mesh [64]. In particular as the equation system changes, to accommodate the altered requirements of the new immersed-boundary positions, it may do so dis-continuously. This is particularly problematic for the continuity equation in which a small discontinuity can illicit a large change in the momentum equation terms. The large change in the momentum equation terms resulting from small discontinuities in the continuity equation occurs as such discontinuities act as a source or sink term. These are significant as a large force is required to 'inject' a finite quantity of mass, or volume, over a short time period, δt. The problem is restricted to particles moving relative to the Eulerian mesh, as without relative motion the problematic transitions, hence discontinuities, do not occur. The oscillations in the pressure field are most significant at low CFL numbers and low Eulerian mesh refinement. The significance of the low refinement is that the continuity equation contributions are weighted by area. Hence, for a given interpolation accuracy, lower refinements give large unphysical 'injections' of mass or volume. The significance of low CFL numbers is that, in such cases, changes to the spacial discretisation (to accommodate spacial changes) occur over shorter time scales than is physical. In contrast, for high CFL numbers, while the spacial discretisation will change between time-instances, such a change is expected and is representative of physical changes.

Modifications to the mirroring immersed-boundary method have been made to reduce these oscillations. As a result good agreement is found between simulations differing only by the relative motion between the immersed-boundary and the Eulerian mesh, demonstrating a non-dependence on this type of motion. Section 4.2.5 shows an example of such a result.

7.1.1. Averaged mirroring

A simple modification, designed to insure that the velocity of boundary-exterior points is similar to that of the immersed-boundary when it is close to the points in question, is developed. The premise of the method is to identify, and use, a list of points on the immersed-boundary, closest to boundary-exterior cells. The motivation for this is that these points will be those used to perform the mirroring after a boundary-exterior to boundary-interior transition. Knowledge of these points can be used to determine how close the immersed-boundary is to such transitions, and accordingly taken into account when enforcing the no slip condition. As a result a smooth transition can be made as these cells transition to become inside the body and used to enforce the no-slip condition.

The modification significantly reduces oscillations in the pressure field during the eclipsing of a fluid cell by the immersed-boundary (a 'dead' cell). The reduction in oscillations occurs because no

'jump' must occur in the velocity at the transitional points during transition. This is because the required velocity at these point, after transition, can be preempted in advance of a near transition. Hence the velocity at these points is no longer subject to discontinuity.

The modification is simple to implement and does not significantly alter the framework of the equation system or the computational cost. Additionally it does not effect the influence of the immersed-boundary in the case the immersed-boundary does not move relative to the Eulerian mesh. Hence this modification has very little negative impact while improving the oscillations in the pressure field, one of the major limiting factors of sharp-interface-resolving immersed-boundary methods.

7.1.2. Cut-cell partial mirroring

The main obstacle to a consistent equation system for the immersed-boundary method, required to avoid spurious pressure oscillations, is a consistent continuity equation. This is also noted by Seo and Mittal [64]. In turn, the principle obstacle to a consistent continuity equation is its conjunction with determining the velocity field for the boundary cells. Here the problem can be seen that, for the cells used to enforce the immersed-boundary condition, three equations governing the three components of velocity already determine the continuity equation, by specification of the relevant variables. Hence maintaining control over the consistency and accuracy of the continuity equation during transition of cell types is problematic as it cannot be governed directly for all cells either side of the transition. In particular those required to enforce the velocity condition at immersed-boundaries.

A novel method is presented to resolve this problem, allowing a consistent and persistent continuity equation to govern the flow, even at cells required to enforce velocity. The premise of the method is the resolution of the velocity field into an out-of-plane and two in-plane components, relative to the nearest immersed-boundary point, using the Gram-Schmidt process. Separate treatment for the in-plane components while allowing continuity to enforce the out-of-plane component of velocity accurately reconstructs the velocity field at the immersed-boundary, while maintaining a continuity equation for each cell. A cut-cell version of the continuity equation is required in order to take advantage of this method. This is developed from previous work and care is taken such that this does not introduce any additional discontinuities. An interpolation scheme is developed to allow the smooth transition of continuity contributions to the cut-cell as the immersed-boundary moves. Section 3.2.3 details the implementation of the cut-cell partial mirroring method. The effects of the method are shown in section 4.2.4.

The advantages of this method are its improved mass conservation and decreased spurious pressure oscillation. The principle disadvantage of the method is the additional associated cost of calculating the intersections of the tessellated surface of the immersed-boundary and the Eulerian mesh, required to solve the cut-cell version of the continuity equation.

7.2. Contact dynamics

Contact dynamics is an integral part of granular dynamics. Often the complexities of collision dynamics for non-spherical particles has limited works to the study of spherical particles, for which the contact dynamics is substantially simpler than for particles with asperities. A contact model that works well for arbitrary shaped particles is developed using the existing immersed-boundary framework. The contact model works by discretising the overlap of two bodies into simply

parametrised collision 'pieces'. From these pieces, the overall force and torque can be determined. For this work a variation of the Hertzian spring-dashpot method, based on the method of Tsuji *et al.* [7] is used however the discretisation scheme does not require this. For refined spheres the results closely match expectations. While the correct force for deformed bodies is not known exactly, predictable trends are followed for rough bodies. This represents a useful advance over spherical-only particle dynamics as particle shape is found to influence granular dynamics (see section 6.1).

7.3. Lubrication force correction

Lubrication force is a significant factor in determining the dynamics of granular flows. An estimation based on the harmonic averaged radius of two spheres is developed, from theoretical results from literature and numerical simulations, to allow the prediction of lubrication forces between arbitrary shaped particles.

