

The copyright of this thesis vests in the author. No quotation from it or information derived from it is to be published without full acknowledgement of the source. The thesis is to be used for private study or non-commercial research purposes only.

Published by the University of Cape Town (UCT) in terms of the non-exclusive license granted to UCT by the author.

A One-way Coupled DEM-CFD Scheme to Model Free-surface Flows in Tumbling Mills

Michael Malahe¹

MLHMIC002

Thesis presented for the degree of
Master of Science

PHY5000W

University of Cape Town

(SUBMITTED 25 SEPTEMBER, 2012; CORRECTIONS 22 NOVEMBER, 2012)

¹E-mail: michael.malahe@uct.ac.za

Acknowledgements

This thesis would not have been possible without the support of my supervisors, Andy Buffler and Indresan Govender. In terms of resources, I am thankful for being given a desk in a great office and access to all the computing power I could possibly need. In terms of supervision, I am thankful for being given enough room to take this project in my own direction, but enough guidance to keep me focused.

I would like to thank Gary Tupper for his honest appraisals of my work and for always managing to find that elusive paper that I couldn't.

I am grateful to Angus Morrison for his help with many of the technical aspects of PEPT and minerals processing and for sharing some of his PEPT analysis code and MATLAB expertise. Additionally, I would like to thank Angus for being a great sounding board for being a great office mate.

I would like to thank Angus Comrie and Marius Hromnik for their help and advice with various C++ issues and Dino Giovannoni for sharing some of his custom L^AT_EX macros with me.

I would like to thank Lawrence Bbosa for allowing me the use of his CAD drawings of the 30 cm mill and for managing all of the EDEM licensing.

I am indebted to the department of Chemical Engineering at the University of Cape Town (UCT) for generously allowing the use of their FLUENT licenses. I would also like to thank Bill Randall and Harish Appa for their help with licensing issues and the technical aspects of FLUENT respectively.

For their (perhaps unknowing) contributions to my thought process, I would like to thank Spencer Wheaton, Matthew Bickell and Maciej Stankiewicz.

During the course of this work I have been fortunate to receive funding from a number of sponsors. I am grateful to the National Research Foundation of South Africa for the award of an NRF Free-standing Innovation/Scarce Skills Bursary over the first two years of my study. From the Centre for Minerals Research (CMR) at the University of Cape Town, I am thankful for the award of a CMR Postgraduate Student Bursary for the same period. Finally, from UCT Sport & Recreation, I appreciate the award of a full sports scholarship for the same period.

Finally, I would like to thank my mother, Lesley Daniel, for being the constant source of positivity that helped me through some of the more trying patches of this work.

Abstract

There is a great need in the minerals processing industry for accurate prediction of slurry flows in tumbling mills. This flow is currently best understood through empirical, mechanistic and computational models. In the case of computational models, a realistic description of the system requires two fluid phases (for slurry and air) and one solid phase (for the charge). Existing 3-phase work makes use of a purely particle-based description, coupling the discrete element method (DEM) for the solids with smoothed particle hydrodynamics (SPH) for the fluids.

This work is the first presentation of a 3-phase approach for tumbling mills that combines a particle description for the solids, DEM, and a continuum description for the fluids, computational fluid dynamics (CFD). In this approach, the phase coupling is only one-way, in the sense that forces are applied from the solids to the fluids, but not the other way around. In the course of developing the approach, some computational geometry algorithms are presented for efficiently converting the particle description of DEM into an appropriate continuum description for CFD.

A set of DEM-only tests is compared with equivalent Positron Emission Particle Tracking (PEPT) experiments of laboratory-scale mills, and excellent agreement is found. Another set of DEM-only tests is compared with PEPT experiments that are identical except for the addition of slurry. It is found that the “dry” DEM-only experiments match the particle trajectories in these “wet” experiments very well. It is concluded that the effect of the slurry on the charge is minimal, suggesting that one-way coupling is a sensible approach. Additionally, some mechanistic modelling assumptions by Taberlet *et al.* are brought into question.

The DEM-CFD approach is implemented, and the effect of using different drag correlations and time-averaging intervals is examined. It is found that the drag correlation presented by Beetsta *et al.* is the most appropriate, and that the choice of time-averaging interval is significant. The results are compared with PEPT experiments and good agreement is found in some cases, although the validity of current PEPT techniques for fluid tracking is examined. Some mechanistic modelling assumptions by Condori *et al.* and Moys are contradicted, while some assumptions by Morrel & Stephenson and results by Cleary *et al.* are supported.

Possible modifications and improvements to the method are discussed, including turbulence modelling, improved non-Newtonian modelling and two-way coupling.

Contents

1	Introduction	7
1.0.1	The Present Work	9
2	Computational and Mechanistic Modelling	10
2.1	The Discrete Element Method	10
2.1.1	Contact Models	11
2.1.2	Time-Stepping	15
2.1.3	DEM in Tumbling Mills	16
2.2	Granular Solids in Tumbling Mills	20
2.2.1	Taberlet <i>et al.</i> - rotating drums	20
2.2.2	Govender <i>et al.</i> - shear rates	21
2.3	Fluid modelling	24
2.3.1	Governing Equations	24
2.3.2	Multiphase Flows	24
2.3.3	Turbulence modelling	25
2.4	Solid-Fluid Coupling	28
2.4.1	Drag Forces	28
2.4.2	DEM-SPH in Tumbling Mills	36
2.5	Slurry in Tumbling Mills	40
2.5.1	Moys - mechanistic modelling	40
2.5.2	Morrell and Stephenson - experimental work	42
2.5.3	Condori <i>et al.</i> - transport regions	42
3	Coupling Implementation	46
3.1	Coupling	46
3.1.1	Scheme	46
3.2	Algorithms	48
3.2.1	Point in Voxel Local Averaging	48

3.2.2	Point in Cell Local Averaging	50
3.3	Verification	56
3.3.1	Pressure drop recovery	56
3.4	Geometry	57
3.4.1	The 12 cm Mill	57
3.4.2	The 30 cm Mill	64
3.4.3	Meshing	64
4	Validation Using PEPT Experiments	73
4.1	The PEPT Method	73
4.1.1	Steady-state Profiles	75
4.2	PEPT Experiments	76
4.2.1	The 12 cm Mill	78
4.2.2	The 30 cm Mill	78
5	Results	81
5.1	Models and Discretisation	81
5.1.1	DEM	81
5.1.2	CFD	81
5.2	Solids Only: The 12 cm Mill	82
5.2.1	Materials	82
5.2.2	Results	82
5.3	The Effect of the Fluid on the Solids: The 30 cm Mill	98
5.4	Fluid Only	102
5.4.1	Simulation Conditions	102
5.4.2	Basic Results	102
5.5	One-Way Coupling	108
5.5.1	Coupling with Slurry	108
5.5.2	PEPT Comparison	111
5.5.3	Full-run Time-averaging	120
6	Discussion and Conclusion	123
6.1	DEM-Only	123
6.1.1	12 cm Mill	123
6.1.2	30 cm Mill	124
6.2	One-way Coupled DEM-CFD	124
6.2.1	Relation to Prior Work	125

6.2.2	Comparison With Experiments	125
6.2.3	Implementation Concerns	126
6.2.4	Future Applications	127

Bibliography	136
---------------------	------------

A Animations	137
---------------------	------------

University of Cape Town

Chapter 1

Introduction

In a government report on South Africa's energy usage and efficiency, a common theme was the extremely high energy usage by the mining and minerals processing industry [1]. The report noted a steady increase in energy use in these sectors, and identified them as having massive potential for increased energy efficiency [1]. A high-value target for such improvements is a piece of minerals processing equipment called a tumbling mill, which alone accounts for approximately 1.8% of South Africa's energy consumption [2] and 60% of the operational costs on a typical minerals processing plant.

The core function of a tumbling mill is the breaking and grinding of ore to expose valuable minerals. The base structure of a mill is a cylindrical drum lined by a series of narrow bars called lifters, which aid the lifting of ore when the mill is turning. When the mill is operating, it is fed in one end with ore and water. The ore is then churned in the mill and reduced in size in two ways. The first way is by breakage, where pieces of ore that are flung into the air by the churning motion strike the mill shell or other pieces of ore and break. The second is by grinding, where slower moving ore particles in the more packed regions of the mill move past each other and grind away at each other by friction. In what are known as semi-autogenous (SAG) mills, steel balls (known as the grinding media) are added to the ore to aid in the breaking and grinding processes.

The result of the breaking and grinding is the production of fine particles mixed with the valuable minerals that have been liberated from the ore. This material is suspended in the water and forms a viscous mineral slurry that exits the mill on the opposite end to where the ore and water were fed in. At this end there are discharge grates that are sufficiently small not to allow through ore particles that are large enough to still encompass large amounts of the target mineral. This slurry is fed into subsequent stages of the minerals processing circuit, which serve to separate the target

mineral from the other slurry particles.

While this process seems fairly straightforward, the volume of slurry present and the slurry dynamics can have a large effect on the efficiency of milling. Of particular importance is the volume of slurry when the mill is operating at steady-state, known as the slurry hold-up. Moys notes that if the hold-up is too low, it causes two key problems [3]. The first is that the rate of grinding is reduced, because the presence of the granular slurry between the ore particles improves the grinding rate [3]. The second is that in SAG mills, the rate of wear on the grinding media increases [3].

On the other hand, if the hold-up is too high, the slurry forms a pool at the base of the mill, which suppresses breakage from impact in that region. In addition to suppressing breakage, the added mass in the mill from this pool increases the power-draw of the mill with no benefit, drastically reducing the efficiency of the mill. Thus, it is key to have good control over the hold-up in the mill to maintain efficiency.

The hold-up depends on a number of factors, including the mill speed, slurry viscosity, inflow rate and discharge grate design [3][4][5][6][7][8][9]. Despite the importance of being able to model the hold-up however, most milling operations rely on empirical models that correlate these (and other) variables to the hold-up. Since taking invasive measurements of an on-plant mill is exceptionally expensive due to long down-times, these correlations often have to rely on a small number of data points and have little predictive power outside a narrow range of operation. A similar modelling approach that has been employed is to run more extensive tests on laboratory-scale mills (see Section 2.5) and generate correlations that hopefully scale up, however Morrell notes that scale-up is frequently a problem with lab-only studies [4].

A more fundamental approach is that of developing mechanistic models that aim to capture the underlying physics of the slurry flow in a mill (see Section 2.5). Such models have the potential to have wide-ranging and easily-applicable predictive power. However, a major difficulty in developing these models is knowing which simplifying assumptions and modelling choices may be valid, and what the extent of their validity is. This arises partly because experimental techniques can't provide sufficiently detailed information about the dynamics in the system.

It is at this point that computational modelling can provide a more complete picture for specific systems and conditions, which can be used to aid the mechanistic modelling process. Such computational modelling has been done extensively for mills, particularly DEM (the Discrete Element Method, see Section 2.1) for modelling the motion of the solids (see Subection 2.1.3).

Investigators have used the simulations directly to predict quantities that are also investigated with experiments, such as power draw [10][11], grinding media wear and lifter wear [12].

1.0.1 The Present Work

For the flow of slurry in tumbling mills, only Cleary has done full-scale computations [13], by coupling SPH (Smoothed Particle Hydrodynamics), a mesh-free computational method for fluids, to DEM (see Subsection 2.4.2). In this work, a similar approach is taken, but the DEM simulations are rather coupled to simulations using CFD (Computational Fluid Dynamics), which is a much more established and well-understood mesh-based method for fluids.

The coupling in this work will be “one-way”, in the sense that information is fed from the DEM simulation to the CFD simulation, but not the other way around (see Section 3.1). This approach is motivated by a massive decrease in the difficulty of implementation and the assertion that the particles have a much larger influence over the motion of the fluid than the fluid has over the particles. This assertion is tested in Section 5.3.

With this one-way process in mind, the aim of this work is to first test the validity of DEM as a method for modelling the solids in the mill, since this information will feed into the fluid dynamics. This validation is done by comparing the results from DEM simulations with experimental results using PEPT (Positron Emission Particle Tracking), which is a non-invasive technique for measuring the solidicity and velocity fields in the mill (see Section 4.1).

Once this is established, the one-way coupled results will be compared with the assumptions in existing empirical, mechanistic and computational models, as well as with experimental results from PEPT measurements.

Chapter 2

Computational and Mechanistic Modelling

With an approach of computational modelling in mind, this section surveys the theoretical descriptions and computational methods that are available for modelling tumbling mills. Since our approach is to model the system as discrete rigid solids interacting with a continuous fluid, the leading methods for those approaches are presented. In addition to this, leading research in the mechanistic modelling of tumbling mills is presented.

2.1 The Discrete Element Method

The discrete element method (DEM) is a computational method for discrete solids developed by Cundall and Strack in 1979 [14]. In its simplest formulation, DEM models the translational and rotational motion of discrete spheres (or discs in 2D). As expected from Newtonian mechanics, the rate of change of the linear and angular velocity of each particle is given by:

$$\begin{aligned}m \frac{d\mathbf{v}}{dt} &= \mathbf{F} \\ I \frac{d\boldsymbol{\omega}}{dt} &= \boldsymbol{\tau}\end{aligned}$$

where \mathbf{F} is the net force on the particle and $\boldsymbol{\tau}$ is the net torque. The rate of change of the position of the particle is given by:

$$\frac{d\mathbf{r}}{dt} = \mathbf{v}. \tag{2.1}$$

Since the particles are spherical, none of the equations depend on their angular orientation, and keeping track of it is not necessary. With these equations in mind, there are only two key decisions to be made in implementing a DEM code. The first, a modelling decision, is to decide how to

determine the net force and torque on a particle at a given moment, and will be covered in Subsection 2.1.1. The second, a numerical decision, is to decide on the time-stepping scheme and the length of time-step used to update the velocities and positions of the particles, and will be covered in Subsection 2.1.2.

2.1.1 Contact Models

The models used in DEM to determine the force and torque on a particle due to interactions with other particles and boundaries are referred to as contact models. Each particle can be in contact with multiple other particles and boundaries and the force due to each contact is calculated separately. The contacts in DEM are treated as non-conforming contacts, that is the initial contact between two spheres is expected to be a point, after which the deformation of the spheres broadens the contact surface. In other words, the boundaries of the spheres do not conform to each other without first deforming.

In DEM, this deformation is not tracked directly, as it would be in the Finite Element Method, but the indented region of contact between the two spheres is tracked by allowing the perfect numerical spheres to overlap. The length of this overlap is then equated to the depth of deformation of the spheres, as illustrated in Figure 2.1. The overlap distance, d or δ , is the difference between the distance between the particles and the sum of their radii.

This overlap distance, as an analogue for the depth of deformation, is used to recover the force that would be required to deform the particles to that depth. The exact way that these are related is determined by the contact model.

Hertz-Mindlin

In 1881, a seminal paper by Hertz [15] outlined the forces acting on elastic bodies in contact. In particular, the forces due to normal loading and relative normal velocities between particles are still used today [16][17].

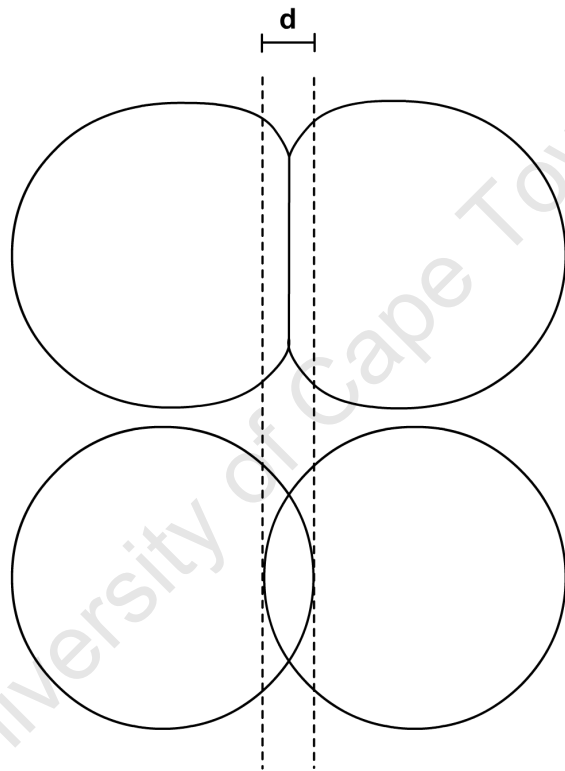


Figure 2.1: The deformation of particles treated as an overlap in DEM, where d is the overlap distance

According to Hertzian theory, the depth of indentation between two elastic spheres, δ , is related to the static loading force between the two spheres, F_n , by [15][16][18]:

$$\delta = \left(\frac{9F^2}{16R^*E^{*2}}E \right) \quad (2.2)$$

where the equivalent Young's Modulus, E^* , and equivalent radius R^* are given by

$$\frac{1}{E^*} = \frac{1 - \nu_1^2}{E_1} + \frac{1 - \nu_2^2}{E_2}$$

$$\frac{1}{R^*} = \frac{1}{R_1} + \frac{1}{R_2}$$

where E_i, ν_i and R_i label the Young's Modulus, Poisson's ratio and radius of sphere i . For the purposes of DEM, we wish to recover the force from the indentation, or overlap, and can rearrange Equation 2.2 to get

$$F_n = \frac{4}{3}E^*\sqrt{R^*}\delta^{\frac{3}{2}}$$

For particles in motion, there is additionally a normal damping force, F_d^n , given by [15][16][18]

$$\mathbf{F}_d^n = -2\sqrt{\frac{5}{6}}\beta\sqrt{S_n m^*} \mathbf{v}_n$$

where \mathbf{v}_n is the normal component of the relative velocity, the equivalent mass, m^* , and the constant β and the normal stiffness, S_n , are given by

$$\frac{1}{m^*} = \frac{1}{m_1} + \frac{1}{m_2}$$

$$\beta = \frac{\ln c_r}{\sqrt{\ln^2 c_r + \pi^2}}$$

$$S_n = 2E^*\sqrt{R^*}\delta_n$$

and c_r is the coefficient of restitution. Hertzian theory is still widely used and continues to be supported by modern experiments [19][20][21][22].

A well-supported [23][22] contact model for forces due to relative tangential velocities was developed from work by Mindlin in 1949 [24] and Mindlin and Deresiewicz in 1953 [25]. The tangential force, F_t is given by [24][16][18]

$$F_t = -S_t\delta_t$$

where δ_n is the tangential overlap and the tangential stiffness, S_t , is given by

$$S_t = 8G^*\sqrt{R^*\delta_n}$$

provided there is no slip between the particles during the contact. The tangential force is bounded above by Coulomb friction [26][18], $F_c = \mu_s F_n$, where μ_s is the coefficient of static friction. When

the tangential force reaches the Coulomb friction bound, there is slip between the particles. In this case, more general forms of the tangential force have been proposed, but the most prominent one, by Thornton and Randall [27], has been shown by Makse *et al.* to have an insignificant effect on results [22].

The tangential damping force, F_t^d , is given by a similar expression to the normal damping:

$$F_t^d = -2\sqrt{\frac{5}{6}}\beta\sqrt{S_t m^*}\mathbf{v}_t$$

Here \mathbf{v}_t is the relative tangential velocity.

Rolling friction is modeled by applying a torque of [28][17][18]

$$\tau = -\mu_r F_n R \hat{\omega}$$

where μ_r is the coefficient of rolling friction and $\hat{\omega}$ is the direction of the object's angular velocity.

The combination of the Hertz description of normal forces, and the Mindlin description of tangential forces, referred to as the Hertz-Mindlin model, is well-supported. In particular, for the contact regimes in which this work falls, the model has been shown to be sound by Makse *et al.*[22]: “We find agreement between our simulations and the experiments, thus confirming the validity of the Hertz-Mindlin contact theory to glass bead aggregates composed of frictional particles.” This is of particular importance to the simulations in this work, because they mirror experiments done with glass bead aggregates.

According to the same paper, the upper pressure limit at which the Hertz-Mindlin model holds is around 10MPa [22]. This limit was discovered by measuring the point at which the predicted $p^{1/3}$ scaling (where p is the pressure) of the elastic moduli gave way to a $p^{5/9}$ scaling [22]. A final limitation on using the Hertz-Mindlin model is that it makes use of quasi-static assumptions that require that the typical relative velocities of the particles are slower than the speed of sound in the solid material [29].

Spring-Dashpot

The spring-dashpot, or damped linear spring model is from the original work by Cundall & Strack [14]. In both the normal and tangential directions, the contact force is modeled with a spring-dashpot system. In the normal direction, the force is given by [14][26][18]:

$$F_n = -k\delta + c\dot{\delta}$$

δ is the overlap, $\dot{\delta}$ is the rate of change of overlap, or overlap velocity and k , the spring coefficient, and c , the damping coefficient are given by:

$$k = \frac{15}{16} \sqrt{R^* E^*} \left(\frac{15 m^* \dot{\delta}^2}{16 \sqrt{R^* E^*}} \right)^{1/5}$$

$$c = \sqrt{\frac{4 m^* k}{1 + \left(\frac{\pi}{\ln c_r} \right)^2}}$$

Similarly, in the tangential direction, the force is given by:

$$F_t = k_t \delta_t + c_t \dot{\delta}_t$$

where δ_t is the tangential overlap, $\dot{\delta}_t$ is the rate of change of tangential overlap and the tangential spring coefficient, k_t , and damping coefficient, c_t , are taken as ratios of the normal versions [18]. As with the Hertz-Mindlin model, the tangential force is bounded by Coulomb friction [30].

2.1.2 Time-Stepping

Stepping Scheme

The leap-frog method is commonly used in DEM simulations [31][32] because it is simple but fairly accurate ($\mathcal{O}(\Delta t^2)$). The update equations are:

$$\begin{aligned} \mathbf{r}_{i+1} &= \mathbf{r}_i + \mathbf{v}_{i+1/2} \Delta t \\ \mathbf{v}_{i+1/2} &= \mathbf{v}_{i-1/2} + (\mathbf{F}_i/m) \Delta t \\ \omega_{i+1/2} &= \omega_{i-1/2} + (\tau_i/I) \Delta t \end{aligned}$$

Size of Time-step

A typically quoted static constraint on the size of the time-step is the Rayleigh time-step, T_R , given by [18]

$$T_R = \frac{\pi R \sqrt{\frac{\rho}{G}}}{0.1631\nu + 0.8766}$$

where R is the particle radius, ρ is the density of the material, G is the shear modulus and ν the Poisson's ratio. This constraint is chosen so that shear waves cannot propagate through the entire extent of a particle in a single time-step. If this were possible, the integration scheme would fail to accurately predict the transfer of momentum from one particle, through another, into a third (eg. in a Newton's cradle). That is, the scheme would break down in situations where a particle is in contact with more than one other particle. Typically, this constraint is reduced by a factor of 3 or more to make the accurate integration of such scenarios possible [18].

A similarly motivated approach, frequently quoted in conjunction with spring-dashpot type contact models, is to base the time-step on the natural vibrational frequency of a system consisting of the contact spring and a single particle [14][33][31][34]. That is, the minimum time-step would be given by:

$$\Delta t = \frac{1}{fn}$$

where f is the vibrational frequency and n is the number of time-steps needed to resolve the contact. The vibrational frequency for a spring of stiffness k and a particle of mass m is given by the usual $f = \frac{1}{2\pi} \sqrt{\frac{k}{m}}$. Lower values of n are chosen to reduce computational time, but also come with a reduction in accuracy. For example, Tsuji *et al.* in 1993 used $n = 5$ [33], but in 2000 Vu-Quoc *et al.* used $n = 40$ [31]. For $n = 2\pi$, this is identical to the stability criterion for the leap-frog time-stepping scheme.

Another constraint on the size of the time-step exists to ensure that the particle overlap, depicted in Figure 2.1, is never greater than the maximum deformation allowed by Hertzian theory. To this end, the time-step should never be large enough so that particles can go from not being in contact to having an overlap greater than this limit in one time-step. That is, the least upper bound on the time-step would be given by

$$v\Delta t < \delta_{\max}$$

where v is the expected maximum velocity in the domain, Δt is the time-step and δ_{\max} is the maximum deformation. Typically, this constraint is reduced by a factor of six or more so that the contact can be better resolved [18].

2.1.3 DEM in Tumbling Mills

The state of the art of DEM in tumbling mills is spearheaded by two research groups, headed by B.K. Mishra and P. Cleary respectively. Before going into their methods and approaches to validation, it would be useful to introduce some terms that are frequently used when referring to mill operation. When a mill's rotational speed is quoted, a commonly used reference speed is the "critical speed", which is the speed at which the charge would centrifuge in a mill without lifters. That is, it is the speed at which the centripetal acceleration of the mill shell matches gravity, i.e.

$$\begin{aligned} \frac{v_{\text{crit}}^2}{R} &= g \\ \Rightarrow v_{\text{crit}} &= \sqrt{gR} \\ \Rightarrow \omega_{\text{crit}} &= \sqrt{g/R} \end{aligned}$$

where v_{crit} is the critical speed, ω_{crit} is the critical angular speed and R is the radius of the mill. The operating load of tumbling mills is typically referenced by the percentage of the total mill

volume occupied by the “charge”, which is a term for the solid material (ore and grinding media) in the mill.

Methods

Both Mishra and Cleary have used a linear spring-dashpot model with the tangential force limited by Coulomb friction [10][11] (see Subsection 2.1.1), although more recent work by Mishra has regarded the Mindlin model for tangential forces to be an improvement on that of the linear spring-dashpot model [32]. In the Mishra implementation, a leap-frog time-stepping scheme is employed [10] and in the Cleary implementation a second-order predictor-corrector scheme is used [35].

Validation

Both Mishra and Cleary have made use of direct comparisons between simulation snapshots and photographs of real-world experiments with matched conditions. Figure 2.2 shows a typical comparison by Mishra for an early DEM code. For these comparisons, good agreement on the motion of particles in the interior of the charge and on the periphery were considered sufficient for validating the character of the overall charge motion [10].

Figure 2.3 shows a comparison by Cleary of an experimental photograph with four different DEM implementations. The “3D Slice” implementation is one where 3D particles were used but were confined to a small axial range, past which a periodic boundary condition was applied [35].

A more quantitative and context-aware comparison that is used by both authors is a comparison with the power draw predicted by the simulations and that derived from torque measurements or direct power input measurements on laboratory-scale and operational mills.

For purely laboratory-scale mills, work by Mishra found the difference between predicted power draw and measured power draw to be approximately 5% on average for three laboratory-scale mills [10]. For a 0.9 m length and 0.15 m diameter mill the comparison was for five sets of operational conditions picked from a matrix of speeds of 30%, 50%, 60% and 70% v_{crit} and fillings of 20% and 30% by volume [10]. For the two other mills, one with 0.292 m length and 0.254 m diameter, one with 0.292 m length and 0.381 m diameter, the comparison was for a speed of 60% v_{crit} and a filling of 50% by volume [10].




Figure 2.2: Comparison of a snapshot of a DEM simulation and a photograph of an experiment by Mishra *et al.* [10]

For a single laboratory scale mill of diameter 0.592 m and length 0.2 m, work by Cleary found that the calculated power draw severely under-predicted the experimental value (by 44%) for low speeds and loads (35% filling by volume and 65% v_{crit}) but matched experiment well (within 3%) for higher speeds and loads (45% filling by volume and 95% v_{crit}) [35]. The authors suggested that the under-prediction in such a small mill was due to a large percentage of the power in the experiment not being used to move the charge, but used and dissipated in other parts of the machinery.

For 7 intermediate-scale mills, ranging from 0.292 m length and 0.254 m diameter to 5.49 m length and 4.11 m diameter, work by Mishra found the average difference between predicted power draw and measured power draw to be approximately 10% [36]. For 15 operational mills, ranging from 3.62 m diameter to 10.8 m diameter, the average difference was again found to be around 10% [36], with the calculated power draw generally under-predicting the measured power draw.