Lubrication forces predicted by the developed method are in close agreement with theory when the bodies are removed from each other by more than a single Eulerian cell. However, given the dependence of the lubrication force on very small lengthscale flows, when particles are in close proximity, at a certain separation distance the prediction of lubrication force becomes inaccurate. This is significant as the lubrication forces between particles in close proximity can be very large.

A sub-grid scale contribution to the lubrication force is proposed to account for the difference in lubrication force, as predicted by theory and as calculated using the immersed-boundary method. A sub-stepping method, where the particles are allowed to move at a frequency different to that of the fluid is developed to allow the solution to the stiff equation system resulting from the combination of lubrication force and collision force.

A low Reynolds number prediction of the lubrication force proposed by Cox and Brenner [66] shows good agreement with the developed method, using the sub-grid scale correction. Justification of the low Reynolds number assumption of Cox and Brenner is found by simulating a range of Reynolds numbers likely to be encountered in granular flows.

7.4. Test cases

In general, good agreement is found with experimental test cases presented in the literature. In particular to the experimental investigations of ten Cate *et al.* [86], where simulated velocity and force are in good agreement with experiments. Good agreement was also found in between different Eulerian mesh refinements, with simulations resulting in very similar flow fields. Larger test cases with multiple particles also demonstrate the sensible behaviour of the developed immersed-boundary method. These larger test cases also show the dependence of the dynamics of a granular material on sub-particle scale variations. In particular the resolution of the pore-scale flow was found to have significant effects on the macroscopic behaviour of the system.

7.5. Future work

7.5.1. Force calculation

A significant limitation on the refinement, and hence tractability, of simulations using the developed method is a dependency of calculation of the velocity gradient at the particle surface, on refinement. This is important, particularly for dynamic cases, as it influences the calculation of force acting on the particle. If this could be resolved it would be a significant step to improving the tractability of the method as it would allow fewer Eulerian grid cells to accurately discretise the flow around a particle. As the fluid field does not vary significantly with refinement it is concluded that the primary cause of the inaccuracy of the velocity gradient is the distance over which the velocity gradient 'fit' occurs. This conclusion is corroborated by the non-dependence of velocity-gradient calculation on the interpolation weightings. Possibilities for reducing the inaccuracy in the prediction of hydro-dynamic interaction force at low refinements include accounting for the domain over which the velocity gradient is calculated, such as:

- A prediction using the far-field value in conjunction with a known 'free-stream' solution when a particle is removed from others.
- A correction factor based on the Eulerian grid resolution, relative to the size of an immersed-boundary.

Alternatively, a method using a per-cell, rather than per triangle, contribution to the force may be used. This could be achieved by evaluating a momentum equation for the cells intersected by an immersed-boundary, after solving the primary equation system. By evaluating the contributions for the fluid phase, and the net force imbalance, the contribution from the immersed-boundary may be deduced.

7.5.2. Lubrication force for arbitrary shaped particles

Currently the overall lubrication force, and hence sub grid contribution, is based on two spherical particles or a particle and wall colliding. In these cases the parametrisation of the collision is simple. It would be possible to use a similar method to that of the collision algorithm, with an expanded or translated surface, to gain data relating to a near collision of arbitrary shaped particles. Hence allowing a similar correction technique to that developed to be used for arbitrary shaped particles. Additionally currently, no account is taken of the interaction of the lubrication force and the deformability, roughness and porosity of the body. This is potentially an issue as some erroneous predictions are made by the theory governing non-deformable perfectly spherical particles and perfectly incompressible fluids. In particular the singular nature of the lubrication force is unphysical in all real-world cases and potentially problematic. The problem is often masked by the discretisation inherent to CFD. This occurs because the singularity is rarely realised in the discretised scheme. This masking does not solve the problem, if the influence of the dynamics of the system around the singularity is of interest, the existing method may not be suitable. However incorporating additional factors may allow the accurate reproduction of the dynamics of close contact. Experimental data, such as that of Joseph and Hunt [17], of particles colliding in liquids exists, and could be used in the development and validation of a combined lubrication and collision model. For example the

gap width used to calculate the lubrication force could be modified by the repulsive force between bodies, along with the contact model.

7.5.3. Tractability

There are three main computational costs associated with the developed method. The solution of the matrix of fluid equations, calculating the forces acting on and incrementing the body (sometimes multiple times per fluid timestep) and calculating geometrical properties of the Eulerian mesh relative to the Lagrangian immersed-boundary particles. The implementation, to optimise computational cost, of the solution of the matrix of fluid equations has been the subject of much, and ongoing research. The incrementation of the body is fully explicit, hence little opportunity for optimisation exits. Additionally the entire cost associated with this operation is small, if not negligible, relative to the overall cost of the method. However there is room for improving the tractability of the developed method by improving the cost of calculating geometrical properties of the Eulerian mesh relative to the Lagrangian immersed-boundary particles. In particular, the tractability of the method could be increased by improving the cell cutting algorithm *e.g.* by:

- Optimising the method for a regular Cartesian meshes, as these are suitable for the immersed-boundary method, in particular for modelling granular materials.
- Improving the search algorithm for triangles near each intersected cell.

Improvements in tractability from these, or similar ideas, would be interesting future work in their own right. Such improvements would also allow the use of the cut-cell partial mirroring method to be used for simulating larger problems.

Bibliography

[1] M. Foster, R. Fell, and M. Spannagle, "The statistics of embankment dam failures and accidents," Canadian Geotechnical Journal, vol. 37, no. 5, pp. 1000–1024, 2000.

[2] M. Dyer, "Performance of flood embankments in England and Wales," Proceedings of the ICE - Water Management, vol. 157, pp. 177–186, Jan. 2004.

[3] S. Garner and J. Sobkowicz, "Internal instability in gap-graded cores and filters," in Proceedings of the Canadian Dam Association Annual Conference, pp. 6—-10, 2002.

[4] R. Fell, C. F. Wan, J. Cyganiewicz, and M. Foster, "Time for Development of Internal Erosion and Piping in Embankment Dams," Journal of Geotechnical and Geoenvironmental Engineering, vol. 129, pp. 307–314, Apr. 2003.

[5] M. Li, Seepage induced instability in widely graded soils. PhD thesis, Nov. 2008.

[6] S. Garner and R. Fannin, "Understanding internal erosion: a decade of research following a sinkhole event," The international journal on hydropower & dams, vol. 17, no. 3, pp. 93–98, 2010.

[7] Y. Tsuji, T. Tanaka, and T. Ishida, "Lagrangian numerical simulation of plug flow of cohesionless particles in a horizontal pipe," Powder Technology, vol. 71, no. 3, pp. 239–250, 1992.

[8] Q. Zheng, H. Zhu, and a.B. Yu, "Finite element analysis of the contact forces between a viscoelastic sphere and rigid plane," Powder Technology, vol. 226, pp. 130–142, Aug. 2012.

[9] H. Kruggel-Emden, S. Wirtz, and V. Scherer, "A study on tangential force laws applicable to the discrete element method (DEM) for materials with viscoelastic or plastic behavior," Chemical Engineering Science, vol. 63, pp. 1523–1541, Mar. 2008.

[10] P. Müller and T. Pöschel, "Collision of viscoelastic spheres: Compact expressions for the coefficient of normal restitution," Physical Review E, vol. 84, pp. 5–8, Aug. 2011.

[11] N.-Q. Nguyen and a. Ladd, "Lubrication corrections for lattice-Boltzmann simulations of particle suspensions," Physical Review E, vol. 66, p. 046708, Oct. 2002.

[12] A. A. Mammoli, "The treatment of lubrication forces in boundary integral equations," Proc. R. Soc. A, vol. 462, pp. 855–881, Mar. 2006.

[13] G. Zhu and M. S. Ingber, "Accurate treatment of lubrication forces between rigid spheres in viscous fluids using a traction-corrected boundary element method," Engineering Analysis with Boundary Elements, vol. 33, no. 4, pp. 467–473, 2009.

[14] R. Glowinski, T. Pan, T. Hesla, D. Joseph, and J. Périaux, "A Fictitious Domain Approach to the Direct Numerical Simulation of Incompressible Viscous Flow past Moving Rigid Bodies: Application to Particulate Flow," Journal of Computational Physics, vol. 169, no. 2, pp. 363–426, 2001.

[15] L. Vu-Quoc and X. Zhang, "An accurate and efficient tangential force-displacement model for elastic frictional contact in particle-flow simulations," vol. 31, pp. 235–269, 1999.

[16] J. Zhang, L. L.-S. Fan, C. Zhu, R. Pfeffer, and D. Qi, "Dynamic behavior of collision of elastic spheres in viscous fluids," Powder Technology, vol. 106, no. 1-2, pp. 98–109, 1999.

[17] G. Joseph and M. Hunt, "Oblique particle-wall collisions in a liquid," Journal of Fluid Mechanics, vol. 510, pp. 71–93, 2004.

[18] C. Hirt, A. Amsden, and J. Cook, "An arbitrary Lagrangian-Eulerian computing method for all flow speeds," Journal of Computational Physics, vol. 14, pp. 227–253, Mar. 1974.

[19] T. Björntegaard and E. M. Rönquist, "Accurate interface-tracking of surfaces in three dimensions for arbitrary Lagrangian–Eulerian schemes," Journal of Computational Physics, vol. 231, pp. 6514–6531, Aug. 2012.

[20] J. Ponthot, "A Method to Reduce Cost of Mesh Deformation in Eulerian-Lagrangian Formulation," in Modelling of Metal Forming Processes (E. Chenot, J.L. and Onate, ed.), pp. 65–74, Springer Netherlands, 1988.

[21] A. Kravchenko, P. Moin, and R. Moser, "Zonal Embedded Grids for Numerical Simulations of Wall-Bounded Turbulent Flows," Journal of Computational Physics, vol. 127, pp. 412–423, Sept. 1996.

[22] P. Eiseman, "Adpative Grid Generation," Computer Methods in Applied Mechanics and Engineering, vol. 64, pp. 321–376, 1987.

[23] X. Li, M. Shephard, and M. Beall, "3D anisotropic mesh adaptation by mesh modification," Computer Methods in Applied Mechanics and Engineering, vol. 194, pp. 4915–4950, Nov. 2005.

[24] L. Remaki and W. Habashi, "3-d mesh adaption on multiple weak discontinuities and boundary layers," SIAM Journal on Scientific Computing, vol. 28, no. 4, pp. 1379–1397, 2006.

[25] A. Johnson and T. Tezduyar, "Methods for 3D computation of fluid-object interactions in spatially periodic flows," Computer Methods in Applied Mechanics and Engineering, vol. 190, pp. 3201–3221, Mar. 2001.

[26] D. Wan and S. Turek, "An efficient multigrid-FEM method for the simulation of solid-liquid two phase flows," Journal of Computational and Applied Mathematics, vol. 203, pp. 561–580, June 2007.