In general, photographic comparisons of DEM and experiment have been favourable, as have comparisons between measured and predicted power draw, which is directly useful for industry. However good photographic comparisons and power draw agreement don't necessarily imply that the flow of the particles is in good agreement. For this work, the solidicity and velocity profiles throughout the mill need to be known with some accuracy for the coupling process, so confidence in those measures is essential. This motivates the use of a measurement technique like PEPT in this work as these measures can be determined directly from experiments.




Figure 2.3: Comparison of snapshots of DEM simulations and a photograph of an experiment by Cleary *et al.* [35]. (a): 2D model with circular particles (b): 2D model with non-circular particles (c): experiment (d): 3D slice model (e): full 3D model

2.2 Granular Solids in Tumbling Mills

2.2.1 Taberlet *et al.* - rotating drums

Work has been done by Taberlet *et al.* to model granular flows in rotating drums that are very much like tumbling mills [37]. In particular, the shape of the free surface is modeled analytically using granular flow theory and numerically using Molecular Dynamics (MD) simulations.

For the granular flow modelling, the authors characterise the system in terms of the following quantities:

- d : The diameter of the particles
- Ω : The rotational speed of the drum
- g : The acceleration due to gravity
- R : The radius of the drum
- L : The axial length of the drum
- $\bar{\rho}$: The average density of the material

The two-dimensional flow is then characterised in terms of the feed rate, $q(x)$, which is the local rate at which material is transferred in the y direction per unit length, and the shear rate, $\dot{\gamma}$.

The authors assume that the feed rate is given by:

$$q(x) = \bar{\rho}\Omega r$$

where r is the distance from the centre of the drum [37]. The authors note that the shear rate depends on the flow properties, but also that the dependence is weak and therefore model the shear rate as a constant. This constant is taken to have the proportionality

$$\dot{\gamma} = \alpha\sqrt{g/d},$$

citing a dimensional analysis argument. The authors choose $\alpha = 1$ simply to have a fixed value, but claim that such fixing has no effect on their qualitative results [37].

With these assumptions in mind, the following differential equation for the height of the free surface as a function of x was derived [37]:

$$\frac{dy_{\text{surf}}}{dx} = \mu_i + \mu_w \left(\frac{\Omega^2 d}{g} \right)^{1/4} \frac{1}{L} \sqrt{R^2 - x^2 - y_{\text{surf}}} \quad (2.3)$$

where μ_i is the effective internal friction and μ_w is the effective wall friction [37].

The authors note further that with the non-dimensionalisations $\tilde{x} = x/R$ and $\tilde{y} = y_{\text{surf}}/R$, Equation 2.3 can be reduced to:

$$\frac{d\tilde{y}}{d\tilde{x}} = \mu_i + \mu_w \Lambda \sqrt{1 - \tilde{x}^2 - \tilde{y}^2} \quad (2.4)$$

where Λ is given by:

$$\Lambda = \left(\frac{\Omega^2 d}{g} \right)^{1/4} \frac{R}{L}$$

Since μ_i and μ_w are constants, the dimensionless shape of the free surface is dependent only on Λ [37].

This dependence of the shape of the free surface on only this one dimensionless parameter was tested with MD simulations, which in this context is very similar to DEM. The authors used a spring-dashpot contact model (see Subsection 2.1.1) and a Verlet time-integration scheme.

An equal number of 8 mm diameter and 6.4 mm diameter particles of density 7460 kg.m⁻³ were used to fill the drum half way. Drum radii in the region 100d and drum lengths between 10d and 200d were used with rotational speeds between 0.05 rad/s and 0.2 rad/s. The authors found that if Λ was kept constant while varying its controllable components (Ω, R, L), the shape of the free surface remained unchanged [37], supporting equation 2.4.

2.2.2 Govender *et al.* - shear rates

A paper by Govender *et al.* details PEPT measurements of the velocities and shear rates in a laboratory-scale tumbling mill [38]. The mill had a diameter of 30 cm and a length of 30.5 cm and is also used in this work.

In their analysis, a line was constructed joining the centre of the mill with the centre of circulation (CoC) of the charge, which is the point about which the charge circulates. This line was then extended in either direction to meet the mill shell. The velocity component perpendicular to this constructed line was calculated referred to as the tangential velocity, v_θ [38]. Positive values of v_θ were taken to be in the clockwise direction. Figure 2.4 illustrates the construction of this line and the velocity vectors at the sample points.

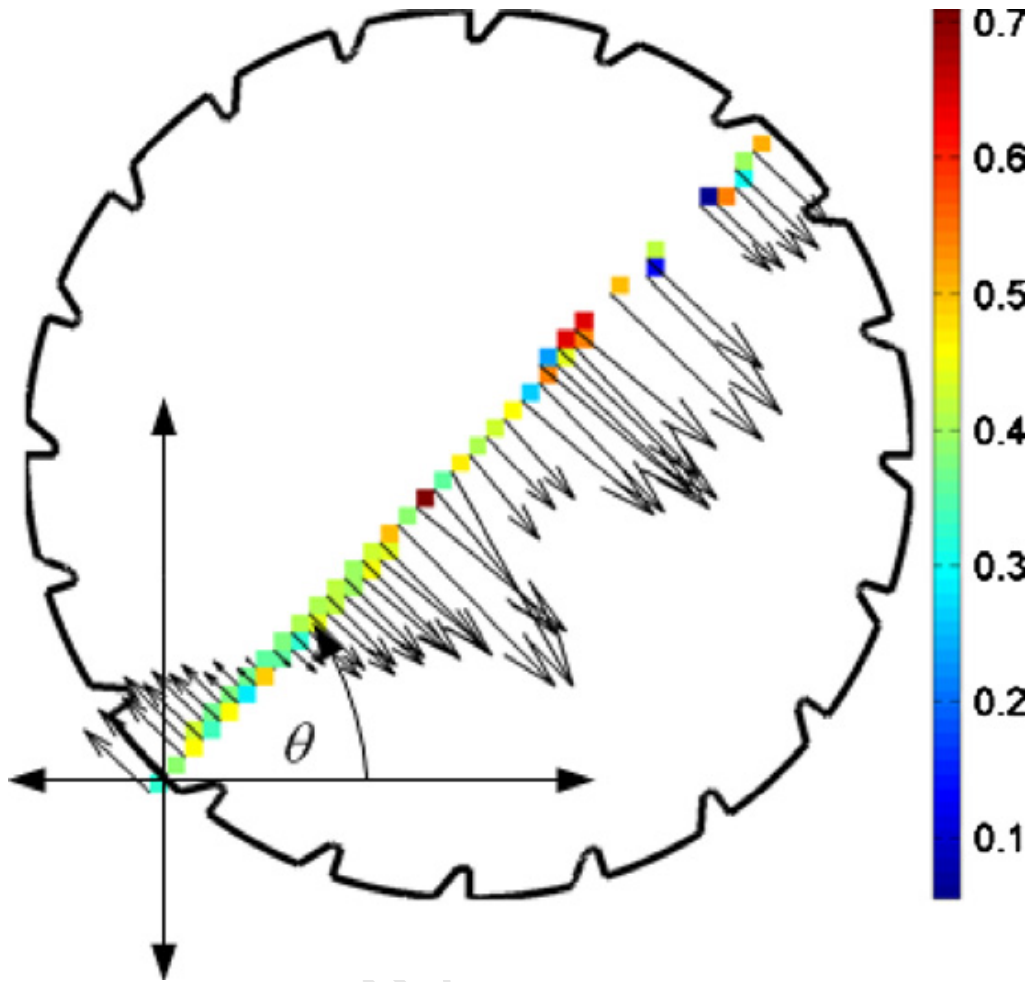


Figure 2.4: Tangential velocity vectors along the radial scanline by Govender *et al.* The data points are coloured by the tangential velocity magnitude in m/s. [38]

A sum of sines of the form $v_{\theta}(r) = a_1 \sin(b_1 r + c_1) + a_2 \sin(b_2 r + c_2)$, where r is measured from bottom left (inside the charge) to top right, was fitted to the tangential velocity data and the shear rate along the radial line, $\frac{dv_{\theta}}{dr} = a_1 b_1 \cos(b_1 r + c_1) + a_2 b_2 \cos(b_2 r + c_2)$, was calculated. Figures 2.5 and 2.6 show the tangential velocity and shear rate profiles respectively.

The shear rate between the mill shell and the free-surface, which will be examined for both DEM and PEPT simulations in this work in Section 5.2, followed a roughly parabolic profile.




Figure 2.5: Tangential velocity along the radial scanline by Govender *et al.* [38]

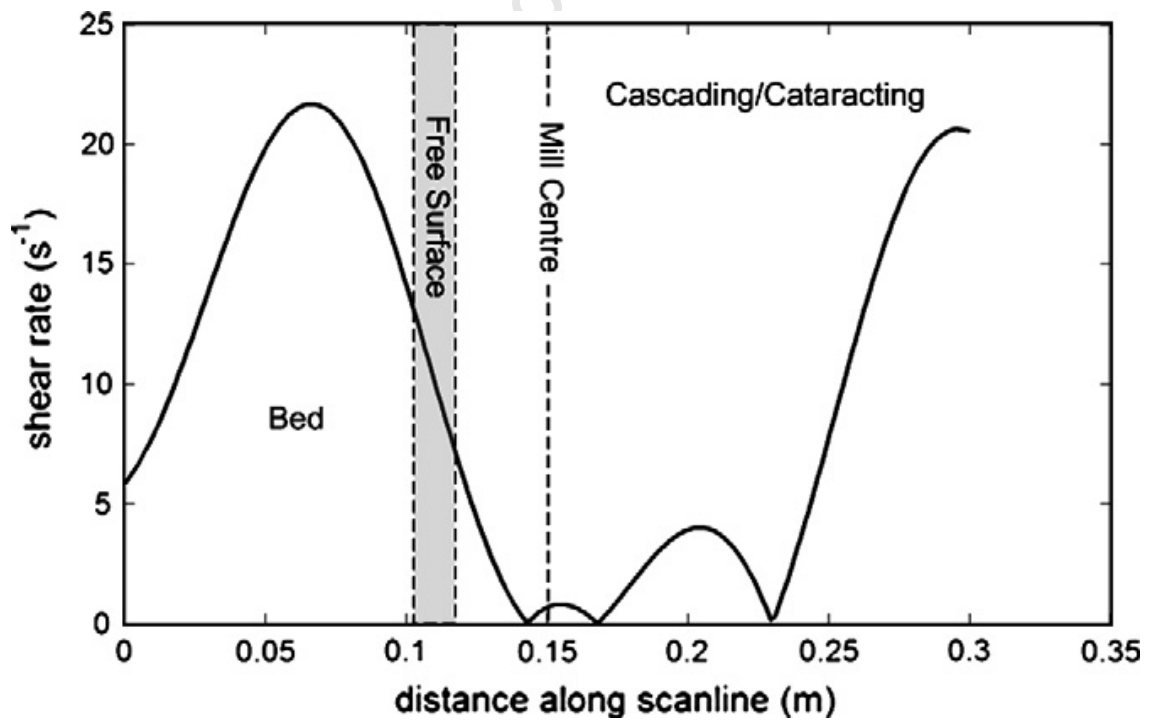


Figure 2.6: Shear rate along the radial scanline by Govender *et al.* [38]

2.3 Fluid modelling

2.3.1 Governing Equations

The conservation of fluid momentum in an inertial reference frame is given by the Navier-Stokes equations [39][40]:

$$\frac{\partial}{\partial t}(\rho \mathbf{u}) + \nabla \cdot (\rho \mathbf{u} \mathbf{u}) = -\nabla p + \nabla \cdot \mathbb{T} + \mathbf{F}, \quad (2.5)$$

where ρ is the fluid density, \mathbf{u} is the fluid velocity, p is the static pressure, \mathbf{F} contains all of the external body forces and \mathbb{T} is the stress tensor [39][40]. Additionally, mass conservation requires that [39][40]

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{u}) = 0. \quad (2.6)$$

In the incompressible case, these reduce to [39][40]:

$$\begin{aligned} \rho \left(\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} \right) &= -\nabla p + \mu \nabla^2 \mathbf{u} + \mathbf{F}, \\ \nabla \cdot \mathbf{u} &= 0, \end{aligned}$$

where μ is the dynamic viscosity.

Non-Newtonian Flows

The viscosity in the Navier-Stokes equation can only be modeled as a constant for certain fluids. The viscosity for the remainder, known as Non-Newtonian fluids, can be modeled in a number of ways, depending on the fluid. In the context of this work, only what is referred to as Bingham plastic behaviour is modeled. Fluids that exhibit this behaviour do not shear at all for shear stresses below a certain threshold, but have a constant viscosity above this threshold [41]. That is, the shear rate is given by:

$$\dot{\gamma} = \begin{cases} 0 & \text{if } \tau < \tau_0 \\ (\tau - \tau_0)/\mu & \text{if } \tau \geq \tau_0 \end{cases}$$

where $\dot{\gamma}$ is the shear rate, τ is the shear stress, τ_0 is the yield stress and μ is the constant viscosity of the fluid above the yield stress threshold.

2.3.2 Multiphase Flows

The VOF (Volume of Fluid) model [42] is a well-established formulation for tracking the volume fractions of two fluids [43][44]. In the context of minerals processing, it has found use mostly in the modelling of hydrocyclones [45][46]. For a two-phase flow, the continuity equations for the volume fractions are [47]:

$$\frac{\partial}{\partial t}(\alpha_q \rho_q) + \nabla \cdot (\alpha_q \rho_q \mathbf{v}_q) = \sum_{p=1}^2 (\dot{m}_{pq} - \dot{m}_{qp}),$$

subject to the constraint

$$\alpha_1 + \alpha_2 = 1,$$

where α_q is the volume fraction of phase q , the flow variables are similarly labelled, and \dot{m}_{pq} is the mass transfer rate from phase p to phase q .

2.3.3 Turbulence modelling

Reynolds Averaging

A standard approach to turbulence modelling is that of Reynolds averaging, in which turbulence is treated as a fluctuation from some mean flow. In flows that are statistically steady, the mean value of some flow variable $\phi(\mathbf{x}, t)$ is given by the time average [48][40]:

$$\bar{\phi}(\mathbf{x}) = \lim_{T \rightarrow \infty} \frac{1}{T} \int_0^T \phi(\mathbf{x}, t) dt$$

For flows that are spatially homogeneous, volume averaging is used [48]:

$$\bar{\phi}(t) = \lim_{V \rightarrow \infty} \frac{1}{V} \int_V \phi(\mathbf{x}, t) d^3x$$

For flows that are neither steady nor spatially homogeneous, ensemble averaging is used [48][40]:

$$\bar{\phi}(t) = \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{n=1}^N \phi_n(\mathbf{x}, t)$$

where each ϕ_n is from an ensemble of flows under apparently identical conditions.

The instantaneous value of the flow variable is then decomposed into its mean and fluctuating component [48][40][47]:

$$\phi(\mathbf{x}, t) = \bar{\phi} + \phi'(\mathbf{x}, t)$$

where ϕ' is the fluctuating component. Writing the velocity, \mathbf{u} , and the pressure, p , in this way, the continuity equation (Equation 2.6) becomes [48][47]

$$\frac{\partial \rho}{\partial t} + \frac{\partial}{\partial x_i} (\rho \bar{u}_i) = 0$$

and the Navier-Stokes equation (Equation 2.5) becomes [48][40][47]

$$\frac{\partial}{\partial t} (\rho \bar{u}_i) + \frac{\partial}{\partial x_j} (\rho \bar{u}_i \bar{u}_j) = -\frac{\partial \bar{p}}{\partial x_i} + \frac{\partial \bar{\mathbb{T}}_{ij}}{\partial x_j} + \frac{\partial}{\partial x_j} (-\rho \overline{u'_i u'_j})$$

The equation that remains is simply the Navier-Stokes equation for the averaged quantities, with the additional terms $\frac{\partial}{\partial x_j} (-\rho \overline{u'_i u'_j})$, referred to as the Reynolds stresses. One approach is to relate these stresses to the mean velocity gradients using the Boussinesq hypothesis [49][40][47]:

$$-\rho \overline{u'_i u'_j} = \mu_t \left(\frac{\partial \bar{u}_i}{\partial x_j} + \frac{\partial \bar{u}_j}{\partial x_i} \right) - \frac{2}{3} \left(\rho k + \mu_t \frac{\partial \bar{u}_k}{\partial x_k} \right) \delta_{ij}$$

where μ_t is referred to as the turbulent viscosity and k is referred to as the turbulent kinetic energy [49][40][47]:

$$k = \frac{1}{2} \overline{u'_i u'_i}$$

Some popular approaches for modelling the turbulent viscosity and turbulent kinetic energy are highlighted in the following sections. In all of the following equations, the mean flow variables, $\bar{\phi}$, are written without the overbar.

Standard k - ϵ Model

The k - ϵ model, proposed in 1974 by Launder and Spalding [50], relates the turbulent viscosity, μ_t , to the turbulent kinetic energy, k , for flows in which the rates of production and destruction of turbulence are approximately equal [50][40]. The relation is given by [50][40][47]:

$$\mu_t = \rho C_\mu \frac{k^2}{\epsilon}$$

where ϵ is the turbulent dissipation rate and C_μ is a model constant. The evolution of k is given by [50][40][47]

$$\frac{\partial}{\partial t}(\rho k) + \frac{\partial}{\partial x_i}(\rho k u_i) = \frac{\partial}{\partial x_j} \left[\left(\mu + \frac{\mu_t}{\sigma_k} \right) \frac{\partial k}{\partial x_j} \right] + G_k - \rho \epsilon \quad (2.7)$$

where σ_k is the turbulent Prandtl number for k , and G_k , the rate of turbulent kinetic generation, is modeled as:

$$G_k = \mu_t S^2$$

where

$$S \equiv \sqrt{\frac{1}{2} \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right) \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right)}$$

The evolution of ϵ is given by [50][40][47]:

$$\frac{\partial}{\partial t}(\rho \epsilon) + \frac{\partial}{\partial x_i}(\rho \epsilon u_i) = \frac{\partial}{\partial x_j} \left[\left(\mu + \frac{\mu_t}{\sigma_\epsilon} \right) \frac{\partial \epsilon}{\partial x_j} \right] + C_{1\epsilon} \frac{\epsilon}{k} G_k - C_{2\epsilon} \rho \frac{\epsilon^2}{k}$$

where $C_{1\epsilon}$ and $C_{2\epsilon}$ are model constants. The typically used values for the model constants, determined from experiments with air and water, are [40][47]

$$C_\mu = 0.09, \quad C_{1\epsilon} = 1.44, \quad C_{2\epsilon} = 1.92, \quad \sigma_k = 1.0, \quad \sigma_\epsilon = 1.3$$

The k - ϵ model is used frequently because of its relatively low computational cost [48][51].

RNG k - ϵ Model

The Renormalization Group k - ϵ Model was proposed by Yakhot and Orszag in 1986 [52]. The evolution of k is given by [52][47]:

$$\frac{\partial}{\partial t}(\rho k) + \frac{\partial}{\partial x_i}(\rho k u_i) = \frac{\partial}{\partial x_j} \left(\alpha_k \mu_{\text{eff}} \frac{\partial k}{\partial x_j} \right) + G_k - \rho \epsilon$$

where α_k is the inverse effective Prandtl number for the turbulent kinetic energy and μ_{eff} is the effective viscosity. The effective viscosity is related to the molecular viscosity by the following differential equation for their ratio, $\hat{\nu} = \mu_{\text{eff}}/\mu$ [52][47]:

$$d \left(\frac{\rho^2 k}{\sqrt{\epsilon \mu}} \right) = 1.72 \frac{\hat{\nu}}{\sqrt{\hat{\nu}^3 - 1 + C_\nu}} d\hat{\nu}$$

where $C_\nu \approx 100$. The evolution of ϵ is given by [52][47]

$$\frac{\partial}{\partial t}(\rho \epsilon) + \frac{\partial}{\partial x_i}(\rho \epsilon u_i) = \frac{\partial}{\partial x_j} \left(\alpha_\epsilon \mu_{\text{eff}} \frac{\partial \epsilon}{\partial x_j} \right) + C_{1\epsilon} \frac{\epsilon}{k} G_k - C_{2\epsilon}^* \rho \frac{\epsilon^2}{k}$$

where α_ϵ is the inverse effective Prandtl number for the turbulent dissipation rate and the modified coefficient $C_{2\epsilon}^*$ is given by [52][47]

$$C_{2\epsilon}^* \equiv C_{2\epsilon} + \frac{C_\mu \eta^3 (1 - \eta/\eta_0)}{1 + \beta \eta^3}$$

where $\eta \equiv \frac{Sk}{\epsilon}$, $\eta_0 = 4.38$ and $\beta = 0.012$.

The RNG theory sets the the inverse effective Prandtl numbers, α_k and α_ϵ equal to one another, and gives them as [47]

$$\left| \frac{\alpha - 1.3929}{\alpha_0 - 1.3929} \right|^{0.6321} \left| \frac{\alpha + 2.3929}{\alpha_0 + 2.3929} \right|^{0.3679} = \frac{\mu}{\mu_{\text{eff}}}$$

where $\alpha_0 = 1.0$. The key model constants are given by [52][47]:

$$C_{1\epsilon} = 1.42 \quad \text{and} \quad C_{1\epsilon} = 1.68.$$

Realisable k - ϵ Model

The Realisable k - ϵ Model was proposed by Shih *et al.* in 1994 [53]. The evolution of k is the same as that in the standard k - ϵ model (Equation 2.7). The evolution of ϵ is given by [53][47]:

$$\frac{\partial}{\partial t}(\rho \epsilon) + \frac{\partial}{\partial x_i}(\rho \epsilon u_i) = \frac{\partial}{\partial x_j} \left[\left(\mu + \frac{\mu_t}{\sigma_\epsilon} \right) \frac{\partial \epsilon}{\partial x_j} \right] + \rho C_{1\epsilon} S - \rho C_{2\epsilon} \frac{\epsilon^2}{k + \sqrt{\nu \epsilon}}$$

where C_2 is a model constant, and

$$C_1 = \max \left[0.43, \frac{\eta}{\eta + 5} \right]$$

In the realisable k - ϵ model, unlike the two above, C_μ is no longer a constant, but it modeled by [53][47]:

$$C_\mu = \left(A_0 + A_s \frac{kU^*}{\epsilon} \right)^{-1}$$

where $A_0 = 4.04$, $A_s = \sqrt{6} \cos \phi$, and

$$U^* \equiv \frac{1}{\sqrt{2}} S$$

for systems in an inertial reference frame. In a single rotating reference frame, the expression for U^* takes on terms from the rate-of-rotation tensor [47].

2.4 Solid-Fluid Coupling

There exist a number of techniques for numerically treating the interaction of solid particles and fluids. These generally involve a hybrid of a technique developed for solids and one developed for fluids. Some common techniques include DEM-CFD, DPM(Discrete Particle Method)-CFD and DEM-SPH(Smoothed Particle Hydrodynamics).

2.4.1 Drag Forces

The calculation of the drag force on the fluid is key to the implementation. In general, a drag model can be written in terms of the drag force on the fluid per unit volume:

$$\mathbf{F}_D = \beta \mathbf{v}$$

where $\mathbf{v} = \mathbf{u}_s - \mathbf{u}_f$ is the difference in velocity between the solid and fluid, and β is known as the interphase momentum exchange coefficient. Equivalently, the drag force is often expressed in the drag force literature in terms of the non-dimensional drag force, F , which is related to the interphase momentum exchange coefficient by:

$$\beta = \frac{18\mu(1-\phi)^2\phi}{d^2} F$$

where μ is the fluid viscosity, d is the particle diameter and ϕ is the solidicity. Typical drag models express this non-dimensional drag force as a function of the solidicity and the Reynolds number, Re , defined as:

$$Re \equiv \frac{(1-\phi)\rho d |\mathbf{v}|}{\mu}$$

where ρ is the fluid density. Many drag models exist with many different levels of applicability in terms of the particle size, Reynolds number and solidicity. One class of models, which will be covered first with the Ergun equation as its most prominent representative, assumes a functional form for F and fits the constants of the model with experimental data. A more modern class of models, which will be covered second, uses fits to data from high-resolution simulations to determine an expression for F .

Ergun and Wen&Yu

For high-solidicity systems and low Reynolds numbers, the following functional form of F was suggested by Blake in 1922 [54][55]:

$$F(\phi, Re) = \frac{a\phi}{18(1-\phi)^2} + \frac{b}{18(1-\phi)^2} Re \quad (2.8)$$

where a and b are unknown constants. Values for these constants were determined by Ergun in 1952 by running a series of experiments on the pressure drop over packed beds under different conditions. The Ergun experiments yielded values of $a = 150$ and $b = 1.75$ [56]. More recent (1987) experiments by Fand *et al.* have yielded values for these constants for different ranges of the Reynolds number [57][55]:

$$\begin{aligned} a = 192, b = 0 & \quad \text{if } Re < 2.3 \\ a = 182, b = 1.92 & \quad \text{if } 5 < Re < 80 \\ a = 225, b = 1.61 & \quad \text{if } Re > 120 \end{aligned}$$

Despite these newer results, the Ergun correlation is still used frequently in the literature. Due to being inapplicable at low solidicities, it is often augmented with a correlation by Wen&Yu for solidicities below 0.2. The interphase momentum exchange coefficient is given by [56][58][59]:

$$\beta = \begin{cases} \frac{150\phi^2\mu}{(1-\phi)d^2} + \frac{1.75\phi\rho|\mathbf{v}|}{d} & \text{if } \phi > 0.2 \\ \frac{3}{4}C_D\rho(1-\phi)^{-1.65}\phi|\mathbf{v}| & \text{if } \phi \leq 0.2 \end{cases}$$

where μ is the fluid viscosity, d is the particle diameter, ρ is the fluid density and the drag coefficient C_D is given by

$$C_D = \begin{cases} \frac{24}{Re} & \text{if } Re \leq 0.5 \\ 24(1.0 + 0.15Re^{0.687})/Re & \text{if } 0.5 < Re \leq 1000 \\ 0.44 & \text{if } Re > 1000 \end{cases}$$

Hill *et al.*

A pair of seminal papers by Hill *et al.* [60][61] describe a low to moderate Reynolds number drag correlation based on data from Lattice-Boltzmann simulations. Their approach was to run a set of fluid simulations on small domains with a limited number of spherical particles at a high enough resolution to resolve flows around and between the particles. Run under different conditions, these simulations are regarded as virtual experiments, where the drag force can be calculated based on the input solidicity and Reynolds number.

The first step in each simulation was to place a set of spheres (totalling between 8 and 64) randomly

in a cubic domain with periodic boundary conditions. The mechanism of this random placement was given particular attention to ensure that the permeability of the domain was isotropic, that is the nature of the flow and the arising drag force would be independent of its direction through the domain. This was done to ensure that the measured drag force was standardised to some class of particle configuration (in this case, one with isotropic permeability), as differences in the placement of the spheres have a large influence on the permeability and hence the drag force [61][55].

The authors characterised the isotropy of the domain in terms of how different the configuration of the spheres was from a crystal-like structure, since crystal structures show a permeability that has a strong dependence on the direction of the flow (as shown in a separate set of experiments in [61]). To this end, Zinchenko's method [62] (which generates arrays of spheres that are almost close-packed, but do not have a crystal-like structure) was used to configure the arrays of spheres. For low solidicities, the spheres were shrunk and additional Monte-Carlo steps were taken. The authors checked the configurations for a crystal-like structure by quantifying such a structure as one where the angle between a given sphere and its two nearest neighbours did not deviate greatly through the configuration (for example there would only be one angle in a hexagonal close packed configuration).