[27] H. Chen, S. Chen, and W. Matthaeus, "Recovery of the Navier-Stokes equations using a lattice-gas Boltzmann method," Physical Review A, vol. 45, pp. 5339–5342, Apr. 1992.

[28] Ladd A.J.C., "Numerical simulations of particulate suspensions via a discretized Boltzmann equation. Part 2. Numerical results," Journal of Fluid Mechanics, vol. 271, pp. 311–339, 1994.

[29] Ladd A.J.C., "Numerical simulations of particulate suspensions via a discretized Boltzmann equation. Part 1. Theoretical foundation," Journal of Fluid Mechanics, vol. 211, pp. 285–309, 1994.

[30] Z. Feng and E. Michaelides, "The immersed boundary-lattice Boltzmann method for solving fluid-particles interaction problems," Journal of Computational Physics, vol. 195, no. 2, pp. 602–628, 2004.

[31] A. Holzer and M. Sommerfeld, "Lattice Boltzmann simulations to determine drag, lift and torque acting on non-spherical particles," Computers & Fluids, vol. 38, pp. 572–589, Mar. 2009.

[32] C. S. Peskin, "Flow patterns around heart valves: a numerical method," Journal of Computational Physics, vol. 10, no. 2, pp. 252–271, 1972.

[33] M. Zastawny, Fundamental Understanding of Turbulent Gas-Solid Flows. PhD thesis, Imperial College London.

[34] R. Mittal and G. Iaccarino, "Immersed Boundary Methods," Annual Review of Fluid Mechanics, vol. 37, pp. 239–261, Jan. 2005.

[35] S. Haeri and J. Shrimpton, "On the application of immersed boundary, fictitious domain and body-conformal mesh methods to many particle multiphase flows," International Journal of Multiphase Flow, vol. 40, pp. 38–55, Apr. 2012.

[36] E. Balaras, "Modeling complex boundaries using an external force field on fixed Cartesian grids in large-eddy simulations," Computers & Fluids, vol. 33, pp. 375–404, Mar. 2004.

[37] J. Yang and E. Balaras, "An embedded-boundary formulation for large-eddy simulation of turbulent flows interacting with moving boundaries," Journal of Computational Physics, vol. 215, pp. 12–40, June 2006.

[38] J. Ferziger and M. Perić, Computational Methods for Fluid Dynamics. Berlin Heidelberg New York: Springer Verlag, 3. ed., 2002.

[39] A. M. Roma, C. S. Peskin, and M. J. Berger, "An Adaptive Version of the Immersed Boundary Method," Journal of Computational Physics, vol. 153, pp. 509–534, Aug. 1999.

[40] M.-C. Lai and C. S. Peskin, "An Immersed Boundary Method with Formal Second-Order Accuracy and Reduced Numerical Viscosity," Journal of Computational Physics, vol. 160, pp. 705–719, May 2000.

[41] C. S. Peskin, "The immersed boundary method," Acta Numerica, vol. 11, pp. 479–517, July 2003.

[42] B. E. Griffith, R. D. Hornung, D. M. McQueen, and C. S. Peskin, "An adaptive, formally second order accurate version of the immersed boundary method," Journal of Computational Physics, vol. 223, pp. 10–49, Apr. 2007.

[43] L. Schneiders, D. Hartmann, M. Meinke, and W. Schröder, "An accurate moving boundary formulation in cut-cell methods," Journal of Computational Physics, vol. 235, pp. 786–809, 2013.

[44] J. Mohd-Yusof, "Combined immersed boundary/B-spline methods for simulation of flow in complex geometries," Center for Turbulence Research Annual Research Briefs, pp. 317–328, 1997.

[45] E. Fadlun, R. Verzicco, P. Orlandi, and J. Mohd-Yusof, "Combined Immersed-Boundary Finite-Difference Methods for Three-Dimensional Complex Flow Simulations," Journal of Computational Physics, vol. 161, pp. 35–60, June 2000.

[46] M. Uhlmann, "An immersed boundary method with direct forcing for the simulation of particulate flows," Journal of Computational Physics, vol. 209, pp. 448–476, Nov. 2005.

[47] A. Gilmanov, F. Sotiropoulos, and E. Balaras, "A general reconstruction algorithm for simulating flows with complex 3D immersed boundaries on Cartesian grids," Journal of Computational Physics, vol. 191, pp. 660–669, Nov. 2003.

[48] A. Gilmanov and F. Sotiropoulos, "A hybrid Cartesian/immersed boundary method for simulating flows with 3D, geometrically complex, moving bodies," Journal of Computational Physics, vol. 207, pp. 457–492, Aug. 2005.

[49] A. Gilmanov and S. Acharya, "A hybrid immersed boundary and material point method for simulating 3D fluid-structure interaction problems," International Journal for Numerical Methods in Fluids, vol. 56, no. 12, pp. 2151–2177, 2008.

[50] J.-I. Choi, R. C. Oberoi, J. R. Edwards, and J. a. Rosati, "An immersed boundary method for complex incompressible flows," Journal of Computational Physics, vol. 224, pp. 757–784, June 2007.

[51] W. Huang and H. Sung, "Improvement of mass source/sink for an immersed boundary method," International journal for numerical methods in fluids, vol. 53, no. 11, pp. 1659–1671, 2007.

[52] J. Kim, D. Kim, and H. Choi, "An immersed-boundary finite-volume method for simulations of flow in complex geometries," Journal of Computational Physics, vol. 171, no. 1, pp. 132–150, 2001.

[53] S. Majumdar, G. Iaccarino, and P. Durbin, "RANS solvers with adaptive structured boundary non-conforming grids," Center for Turbulence Research Annual Research Briefs, pp. 353–366, 2001.

[54] R. Ghias, R. Mittal, and T. Lund, "A non-body conformal grid method for simulation of compressible flows with complex immersed boundaries," AIAA paper, vol. 80, p. 2004, 2004.

[55] Y.-H. Tseng and J. H. Ferziger, "A ghost-cell immersed boundary method for flow in complex geometry," Journal of Computational Physics, vol. 192, pp. 593–623, Dec. 2003.

[56] R. Ghias, R. Mittal, and H. Dong, "A sharp interface immersed boundary method for compressible viscous flows," Journal of Computational Physics, vol. 225, pp. 528–553, July 2007.

[57] R. Mittal, H. Dong, M. Bozkurttas, F. M. Najjar, A. Vargas, and A. von Loebbecke, "A Versatile Sharp Interface Immersed Boundary Method for Incompressible Flows With Complex Boundaries.," Journal of computational physics, vol. 227, pp. 4825–4852, Jan. 2008.

[58] A. Mark and B. van Wachem, "Derivation and validation of a novel implicit second-order accurate immersed boundary method," Journal of Computational Physics, vol. 227, pp. 6660–6680, 2008.

[59] S. Shin, S. Bae, I. Kim, Y. Kim, and J. Goo, "Computations of flow over a flexible plate using the hybrid Cartesian/immersed boundary method," International Journal for Numerical Methods in Fluids, vol. 55, no. 3, pp. 263–282, 2007.

[60] J. Oliveira, A. Mark, and B. van Wachem, "A Fully Coupled and Implicit Immersed Boundary Method for Multiphase Flows," in ICMF, no. 2003, pp. 1–9, 2007.

[61] M. Zastawny, B. van Wachem, and J. Oliveira, "An Immersed Boundary Method for Interacting Particles," in 7th International Conference on Multiphase Flow, ICMF 2010, Tampa, FL, May 30 - June 4, no. 4, pp. 1–8, 2010.

[62] A. Mark, R. Rundqvist, and F. Edelvik, "Comparison Between Different Immersed Boundary Conditions for Simulation of Complex Fluid Flows," in ICMF, pp. 1–9, 2010.

[63] T. Ye, R. Mittal, and H. Udaykumar, "An accurate Cartesian grid method for viscous incompressible flows with complex immersed boundaries," Journal of Computational Physics, vol. 156, pp. 209–240, Dec. 1999.

[64] J. H. Seo and R. Mittal, "A Sharp-Interface Immersed Boundary Method with Improved Mass Conservation and Reduced Spurious Pressure Oscillations.," Journal of computational physics, vol. 230, pp. 7347–7363, Aug. 2011.

[65] H. Brenner, "The slow motion of a sphere through a viscous fluid towards a plane surface," Chemical Engineering Science, vol. 16, no. 3-4, pp. 242–251, 1961.

[66] R. Cox and H. Brenner, "The slow motion of a sphere through a viscous fluid towards a plane surfaceII Small gap widths, including inertial effects," Chemical Engineering Science, vol. 22, no. 4, pp. 637–651, 1967.

[67] M. Adams and V. Perchard, "The cohesive forces between particles with interstitial liquid," Inst. Chem. Eng. Symp. Ser, vol. 91, p. 147, 1985.

[68] G. Joseph and R. Zenit, "Particle-wall collisions in a viscous fluid," Journal of Fluid Mechanics, vol. 433, pp. 329–346, 2001.

[69] A. Mammoli, "Towards a reliable method for predicting the rheological properties of multiphase fluids," Engineering Analysis with Boundary Elements, vol. 26, no. 10, pp. 861–869, 2002.

[70] S. Benyahia, M. Syamlal, and T. O'Brien, "Evaluation of boundary conditions used to model dilute, turbulent gas/solids flows in a pipe," Powder Technology, vol. 156, no. 2, pp. 62–72, 2005.

[71] A. Kezdi, Stabilized Earth Roads. Elsevier, 1979.

[72] T. Kenney and D. Lau, "Internal stability of granular filters," Canadian Geotechnical Journal, vol. 22, no. 2, pp. 215–225, 1985.

[73] J. Lee and D. You, "An implicit ghost-cell immersed boundary method for simulations of moving body problems with control of spurious force oscillations," Journal of Computational Physics, vol. 233, pp. 295–314, Jan. 2013.

[74] H. Rönnqvist, Predicting surfacing internal erosion in moraine core dams. PhD thesis, 2010.

[75] W. C. Sherman, "Filter experiments and design criteria.," tech. rep., USACE, 1953.

[76] P. Vaughan, "Criteria for the use of weak and weathered rock for embankment fill and its compaction control," in Proceedings of the international conference on soil mechanincs and foundation engineering., pp. 195–212, 1994.

[77] T. Shire and C. O'Sullivan, "Micromechanical assessment of an internal stability criterion," Acta Geotechnica, vol. 8, no. 1, pp. 81–90, 2013.

[78] M. Li and R. Fannin, "Comparison of two criteria for internal stability of granular soil," Canadian Geotechnical Journal, vol. 45, no. 9, pp. 1303–1309, 2008.

[79] C. Wan and R. Fell, "Experimental investigation of internal instability of soils in embankment dams and their foundations," Journal of geotechnical and geoenvironmental engineering, vol. 134, no. 3, pp. 401–407, 2004.