Once set up with a given ϕ , each simulation was run to steady state for a given Re and the drag force maintaining the equilibrium was calculated. This drag force for the given ϕ and Re was then used as a data point for fits to a given functional form of the non-dimensional drag force, $F(\phi, Re)$. The Reynolds numbers were kept relatively low (below 80) and it was found that the flows in the simulations were persistently unsteady for Reynolds numbers in the range 35-70 [61], where the Reynolds number here is defined as $Re = \frac{\rho(1-\phi)|\mathbf{v}|d}{2\mu}$.

The simulation results and their functional fits were separated into two regimes. The first set of simulations (presented in [60]), for Reynolds numbers below 20 were done for 19 values of ϕ in the range 1.04×10^{-3} - 6.41×10^{-1} and were fit to the following functional form:

$$F(\phi, Re) = F_0(\phi) + F_1(\phi)Re^2 \quad (2.9)$$

The fits gave

$$F_0 = \begin{cases} \frac{(1+3\sqrt{\phi/2}+(135/64)\phi \ln(\phi)+17.14\phi)}{1+0.681\phi-8.48\phi^2+8.16\phi^3} & \text{if } \phi < 0.4 \\ 10 \frac{\phi}{(1-\phi)^3} & \text{if } \phi \geq 0.4 \end{cases}$$

and

$$F_1 = \begin{cases} \sqrt{\frac{2}{\phi}}/40 & \text{if } 0.01 \leq \phi \leq 0.1 \\ 0.11 + 0.00051e^{11.6\phi} & \text{if } \phi > 0.1 \end{cases}$$

The second set of simulations (presented in [61]) for Reynolds numbers above 20, were done for 19 values of ϕ in the range 0.0953-0.641 and were fitted to the following functional form:

$$F(\phi, Re) = F_2(\phi) + F_3(\phi)Re$$

The fits gave

$$F_3 = 0.0673 + 0.212\phi + \frac{0.0232}{(1-\phi)^5}$$

The expression for F_2 was taken to be the same as that for F_0 in spite of different (though marginally so) data for F_2 . Beetstra *et al.* claim that the differences are only insignificant above $Re > 200$ [55]. Their expression is described later.

In a paper by Bokkers *et al.* [63], the Hill drag correlation was implemented for a gas-solid fluidised bed, and was found to match experimental data better than the Ergun, Wen & Yu combination. The implementation was exactly as described above, but with the high-solidicity reduction of:

$$F_1 = 0.110 + 5.10 \times 10^{-4} e^{11.6\phi}.$$

Benyahia *et al.*

A modification and extension of the Hill-Koch-Ladd (HKL) formulas to a larger range of Reynolds numbers and solidicities was done by Benyahia, Syamlal and O'Brien in 2006 [64]. The modifications were primarily done to smooth discontinuities in the HKL formulas and ensure that they matched the limiting values for single-particle drag. The extension was primarily done to address the fact, as raised by Hill *et al.*, that the expression for F_3 loses validity at low solidicities [61].

The modifications and extensions are manifested in the drag coefficients and the non-dimensional drag force as follows [64]:

$$\begin{aligned}
 F_0 &= \begin{cases} (1-w) \frac{(1+3\sqrt{\phi/2}+(135/64)\phi \ln(\phi)+17.14\phi)}{1+0.681\phi-8.48\phi^2+8.16\phi^3} + w \left(\frac{10\phi}{(1-\phi)^3} \right) & \text{if } 0.01 < \phi < 0.4 \\ 10 \frac{\phi}{(1-\phi)^3} & \text{if } \phi \geq 0.4 \end{cases} \\
 F_1 &= \begin{cases} \sqrt{\frac{2}{\phi}}/40 & \text{if } 0.01 < \phi \leq 0.1 \\ 0.11 + 0.00051e^{11.6\phi} & \text{if } \phi > 0.1 \end{cases} \\
 F_2 &= \begin{cases} (1-w) \frac{(1+3\sqrt{\phi/2}+(135/64)\phi \ln(\phi)+17.89\phi)}{1+0.681\phi-11.03\phi^2+15.41\phi^3} + w \left(\frac{10\phi}{(1-\phi)^3} \right) & \text{if } \phi < 0.4 \\ 10 \frac{\phi}{(1-\phi)^3} & \text{if } \phi \geq 0.4 \end{cases} \\
 F_3 &= \begin{cases} 0.9351\phi + 0.03667 & \text{if } \phi < 0.0953 \\ 0.0673 + 0.212\phi + 0.0232/(1-\phi)^5 & \text{if } \phi \geq 0.0953 \end{cases}
 \end{aligned}$$

where $w = e^{-10(0.4-\phi)/\phi}$. The expression for the non-dimensional drag force is divided into four regions in (Re, ϕ) space. For very low solidicities, $\phi \leq 0.01$, the force is [64]:

$$F = \begin{cases} 1 + \frac{3}{8}Re & \text{if } Re \leq \frac{F_2-1}{3/8-F_3} \\ F_2 + F_3Re & \text{if } Re > \frac{F_2-1}{3/8-F_3} \end{cases}$$

For typical solidicities, $\phi > 0.01$, the force is [64]:

$$F = \begin{cases} F_0 + F_1Re^2 & \text{if } Re \leq F_3 + \frac{\sqrt{F_3^2 - 4F_1(F_0 - F_2)}}{2F_1} \\ F_2 + F_3Re & \text{if } Re > F_3 + \frac{\sqrt{F_3^2 - 4F_1(F_0 - F_2)}}{2F_1} \end{cases}$$

Beetstra

Work by Beetstra *et al.* in 2007 [55] also used data from Lattice-Boltzmann simulations to derive a drag correlation using methods very similar to those used by Hill *et al.* [61]. For each simulation, exactly 54 particles were randomly distributed into a cubic domain with periodic boundary conditions. The simulation was then run to equilibrium and a time average of the drag force was calculated for unsteady flows. Contrary to the Hill approach of carefully creating a domain with isotropic permeability, simulations were simply done with 10-30 different combinations of particle configuration and flow direction. It was found (as found by Hill *et al.* [61]) that these variations caused large fluctuations in the drag force, and the mean over the different configurations was used. Values that were more than 2.5σ away from the mean were discarded.

In contrast to the work by Hill, which was directed at a fundamental understanding of flow for low and moderate Reynolds numbers, the aim of Beetstra's correlation is to provide a drag force expression that is valid at high Reynolds numbers and extrapolates well to accepted values at very high Reynolds numbers [55]. In particular, their correlation is constrained to match the expression $F(0, Re) = 0.413Re/24$ in the $Re \rightarrow \infty$ limit. The authors admit that in pursuing their goal of high Reynolds number validity and having a simple functional form for the expression, they do not make use of the fact that the drag force scales as Re^2 for very small Reynolds numbers. However, the expression is still constrained at $Re = 0$ to match the Carman equation:

$$F(\phi, 0) = \frac{180\phi}{18(1-\phi)} + (1-\phi)(1 + 1.5\sqrt{\phi}), \quad (2.10)$$

which is an expansion of the $Re = 0$ expression for the Blake Equation (2.8) to arbitrary solidicities.

The simulations were carried out for 35 different (Re, ϕ) pairs, taken from Re values in the set $\{21.0, 30.5, 105.0, 153.0, 210.0, 420.0, 612.5, 1049.4\}$ and ϕ values in the set $\{0.100, 0.200, 0.300, 0.400, 0.450, 0.500, 0.550, 0.600\}$. All of the Re values were used for $\phi = 0.600$ and all of the ϕ values were used for $Re = 1049.4$. The functional form of the non-dimensional drag force was

assumed to be:

$$F(\phi, Re) = F(\phi, 0) + \alpha(\phi, Re)Re$$

where $F(\phi, 0)$ is given by the Carman expression (Equation 2.10). The function $\alpha(\phi, Re)$ was then fitted to the data in order to obtain the following expression for the non-dimensional drag force [55]:

$$F(\phi, Re) = \frac{10\phi}{1-\phi} + (1-\phi) \left(1 + 1.5\sqrt{\phi}\right) + \frac{0.413Re}{24(1-\phi)^3} \left[\frac{(1-\phi)^{-1} + 3\phi(1-\phi) + 8.4Re^{-0.343}}{1 + 10^3\phi Re^{-(1+4\phi)/2}} \right],$$

The result extends earlier low Re results by the group (in [65]) and is given notably free of conditions on Re and ϕ . The authors note that the maximum deviation of this function from any of the data points (the simulation results) is 8% and the average deviation is 3%. This expression is believed by Deen [58] to be of greater practical value than that of Hill *et al.* [61].

Drag model comparison

The drag correlations for $\phi = 0$, i.e. an isolated particle in a large volume of fluid, are compared for a large range of Reynolds numbers in Figure 2.7. The Beetstra and Wen & Yu correlations can be seen matching the same high-Reynolds number limit.

Figure 2.8 compares the drag correlations in the regime that will be encountered most frequently in this work, high solidities (≈ 0.5) and moderate Reynolds numbers. The Beetstra and Benyahia correlations can be seen matching the Carman equation value (Equation 2.10) for $Re = 0$. While these two correlations produce similar values, the Ergun correlation predicts a lower drag for Reynolds numbers below 50 and predicts an increasingly greater drag for Reynolds numbers above 50. The effect of these differences in the drag correlations is examined in Subsection 5.5.1.

Non-Static Configurations

All of the above correlations have been derived for static configurations of solid particles, but for the systems under investigation in this work, the solid configuration is manifestly not static. Accurate correlations that take granular temperature into account remain an open problem in the field [55] but many investigators have achieved good experimental agreement in highly dynamic systems using the above correlations [58].




Figure 2.7: Comparison of drag models for an isolated particle




Figure 2.8: Comparison of drag models at $\phi = 0.5$ for moderate Reynolds numbers

2.4.2 DEM-SPH in Tumbling Mills

SPH (Smoothed Particle Hydrodynamics) is a particle-based method for solving PDEs, which has been combined with DEM to solve coupled solid-fluid problems, where the solid motion is modeled using DEM and the fluid motion modeled using SPH [66][13].

Cleary *et al.*

One-way coupled simulations of a full-scale tumbling mill were done by Cleary *et al.* [13] using a DEM-SPH method. The basic approach was to first do a “dry” DEM simulation with only the solid particles and then use the information from that simulation to determine the drag force in the SPH simulation. The dynamics from the SPH simulation were not fed back into the DEM simulation.

The drag force of the DEM particles on the fluid was calculated using a Darcy drag law [13]:

$$\mathbf{F}_D = \frac{(1 - \phi)^2 \mu \mathbf{v}}{\rho K}$$

where K is the permeability of the solids, given by the Kozeny-Carman equation [13]:

$$K = \frac{(1 - \phi)^3}{CT\phi^2 S^2},$$

where C is the shape factor of the particles, T is the tortuosity of the fluid pathways through the particles and S is the ratio of surface area to volume for the particles. For the case of spheres ($S = \frac{6}{d}$, $C = 1$), this reduces to:

$$K = \frac{(1 - \phi)^3 d^2}{36T\phi^2}.$$

The simulations were done in 2D and values for the particle solidicity and velocity were obtained by averaging the solidicities and velocities of particles over a cylindrical grid with 30 subdivisions in the radial direction and 40 subdivisions in the angular direction. In addition to this spatial averaging, there was also time averaging of the particle quantities over the entire length of the simulation. The drag force on each SPH particle was then determined based on its velocity and the particle velocity and solidicity at that point, interpolated from the cylindrical grid. This one-way coupling and time-averaging was justified by noting that the flow of the particles is in steady state and that the effect of the slurry on the particles would be minimal [13].

For the DEM, a linear spring-dashpot contact model (Subsection 2.1.1) was used to resolve contact forces and time-stepping was done using a second-order predictor-corrector scheme with 15-25

timesteps per collision [13]. For the SPH, a quasi-compressible (with a maximum density variation of 1%) Navier-Stokes equation was used [13].

The mill used in the work was a 10.86 m diameter mill with 72 lifters with a height of 20 cm. The solids consisted of steel balls in the size range 31-125 mm in diameter and rocks in the size range 35-200mm in diameter. The area of slurry introduced to the mill was 7.2 m², which is 80% of the area needed to completely fill the interstices between the solids when the mill is stationary.

The spring-dashpot friction coefficient was kept constant at 0.5, the coefficients of restitution were 0.3, 0.5 and 0.8 for rock-rock, rock-steel and steel-steel respectively and 10,000 particles were used in the DEM simulation. The slurry was given a viscosity of 0.1 Pa s and 10,300 particles of radius 3 cm were used in the SPH simulation [13].

Figure 2.9 shows the time-averaged and spatially-binned results from the DEM simulation and Figure 2.10 shows the slurry profile throughout the simulation. The slurry profile has roughly the same shape as the charge and for a short time there is a dry core near the centre of circulation.




Figure 2.9: Time-averaged and spatially-binned solid quantities by Cleary *et al.* [13]: (a) solid fraction distribution; (b) velocity distribution of the charge; (c) spatial distribution of the x component of velocity and (d) spatial distribution of the y component of velocity [13]

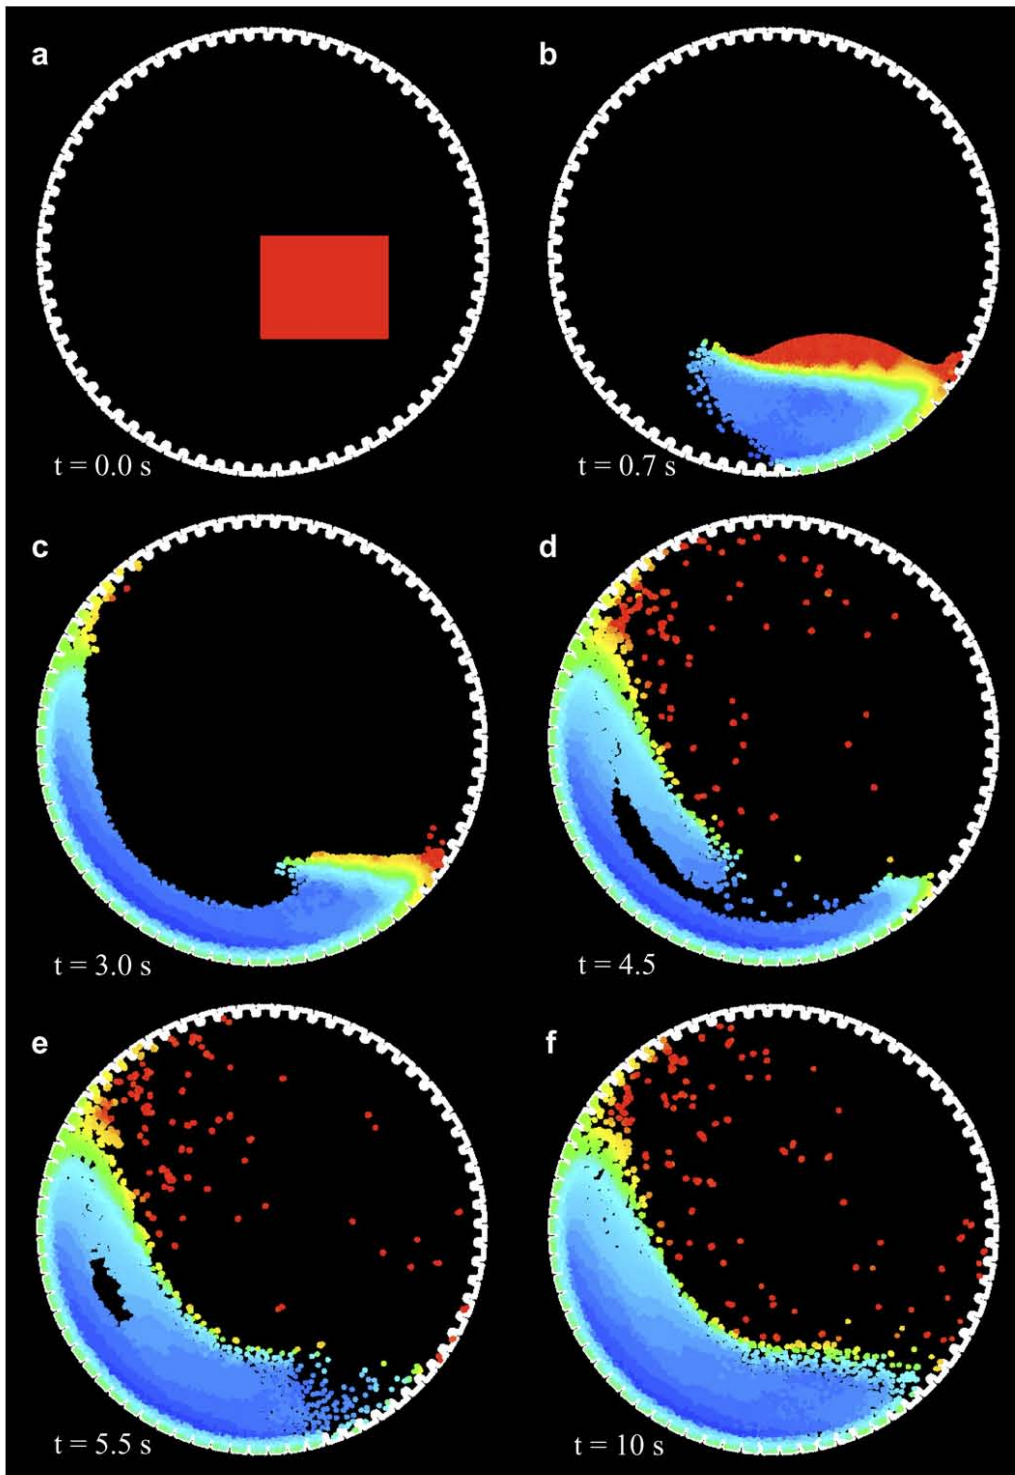


Figure 2.10: Slurry distribution throughout the simulation by Cleary *et al.* [13]. The slurry is coloured by its volume fraction, where red = 1.0 and dark blue = 0.35

2.5 Slurry in Tumbling Mills

The work reviewed in this section represents a few of the existing approaches to modelling slurry transport in tumbling mills. The features of the work that allow it to be compared with this work are highlighted and will be referenced as comparisons in Sections 5.4 and 5.5.

2.5.1 Moys - mechanistic modelling

A mechanistic model for the axial transport of slurry has been developed by Moys [3] to relate slurry hold-up to the flowrate and viscosity. The model is conceptually quite simple, and treats the slurry in the mill as having a uniform level in the transverse direction and a decreasing level across the axis of the mill [3]. This conjectured axial slurry profile is shown in Figure 2.11. The model suggests that the gravitational pressure gradient caused by this difference in slurry level is the key force in driving axial slurry transport [3].

The model also assumes that the Reynolds numbers in the flow, defined in the same way as earlier, $Re = \frac{\rho(1-\phi)|v|d}{\mu}$, would be below 30 and that the flow would be laminar in nature. The author notes that the observed Reynolds numbers in their pilot mill were below 20 [3]. With this assumption in mind, the author models the axial pressure drop with the Blake-Kozeny equation [3]:

$$\frac{dp}{dz} = \frac{-72\mu v \phi^2}{(1-\phi)^2 d^2}$$

where v is the physical velocity of the fluid. The author acknowledges that a drag model with a v^2 dependence would be necessary for Reynolds numbers above 30.

With these assumptions in mind, the author finds that at a constant feed rate Q , the axial profile of the height of the slurry would be given by:

$$h^{2.5} = h_L^{2.5} + 2.5\beta Q(L - z),$$

where Q is the feed rate, L is the length of the mill, z is the axial position (starting with zero at the feed end of the mill), h_L is the height of the slurry at the discharge end of the mill, and β is given by:

$$\beta = \frac{1.463}{R^{0.5} \rho k_Q g},$$

where R is the mill radius, ρ is the fluid density, g is the acceleration due to gravity, and k_Q is assumed to have the following proportionality:

$$k_Q \propto \frac{d^2(1-\phi)^3}{\mu\phi^2}.$$




Figure 2.11: The axial slurry profile conjectured by Moys [3]

By treating d and ϕ as constants and allowing the proportionality of μ to be flexible, the author models k_Q as:

$$k_Q = k'_Q \mu^{-m_Q}$$

where k'_Q and m_Q are parameters to be fitted, with m_Q expected to be approximately one. Finally, the holdup volume is given by [3]:

$$V = \pi(1 - \phi)R^{0.5} ((h_L^{2.5} + 2.5\beta QL)^{1.6} - h_L^4) / 8\beta Q$$

The parameters k'_Q and m_Q were determined by fitting the volume model to data from 19 experiments using a pilot-scale mill. The mill had a diameter of 0.434 m and a length of 0.83 m and was filled with a 0.43:0.25:0.32 ratio of 30 mm diameter steel balls, 22 mm and 16 mm cylpebs (cylindrical pebbles) respectively to 30% of the available volume. The mill was operated at variable speeds and was fed with different concentrations of a Witwatersrand quartzite slurry.

The fits to the experimental data gave values of $k'_Q = 536$ and $m_Q = 0.7$, which was lower than expected. However, the standard deviation of the residuals of the fit was 7% of the mean, which gave the author confidence in the model [3].

2.5.2 Morrell and Stephenson - experimental work

Experimental work by Morrell and Stephenson sought to examine the effect of mill operating parameters and grate design on discharge and hold-up in mills [4]. The laboratory-scale mill used in their work had a diameter of 0.3 m and a length of 0.15 m. Figure 2.12 shows the experimental setup. The mill was filled to 30% by volume with steel balls of diameter 6 mm [4] and was operated for a series of flowrates with different grate designs.

For each setup, the mill was fed with water at the given flowrate until it was judged by eye that the water level in the mill had stabilised. The inflow was then stopped and the water present in the mill was allowed to drain into a bucket. With this method, the water volume in the mill at steady-state, i.e. the hold-up, was determined.

It was found (as one might expect), that the hold-up increased linearly with the flowrate. It was also found that the hold-up increased with the mill rotation speed. These relationships are shown in Figure 2.13.

Their explanation for the increase of hold-up with an increase in the speed of the mill is based on the fact that the shape of the water free surface tended to match the shape of the charge [4]. At higher speeds, the charge was raised, and the height of the water free surface in the region with the highest discharge, the bottom of the mill, was reduced [4]. Consequently, the discharge rate in that region was reduced, and hold-up increased. This explanation was supported by their observations shown in Figure 2.14.

The notion of the holdup being dependent on the mill speed is explored further in work by Latchireddi and Morrell, which shows that the hold-up is approximately linear in the mill speed [7]. They also found that the effect was more pronounced for lower mill loads [7], which gives credence to the claim that the lower discharge rate has to do with the charge having a higher position, since the change in the position of the charge is more pronounced for lower loads.

2.5.3 Condori *et al.* - transport regions

Similarly to Morrell and Stephenson, Condori *et al.* suggest that the the majority of transport and discharge occurs at the bottom of the mill [67]. Figure 2.15 illustrates this schematically. The right-hand image shows the conjectured axial profile of the fluid height in red and the path of a typical packet of slurry in blue [67].




Figure 2.12: The experimental setup used by Morrell and Stephenson [4]

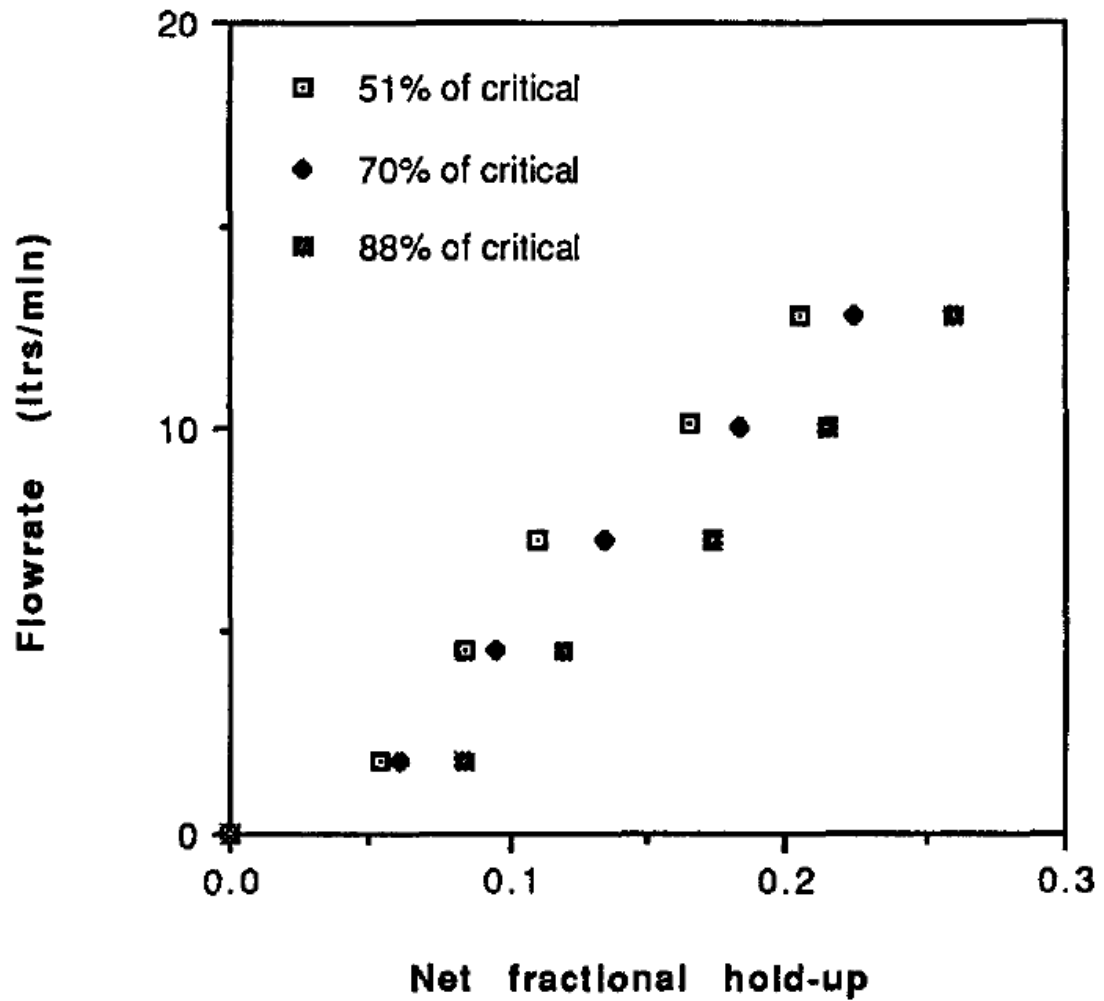


Figure 2.13: The effect of flowrate and mill speed on hold-up as determined by Morrell and Stephenson [4]




Figure 2.14: The discharge flow observed by Morrell and Stephenson for different mill speeds [4]




Figure 2.15: Left: The two zones of slurry motion suggested by Condori et al.. Right: The conjectured axial profile of the fluid height in red and path of a typical packet of slurry in blue by Condori *et al.* [67]

Chapter 3

Coupling Implementation

3.1 Coupling

3.1.1 Scheme

The drag models mentioned in Subsection 2.4.1 are all predicated on the assumption that the flow is steady. Thus, in order for them to be correctly implemented, the flow variables must be constant over time at least within some statistical bounds. For this to be the case, the force applied by the particles on the fluid should also be steady, which requires the same of the velocity of the particles and the local solidicity. Since the particles are in motion, this will never hold for a given point in the charge, and can only be true over some averaging volume and time. Thus, the drag force must be computed over some averaging time and volume. The minimum size of the averaging volume is informed by the conditions under which the given drag model was derived. In the case of the models derived from Lattice-Boltzmann simulations, the simulations were typically done with 8-64 particles in the computational domain, so an averaging volume that could contain that many particles is appropriate.

The length of the averaging time depends on how strictly one defines the flow to be steady. In this work, the relaxation time of the fluid is taken as the expected interval over which the flow is steady. Work by Blakey and James [68] has shown that for Laterite slurries under industrial-scale shear rates, one can expect near-instant relaxation times (reported as zero) for slurries with Bingham viscosities below 20 mPa s and Bingham yield stresses below 7 Pa (which for Laterite slurries corresponded to a solids concentration of 9 %). The slurries in this work fall well below these bounds, so the majority of the simulations use a short time-averaging interval. However, since there is no reliable data available for the relaxation times of the El Soldado slurry used in this work, some simulations are also done with much longer time-averaging intervals. The effect of this choice of averaging time is examined in Subsection 5.5.3.

In this implementation, volume averages are obtained by distributing every particle's mass throughout its own volume and averaging the contributions from all the particles over either CFD cells or axis-aligned voxels. Averaging over individual CFD cells is used in the case where the cell-size is sufficiently large, otherwise axis aligned voxels are used. In this case, the voxels are chosen such that their volume satisfies the requirements mentioned above, and the drag correlations hold on a per-voxel basis. These methods are described in more detail in Section 3.2.

The coupling scheme itself is one-way, in the sense that a full DEM simulation is run without any input from the fluid phase. The data from the DEM simulation is then used to apply a drag force to the fluid simulation. The scheme, for each averaging volume V and averaging time interval $[t, t + T)$, is as follows:

1. For each time-step in the interval for which there is DEM data, $t_i \in [t, t + T)$
 - (a) The volume average of the solidicity and velocity are obtained: $\bar{\phi}(t_i) = \iiint_V \phi(\mathbf{x}, t_i) dV$
and $\bar{\mathbf{v}}(t_i) = \iiint_V \mathbf{v}(\mathbf{x}, t_i) dV$
 - (b) The new volume averages are added to a running time average for the time interval.
2. The CFD simulation is run from time t until time $t + T$ using the drag force calculated from the now volume- and time-averaged particle data and the fluid data at each time-step.

The time-averaging aspect of the scheme is illustrated in Figure 3.1.




Figure 3.1: The time-averaging of DEM timesteps feeding into the CFD simulation for the one-way coupling scheme

3.2 Algorithms

The following two sections describe the local averaging algorithms used in this work for both the coupling process and the presentation of results.

3.2.1 Point in Voxel Local Averaging

The simplest local averaging procedure developed for this work is one that determines the solidity and locally-averaged particle properties in a given set of 3D voxels using data from spherical DEM particles.

Sample point distribution

The first step of the algorithm is to distribute sample points evenly throughout each particle. However, determining the sample points for each particle individually would consume a large amount of computational time, so a procedure was developed to take a single distribution of points, which is repositioned and scaled depending on the particle. The first step in this procedure is creating this single distribution, which will be referred to as the template distribution.

The template distribution is created as follows:

1. An evenly-spaced $n_r \times n_r \times n_r$ grid of points is created in the domain $[-1, 1] \times [-1, 1] \times [-1, 1]$, where n_r is referred to as the template resolution.
2. A set of these template points is created, but any points that fall outside of a sphere of radius 1 centred at the origin are discarded. That is, all points for which $x^2 + y^2 + z^2 > 1$ are discarded.

The points that remain are an even distribution of points throughout the volume of a sphere of radius 1 centered on the origin.

With this template distribution created, determining the sample points for a given particle is simple. If the particle has position \mathbf{x} and radius r , each sample point is generated from a point in the template distribution as follows:

$$\mathbf{P}_S = \mathbf{x} + r\mathbf{P}_T$$

where \mathbf{P}_S is the sample point and \mathbf{P}_T is the point in the template distribution. The template sphere is simply scaled up and repositioned.

Averaging

Once sample points are distributed throughout each of the particles like this, the contribution from each sample point must be added to the voxel containing it. The particle volume contained within a given voxel j , $V_{p,j}$, is simply:

$$V_{p,j} = \sum_i \frac{V_i N_{i,j}}{N_i}$$

where V_i is the volume of particle i , $N_{i,j}$ is the number of samples of particle i that fall in voxel j and N_i is the total number of samples taken of particle i . Figure 3.2 illustrates the process for a low number of samples.

In a similar manner, the particle velocity averaged over a voxel, $\mathbf{v}_{\text{avg},j}$, can also be determined:

$$\mathbf{v}_{\text{avg},j} = \left(\sum_i \frac{V_i N_{i,j} \mathbf{v}_i}{N_i} \right) / V_{p,j}$$

where \mathbf{v}_i is the velocity of particle i . Any other quantities that are weighted with the particle volume can be determined in exactly the same way.

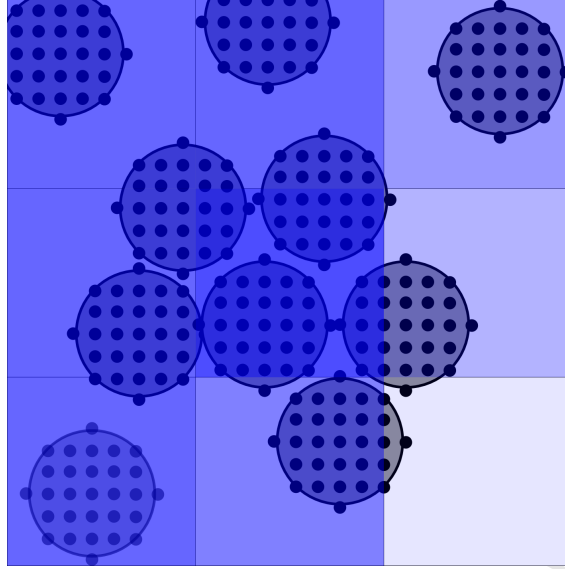


Figure 3.2: Sampling of particles. Darker shaded voxels will have a higher solidicity.

Determining which voxel a sample point falls in

Suppose a grid of $m_x \times m_y \times m_z$ voxels is used to cover the domain $[x_l, x_r] \times [y_l, y_r] \times [z_l, z_r]$. The voxels can then be labeled by indices running from 0 in each dimension such that voxel v_{ijk} has lower bounding co-ordinates $\{x_l + i \frac{x_r - x_l}{m_x}, y_l + j \frac{y_r - y_l}{m_y}, z_l + k \frac{z_r - z_l}{m_z}\}$ and upper bounding co-ordinates $\{x_l + (i + 1) \frac{x_r - x_l}{m_x}, y_l + (j + 1) \frac{y_r - y_l}{m_y}, z_l + (k + 1) \frac{z_r - z_l}{m_z}\}$. The indices of the voxel containing a given sample point are then simply

$$\{i, j, k\} = \left\{ \left\lfloor \frac{x - x_l}{m_x} \right\rfloor, \left\lfloor \frac{y - y_l}{m_y} \right\rfloor, \left\lfloor \frac{z - z_l}{m_z} \right\rfloor \right\}$$

where x, y, z are the co-ordinates of the sample points.

3.2.2 Point in Cell Local Averaging

A much more complex and computationally intensive procedure is that of conducting a similar averaging process to the above, but where the particle properties are averaged over the cells from a CFD mesh, which in this implementation are arbitrary hexahedrons. This local-averaging can be used for applying drag correlations on a per-cell basis and an adaptation of this algorithm is also used to average cell-based CFD data over axis-aligned voxels.

The first stage of the voxel-based algorithm, the distribution of the sample points, is kept as is. Also, the averaging process remains the same, but instead of voxels, the averaging is over cells. Key to this averaging process though, is determining which sample points fall within which cells. When the cells are an arbitrary collection of hexahedrons, a brute-force check of every sample

with every cell would have an asymptotic complexity of $\mathcal{O}(MN)$ where M is the number of cells and N is the number of samples. For this implementation one could reasonably expect roughly a million samples and approximately 50 000 cells, so such a method over thousands of timesteps would be computationally intractable.

An algorithm that drastically reduces this requirement, but that relies on meshes that do not deform (but may rotate and translate) is presented in the following section. Following that, a slower method that still scales well is presented for meshes that can deform arbitrarily.

Non-deforming meshes

For the more computationally inexpensive case of a non-deforming mesh, neither the mesh as a whole, nor any of the cells undergo any kind of deformation over the course of the simulation. In this case, the mesh is a rigid body, and its motion can be entirely described by a sequence of translations and rotations. One can thus describe the position of every vertex in the mesh at every point in time in terms of an appropriate linear transformation applied to the original position of the vertex. If we label the initial position of the vertex $\mathbf{x}(0)$, then the position of the vertex at time t , $\mathbf{x}(t)$, is:

$$\mathbf{x}(t) = T(\mathbf{x}(0), t)$$

where T is some combination of time-dependent translations and rotations. For example, a linear rotation about the z-axis with period 1 would be:

$$T(\mathbf{x}, t) = \{\cos(2\pi t)x - \sin(2\pi t)y, \sin(2\pi t)x + \cos(2\pi t)y, z\}$$

Such transformations will always have an inverse:

$$\mathbf{x}(0) = T^{-1}(\mathbf{x}(t), t)$$

In the case of the linear rotation about the z-axis, the inverse would be:

$$T^{-1}(\mathbf{x}, t) = \{\cos(2\pi t)x + \sin(2\pi t)y, -\sin(2\pi t)x + \cos(2\pi t)y, z\}$$

As will be shown later, the algorithm exploits the fact that inverses exist for this class of transformations.

The first two stages of the algorithm are once-off processes that establish the spatial relationships between the cells and discretise the domain. These operations are performed once only on the mesh in its position at $t = 0$. The first of these two stages is dedicated to arranging the cells in the mesh into coarse spatial groups based on their initial positions. A grid with dimensions $m_x \times m_y \times m_z$ is

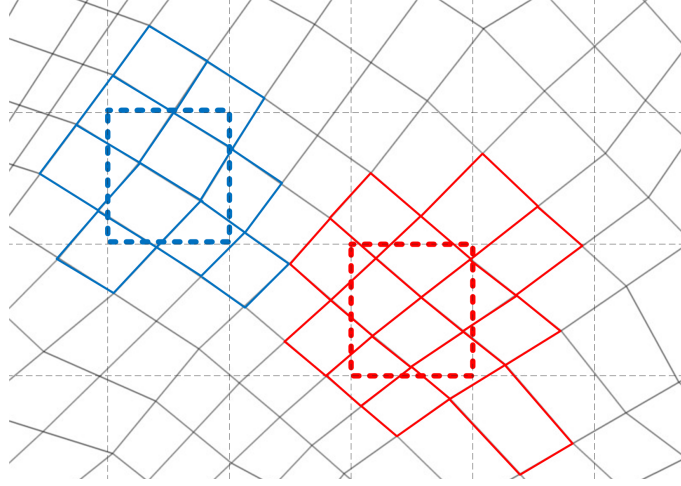


Figure 3.3: Association of CFD cells with voxels for the cell-based local averaging algorithm. The blue cells are associated with the blue voxel and the red cells are associated with the red voxel.

imposed on the computational domain, which we will label $[x_l, x_r] \times [y_l, y_r] \times [z_l, z_r]$. It is important to note at this stage that this grid is a device for the algorithm, not a region over which any averaging is done. The voxels of the grid are labeled by indices running from 0 in each dimension such that voxel v_{ijk} has lower bounding co-ordinates $\{x_l + i \frac{x_r - x_l}{m_x}, y_l + j \frac{y_r - y_l}{m_y}, z_l + k \frac{z_r - z_l}{m_z}\}$ and upper bounding co-ordinates $\{x_l + (i + 1) \frac{x_r - x_l}{m_x}, y_l + (j + 1) \frac{y_r - y_l}{m_y}, z_l + (k + 1) \frac{z_r - z_l}{m_z}\}$. For each voxel in the grid, it is determined which cells have bounding boxes that intersect that voxel. i.e. a cell intersects a voxel if for the set of vertex positions of the voxel, V , and for the set of vertex positions of the cell, C ,

$$\exists i \in \{0, 1, 2\} : \min_{\mathbf{v} \in C} \{v_i\} > \max_{\mathbf{w} \in V} \{w_i\} \vee \min_{\mathbf{w} \in V} \{w_i\} > \max_{\mathbf{v} \in C} \{v_i\}.$$

Figure 3.3 illustrates how cells are associated with the two highlighted voxels. Notice that there are two cells that don't intersect the voxels, but whose bounding boxes do. At the end of this process, every voxel contains an array of all the unique integer IDs of the cells whose bounding boxes overlap with it. A given cell can overlap with many voxels.

The second stage of the algorithm is concerned with creating a much finer grid with dimensions $n_x \times n_y \times n_z$ to quickly determine which cells could potentially contain a given sample point. For each vertex with position $\{x, y, z\}$ in the finer grid, the following process is followed:

1. Determine which voxel in the coarse grid contains the vertex. The indices of the voxel are $\{i, j, k\} = \left\{ \left\lfloor \frac{x - x_l}{m_x} \right\rfloor, \left\lfloor \frac{y - y_l}{m_y} \right\rfloor, \left\lfloor \frac{z - z_l}{m_z} \right\rfloor \right\}$.
2. For each cell associated with that voxel (from the coarse grid), determine if the vertex (from the fine grid) is contained in the cell.

3. For the single cell found (since the cells don't overlap, only one cell can contain the vertex):
 - (a) Get the cell ID.
 - (b) Associate that cell ID with every voxel in the fine grid that contains the vertex.

Figure 3.4 illustrates this process. The first image shows the vertex in question. The second shows the voxel in the coarse grid that contains the vertex. The third shows the cells that are associated with the coarse grid voxel. The fourth image shows the unique cell that contains the fine grid vertex. The final image shows the voxels in the fine grid that will have that cell associated with them. At the end of the process, every voxel in the fine grid contains an array of all the unique integer IDs of the cells that intersect it. The third stage of the algorithm is implemented every time-step and uses the information from the previous two stages to quickly determine which cell will contain an arbitrary sample point. For a given particle, the process is as follows:

1. Determine the current position of the particle, \mathbf{r} .
2. Determine the position of the particle in the reference frame of the mesh, \mathbf{r}' . That is, if the rigid motion of the mesh is governed by the time-dependent transformation T , then $\mathbf{r}' = T^{-1}(\mathbf{r}, t)$.
3. Take a pre-defined set of samples (using the template scaling and re-positioning procedure described in Subsection 3.2.1) through the volume of the particle in the reference frame of the mesh.¹ For each sample point \mathbf{x} :
 - (a) Determine which fine grid voxel the point \mathbf{x} lies in. The voxel indices are $\{i, j, k\} = \left\{ \left\lfloor \frac{x-x_l}{n_x} \right\rfloor, \left\lfloor \frac{y-y_l}{n_y} \right\rfloor, \left\lfloor \frac{z-z_l}{n_z} \right\rfloor \right\}$.
 - (b) Get the pre-computed array of cells that intersect the voxel.
 - (c) Still in the reference frame of the mesh, check each of these cells to determine which one contains \mathbf{x} .

This algorithm for averaging the particle properties over cells has an asymptotic complexity of $\mathcal{O}(N)$, where N is the number of samples taken.

Determining if a point is in a concave hexahedral cell

In the description above, the method for determining whether a point is in a given cell is not described, as it has no effect on the asymptotic complexity of the algorithm. However, the employed method is described here for completeness.

¹If the particle were not spherical, this step would require the pre-defined sampling space for the particle to also be rotated.

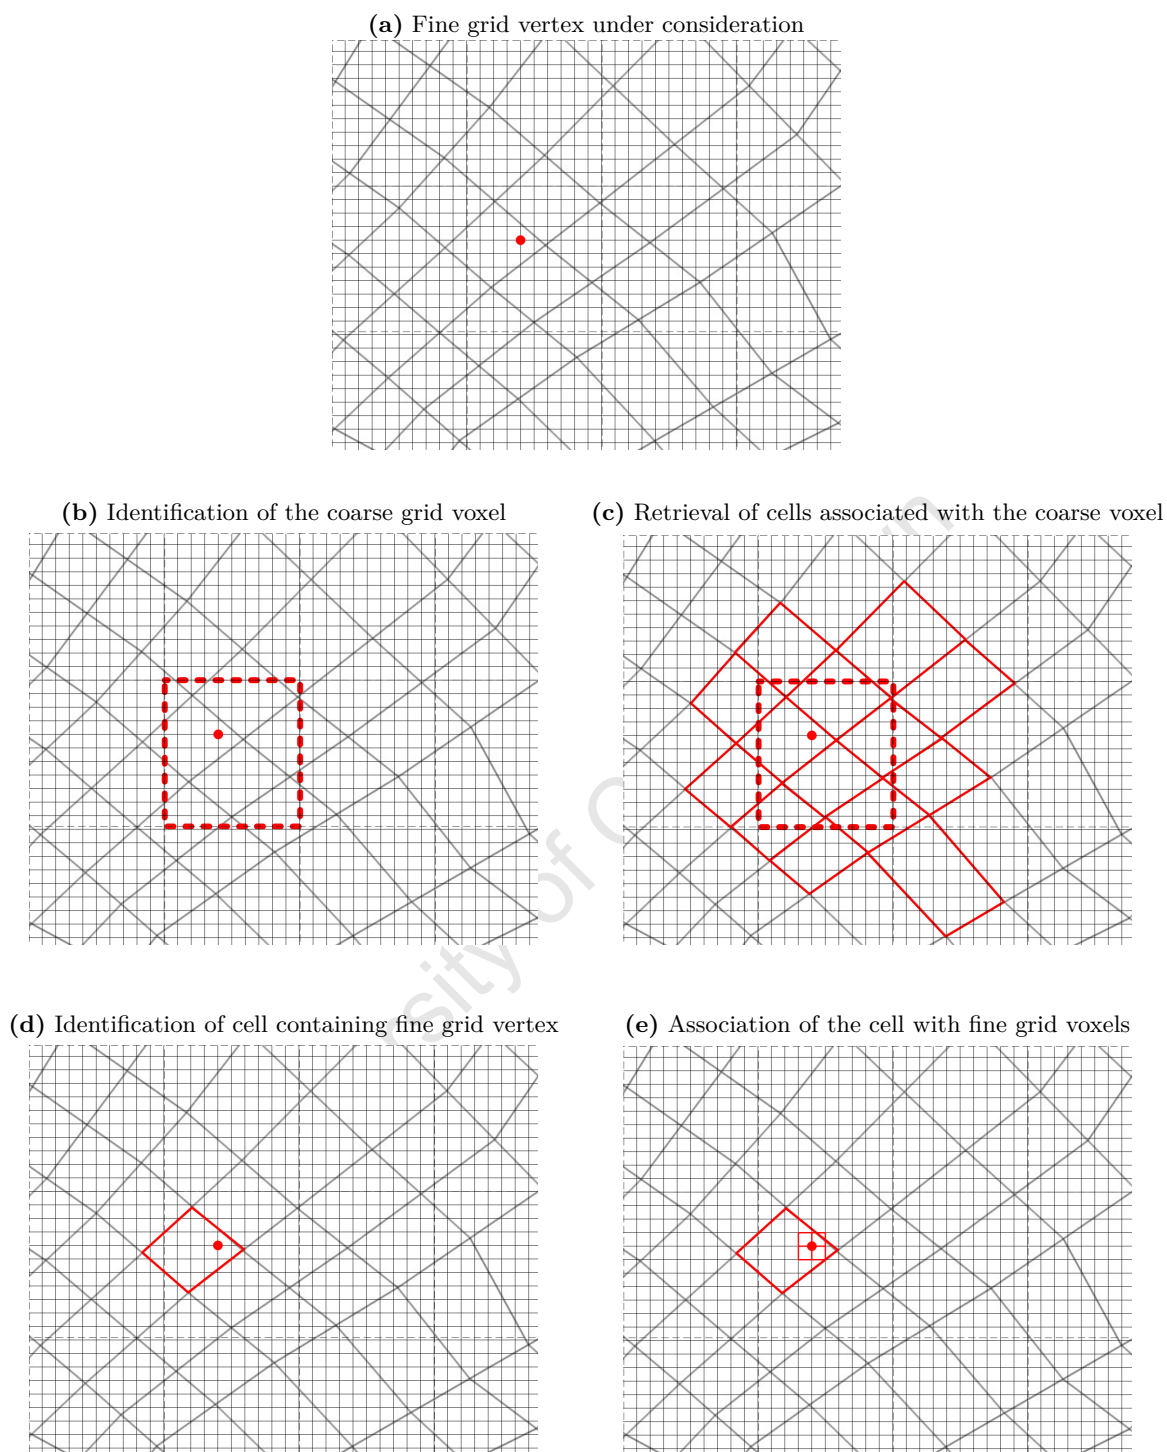


Figure 3.4: The fine grid creation process for cell-based local-averaging. (a): The red dot is the fine grid vertex under consideration. (b): The coarse grid voxel containing the fine grid vertex is identified and is shown as the dashed red rectangle. (c): The cells associated with the coarse voxel, shown as solid red lines, are retrieved. (d): From these cells, the cell containing the fine grid vertex is found. (e): The found cell is identified with all of the fine grid voxels that share the fine grid vertex under consideration.

For each of the faces in the CFD cell, the four vertices are given in a specific order $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3, \mathbf{v}_4$. The face normal could be regarded as pointing one of two ways, so a canonical (but not unit) normal is defined as:

$$\mathbf{n} = (\mathbf{v}_2 - \mathbf{v}_1) \times (\mathbf{v}_3 - \mathbf{v}_1) \quad (3.1)$$

For each face it is also known whether this normal points into or out of the CFD cell. If the normal points into of the cell, one can check if an arbitrary point, x , is on the “cell-side” of the face by examining the dot product $(\mathbf{x} - \mathbf{v}_1) \cdot \mathbf{n}$. If it is positive, then the point is on the cell-side of the face and if it is negative the point is on the opposite side of the face. Similarly, if the face in question has its normal pointing out of the cell, a negative dot product would imply that the point is on the cell-side of the face and a positive dot product would imply that it is on the opposite side.

Such checks are done for every face and as soon as the point is found to be on the wrong side of one of the faces, it is discarded. If the point passes all the checks, it is known that it is inside the cell.

Application to voxel-averaging CFD data

The algorithm for sampling non-deforming meshes can also be used to average data from CFD simulations into voxels for analysis. The algorithm proceeds exactly as outlined above, but instead of the sample points being generated from particles, they are simply taken in an even grid over the domain, which has a finer resolution than the desired voxel resolution. For each of samples it is known which cell contains it, and the properties of that cell are contributed to the voxel that the sample point falls in. The contributions are done exactly the same way as for the particle-property averaging, i.e. they are averaged by volume.

Deforming meshes

For the case of deforming meshes, the spatial relationships between vertices of the mesh are not constant. There is still the need to reduce the computational complexity from the brute force method, but a method is needed that can cope with truly arbitrary deformations, where the distance between two given vertices can change arbitrarily (within the domain of the simulation) and techniques exploiting rigid motion cannot be used.

A standard method for spatial searching problems is binary space partitioning [69], which can be used to quickly determine the overlap between cubes. In this implementation, the method is used to determine which cells have potential intersections with a given particle, after which all of

the samples taken in that particle are tested against all of the potential cells. The scheme is as follows for each time-step:

1. Initialise the binary search tree.
2. Determine the axis-aligned bounding cubes of the particles and cells based on their current positions.
3. Add the bounding cubes of the particles and cells to the binary search tree.
4. For each particle:
 - (a) Traverse the tree to create a list of cells whose bounding cubes overlap the particle's.
 - (b) Take the pre-defined set of sample points through the volume of the particle.
 - (c) For each of the sample points, find which of the cells in the list contains the sample.

The method has an asymptotic complexity of $\mathcal{O}(N \log M)$, where N is the number of samples taken and M is the sum of the number of cells and particles.

3.3 Verification

3.3.1 Pressure drop recovery

The CFD step in the coupling scheme (described in Subsection 3.1.1) is implemented in Ansys Inc.'s FLUENT 12.1 [47]. This verification section is presented in order to check that the drag forces are correctly implemented using FLUENT's user-defined functions and cell-wise momentum sources. Two domains were used, one large and one small. The large domain had dimensions $8 \times 2 \times 2$ m and the small domain had dimensions $8 \times 2 \times 2$ cm. Each domain was discretised using a grid of $32 \times 8 \times 8$ cells for a total of 2048 cubic cells. Figure 3.5 shows a diagram of the smaller domain.

For each domain the velocity at the left and right ends was fixed to zero (in all directions) and the remaining boundaries were treated as no-slip walls. Next, a given solidicity and particle diameter were defined. Then, a series of short runs was performed, each to convergence, with the coupling applied to a range of bounds in the x-direction and a range of particle velocities in the x-direction. The range of x-bounds represents a series of fixed beds of the given length with a uniform solidicity and particle diameter. Since the number of cells is fixed, a reduction in the bed length reduces the number of cells in which the momentum sources are applied. The range of particle velocities was done to check that the implementation is sound over a range of Reynolds numbers that traverses all conditional statements in the drag correlations (see Subsection 2.4.1). The results for the Ergun




Figure 3.5: The smaller domain used for pressure drop comparisons. The shaded area shows a set of x-bounds.

drag correlation are presented here.

Two simple comparisons are used for verification purposes. The first is the absolute difference between the expected pressure drop from the drag correlation and the numerical pressure drop, $|\Delta p_{\text{analytical}} - \Delta p_{\text{numerical}}|$, and the second is the relative error in the numerical pressure drop $\frac{|\Delta p_{\text{analytical}} - \Delta p_{\text{numerical}}|}{\Delta p_{\text{analytical}}}$. The numerical pressure drop was calculated by literally recording the pressure on either side of the domain and taking the difference. Figure 3.6 shows in order the pressure drop from the drag correlation, the numerical pressure drop, the absolute difference and the relative error for particles of diameter 5mm and a solidity 0.15. The comparison shows that the pressure drop was recovered correctly. The tests in the larger domain and tests with a solidity of 0.5 all showed the pressure drop being recovered correctly.

3.4 Geometry

Two laboratory-scale mills are used in this work. The first, a mill with a diameter of 12 cm, is used primarily for DEM-only simulations as a large number of dry PEPT experiments exist for this mill. The second, a mill with a diameter of 30 cm, is used for the one-way coupled runs as there are a number of PEPT experiments that have water or slurry present for this mill. Details of the internal dimensions, the surface mesh (for DEM) and the volume mesh (for CFD) for each mill are given in the following two subsections. The volume meshes were all generated using the process described in Subsection 3.4.3.