[80] V. Burenkova, "Assessment of suffusion in non-cohesive and graded soils," Filters in geotechnical and hydraulic engineering. Balkema, Rotterdam., pp. 357–360, 1993.

[81] A. Skempton and J. Brogan, "Experiments on piping in sandy gravels," Geotechnique, vol. 44, no. 3, pp. 449–460, 1994.

[82] J. Lafleur, J. Mlynarek, and A. Rollin, "Filtration of broadly graded cohesionless soils," Journal of Geotechnical Engineering, vol. 115, no. 12, pp. 1747–1768, 1989.

[83] R. Chapuis, "Similarity of internal stability criteria for granular soils," Canadian Geotechnical Journal, vol. 29, no. 4, pp. 711–713, 1992.

[84] C. M. Rhie and W. L. Chow, "Numerical study of the turbulent flow past an airfoil with trailing edge separation," AIAA Journal, vol. 21, pp. 1525–1532, Nov. 1983.

[85] E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, and D. Sorensen, LAPACK Users's Guide. Philadelphia, PA: Society for Industrial and Applied Mathematics, third ed., 1999.

[86] A. ten Cate, C. H. Nieuwstad, J. J. Derksen, and H. van den Akker, "Particle imaging velocimetry experiments and lattice-Boltzmann simulations on a single sphere settling under gravity," Physics of Fluids, vol. 14, no. 11, p. 4012, 2002.

[87] F. F. Abraham, "Functional Dependence of Drag Coefficient of a Sphere on Reynolds Number," Physics of Fluids, vol. 13, p. 2194, Aug. 1970.

[88] A. Henderson, ParaView Guide. Kitware, Inc., 2012.

[89] J. Ahrens, B. Geveci, and C. Law, Visualization Handbook. Elsevier, 2005.

[90] R. Mei, "Flow due to an oscillating sphere and an expression for unsteady drag on the sphere at finite Reynolds number," Journal of Fluid Mechanics, vol. 270, pp. 133– 174, Apr. 1994.

[91] J. D. Franklin and J. S. Lee, "A high quality interpolation method for colocated polyhedral/polygonal control volume methods," Computers & Fluids, vol. 39, pp. 1012–1021, June 2010.

Appendix

A. Appendix

A.1. Interpolation

Interpolation, the process via which information is extrapolated from discrete points to a continua, is vital to CFD, as information is not always available at the exact position it is needed. Hence appropriate interpolation is a critical part of an accurate solution of the Navier-Stokes equations in fluid-solid flows. Details of the methods used, there motivation and implementation explicitly and implicitly is given in this section.
When interpolating data the objective is accuracy, at the cost of computational expense. However many different functions will intersect precisely any given set of points, even when restricted to $\mathbf{C}_\infty$ functions. Hence there is no canonical way of interpolating a field, there are many different 'correct' interpolated fields. There are however certain desirable properties, such as continuity. Indeed most physical fields are continuous hence the restriction to continuous interpolated fields is sensible. Further more, if the gradient is required, a continuous field is necessary. In the context of computational fluid dynamics, the choice often depends on the regularity of the grid over which to interpolate. Often choice of interpolation scheme often depends on the nature of the Eulerian grid, in particular on whether it is regular or not.

Interpolation over a uniform grid For a uniform grid interpolation is a lot simpler than the general case. Indeed only the fixed case need be considered. As the immersed-boundary method is conducive to using a uniform grid, it is expected this will be used for the Eulerian mesh. This is the case for the present work. Interpolations schemes for a fixed regular grid are well documented in the literature.

Interpolation over a non-uniform grid For a non-uniform grid enforcing continuity is significantly more complex. This is mainly as a result of much wider range of cases that need be considered. This is of interest for immersed-boundary method as interpolation between an immersed-boundary and the Eulerian mesh and for cut-cells, requires it. A flexible method, commonly used in CFD, Shepard's algorithm, is not ideal for a number of reasons:

- It does not account for the significant non-linear variation over the interpolation stencil. This is particularly problematic quadratic terms significant over the size of the fluid cells. This is often the case near an immersed-boundary.
- The method does not enforce continuity as the stencil used changes.

Franklin and Lee [91] suggests a number of desirable properties:

Comprehensive The method must offer reconstructed interpolation data throughout the entire volume of the control volume mesh.

Continuous The reconstructed data must be continuous at all points throughout the control volume mesh (Co continuity).

Reconstruct-ability The reconstruction method must reproduce the discrete data at the known data points.

Linear The variation between known data points should follow an approximate linear variation.

Geometry only The interpolation/reconstruction algorithm must be data independent, *i.e.*, only a function of the spatial locations of the discrete data.

Strong Connectivity The method should follow the strong connectivity paradigm implied by the underlying control volume discretisation practises.

Low operation The operation count for each interpolation should be minimised.

Connectivity structures The interpolation data structures should also provide data connectivity for efficient particle locating methods during Lagrangian particle tracking calculations.

Poly-capable The interpolation method should be applicable to polyhedral/polygonal mesh topologies.