3.4.1 The 12 cm Mill

The 12 cm mill is a simple laboratory-scale mill that is effectively a cylinder of length 12 cm and diameter 12 cm, lined with 12 lifters. The dimensions of the mill and the lifters are shown in

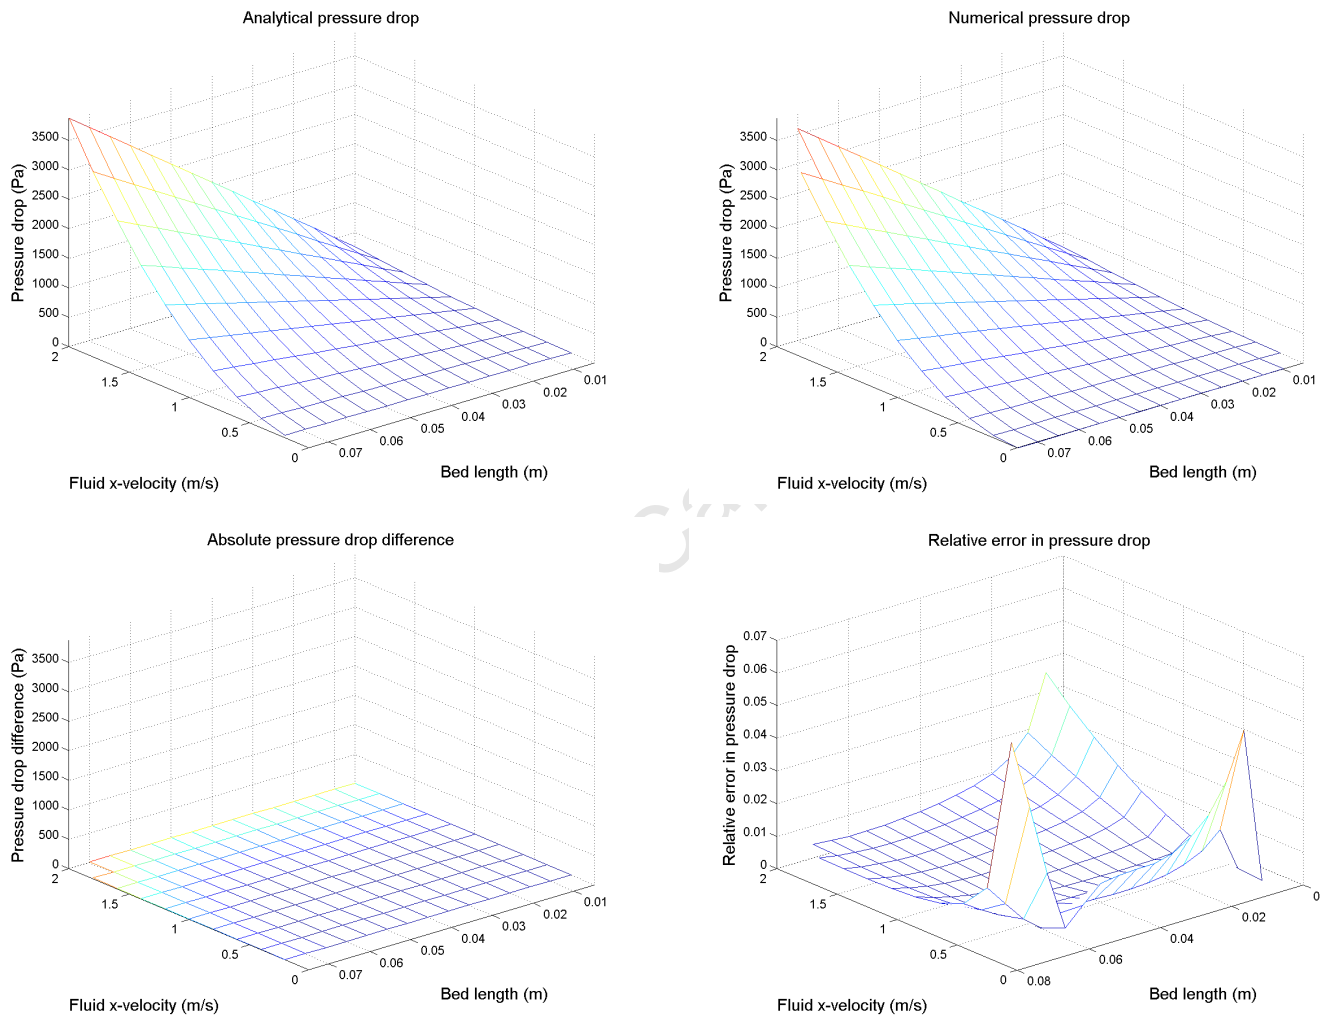


Figure 3.6: Comparison between analytical and numerically calculated pressure drop in the small domain with 5mm diameter particles and $\phi = 0.15$. The “Fluid x-velocity” axis refers to the chosen relative velocity between the particles and the fluid.




Figure 3.7: Dimensions of the 12 cm mill

Figure 3.7. The mill is uniform in the axial direction.

Isometric views of the surface mesh with and without the front wall are shown in Figure 3.8 and Figure 3.9 respectively.

Two core volume meshes were created, a low-resolution version with 3948 faces describing the end walls, and a high-resolution version with 17328 faces describing the end walls. The end walls for each mesh are shown respectively in Figure 3.10 and Figure 3.11. All of the faces are quadrilaterals and a sequence of hexahedral cells was created by simply extruding the the quadrilaterals in the axial direction. The number of times this extrusion was done determined the axial resolution of each mesh, and meshes with 1 cell in the axial direction as well as meshes with 12 cells in the axial direction were created.

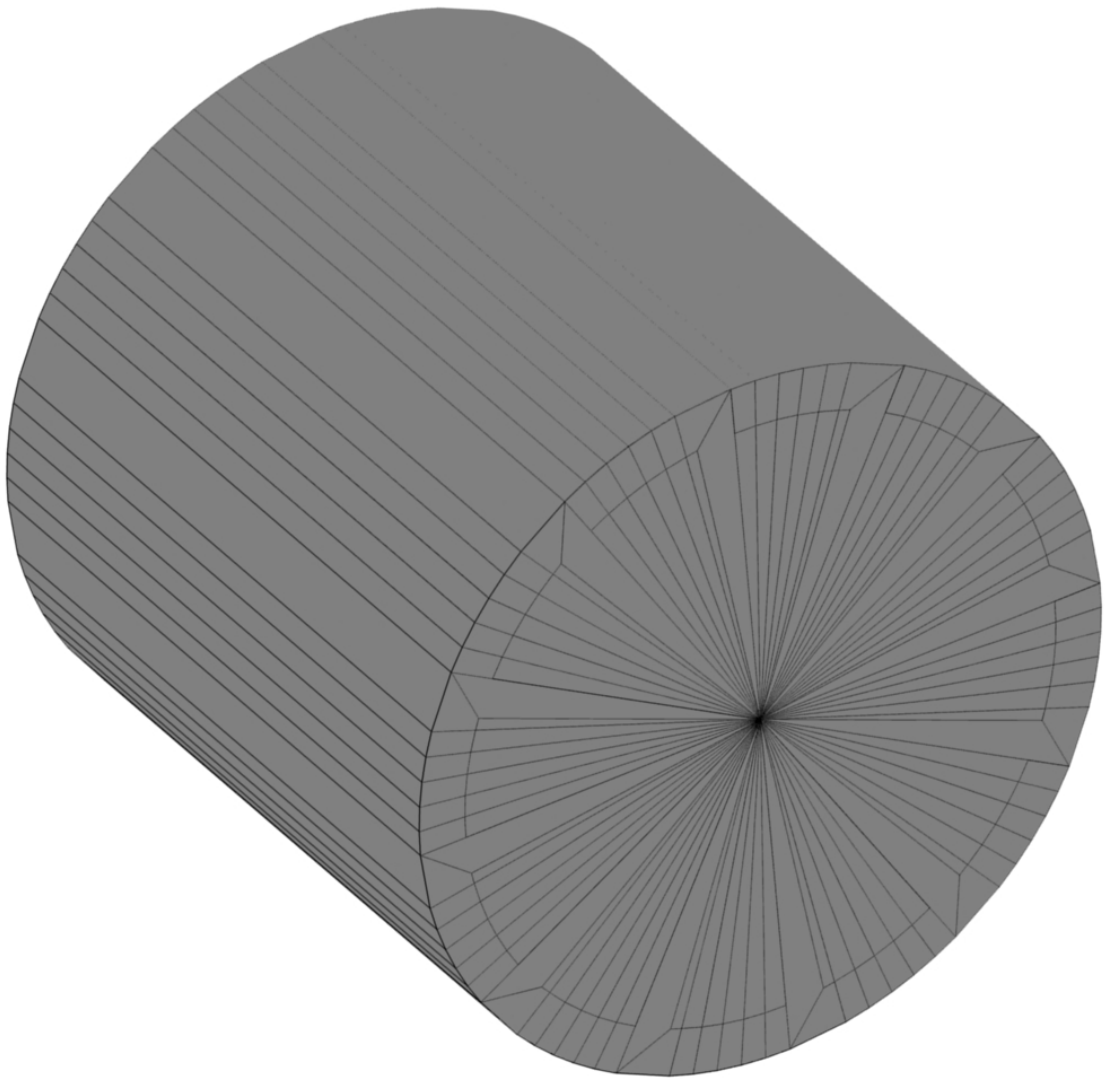


Figure 3.8: Isometric view of the 12 cm mill surface mesh




Figure 3.9: Isometric view of the 12 cm mill surface mesh without the front wall




Figure 3.10: End wall of the low-resolution volume mesh for the 12 cm mill




Figure 3.11: End wall of the high-resolution volume mesh for the 12 cm mill

3.4.2 The 30 cm Mill

The core of the 30 cm mill is a cylinder of diameter 300 mm and length 305 mm and the mill is lined by 20 equally-spaced lifters. On the discharge end, there are 12 equally-spaced pulp lifters, each with a discharge grate composed of holes 2 mm in diameter. The discharge cone and other features outside the main body of the mill are not included in this work. On the feed end, there are four observation windows and a feed tube with a diameter of 64 mm. The dimensions of the mill are shown for the discharge end, feed end and side of the mill respectively in Figures 3.12 - 3.13.

Isometric views of the discharge end of the surface mesh with and without the front wall are shown in Figure 3.15 and Figure 3.16 respectively. An isometric view of the feed end of the surface mesh is shown in Figure 3.17.

One core volume mesh was created, with 1616 faces describing the end walls. The end walls on the discharge end are shown in Figure 3.18. The green faces are those that were used to describe the discharge grates, as detailed later in Section 5.1. The end walls on the feed end are shown in Figure 3.19. The red faces are those that were used to describe the fluid inflow, as detailed later in Section 5.1. Two versions of the mesh were created, one with a depth of 9 cells in the axial direction and one with a depth of 33 cells in the axial direction.

3.4.3 Meshing

In a typical CFD pipeline, the volume mesh is usually generated automatically (using one of a host of methods [70]) from a surface mesh. This process saves the user a lot of time, but is unfortunately fairly imprecise, and often creates meshes that require further refinement [70]. For this implementation, complete control is needed over the exact structure of the CFD mesh. This is firstly to ensure that complex inputs, such as flow through grates, can be adequately handled and secondly so that the sizes of cells can be directly chosen. Direct control of the cell size is useful when the locally-averaged particle properties are calculated on a per-cell basis, as mentioned in Subsection 3.1.1, in which case the cell size must at the very least be greater than the largest particle size. Additionally, for the domains being meshed, the mills, there is discrete axial symmetry, so creating the basis of the mesh by hand is very manageable.

With these things in mind, the first step in this implementation is to create by hand a mesh that has all the vertices, edges and faces of the final volume mesh. This mesh, with zero information about cells or cell connectivity, is then processed to determine which collections of faces enclose cells and determine how those cells are connected to generate the completed volume mesh.




Figure 3.12: Dimensions of the 30 cm mill: discharge end




Figure 3.13: Dimensions of the 30 cm mill: feed end




Figure 3.14: Dimensions of the 30 cm mill: side




Figure 3.15: Isometric view of the discharge end of the 30 cm mill surface mesh




Figure 3.16: Isometric view of the discharge end of the 30 cm mill surface mesh without the end wall




Figure 3.17: Isometric view of the feed end of the 30 cm mill surface mesh

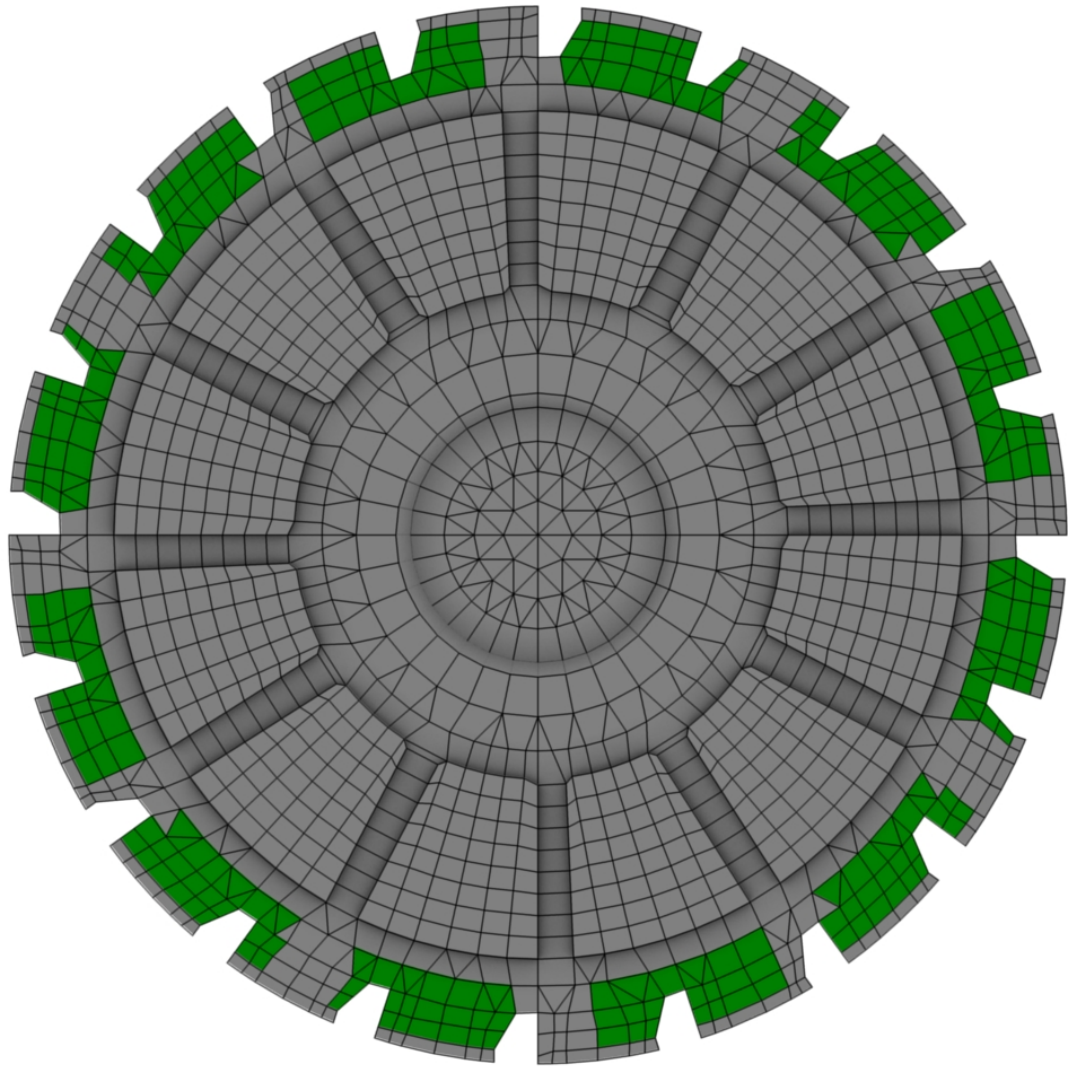


Figure 3.18: End walls on the discharge end of the 30 cm mill volume mesh




Figure 3.19: End walls on the feed end of the 30 cm mill volume mesh

Chapter 4

Validation Using PEPT

Experiments

4.1 The PEPT Method

Positron Emission Particle Tracking (PEPT) is a modification of Positron Emission Tomography, (PET), which is a 3D medical imaging technique. PET uses positron-emitting radionuclides to determine a time-averaged spatial concentration of tracer particles containing these radionuclides. In PEPT, this process is modified by tracking the position of only one tracer particle over time.

A typical positron emitted by such a particle collides with an electron within a short distance and the particles annihilate each other to produce two effectively back-to-back gamma rays, each with an energy of 511 keV. These gamma rays then strike gamma ray detectors in the PEPT camera simultaneously. A PEPT camera will have many such detectors and if any two detectors are triggered by gamma rays within a sufficiently short period of time, it is assumed that the gamma rays came from the same annihilation event. If this is the case, it is known that the annihilation event occurred somewhere along the line joining the two points at which the gamma rays were detected, since they were emitted back-to-back.

Over a very short interval of the order milliseconds, many such “lines of response” are recorded. If there is only one source of positron emitters in the camera’s field of view, the lines must all pass through the same small region of space, which corresponds to the position of the source. A typical collection of these lines of response is shown diagrammatically in Figure 4.1.




Figure 4.1: The lines of response from a PEPT tracer

A standard method for determining the particle's position from the lines of response is to find the point that minimises sum of the perpendicular distances between the point and the lines of response [71][72]. In practice, there are often lines of response that don't correspond to annihilation events but are from background radiation. Additionally, there false coincidences, which are nearly simultaneous gamma ray detections that originate from different events and generate a false line of response. The presence of these spurious lines of response affects the process for determining the particle's position.

A typical triangulation process starts with a collection of 100-500 lines of response and the best estimate of the particle's position using those lines is determined. Then, a large number (between 50% and 70%) of the lines furthest from this point are discarded. Figure 4.2 shows the lines of response from Figure 4.1 after this discarding process. The best position for the particle is then recalculated from the remaining lines and is accurate to within 1-2 mm for typical particle speeds.




Figure 4.2: The lines of response from a PEPT tracer after the spurious lines have been removed

4.1.1 Steady-state Profiles

Over time, the time-sequence of points determined by the triangulation process traces out a trajectory for the particle being tracked. However, for the systems under investigation the quantities of interest involve the average motion of the entire ensemble of particles. It is possible to approximate the ensemble average by a time average of the motion of a single particle, using Boltzmann’s Ergodic Hypothesis, which is paraphrased by Szasz as follows [73]: “For large systems of interacting particles in equilibrium, time averages are close to the ensemble, or equilibrium average.”

The process of constructing this approximated ensemble average from the trajectory of a single particle begins with dividing the domain into a collection of 2D voxels. The voxels are 2D because in this implementation there is spatial averaging in z in addition to the ensemble averaging.

As the particle moves through the domain, it enters and exits these voxels. Every time the particle enters a voxel, the time it spends in that voxel is calculated and added to a running total of time spent in that voxel. At the end of the analysis, each voxel, labelled with indices i, j , has the total time that the particle spent in that voxel, τ_{ij} . Figure 4.3 shows these τ_{ij} , known as the residence times, being determined from the particle trajectory (images courtesy of Angus

Morrison). The proportion of time the particle spent in that voxel is then

$$\alpha_{ij} = \tau_{ij} / \sum_{kl} \tau_{kl}$$

According to the ergodic hypothesis, this time- and spatially-averaged (where the spatial average is over the voxel) residence fraction will be approximately equal to the ensemble- and spatially-averaged residence fraction. Thus, if the total volume of the particles, V , is known, then the average particle volume in a given voxel, V_{ij} can be calculated as

$$V_{ij} = \alpha_{ij} V.$$

Finally, the steady-state solidicity in the voxel is simply this volume divided by the voxel volume, $\phi_{ij} = V_{ij} / V_{\text{vox}}$.

The ensemble- and spatially-averaged velocity can be calculated in a similar way. First, a lagrangian interpolation function is constructed from the position vs. time data in small slices (usually 3 or 5 points). Then, the velocity at these points is calculated as the first derivative of this function evaluated at these points. For a given transit through a voxel, the velocity is time-averaged and this time-average, weighted with the residence time in the voxel is added to a running total. At the end of the analysis, this total is:

$$\mathbf{v}_{ij,\text{tot}} = \sum_k \tau_{ijk} v_{ijk}$$

where τ_{ijk} is the time spent in the k_{th} crossing through voxel i, j and \mathbf{v}_{ijk} is the time-averaged velocity for the same crossing. The steady-state velocity in the voxel is then:

$$\mathbf{v}_{ij} = \mathbf{v}_{ij,\text{tot}} / \sum_k \tau_{ijk}$$

It is these steady-state profiles for the solidicity and velocity that are presented alongside the simulation results in Chapter 5.

4.2 PEPT Experiments

The PEPT experiments used for comparison in this work were done at the University of Birmingham Positron Imaging Centre, using the Birmingham MC40 cyclotron to activate the tracers and an ADAC Forte PET scanner [74] was used to detect annihilation events. The camera contains two parallel plates of detectors with a variable separation distance.

Sets of experiments were done using the two mills described in Section 3.4. The experimental setup and operational parameters for each set of experiments is detailed in the next two subsections. Figure 4.4 shows the 30 cm mill between the parallel plates of the PEPT camera.




Figure 4.3: PEPT binning process. Left: The path of the tracer. Right: The residence time for each voxel. (Images courtesy of Angus Morrison)




Figure 4.4: The 30 cm mill between the parallel plates of the Birmingham PEPT camera

Particle Diameter (mm)	12.5% load (kg)	25% load (kg)	37.5% load (kg)	50% load (kg)
2	0.038	0.076	0.114	0.152
3	0.054	0.108	0.161	0.215
4	0.056	0.113	0.169	0.225
5	0.052	0.103	0.155	0.206
6	0.044	0.089	0.133	0.178
8	0.030	0.061	0.091	0.122
Total	0.275	0.549	0.824	1.098

Table 4.1: Mass distribution between particle size classes for the 12 cm mill PEPT experiments and DEM simulations

% v_{crit}	40	50	60	70	80	90	100	110
rad.s ⁻¹	5.116	6.395	7.674	8.953	10.23	11.51	12.79	14.069

Table 4.2: Rotational speeds for the 12 cm mill PEPT experiments

4.2.1 The 12 cm Mill

All of the experiments with the 12 cm mill were done with only glass beads present in the mill for a matrix of particle loads and rotational speeds. The rotational speeds were in the range 40% - 110% of the critical speed of the mill, in increments of 10%, as tabulated in Table 4.2. The range of particle loads was 12.5%, 25%, 37.5% and 50% of the available volume. Glass beads of diameter 2 mm, 3 mm, 4 mm, 5 mm, 6 mm and 8 mm were used to fill the mill. The particle size distribution was a truncated Rosin-Rammler distribution [75]. Table 4.1 shows the mass of each particle size in the overall load.

For each experiment, a single glass bead was placed in a 36 MeV ³He beam from the Birmingham cyclotron, which converted a small fraction of the naturally occurring ¹⁶O atoms to ¹⁸F, which is a positron emitter with a half life of 110 minutes. For each set of operating conditions, this process was carried out for each of the bead sizes, so that trajectories for each of the size classes could be determined. In the results presented in Subsection 5.2.2, the data for each of the size classes is averaged by mass to provide data for the overall charge.

4.2.2 The 30 cm Mill

Three sets of PEPT experiments were performed with the 30 cm mill. In the first set, the mill was filled only with glass beads and was kept dry. In the second set, there were glass beads and water, and in the third set there were glass beads and slurry. These experiments are detailed here.

Dry

The dry PEPT experiments of the 30 cm mill were done with two rotational speeds: 60% and 75% v_{crit} . The experiments were done with mono-sized distributions of first 5 mm then 3 mm glass particles at a filling of 31.25% of the available volume. In both cases this corresponded to a charge mass of 9.662 kg. The glass particles were activated using the same method as in the experiments with the 12 cm mill.

30 cm mill filled with water

For the experiments with water, a pump was connected to a hose that was inserted into the feed end of the mill. Before running the main experiments, some steps were taken to choose an appropriate fill-rate for the experiments. In the minerals processing industry, a desirable slurry volume is one for which the slurry will only just cover the solid charge when the mill is stationary. Thus, a fill rate for which the slurry volume is held constant at this value is desirable.

In order to determine this fill rate, the mill was run with the charge present while water was pumped into the mill. When the water volume stabilised the mill was stopped abruptly and the water level was measured. This process was repeated for different pump strengths until the measured water level was exactly in line with the top of the stationary charge. This pump setting was then used for the main experiments. The volumetric fill-rate for that pump setting was 50.35 ml/s.

Experiments were run at 60% and 75% v_{crit} with 5 mm mono-sized glass beads filling 31.25% of the available volume and the tracer creation and placement was identical to that in the dry runs.

30 cm mill filled with slurry

For the experiments with slurry, the same process was used to determine the correct fill-rate as was used in the water experiments. The fill-rates for the various conditions are summarised in Table 4.3.

The slurry used in the experiments was prepared by mixing water and finely ground particles from an El Soldado ore in the size range 53-75 μm . Slurries with a range of solids concentrations were prepared for the experiments. The viscous properties of these slurries are summarised in Table 4.4 (see Subsection 2.3.1 for definitions).

Solids concentration % / Speed (% v_{crit})	60	75
20	72.2	65.3
30	83.1	93.2
40	72.2	93.2

Table 4.3: Volumetric fill rates (in millilitres per second) for PEPT experiments with slurry

Solids concentration % (by mass)	Yield Stress τ_0 (Pa)	Yield Viscosity μ (Pa s)	Density ρ (kg m^{-3})
20	0.021	0.0013	1148
30	0.16	0.0014	1239
40	0.38	0.0022	1346
50	0.87	0.0035	1474
60	2.7	0.0065	1628

Table 4.4: Yield stress, yield viscosity and density for the different slurry compositions used in this work

Experiments were run at 60% and 75% v_{crit} with 5 mm mono-sized glass beads filling 31.25% of the available volume (excluding the slurry) for slurries with solids concentrations of 20%, 30% and 40% by mass. For all but the 60% v_{crit} and 40% slurry concentration condition, a 5 mm glass bead was tracked, after applying the same activation process as in the dry experiments. Additionally, and in a separate set of runs for all 6 conditions, a 1 mm blue stone was tracked.

The purpose of tracking the blue stone particle was to track a particle that would by some measure be representative of a particle in the slurry. At the time of the experiments, particles smaller than 1 mm in diameter either broke or were difficult to activate, so the blue stone particles were the smallest trackable particles. The extent to which these particles represent particles in the slurry will be examined in Section 5.5.

Since the blue stone particles were small enough to fit through the discharge grates, they were frequently discharged from the mill. When they were, they were simply placed back in the feed end of the mill. This happened many times over the course of a typical experiment.

Chapter 5

Results

5.1 Models and Discretisation

5.1.1 DEM

All of the discrete element modelling was done using EDEM 2.3, developed by DEM Solutions [18]. Particle-particle and particle-wall contacts were modeled using the Hertz-Mindlin contact model (described in Subsection 2.1.1). The time-step used for all simulations was 1×10^{-6} s.

5.1.2 CFD

All of the fluid modelling, with the exception of the fluid-particle drag force, was done using the built-in models in FLUENT 12.1, developed by Ansys Inc. [47]. The two-phase system of water or slurry and air was modeled with each as an incompressible fluid, with the multiphase dynamics modeled using an explicit Volume of Fluid method. The viscosity of the slurry was modeled using a set of parameters in a Herschel-Bulkley model that corresponded to a Bingham plastic model. These models are described in Section 2.3. Turbulence was not modelled for the slurry in this work due to incompatibility between FLUENT's turbulence and non-Newtonian models.

Regarding the computation itself, pressure-velocity coupling was done using SIMPLE (Semi-Implicit Method for Pressure Linked Equations) [76][77]. The momentum and volume-fraction equations were discretised using second-order upwinding and the time-stepping was done using a first-order implicit scheme.

Interaction	Coefficient of restitution	Coefficient of static friction	Coefficient of rolling friction
Glass-Glass	0.8	0.94	0.01
Glass-Aluminium	0.85	0.306	0.01
Glass-HDPE	0.703	0.287	0.01
Glass-Nylon	0.279	0.532	0.01
Glass-Perspex	0.5	0.424	0.01

Table 5.1: Contact coefficients for the interactions involved in the DEM simulations

5.2 Solids Only: The 12 cm Mill

5.2.1 Materials

For all of the solids-only simulations, the charge material consisted of glass beads. The density of the glass was $2700 \text{ kg}\cdot\text{m}^{-3}$, the shear modulus was 26.42 GPa and the Poisson's ratio was 0.23. The remaining DEM material parameters are the contact coefficients of the glass with the materials of the mill geometry. These are detailed in Table 5.1.

5.2.2 Results

DEM simulations were performed for the 12 cm mill filled with only glass beads for a matrix of filling conditions and rotational speeds identical to that in the equivalent PEPT experiments. These are detailed in Subsection 4.2.1 and in Tables 4.2 and 4.1.

At the beginning of each simulation, the particles were initialised evenly throughout a cylindrical volume of length 11.8 cm and radius 5 cm centred on the centre of the mill. The particles were initialised in order of largest to smallest to prevent space availability issues with randomly placing the particles. All of the particles were initialised with zero velocity.

For all of the following time-averaged results for the 12 cm mill, the first one second of data has been omitted to remove the influence of startup effects before the system reaches steady state. In addition to time-averaging, all of the results in this section are locally averaged into a 65×65 grid of voxels, to make direct, per-voxel comparisons possible.

First, the DEM simulations and PEPT experiments for 37.5% loading and 50% v_{crit} are presented in detail, as this set of parameters represents a balance between typical milling conditions and regimes accesible to granular flow theories. This will allow the results to be compared with these theories. The detailed results and comparisons are followed by summaries of the results for all 32 DEM simulations.

Solidicity comparisons for 37.5% loading and 50% v_{crit}

Figure 5.1 shows a comparison of the solidicity obtained from the DEM simulations and the PEPT experiments. The “percentage difference” plot shows the quantity $100\|\phi_{\text{DEM}} - \phi_{\text{PEPT}}\|/0.634$. This is the solidicity difference expressed as a percentage and normalised to the theoretical maximum difference, which is the solidicity for random close packing, $\phi = 0.634$ [78].

It is clear that the overall solidicity profile is in good agreement, but there are three main localised areas with discrepancies. The largest is roughly in the centre of the charge, where the PEPT experiment indicates extremely high solidicities (above random close packing, which suggests they aren’t physical) that aren’t present in the DEM simulation. The next largest area of difference is in the toe region of the charge, where the DEM simulation again underpredicts the solidicity, seen as a teal band in the percentage difference plot. Finally, there is a band of slight differences along the free surface where the DEM simulation over-predicts the solidicity, seen as a yellow band on the difference plot.

It is worth noting that the PEPT experiments for the 12 cm mill were each run for approximately 5 minutes, which is a fairly short time to built up good time-averaged statistics from a single particle. This may contribute to the lack of smoothness seen in the PEPT solidicity profile.

The overall profile of the per-voxel differences in solidicity between the DEM simulation and the PEPT measurement is summarised in Figure 5.2, which is a histogram of the normalised solidicity difference, $\|\phi_{\text{DEM}} - \phi_{\text{PEPT}}\|/0.634$, for all the voxels. In this diagram it can be seen that a large number of voxels have differences below 2% and that remainder of the voxels reach a maximum difference at around 43%.

Velocity comparisons for 37.5% loading and 50% v_{crit}

Figures 5.3 - 5.5 show comparisons of the x-velocity, y-velocity and velocity magnitude respectively. The data has been truncated to show only velocity data for areas with a solidicity above $\phi = 0.01$. This was done so that PEPT data with very little statistical support would not bias comparisons. The “percentage difference” plots, in the case of the velocity components, show the quantity $100\|v_{\text{DEM}} - v_{\text{PEPT}}\|/v_{\text{max}}$, where v_{max} is the maximum for the absolute value of that velocity component in the domain for the DEM simulation. In the case of the velocity magnitude, v_{max} is simply the largest velocity magnitude in the domain.

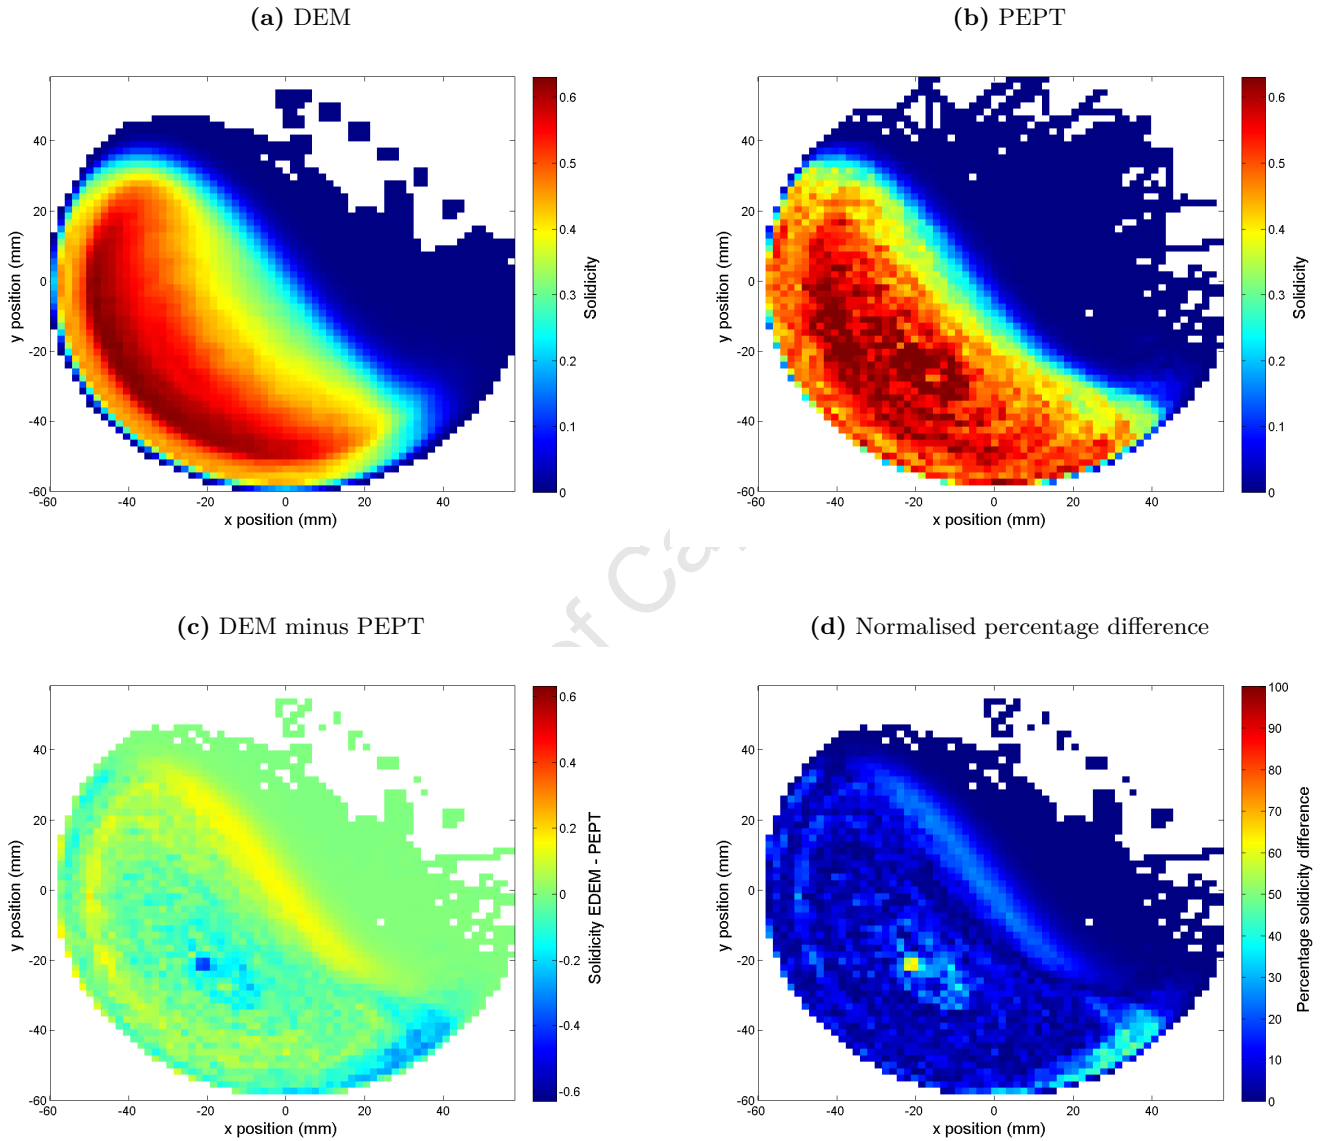


Figure 5.1: Comparison of DEM and PEPT solidicity in the 12 cm mill at 50% v_{crit} and 37.5% load. (a): DEM simulation results (b): PEPT experiment results (c): ϕ_{PEPT} subtracted from ϕ_{DEM} (d): $100\|\phi_{DEM} - \phi_{PEPT}\|/0.634$




Figure 5.2: Histogram of relative solidity differences between DEM and PEPT data in the 12 cm mill at 50% critical speed and 37.5% load

The agreement for the velocities is clearly very good. The only point at which the agreement breaks down is at the very edge of the free surface, where the values from the PEPT experiment become erratic due to very little sampling of the region outside of the free surface. A more quantitative comparison of the velocity field can be made by isolating a line along which to make the comparison.

Tangential velocity comparisons for 37.5% loading and 50% v_{crit}

In the following comparisons, similarly to that of Govender *et al.* [38] (See Subsection 2.2.2), a line was constructed joining the centre of the mill with the centre of circulation (CoC) of the charge. This line was truncated at the point where it intersected the free surface and extended to the point just before the beginning of the lifters (i.e. 8 mm from the mill shell). Figure 5.6 illustrates the positioning of the line in relation to the free surface and the CoC along with the tangential velocity vectors.

The velocity component perpendicular to this constructed line was calculated and will be referred to as the tangential velocity, v_{θ} . Positive values of v_{θ} are taken to be in the clockwise direction. A cubic polynomial of the form $v_{\theta}(r) = a + br + cr^2 + dr^3$ was fitted to the tangential velocity data, where r is the distance from the mill centre. Finally, the shear rate along the line, $\frac{dv_{\theta}}{dr} = b + 2cr + 3dr^2$, was calculated.

Figure 5.7 shows a comparison of the tangential velocities at 50% v_{crit} and 37.5% load. The profiles are in reasonably good agreement, with the DEM simulation under-predicting near the free surface and over-predicting slightly close to the shell.

Figure 5.8 shows a comparison of the shear rates at 50% v_{crit} and 37.5% load. The shear rates follow a similar parabolic profile to that measured by Govender *et al.* [38], with the minimum (in magnitude) shear rate occurring at the shell and a turning point in the shear rate profile occurring near the free surface. This similarity is notable because the shear rates are for mills of significantly different sizes.

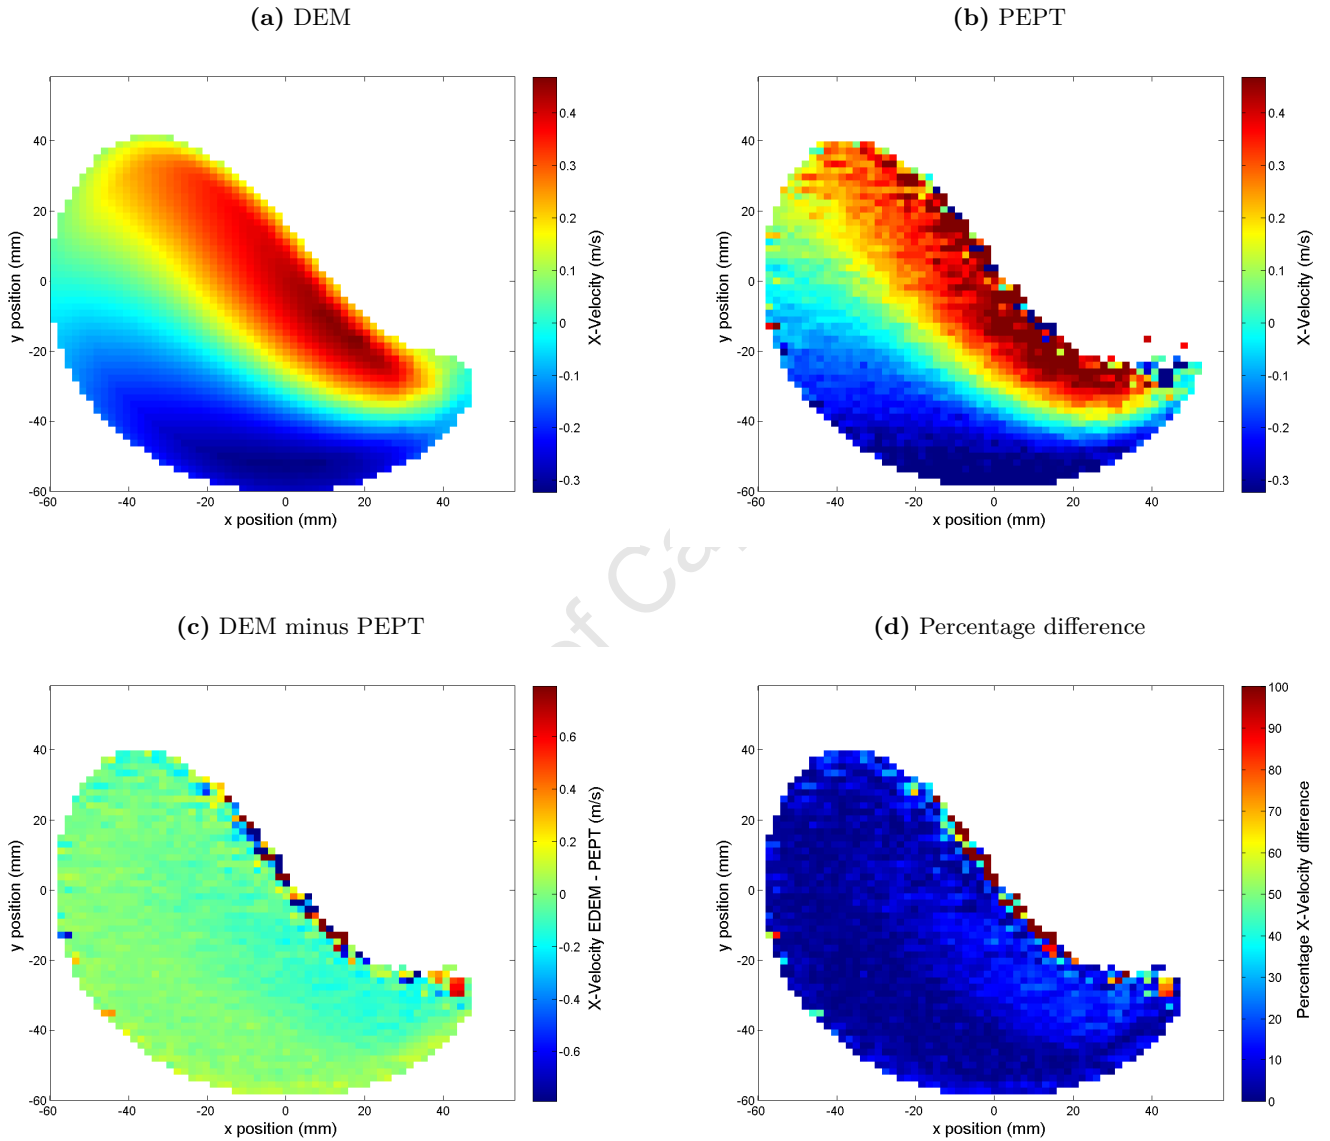


Figure 5.3: Comparison of DEM and PEPT x-velocity in the 12 cm mill at 50% v_{crit} and 37.5% load. (a): DEM simulation results (b): PEPT experiment results (c): v_{PEPT} subtracted from v_{DEM} (d): $100\|v_{DEM} - v_{PEPT}\|/v_{max}$




Figure 5.4: Comparison of DEM and PEPT y-velocity in the 12 cm mill at 50% v_{crit} and 37.5% load. (a): DEM simulation results (b): PEPT experiment results (c): v_{PEPT} subtracted from v_{DEM} (d): $100||v_{DEM} - v_{PEPT}||/v_{max}$

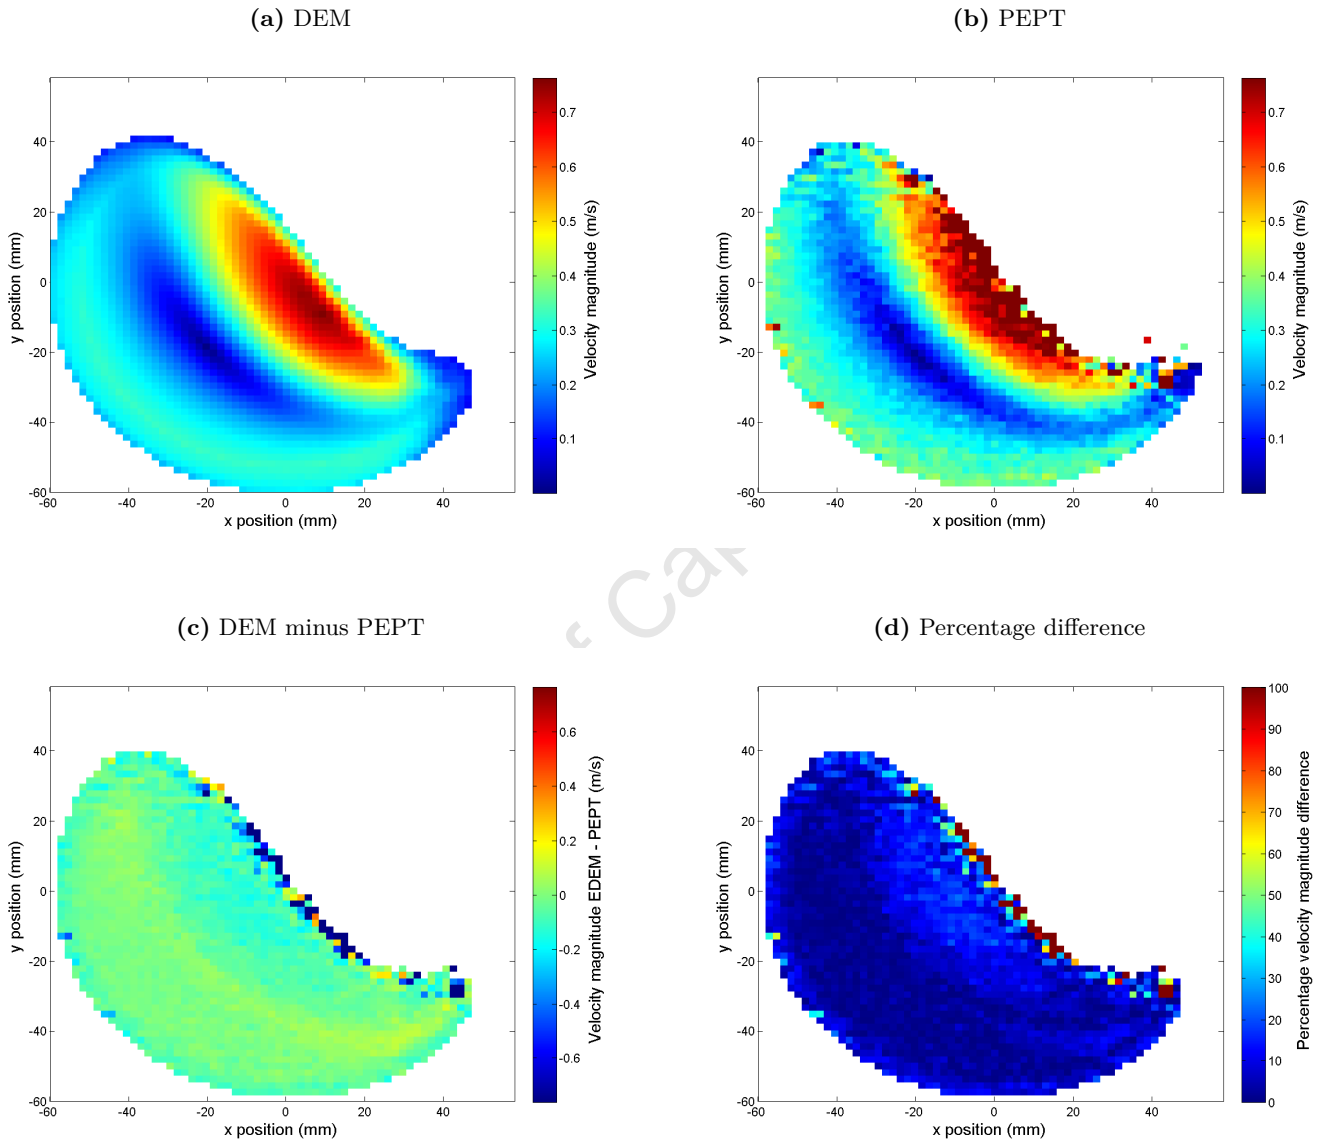


Figure 5.5: Comparison of DEM and PEPT velocity magnitude in the 12 cm mill at 50% v_{crit} and 37.5% load. (a): DEM simulation results (b): PEPT experiment results (c): v_{PEPT} subtracted from v_{DEM} (d): $100\|v_{DEM} - v_{PEPT}\|/v_{max}$




Figure 5.6: The analysis geometry for the tangential velocity and shear rate. The black curve is the free surface, the red + is the position of the CoC and the straight line is the line along which the tangential velocity was taken. The vector arrows represent (to no particular scale) the tangential velocity component.

Also worth noting is that the shear rate is definitively not constant. Consequently, the assumption of a constant shear rate, as used in leading granular flow modelling theories such as that by Taberlet [37] *et al.* (see Subsection 2.2.1) may not be warranted. However, a key difference in this system to the smooth drums of granular flow modelling is the presence of lifters, which may contribute to this shear rate profile. An additional consideration is that this system is polydisperse, while the system for Taberlet's modelling was monodisperse.




Figure 5.7: Comparison of tangential velocities between DEM and PEPT data in the 12 cm mill at 50% critical speed and 37.5% load




Figure 5.8: Comparison of shear rates between DEM and PEPT data in the 12 cm mill at 50% critical speed and 37.5% load




Figure 5.9: The free surfaces profiles for the PEPT data and DEM simulations of the 12cm with 37.5% loading and at 50% v_{crit}

Free surface comparison for 37.5% loading and 50% of the critical speed

For the DEM and PEPT solidicity profiles in Figure 5.1, the free surface profile was determined using the following process:

- A gaussian filter was applied to the solidicity profile to smooth it to within 3 voxels.
- Canny edge detection [79] was used to create a list of voxels in which there was a sharp change in the solidicity.
- The voxels that corresponded to the change from the interior to the exterior of the mill were removed.
- A generalised logistic function, $\frac{a}{1+be^{-cx}} + d$, was fitted to the points.

Figure 5.9 shows the free surface profiles for the DEM simulation and PEPT experiment.

Rotational Speed (% v_{crit})	12.5% Load	25% Load	37.5% Load	50% Load
40	0.0426	0.0578	0.0668	0.0701
50	0.0473	0.0596	0.0592	0.0672
60	0.0491	0.057	0.0678	0.0701
70	0.0563	0.0604	0.0676	0.0858
80	0.0571	0.0798	0.0591	0.104
90	0.0732	0.0919	0.101	0.067
100	0.0703	0.0855	0.11	0.0772
110	0.0498	0.0783	0.0826	0.098

Table 5.2: Mean normalised difference in the solidicity between DEM simulations and PEPT experiments

Summary of basic comparisons

The tables in this section summarise bulk comparisons of the 32 sets of DEM simulations and PEPT experiments. Table 5.2 shows the mean normalised solidicity difference, $\sum_{i=1}^N \|\phi_{\text{DEM}} - \phi_{\text{PEPT}}\| / 0.634N$, where N is the number voxels, for each simulation and experiment pair. The mean differences range between 4% and 11%, showing generally good agreement.

For the bulk velocity comparisons, the mean is weighted with the average of the solidicity between the simulation and the experiment, so that voxels with better statistics from the PEPT experiments are weighted higher. The mean is calculated as:

$$\bar{v} = \sum_{i=1}^N ((\phi_{\text{DEM},i} + \phi_{\text{PEPT},i})/2) |v_{\text{DEM},i} - v_{\text{PEPT},i}| / \sum_{i=1}^N ((\phi_{\text{DEM},i} + \phi_{\text{PEPT},i})/2)$$

The velocity mean is then normalised to the average of the maximum velocity in the DEM simulation and the maximum velocity in the PEPT experiment:

$$\bar{v}_{\text{norm}} = \bar{v} / ((v_{\text{max,DEM}} + v_{\text{max,PEPT}})/2)$$

The normalised mean velocity differences are shown in Tables 5.3 - 5.5.

Tangential velocity summary

Bulk comparisons of the tangential velocity won't add any new information, since the comparisons for the x and y components of velocity are already known. Here we summarise the appropriateness of the cubic fit to the tangential velocity profile. The Root Mean Squared Deviation (RMSD) of the nonlinear fits to the tangential velocities in the DEM simulations are summarised in Table 5.6.

Rotational Speed (% v_{crit})	12.5% Load	25% Load	37.5% Load	50% Load
40	0.0102	0.00857	0.0199	0.0431
50	0.00691	0.00673	0.00941	0.0383
60	0.011	0.012	0.00731	0.0217
70	0.00728	0.0124	0.0134	0.0182
80	0.00781	0.0157	0.0226	0.0437
90	0.0262	0.0238	0.0344	0.0523
100	0.0296	0.0528	0.0388	0.0435
110	0.0559	0.0423	0.0561	0.0652

Table 5.3: Mean normalised difference in the x velocity between DEM simulations and PEPT experiments

Rotational Speed (% v_{crit})	12.5% Load	25% Load	37.5% Load	50% Load
40	0.0116	0.00995	0.0248	0.0467
50	0.0186	0.0151	0.0137	0.0869
60	0.0159	0.0186	0.0316	0.0344
70	0.0258	0.0345	0.0169	0.0273
80	0.0223	0.0407	0.0247	0.047
90	0.061	0.0676	0.0569	0.043
100	0.0299	0.0602	0.0696	0.0316
110	0.069	0.0394	0.0861	0.0652

Table 5.4: Mean normalised difference in the y velocity between DEM simulations and PEPT experiments

Rotational Speed (% v_{crit})	12.5% Load	25% Load	37.5% Load	50% Load
40	0.0149	0.0132	0.032	0.0594
50	0.0224	0.0176	0.0168	0.128
60	0.0184	0.0211	0.0352	0.0421
70	0.0292	0.0357	0.019	0.0295
80	0.0258	0.0466	0.0281	0.0585
90	0.0741	0.0844	0.0682	0.0516
100	0.0411	0.0783	0.0846	0.0395
110	0.0987	0.0573	0.107	0.0972

Table 5.5: Mean normalised difference in the velocity magnitude between DEM simulations and PEPT experiments

Rotational Speed (% v_{crit})	12.5% Load	25% Load	37.5% Load	50% Load
40	0.0218	0.0451	0.329	0.283
50	0.0147	0.0269	0.0293	0.398
60	0.0327	0.0241	0.0356	0.163
70	0.0393	0.0323	0.0331	0.0437
80	0.0966	0.053	0.052	0.0605
90	0.147	0.14	0.0449	0.0553
100	0.404	0.17	0.558	0.0643
110	0.18	0.202	0.0648	0.168

Table 5.6: RMSD in $\text{m}\cdot\text{s}^{-1}$ of the cubic polynomial fits to the DEM tangential velocity data

The fit is clearly not appropriate for very high speeds and low loads as seen by RMSD values above $0.1 \text{ m}\cdot\text{s}^{-1}$ for the 6 simulations with speeds 90%, 100% and 110% v_{crit} and loads 12.5% and 25%. It is also not appropriate for very low speeds and high loads, again as seen by RMSD values above $0.1 \text{ m}\cdot\text{s}^{-1}$ for 40%, 50% and 60% v_{crit} with 50% load and 40% v_{crit} with 37.5% load.