Of particular interest to this work is interpolation from a moving set of points. One of the limitations of interpolation using a discrete number of points, which change abruptly given some condition is discontinuity. That is that there is likely, though not necessary to be a discontinuity in some or all of the interpolation coefficients at the when this change of points used occurs. A necessary, but not sufficient, condition for smoothness is that the transitional points (old ones being removed and the new ones added) have $w = 0$. Hence the choice of interpolation scheme is not trivial. There are three common responses to this problem:

1. Use a method of interpolation that gives the same result on a 'point transition boundary' no matter which set of points is used. This is a common approach on regular grids if only a small interpolation stencil is needed as in this case is easy to implement and cheap to run However this method does have problems: If done based on a small number of points linear interpolation is often the only feasible option However this is problematic as linear interpolation is not always accurate and can cause a number of inaccurate trends. If a larger (regular) stencil is available higher order interpolation is feasible. The same is true in theory of irregular stencils Although for technical reasons this is difficult to implement. However doing so has potential as interesting future work. The basis of the idea is while a change of points can lead to a discontinuity, if careful attention is paid it can be avoided. In particular by governing and preempting how the new points are integrated (and the old ones discarded) a discontinuity can be avoided. An additional problem with this method is it requires points in the set of used ones that are limited in the way in which they can be used. This means it requires more points in each direction which may negatively effect performance for CFD.

2. Over specify and weight the interpolation method. If a weighted least squares fit method is used one can remove and add points without discontinuity by insuring their weight varies continuously. In particular that their weight is zero on introduction and removal varies

smoothly when no transition is made. This method is used for this work as despite the increased cost of fitting the results are good and a regular mesh is not required.

3. Live with a small discontinuity. Tolerating a small discontinuity allows a fast efficient and accurate interpolation. The down side of course is the discontinuity. However if the discontinuity can be kept very small it is not necessarily a major problem. This type of method is used for interpolation of the fluid properties to the boundary when a small discontinuity particularly problematic.

A.1.1. Interpolation coefficients

For a given arbitrary point, amongst 'known' points (points from which the field is specified) it is required to determine the interpolated value at this point. Often in addition to this it is also required that coefficients of each of the points be known such that the interpolated value of any field can be computed as a linear sum of these coefficients and the values of the field at their respective known points. This requirement can be be a computational requirements or in order to linearise an equation system. The equation in this can be written as:

$$w_i \text{ s.t.: } \sum_{i=1}^{n} \phi_i \cdot w_i = \phi_p \tag{A.1}$$

Where:

$$\begin{aligned} n &= \text{number of 'known' points} \\ w_i &= \text{required coefficient of } i_{th} \text{ 'known' point} \\ \phi &= \text{the field that is being studied} \\ \phi_i &= \text{value of the field at the } i_{th} \text{ 'known' point} \\ \phi_p &= \text{value of the field at the point which is being interpolated to} \end{aligned}$$

In order to calculate these coefficients the choice function on each of the stencil points (i) is interpolated to the fixed point (p) The resulting values are equivalent to the interpolation technique used to interpolate to each of the the choice functions.

A.1.1.1. Choice function

For any given set, the choice functions of this set are functions that map elements of the set to either 1 or 0 *i.e.*:

$$C_S(i,j) = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{otherwise} \end{cases} \tag{A.2}$$

for a set S, $i \in S$, $j \in S$. In the context of this document, S is the set of known points (cell centres). For an arbitrary fixed point j, $C_S(j,i)$ will be written as $C_j(i)$ and referred to the choice function of j. This is relevant to the determination of w_j, for an arbitrary fixed point j, as it is necessarily

equal to the interpolation of the choice function (C_j) to the j^{th} point. Indeed, if $\phi_i = C_j(i)$ then:

$$\begin{aligned}
\phi_p &= \sum_{i=1}^{n} \phi_i \cdot w_i && \text{required} \\
&= \sum_{i=1}^{n} C_j(i) \cdot w_i && \text{substitution for chosen } \phi_i \\
&= w_j && \text{simplification of the above}
\end{aligned}$$

Hence the weight of the i^{th} known point (w_j, the value we wish to determine) is the interpolation of the choice function associated with the j^{th} known point.

A.1.2. Interpolation schemes used

For this particular application, it has been decided to:

- Compute a least squares fit of the second order polynomial equation in the polynomial case.
- Compute a weighted least squares regression space for the linear case.

A.1.2.1. Polynomial interpolation

From the (known) interpolation points, the ten coefficients of the second order polynomial over three dimensions will be found *i.e.* $a_0 \dots a_9$ where:
$f(x, y, z)$, the interpolating function is:
$a_0 + a_1 \cdot x + a_2 \cdot y + a_3 \cdot z + a_4 \cdot x \cdot x + a_5 \cdot y \cdot x + a_6 \cdot z \cdot x + a_7 \cdot y \cdot y + a_8 \cdot y \cdot z + a_9 \cdot z \cdot z$
This will be done to minimise error, here error is measured using the weighted L2 norm:

$$\sum_{i=1}^{k} \left[h_i (f(x_i, y_i, z_i) - \phi_i)^2 \right]$$

Where: h_i, the weight of the i^{th} point is chosen to be large if the i^{th} point is close to the interpolation point. x_i, y_i, z_i are components of position of the i^{th} point relative to the interpolation point.
One over the square of the distance is taken for these weighting coefficients h_i. $h_i = ||(x_i, y_i, z_i)||^{-1}$
The rapidly increasing nature of h_i ensures the exactness of the interpolation.
The solution to the L2 norm minimum solution is found using LAPACK routine `dgels` [85] While this is well defined for all lengthscales as is, an adjustment is made to improve the numerics. The input variables (x_i, y_i, z_i) are scaled proportional to the grid lengthscale. *I.e.*:

$$\sum_{i=1}^{k} \left[\frac{s}{||(x_i, y_i, z_i)||} (f(\frac{x_i}{s}, \frac{y_i}{s}, \frac{z_i}{s}) - \phi_i)^2 \right]$$

is minimised, where s is the cell spacing.