Free surface summary

For the free-surface acquisition, only the free surfaces for low loads (excluding 12.5% load) and low speeds (40% - 60% v_{crit}) could be acquired consistently. For all other operating conditions, the free-surface that is extracted is highly sensitive to the method of edge detection used, including any pre-filtering or blurring.

Figure 5.10 shows the free surface extracted from DEM simulations of the aforementioned conditions along with the corresponding position of the centre of circulation.

The validity of the generalised logistic fit to the free surface can be seen in Table 5.7, which summarises the RMSD values for the fits. The values are surprisingly high, indicating that the generalised logistic may not be the best possible function to fit the free surface.

The differences in free surface position between DEM simulations and PEPT are compared by looking at the mean difference in predicted free surface height. A large number of points was evenly taken in x , and the absolute difference in height between the DEM free surface and the PEPT free surface at each of these points in x was recorded. The mean and standard deviation of this set of height differences is summarised in Table 5.7.




Figure 5.10: The free surface and CoC positions for a selection of DEM simulations of the 12 cm mill. The black features are for 25% loading, red for 37.5% loading and blue for 50% loading. The circles and solid lines are the CoC and free surface respectively for 40% speed. The plus signs and dotted lines are for 50% speed and the crosses and dashed lines are for 60% speed

Rotational Speed (% v_{crit})	25% Load	37.5% Load	50% Load
40	12.1	7.95	9.1
50	14.6	8.69	9.33
60	18.9	14.7	10.5

Table 5.7: RMSD in mm of the generalised logistic fits to the DEM free surface data

Rotational Speed (% v_{crit})	25% Load	37.5% Load	50% Load
40	2.87± 2.44	3.13± 2.57	2.89± 4.36
50	5.49± 3.65	6.58± 3.78	3.05± 3.61
60	8.59± 6.44	6.3± 4.47	5.14± 3.5

Table 5.8: Mean and standard deviation of absolute differences (in mm) in free surface height for DEM simulations and PEPT experiments with the 12 cm mill

5.3 The Effect of the Fluid on the Solids: The 30 cm Mill

The previous section has established that DEM is an effective technique for simulating the interactions of rigid particles in mills. In this section, the assumption that a dry DEM simulation will provide a good approximation to the motion of particles surrounded by a fluid is examined. This assumption is invoked in the one-way coupled simulations (Section 5.5), so it is of particular interest to run simulations with the same fluid conditions as those that are used in those simulations.

Thus the simulations were done with two rotational speeds, 60% and 75% of the critical speed, with 5 mm glass particles filling of 31.25% of the available volume. As with the 12 cm simulations, the particles were initialised evenly with zero velocity throughout a cylindrical volume centred on the mill centre. In this case the length of the cylinder was 25 cm and its radius was 12.5 cm. For all of the following time-averaged results for the 30 cm mill, the first one second of data has been omitted to remove the influence of startup effects on the average.

The solids motion in the dry simulations is compared with the solids motion in the equivalent wet PEPT experiments, which include slurry solids concentrations of 20%, 30% and 40%. The solidicities are shown in Figure 5.11, while the x- and y-velocity profiles are shown in Figures 5.12 and 5.13 respectively. The PEPT solidicity profiles, though streaky, seem to show that the charge is dilated somewhat. This effect isn't particularly strong though, and the good matching of the velocity profiles suggests that the presence of the fluid has a minimal effect on the motion of the solids.

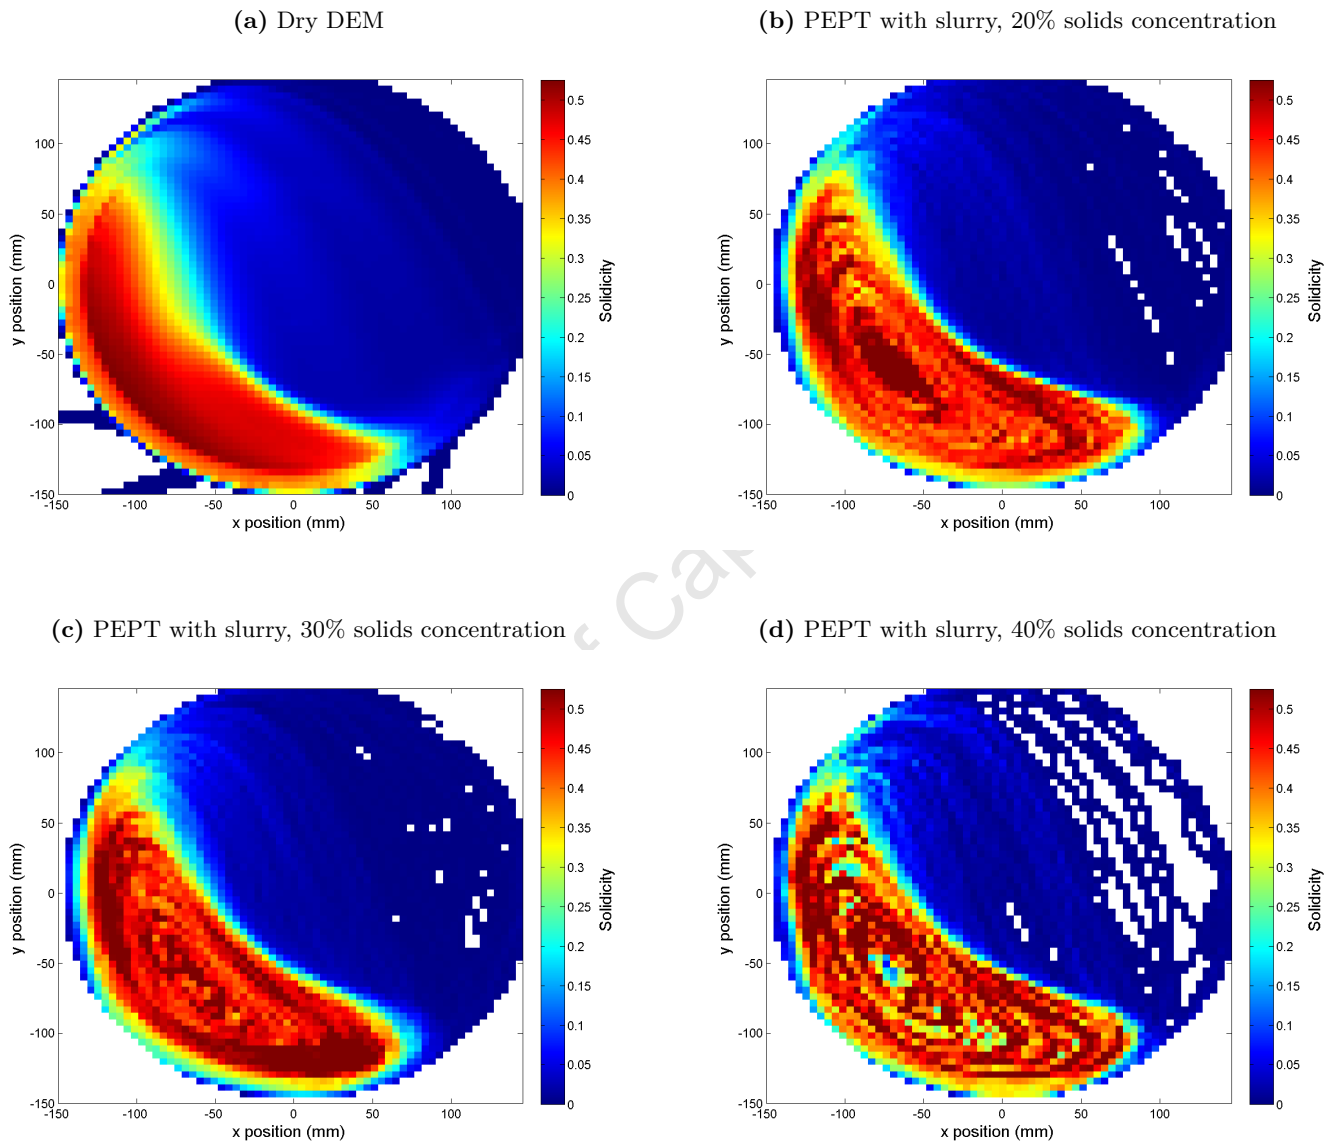


Figure 5.11: Comparison of dry DEM and PEPT with slurry solidicity profiles in the 30 cm mill at 75% v_{crit} and 31.25% load. (a): Dry DEM results (b): PEPT with slurry, 20% solids concentration (c): PEPT with slurry, 30% solids concentration (d): PEPT with slurry, 40% solids concentration

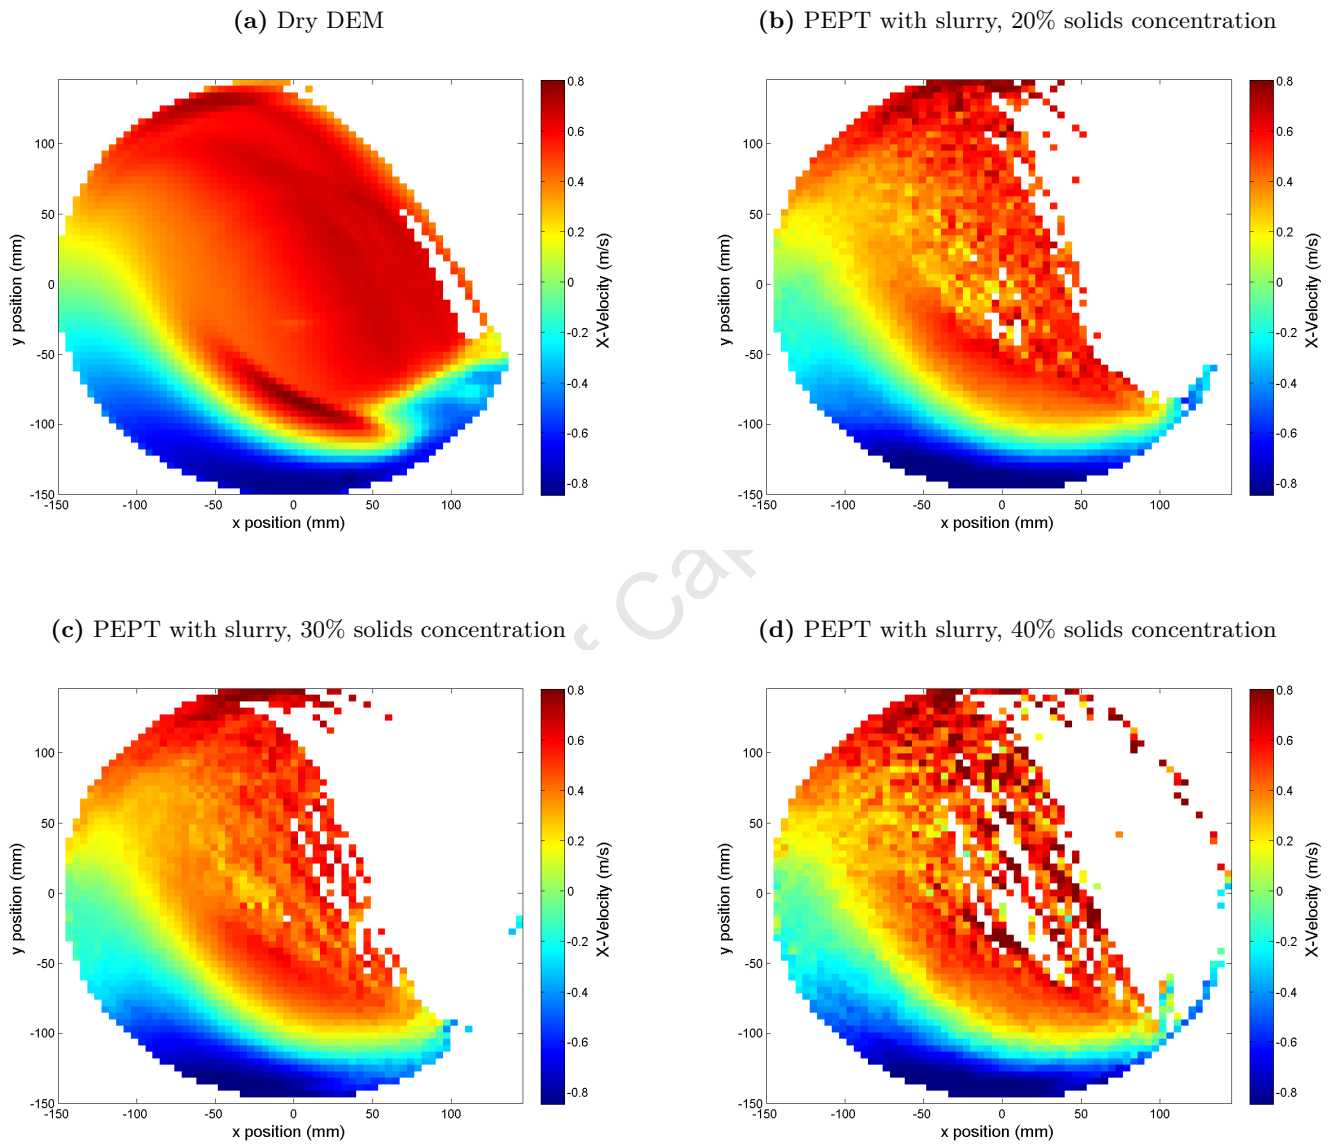


Figure 5.12: Comparison of dry DEM and PEPT with slurry x-velocity profiles in the 30 cm mill at 75% v_{crit} and 31.25% load. (a): Dry DEM results (b): PEPT with slurry, 20% solids concentration (c): PEPT with slurry, 30% solids concentration (d): PEPT with slurry, 40% solids concentration

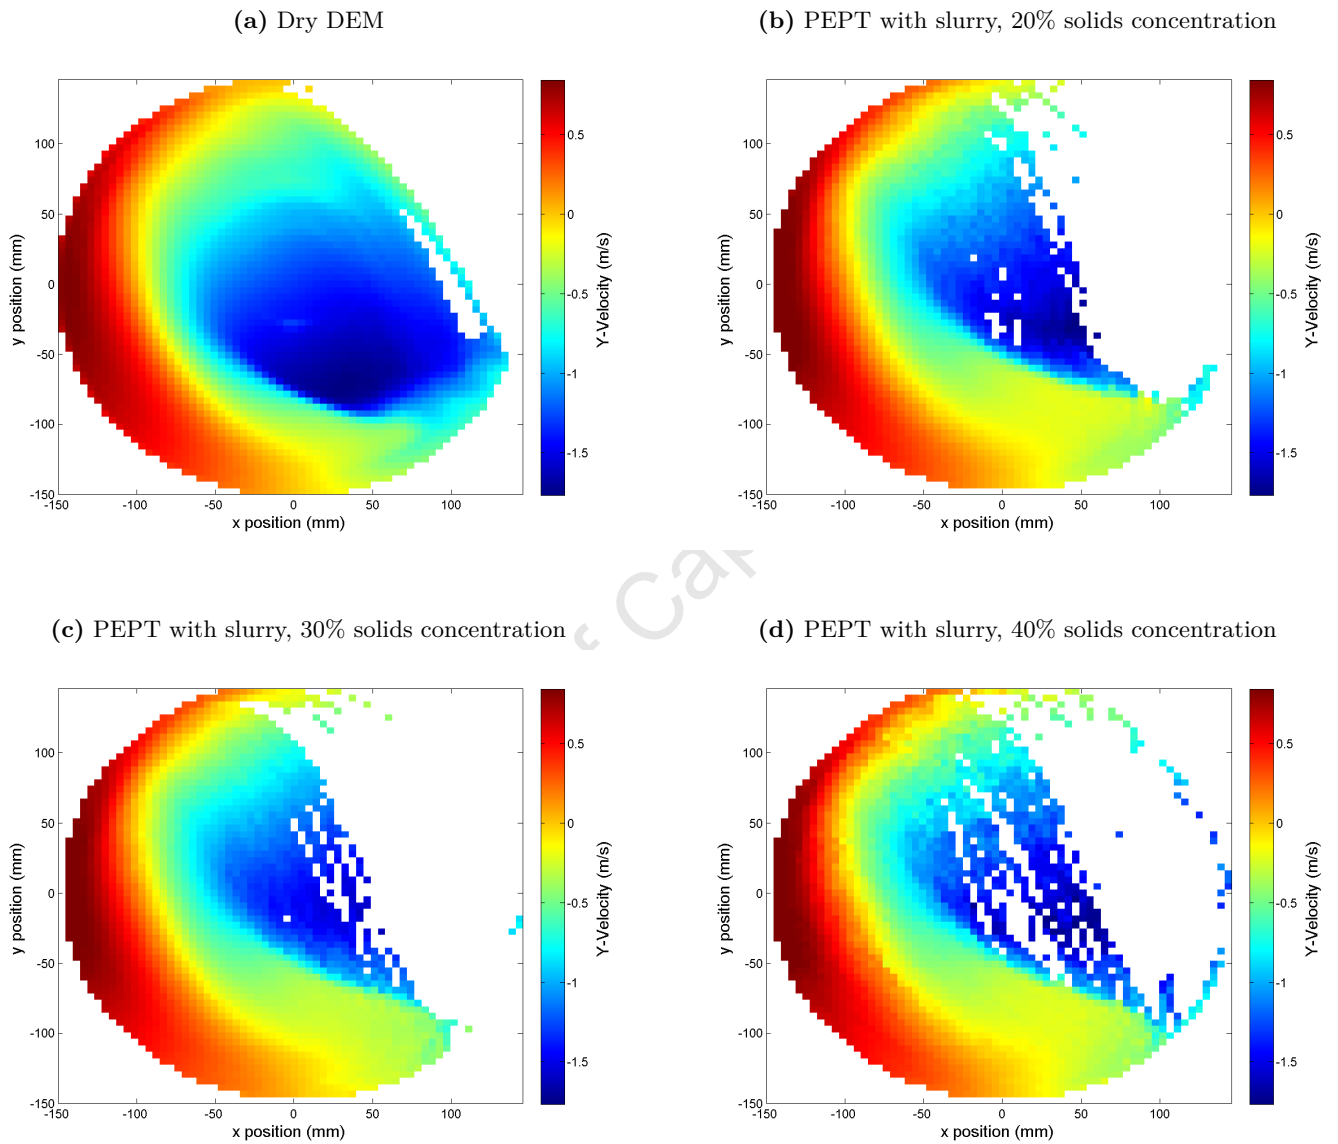


Figure 5.13: Comparison of dry DEM and PEPT with slurry y-velocity profiles in the 30 cm mill at 75% v_{crit} and 31.25% load. (a): Dry DEM results (b): PEPT with slurry, 20% solids concentration (c): PEPT with slurry, 30% solids concentration (d): PEPT with slurry, 40% solids concentration

5.4 Fluid Only

A set of simulations with fluid only were done as initial tests and to create a reference point for the coupled simulations. Since no PEPT experiments were done with fluid only, the operating conditions, such as the fill rate and rotational speeds, were generally matched to the equivalent PEPT experiments with glass beads present.

5.4.1 Simulation Conditions

Materials

The viscous properties of the slurry were matched to those for the PEPT experiments (see Table 4.4). The surface tension of the slurry was taken to be the same as that of water, which is 0.07 N/m. This is in agreement with surface tensions for a variety of mineral slurries measured by Ozcan *et al.*, which were in the range 0.0669-0.0747 N/m [80].

Operating Conditions

For all of the fluid-only runs, the mill was filled to 4.67 cm below its centre with slurry, to match the filling in the PEPT experiments (see Subsection 4.2.2). The mill was then rotated clockwise at either 60% or 75% v_{crit} . In the case of 60% v_{crit} , the simulation was run for at least 6.25 s and in the case of 75% v_{crit} , it was run for at least 5 s. For both cases, this corresponds to about 5 mill revolutions. The results data presented here are from time-averages accumulated beginning 2 s into the simulation to eliminate startup effects. For each of the two speeds, three slurry concentrations were used: 20% solids concentration by mass, 30% concentration and 40% concentration for a total of 6 simulations.

The flow into the mill was matched to the volumetric fill rates measured in the equivalent PEPT experiments that included solids. The fill-rates are summarised in Table 4.3. For each fill rate, the boundary velocity of the fluid entering the mill was simply calculated as the volumetric fill rate divided by the inlet area (see Figure 3.19), which was 128.7 cm², giving inflow speeds of the order of 5 mm s⁻¹.

5.4.2 Basic Results

Transverse volume fraction and velocities

Figure 5.14 shows the axially-averaged and time-averaged flow quantities for slurry only with 30% solids concentration by mass and no coupling in the 30 cm mill rotating at 75% v_{crit} . In the transverse direction, the bulk of the slurry remains almost stationary in the mill, with some slurry

being dragged around the shell in a thin layer by the lifters. In a similar manner, small amounts of air are dragged under the slurry layer by the lifters on the right-hand side.

In the z direction, the majority of the flow towards the discharge end of the mill (the flow in the positive z direction) occurs in a band in the left half of the slurry, while flow back to the feed end of the mill occurs in a pocket on the far right of the mill. These features of the slurry position and velocity profile hold over all of the 6 simulations.

Axial centre of mass

In these fluid-only situation, the idea in the Moys model (See Section 2.5) of there being a flat slurry profile in the transverse direction is supported. However, the notion of there being a purely linear profile in the axial direction is questionable. Figure 5.15 shows an overlay of all the axial centre of mass (CoM) profiles over time for a simulation of the 30 cm mill at 60% v_{crit} and 20% slurry solids concentration. At the far left of the figure is the highest point, which is at the inlet. A wave-like structure with an amplitude of approximately 1.5 cm follows, ending in the lowest point at the discharge end of the mill.

The wave-like structure is revealed more directly in a time-averaged plot of the axial CoM profile, shown in Figure 5.16, with Figure 5.17 showing the same plot in the context of the mill's dimensions. A consistent speed-dependent effect on this structure is illustrated in Figure 5.18, which shows the time-averaged axial CoM profile of the slurry for all of the simulations. In each of the plots, the leftmost three points are discarded to remove the data from the slurry initially trickling into the mill. The rightmost point is also removed, as it is artificially raised due to slurry being caught in the slurry lifters (the protrusions from the mill shell seen in Figure 3.15). To aid comparisons, the profiles are fitted to a sum of two Gaussians with a constant shift:

$$y(z) = a_0 e^{-(z-b_0)^2/2c_0^2} + a_1 e^{-(z-b_1)^2/2c_1^2} + d;$$

In all cases, there is a double-humped structure to the profile (which motivated the above fit), which is consistently broader at 75% v_{crit} than at 60% v_{crit} . This consistent difference is tabulated in Table 5.9, where it can be seen that the distance between the peaks is approximately double at 75% v_{crit} compared to 60% v_{crit} .




Figure 5.14: Axially-averaged and time-averaged flow results for slurry only with 30% solids concentration by mass and no coupling in the 30 cm mill rotating at 75% v_{crit} . (a): Slurry volume fraction. (b): Slurry x velocity. (c): Slurry y velocity. (d): Slurry z velocity.

Solids concentration % / Speed (% v_{crit})	60	75
20	95.9	173
30	96.4	179
40	97.0	174

Table 5.9: Distance between the peaks (in mm) in the axial CoM profiles for fluid-only simulations




Figure 5.15: An overlay plot of the axial centre of mass profile over time for a fluid-only simulation of the 30 cm mill at 60% v_{crit} with a slurry solids fraction of 20% by mass




Figure 5.16: A plot of the time-averaged axial centre of mass profile for a fluid-only simulation of the 30 cm mill at 60% v_{crit} with a slurry solids fraction of 20% by mass






Figure 5.17: A plot of the time-averaged axial centre of mass profile along with the mill bounding box for a fluid-only simulation of the 30 cm mill at 60% v_{crit} with a slurry solids fraction of 20% by mass


(a) 60% v_{crit} , 20% solids concentration




(b) 75% v_{crit} , 20% solids concentration




(c) 60% v_{crit} , 30% solids concentration



(d) 75% v_{crit} , 30% solids concentration



(e) 60% v_{crit} , 40% solids concentration



(f) 75% v_{crit} , 40% solids concentration




Figure 5.18: The axial profile of the height of the centre of mass for each of the fluid-only simulations.

5.5 One-Way Coupling

Coupling

The material properties for the one-way coupled simulations were identical for those of the fluid-only simulations. Since experiments with slurry or water were only run in the 30 cm mill with 5 mm glass beads, only two different DEM simulations were used for the coupling process, one with a mill speed of 60% v_{crit} and one at 75% v_{crit} . For 60% v_{crit} , 6.25 s of DEM data was used, and for 75% v_{crit} , 5 s of data was used, corresponding to approximately 5 mill rotations.

Data from these simulations was spatially averaged into a $33 \times 33 \times 33$ grid of dimension $300 \times 300 \times 338$ mm using the algorithm described in Subsection 3.2.2. Additionally, the data was time-averaged into 0.05 s intervals. Since the DEM data was recorded in intervals of 0.002 s, this corresponds to 25 data points per time average.

The fluid simulations were then run independently, updating the particle-fluid forces from the pre-averaged DEM data using the process described in Section 3.1.

Simulations

The one-way coupled simulations that were run can be divided into 2 sets. The first set contains simulations with slurry for conditions that match those in the PEPT experiments. The effect of the different drag correlations on the results is analysed for this set, and comparisons are made between PEPT experiments and the results for Beetstra drag. The second set is a repeat of the Beetstra drag simulations from the first set, but with the input from the DEM simulations time-averaged over the entire simulation, in a manner similar to that employed by Cleary (see Subsection 2.4.2).

5.5.1 Coupling with Slurry

In this section, simulations with both the Beetstra and Ergun, Wen&Yu drag models are examined. Detailed comparisons are done with data from simulations at 75% v_{crit} and 31.25% load only, since the observations from the comparisons are similar for the simulations at 60% v_{crit} .

Transverse volume fraction

Figure 5.19 shows the volume fraction (slurry vs. air) profiles for Beetstra and Ergun, Wen & Yu drag at 75% v_{crit} and 31.25% load. The profiles for each drag model are remarkably similar, with the Ergun drag predicting a slurry profile that is slightly more raised in the mill. The most likely

explanation for this is that the Ergun correlation over-predicts the drag force at high Reynolds numbers, as the correlation was generated with low Reynolds number data. This is supported by the observation that when the slurry solids concentration is increased from 20% to 40%, increasing the viscosity and reducing the Reynolds numbers, the differences are decreased significantly. The Reynolds numbers for 20% concentration slurry are mostly under 150, while for the 40% concentration slurry they are mostly under 90. It is at Reynolds numbers of around 90 that the two drag correlations diverge (see Figure 2.8 in Subsection 2.4.1), so this seems like the most likely explanation.

For both drag correlations, a higher solids concentration in the slurry led to a more raised profile. These small differences are summarised in Figure 5.20.

While the simulations with only fluid supported the Moys picture (Subsection 2.5.1) of there being a flat slurry profile in the transverse plane, these coupled simulations clearly have profiles that are not flat. Rather, the profiles match the assertion by Morrell and Stephenson (Subsection 2.5.2) that the free surface tends to match the shape of the charge. In this respect they also agree qualitatively with the profiles computed by Cleary *et al.* [13] (Subsection 2.4.2).

Axial velocity

A larger difference between the drag correlations is evident in plots of the z-velocity, which are shown in Figure 5.21. The positive z-direction is towards the discharge end of the mill. Both correlations agree with the observation by Morrell and Stephenson that the majority of the fluid transport occurs at the bottom of the mill.

In this respect, the simulations also agree with Condori *et al.* [67] (Subsection 2.5.3). However, the suggestion by Condori *et al.* that the fluid is stationary in the upper part of the charge (what they called the “suspension zone”) is contradicted. The simulations have the fluid stationary in the middle of the charge, but at the very top they show the fluid flowing back towards the inflow end of the mill.