A.1.2.2. Linear interpolation

From the interpolation points the 4 coefficients of the first order polynomial are fitted. *I.e.* $a_0 \dots a_3$ where:

$f(x, y, z)$, the interpolating function is:
$a_0 + a_1 \cdot x + a_2 \cdot y + a_3 \cdot z +$
This will be done to minimise error, here error is measured using the weighted L2 norm again:

$$\sum_{i=1}^{k} \left[h_i (f(x_i, y_i, z_i) - \phi_i)^2 \right]$$

Unlike the second order polynomial case an analytic solution to this can be found. This is done using the using the correlation coefficients of ϕ_i with position.

A.2. Non sphericity

The particles of which an embankment structure is comprised are invariably not completely spherical. In some cases this may be of relevance to the macroscopic behaviour. In particular the inhibition of rolling and orientation dependence of non-spherical particles may be pertinent. In order to investigate the influence of non-sphericity to the macroscopic behaviour, chaotically generated non-spherical particles are employed. The investigation is limited to predominantly spherical particles as this is physical and allows easy comparison with the spherical case.
Non-sphericity are added by adding randomised asperities. Each of the asperities is characterised by the following:

- Position on the surface p
- Radius r
- Height (which can be negative) h

There is no obvious relation between these quantities, both for an individual asperity and the set. However, the following guidelines were used:

- The position is evenly distributed around the surface of the body
- Each asperity is independent of all the other asperities of that body
- Each asperities size, both height and radius, is independent of its position
- Each lengthscale of roughness is self similar

The central position on the surface required for each of the asperities is calculated as follows:

- Create a random point inside a unit diameter sphere, with constant volumetric probability density. This is achieved via:
    1. Create a point inside a unit cube, centred at $(0, 0)$ using three random numbers between -0.5 and $+0.5$ for each of the three coordinate directions.
    2. Reject the point if it lies outside the sphere centred at $(0, 0)$ with radius 0.5 and go back to the previous step, otherwise proceed.
- Project this point onto the surface of the particle in question.

In order to recreate the mix of lengthscales generated by physical processes a distribution of radii is chosen. The distribution is chosen to such that there are more smaller asperities. In particular, an inverse area weighted distribution was chosen.

$$r = \sqrt{a \times b} \times \left[\sqrt{\frac{a}{b}}\right]^n \tag{A.3}$$

Where a is the maximum asperity radius, b the minimum asperity radius. n is a random variable, with $-1 < n < 1$, used to determine where in the permissible range the added asperity radius should by. n is given by:

$$n = 2\left[\int_{-1}^{l_a} \frac{b}{a}^x \, \mathrm{x}x / \int_{-1}^{1} \frac{b}{a}^x \, \mathrm{x}x\right] - 1 \tag{A.4}$$

This simplifies to:

$$2\left(\frac{\frac{1}{\ln b/a} \times \left[\frac{b}{a}^{l_a} - \frac{a}{b}\right]}{\frac{1}{\ln b/a} \times \left[\frac{b}{a} - \frac{a}{b}\right]}\right) - 1 \tag{A.5}$$

subsequently:

$$2\left(\frac{\left[\frac{b}{a}^{l_a} - \frac{a}{b}\right]}{\left[\frac{b}{a} - \frac{a}{b}\right]}\right) - 1 \tag{A.6}$$

With l_a a random number between -1 and 1.
Non-spherical particles are generated using a random number generator, to determine l. The distribution of l_a was chosen to be a binomial distribution as a this distribution roughly approximates a normal distribution while maintaining boundedness. The generator that produces l_a can either be made to be repeatable in order to recreate specific cases or allowed to vary to compare cases.
The height of the h of each of the asperities is calculated using a second random variable, l_b, scaled relative to the radius of the asperity.

$$h = c(l_b + d)r \tag{A.7}$$

For constants c and d and l_b a random number between -1 and 1. Again a binomial distribution was chosen for l_b. Constants c and d have a physical interpretation:

- c represents the range of heights permissible for a given asperity, relative to its radius.
- d represents the offset of this range of heights, also relative to the asperity's radius. *E.g.* $d = -1$ corresponds to only permitting indentations and $d = +1$ corresponds to only permitting protrusions.

With position on the surface (p), radius (r) and height (h) an asperity is added by modifying the distance away from the centre of the particle of each of the nodes used to tessellate the particles surface. The change in height, δh, is given by:

$$\delta h = h\left(\frac{(r+a)(r-a)}{r^2}\right) \tag{A.8}$$

Where a is the arc length from the centre point of the asperity, p.

A.3. Permiameter type simulation results

A.3.1. Greater number of spheres

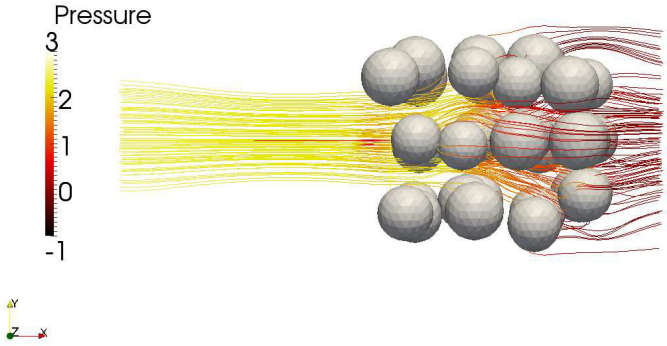


Figure A.1.: Test case showing the free moving particles in a larger test case similar to those used in section 6.1. 36 spherical particles are use, all of the same time. Streamlines of the fluid, coloured by pressure are shown. A static mesh of particles restraining the flow is present by not shown.