Pressure drop

Figure 5.22 shows the pressure drop profile for the two drag correlations. Clearly the pressure drop is significantly higher for the Ergun, Wen & Yu drag, which stems from again from the fact that for the high solidities and high Reynolds numbers in these simulations, the Ergun correlation predicts a significantly higher drag force than the Beetstra correlation.

Discharge rates

Table 5.10 shows the discharge rates for the Beetstra simulations. These simulations can be compared with the experimental hold-up data gathered by Morrell and Stephenson, but it must be noted that the simulations could not be run for long enough to reach a steady fluid volume in the mill. Thus, the net rate of volume change in the mill will be used as a proxy for the expected eventual hold-up. That is, a negative rate of net volume change suggests an eventual hold-up that is smaller and a positive rate of net volume change suggests an eventual hold-up that is larger. However, due to the very short length of the simulations, it must be noted that this inference is tenuous.

With this in mind, the simulation results seem to agree with the experimental work by Morrell and Stephenson that show that the hold-up increases with higher inflow rates. Additionally, the results suggest that the hold-up increases with increased slurry viscosity. This is to be expected, since the drag force increases with viscosity and one would expect this to decrease the axial flow rate through the mill.

Axial centre of mass

With the presence of solids completely changing the transverse slurry profile of the Moys picture, it is interesting to see, as shown in Figure 5.23, that the axial CoM profile for both drag correlations maintains the wave-like structure seen in the fluid-only simulations. Unlike the fluid-only simulations however, the structure appeared to have three peaks for both 20% solids concentration simulations as well as the Beetstra 30% solids concentration simulation. The structure for the remaining three simulations seemed to have two peaks, but not as definitively as for the fluid-only simulations.

Solids concentration %	Inflow rate (ml/s)	Beetstra outflow rate (ml/s)	Net rate of volume change (ml/s)
20	65.3	113	-47.7
30	93.2	60.2	33.0
40	93.2	42.0	51.2

Table 5.10: Discharge rates at 75% v_{crit} for Beetstra drag

For all of the plots, the points are fitted to a sum of three Gaussians with a constant shift:

$$y(z) = a_0 e^{-(z-b_0)^2/2c_0^2} + a_1 e^{-(z-b_1)^2/2c_1^2} + a_2 e^{-(z-b_2)^2/2c_2^2} + d;$$

5.5.2 PEPT Comparison


In this section the simulation results for Beetstra drag are compared to the equivalent PEPT experiments with the 30 cm mill. As described in Subsection 4.2.2, a 1 mm blue stone was tracked with the assumption that it would be somewhat representative of a parcel of fluid, due to being much smaller than the rest of the charge. The upper and lower cusps of the charge need to be distinguished in this section, and are referred to by the standard terms in the milling literature, which are the “shoulder” and “toe” regions respectively.

Transverse volume fraction


Figure 5.24 shows the transverse volume fraction profiles for the PEPT experiments. They are qualitatively similar to the results with Beetstra drag in the sense that the profile follows the shape of the charge, but the centre of mass of the profile is much lower. The simplest explanation for this is that in the computational case, near the shoulder of the charge, it takes longer for the fluid trapped between the lifters to flow out than it would for a small particle. In this particular region, the particle is not representative of the fluid.

An additional feature that is not present in the simulation results is a low volume fraction region in the centre of circulation of the charge. This is a phenomenon also noted in SPH simulations by Cleary *et al.* [13] (Figure 2.10). The absence of this phenomenon in these simulations is most likely a consequence of the way the fluid is initialised, which was as a continuous bulk in the bottom half of the mill. With this kind of initialisation it would take a long time for a dry patch to develop (if it develops at all) in the centre of circulation.


(a) Ergun, Wen & Yu drag: 20% solids concentration




(b) Beetstra drag: 20% solids concentration




(c) Ergun, Wen & Yu drag: 30% solids concentration



(d) Beetstra drag: 30% solids concentration



(e) Ergun, Wen & Yu drag: 40% solids concentration



(f) Beetstra drag: 40% solids concentration




Figure 5.19: A comparison of the volume fraction for one-way coupled simulations at $75\% v_{crit}$ and 31.25% load. (a): Using Ergun, Wen & Yu drag and slurry with 20% solids concentration by mass. (b): Beetstra drag with 20% solids concentration. (c): Ergun, Wen & Yu drag with 30% solids concentration. (d): Beetstra drag with 30% solids concentration. (e): Ergun, Wen & Yu drag with 40% solids concentration. (f): Beetstra drag with 40% solids concentration.






Figure 5.20: The y position of the centre of mass of the slurry at 75% v_{crit} for both Beetstra and Ergun, Wen & Yu drag correlations for slurry solids concentrations of 20%, 30% and 40%


(a) Ergun, Wen & Yu drag: 20% solids concentration




(b) Beetstra drag: 20% solids concentration




(c) Ergun, Wen & Yu drag: 30% solids concentration



(d) Beetstra drag: 30% solids concentration



(e) Ergun, Wen & Yu drag: 50% solids concentration



(f) Beetstra drag: 40% solids concentration






Figure 5.21: A comparison of the z velocity for one-way coupled simulations at 75% v_{crit} and 31.25% load. (a): Using Ergun, Wen & Yu drag and slurry with 20% solids concentration by mass. (b): Beetstra drag with 20% solids concentration. (c): Ergun, Wen & Yu drag with 30% solids concentration. (d): Beetstra drag with 30% solids concentration. (e): Ergun, Wen & Yu drag with 40% solids concentration. (f): Beetstra drag with 40% solids concentration. Voxels with a slurry volume fraction below 0.5 are excluded.


(a) Ergun, Wen & Yu drag: 20% solids concentration




(b) Beetstra drag: 20% solids concentration




(c) Ergun, Wen & Yu drag: 30% solids concentration



(d) Beetstra drag: 30% solids concentration



(e) Ergun, Wen & Yu drag: 40% solids concentration



(f) Beetstra drag: 40% solids concentration






Figure 5.22: A comparison of the axial pressure drop for one-way coupled simulations at 75% v_{crit} and 31.25% load. (a): Using Ergun, Wen & Yu drag and slurry with 20% solids concentration by mass. (b): Beetstra drag with 20% solids concentration. (c): Ergun, Wen & Yu drag with 30% solids concentration. (d): Beetstra drag with 30% solids concentration. (e): Ergun, Wen & Yu drag with 40% solids concentration. (f): Beetstra drag with 40% solids concentration.


(a) Ergun, Wen & Yu drag: 20% solids concentration




(b) Beetstra drag: 20% solids concentration




(c) Ergun, Wen & Yu drag: 30% solids concentration



(d) Beetstra drag: 30% solids concentration



(e) Ergun, Wen & Yu drag: 40% solids concentration



(f) Beetstra drag: 40% solids concentration




Figure 5.23: A comparison of the axial profiles of the height of the centre of mass for one-way coupled simulations at 75% v_{crit} and 31.25% load. (a): Using Ergun, Wen & Yu drag and slurry with 20% solids concentration by mass. (b): Beetstra drag with 20% solids concentration. (c): Ergun, Wen & Yu drag with 30% solids concentration. (d): Beetstra drag with 30% solids concentration. (e): Ergun, Wen & Yu drag with 40% solids concentration. (f): Beetstra drag with 40% solids concentration.

Axial velocity

Figure 5.25 shows plots of the axially-averaged z-velocity profiles for the PEPT experiments. As with the simulations, there is clearly a zone in which there is transport towards the discharge end of the mill and one in which there is transport back towards the feed end.

The location of the flow towards the feed end is the same, around the shoulder of the charge. However, the flow towards the discharge end is localised near the surface of the charge, whereas in the simulations this zone was the whole toe region.

The reason for this discrepancy was noted during the running of the experiments. When the blue stone was placed in the feed end (described in Subsection 4.2.2), it would often not penetrate the charge at all and simply bounce along the surface of the charge until it reached the discharge end. In that particular regime of motion it was manifestly not representative of a fluid parcel and this difference in the zone locations is most likely a result of that.

University of Cape Town




Figure 5.24: The volume fraction profiles for the PEPT experiments with slurry. (a): 60% v_{crit} , 20% slurry solids concentration by mass. (b): 75% v_{crit} , 20% slurry solids concentration by mass. (c): 60% v_{crit} , 30% slurry solids concentration by mass. (d): 75% v_{crit} , 30% slurry solids concentration by mass. (e): 60% v_{crit} , 40% slurry solids concentration by mass. (f): 75% v_{crit} , 40% slurry solids concentration by mass.

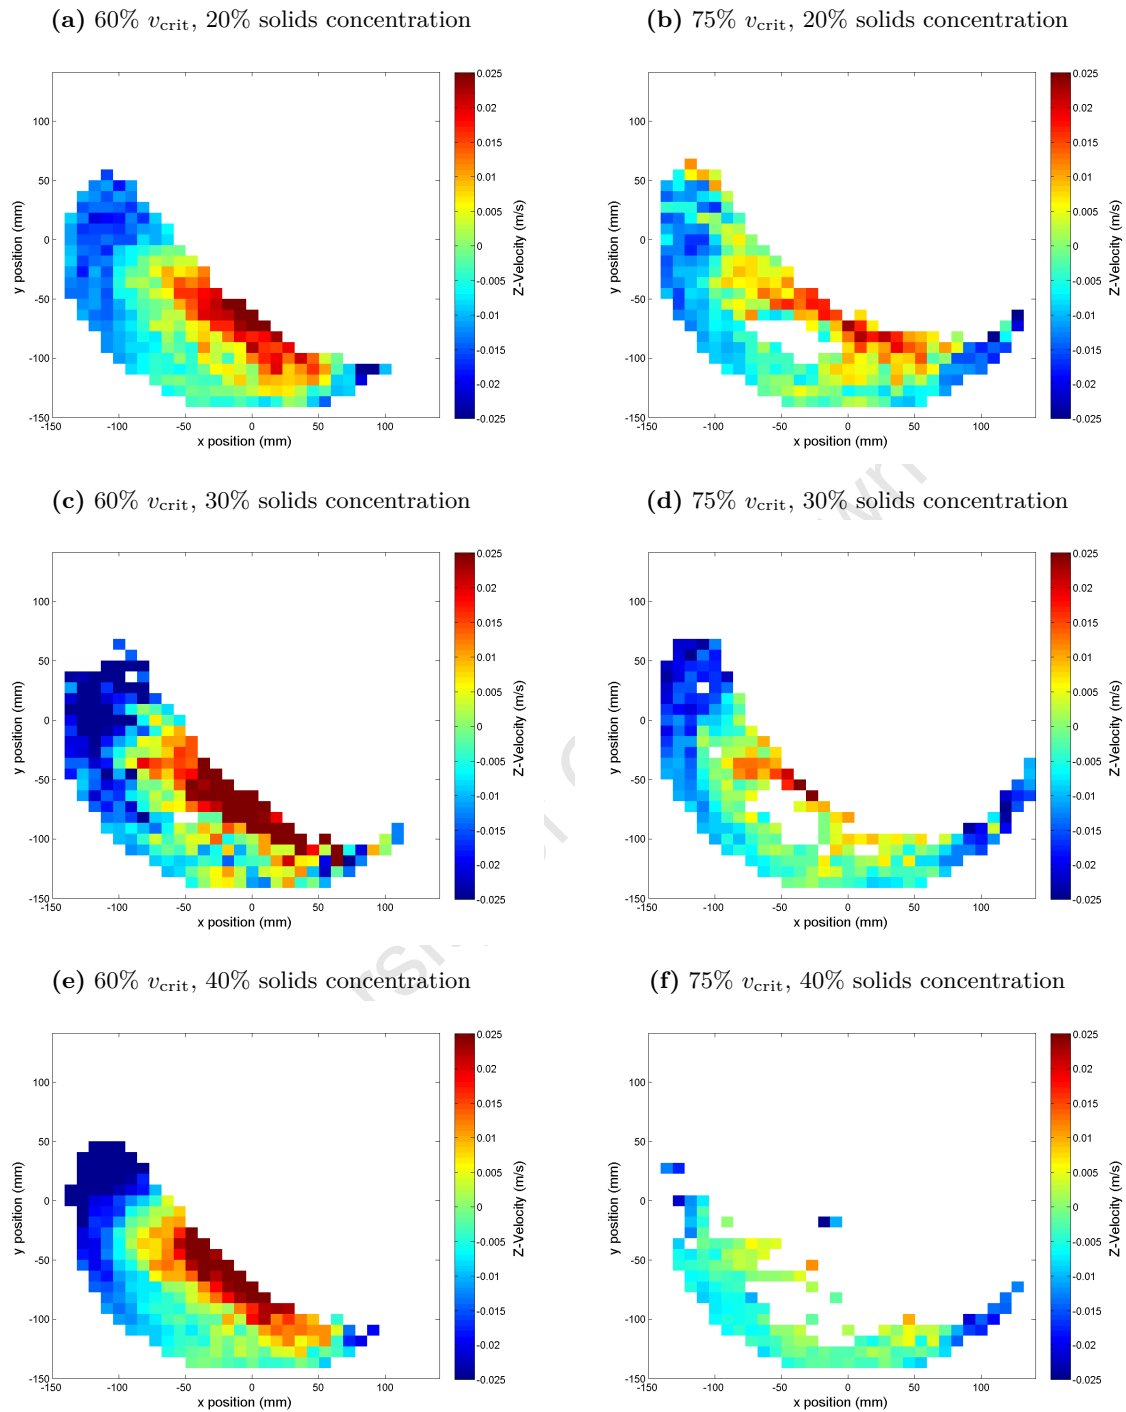


Figure 5.25: The z velocity profiles for the PEPT experiments with slurry. (a): 60% v_{crit} , 20% slurry solids concentration by mass. (b): 75% v_{crit} , 20% slurry solids concentration by mass. (c): 60% v_{crit} , 30% slurry solids concentration by mass. (d): 75% v_{crit} , 30% slurry solids concentration by mass. (e): 60% v_{crit} , 40% slurry solids concentration by mass. (f): 75% v_{crit} , 40% slurry solids concentration by mass.

Solids concentration %	20	30	40
Inflow rate (ml/s)	65.3	93.2	93.2
Full-run net rate (ml/s)	90.7	53.8	59.8
Averaging over 0.05 s net rate (ml/s)	-47.7	33.0	51.2

Table 5.11: Comparison of discharge rates at 75% v_{crit} between full-run time-averaging results and results from time-averaging over 0.05 s

5.5.3 Full-run Time-averaging

In this section, the effect of changing the time-averaging interval for the coupling data is investigated, specifically changing the time-averaging to cover the entire DEM simulation in the manner of Cleary *et al.* [13].

Transverse Volume Fraction



Figure 5.26 shows the volume fraction profiles for each time-averaging approach. The implementation of full-run averaging seems to have the effect of lowering the centre-of-mass of the fluid bulk. This effect becomes more pronounced at higher slurry solids fractions and points towards the relaxation time issues, discussed in Subsection 3.1.1, that manifest themselves at higher viscosities. That is, at higher viscosities the relaxation time of the fluid may be much longer than 0.05 s and the steady flow assumptions behind the drag models are no longer valid over such a short time interval.

Axial Velocity and Discharge Rate

Figure 5.27 shows the z-velocity profiles for each time-averaging approach. Again the greatest differences are seen at higher viscosities.



An even greater difference in the two approaches is seen in the discharge rates, shown in Table 5.11, but again the runs are too short to confidently infer too much from that data.

(a) Full-run time-averaging: 20% solids concentration (b) Averaging over 0.05 s: 20% solids concentration



(c) Full-run averaging: 30% solids concentration

(d) Averaging over 0.05 s: 30% solids concentration



(e) Full-run averaging: 40% solids concentration

(f) Averaging over 0.05 s: 40% solids concentration








Figure 5.26: A comparison of the volume fraction for one-way coupled simulations at 75% v_{crit} and 31.25% load with Beetstra drag (a): Using full-run time-averaging and slurry with 20% solids concentration by mass. (b): Time-averaging over 0.05 s with 20% solids concentration. (c): Using full-run time-averaging with 30% solids concentration. (d): Time-averaging over 0.05 s with 30% solids concentration. (e): Using full-run time-averaging with 40% solids concentration. (f): Time-averaging over 0.05 s with 40% solids concentration.


(a) Full-run time-averaging: 20% solids concentration (b) Time-averaging over 0.05 s: 20% solids concentration




(c) Full-run averaging: 30% solids concentration



(d) Averaging over 0.05 s: 30% solids concentration



(e) Full-run averaging: 40% solids concentration



(f) Averaging over 0.05 s: 40% solids concentration




Figure 5.27: A comparison of the z velocity for one-way coupled simulations at $75\% v_{crit}$ and 31.25% load with Beetstra drag (a): Using full-run time-averaging and slurry with 20% solids concentration by mass. (b): Time-averaging over 0.05 s with 20% solids concentration. (c): Using full-run time-averaging with 30% solids concentration. (d): Time-averaging over 0.05 s with 30% solids concentration. (e): Using full-run time-averaging with 40% solids concentration. (f): Time-averaging over 0.05 s with 40% solids concentration.

Chapter 6

Discussion and Conclusion

In this work, two main sets of simulations were designed and conducted. The first set, which employed DEM only, sought chiefly to establish the validity of DEM as a rigid solids computational technique for tumbling mills. Additionally, information about the validity of the one-way approach to coupling was garnered through comparisons of DEM only simulations with PEPT experiments. The second set, which employed one-way coupled DEM-CFD, sought to investigate the validity of the implementation and to assess the effects of different implementation choices.

In this final section, the main findings from these two sets of simulations will be revisited. Specifically, their agreement or disagreement with prior work, their agreement or disagreement with experiments and their apparent limitations will be discussed. With this information in mind, potential directions for development or modification of the one-way coupling process will be explored.

6.1 DEM-Only

6.1.1 12 cm Mill

For the DEM-only simulations of the 12 cm mill, it was found that the computational solidicity profile and particle velocity field matched those from the equivalent PEPT experiments exceptionally well. However, it was found that for simulations with high speeds and low loads this agreement was not as good. Particularly, derived data such as the free surface and shear rates did not match the PEPT experiments as well for that operating regime.

At high loads and low speeds, which are the domain of both mechanistic granular flow modelling and industrial mill operation, it was found that there was still good agreement with the PEPT experiments. Data from both simulations and experiments contradicted the assumption

by Taberlet *et al.* [37] that there is a constant shear rate in the mill. Rather, the data showed that there was a variation in the shear rate going from the centre of the mill to the shell. This variation was well-fitted with a quadratic function, in agreement with the experimental findings of Govender *et al.* [38].

One potential avenue for future work in this area might be to focus on the inputs for the DEM modelling. A particularly straightforward line of investigation could be to examine the material properties and collision parameters for the glass beads, which were not known to great accuracy and relied on somewhat crude laboratory measurements.

Another potential avenue for investigation is the effect of using different DEM contact models, either a simpler model such as spring-dashpot, or a novel model, such as that proposed by An and Tannant [81]. This could be examined specifically in the high speed, low load regime, where the correct resolution of extreme collisions is critical.

6.1.2 30 cm Mill

For the 30 cm mill, DEM-only simulations were compared with PEPT experiments that had the same solids filling but also contained slurry. It was found that the solidicity profiles were in reasonable agreement, but that the experimental profile was somewhat dilated. The velocity fields however were in excellent agreement. This data supported the assumption that for typical operational fluid levels, it is a reasonable approximation to ignore the effect of the fluid on the solids.

Further work in this direction might investigate if this approximation holds up in situations where there is a large volume of fluid in the mill compared to the volume of solids. This is of industrial interest for when operating conditions lead to such non-optimal regimes with high holdup in the mill.

6.2 One-way Coupled DEM-CFD

For the one-way coupled simulations, it was found that the choice of drag model has a significant effect on the results. The solidicity profiles and velocity fields for simulations using Beetstra drag were significantly different from those using Ergun drag. It was found that the key to this difference was the fact that the Ergun correlation over-predicts the drag force at the high Reynolds numbers present in some regions of the mill. It is thus clear that the use of the Ergun correlation and other correlations developed for low Reynolds number flows is inappropriate for the operating conditions used in this work. In fact, in industrial scale mills such modelling would be even less

appropriate, as the higher fluid velocities would lead to even higher Reynolds numbers (provided the fluid density, fluid viscosity and particle size remain constant).

6.2.1 Relation to Prior Work

In terms of the mechanistic modelling of axial transport, two assumptions about the slurry profile in the Moys [3] model were contradicted by this work. The first assumption, about the transverse profile being flat, was clearly contradicted, although the author was aware that this was a major simplification. The transverse profile rather followed the shape of the charge, as also seen by Cleary *et al.* [13] and Morell and Stephenson [4]. The second assumption, which was more subtle, suggested that the axial profile of the slurry would be a smooth slope. This too was contradicted by the simulations in this work, which showed a wave-like structure in the axial profile.

Regarding the regions in which transport occurs, the model presented by Condori *et al.* [67] is contradicted by this work. In their model, transport towards the discharge end of the mill occurs exclusively in the toe region, while higher up in the mill the slurry is completely stationary in the axial direction. In the simulations in the present work, there is also flow towards the discharge end in the toe region, but in the shoulder region there is in fact some flow back towards the feed end. This phenomenon is also observed in the PEPT experiments presented in this work.

6.2.2 Comparison With Experiments

While both the simulations and PEPT experiments agreed on the transverse slurry profile roughly matching the shape of the charge, the PEPT experiments showed a dry region near the centre of circulation. This phenomenon was also seen in simulations by Cleary *et al.* [13], and highlights a limitation of the implementation in this work. The limitation is that the fluid volume was initialised as a continuous bulk in the bottom half of the mill, while in the work by Cleary *et al.* [13] it was initialised as a bulk above the charge, and allowed to splash down. This, combined the fact that a continuum description (CFD) is used, means that the development of a dry region in the centre of the mill is impossible, at least over short time-scales.

This limitation motivates two potential improvements for future work. The first would be to initialise the fluid volume in the same manner as Cleary *et al.* [13], as a bulk above the charge. The second would be to introduce the fluid in the same way it is introduced in the industrial context, which is as a steady stream from the feed end of the mill. This approach would be very computationally expensive due to the additional time needed to ramp the simulation up to an operating state, but it has the potential to capture phenomena that the other two approaches cannot.

Another major difference between the simulations and PEPT experiments was that the flow towards the feed end in the PEPT experiments, while still at the bottom of the mill, was concentrated closer to the free surface. This is most probably due to situations in which the tracer particle did not penetrate the charge at all, and instead skipped along the free surface, as observed when the experiments were run. In this case the tracer particle is distinctly not representative of the fluid, and possible future work might involve smaller tracer particles or potentially a PET analysis of a radioactively labelled fluid.

6.2.3 Implementation Concerns

In the main body of simulations, it was assumed that the relaxation time of the fluid was very short, allowing for a time-averaged implementation with intervals of 0.05 s. However, work by Blakey *et al.* [68] has shown that relaxation times for highly viscous slurries can be on the order of 10 s. In such cases, much longer time-averaging intervals would be necessary and in this work it was shown that such implementations may even be necessary even at moderate viscosities.

Turbulence was not modelled for the slurry in this work due to incompatibility between FLUENT's turbulence and non-Newtonian models. However, the drag correlations that were employed include the effect of turbulence on drag at least on the scale of the averaging volumes. Potential future work might include a user-defined viscosity description combined with a standard turbulence model (some candidates are surveyed in Subsection 2.3.3).

Regarding the drag models used in this work, both rest on the assumption that the granular temperature of the solids is zero. This is clearly not true in mills, but it is used as an approximation. The effect of this approximation may be difficult to explore as there are no fully-developed drag models that take granular temperature into account for a large range of Reynolds numbers.

It was shown in this work that one-way coupling is a reasonable approximation, but there is potential benefit in exploring two-way coupling, especially in regimes with high fluid volumes. Two-way coupling could be implemented in at least three ways, which will be discussed in order of increasing complexity.

The first way, sometimes called “weak coupling”, would be to capture the effect of the fluid on the solids through modified coefficients of friction and restitution in a dry DEM simulation. This DEM simulation would then be used to calculate drag forces on the fluid in the one-way manner presented in this work.

The second way, an iterated approach, would rely on the system in question reaching some steady state. In this approach, a dry DEM simulation would be run, then used as input for a CFD simulation. The data from the CFD simulation would then be used to calculate fluid forces on a new DEM simulation, and the process would be repeated. The repetition would be continued until some criterion for determining that the system is in steady state is met.

The final way, which is the most computationally expensive, is to have genuine two-way coupling, where forces are exchanged between the solids and the fluids either at every time-step or at the end of each time-averaged interval.

6.2.4 Future Applications

In this work, results of very short simulations on the order of 5 s have been presented. In the industrial context, especially with reliable measurements of the fluid hold-up in the mill, significantly longer simulations will be required. With these long simulations in mind, engineers might like to explore a broad range of design and operation parameters, from the area of grates on the discharge end of the mill, to fill-rates, operating loads and speeds. The cost of exploring this entire parameter space for optima would be immense, which is why the value of such work is not in running simulations for every imaginable case.

The true need in the minerals processing industry is for simply-implemented models that can quickly provide accurate predictions for a broad range of parameters. Such a need is best met by mechanistic modelling approaches, but these approaches have difficulties of their own. Chief among these difficulties is the fact that a number of simplifying assumptions need to be made, and assessing their validity directly through experiment may be difficult or impossible. Comparing their assumptions to computational data however is extremely simple, and the simulations in this work have shown that some assumptions in the literature range from spot on to patently false. It is the ability to test these assumptions that is the true strength of this computational approach.

Bibliography

- [1] P. Mlamo-Ngcuka, “Draft energy efficiency strategy of the republic of south africa,” tech. rep., South African Department of Minerals and Energy, 2004.
- [2] D. Tromans, “Mineral comminution: Energy efficiency considerations,” *Minerals Engineering*, vol. 21, pp. 613–620, 2008.
- [3] M. H. Moys, “The effect of grate design on the behaviour of grate-discharge grinding mills,” *International Journal of Mineral Processing*, vol. 18, pp. 85–105, 1986.
- [4] S. Morrell and I. Stephenson, “Slurry discharge capacity of autogenous and semi-autogenous mills and the effect of grate design,” *International Journal of Mineral Processing*, vol. 46, pp. 53–72, 1996.
- [5] S. Powell, S. Morrel, and S. R. Latchireddi, “Developments in the understanding of south african style sag mills,” *Minerals Engineering*, vol. 14, pp. 1143–1153, 2001.
- [6] S. R. Latchireddi and S. Morrel, “A laboratory study of the performance characteristics of mill pulp lifters,” *Minerals Engineering*, vol. 10, pp. 1233–1244, 1997.
- [7] S. R. Latchireddi and S. Morrel, “Slurry flow in mills: grate-only discharge mechanism (part-1),” *Minerals Engineering*, vol. 16, pp. 625–633, 2003.
- [8] S. R. Latchireddi and S. Morrel, “Slurry flow in mills: grate-pulp lifter discharge systems (part 2),” *Minerals Engineering*, vol. 16, pp. 635–642, 2003.
- [9] S. R. Latchireddi and S. Morrel, “Slurry flow in mills with tcpl - an efficient pulp lifter for ag/sag mills,” *International Journal of Mineral Processing*, vol. 79, pp. 174–187, 2006.
- [10] R. K. Rajamani, B. K. Mishra, R. Venugopal, and A. Datta, “Discrete element analysis of tumbling mills,” *Powder Technology*, vol. 109, pp. 105–112, 2000.
- [11] P. W. Cleary, “Modelling comminution devices using dem,” *International Journal For Numerical And Analytical Methods In Geomechanics*, vol. 25, pp. 83–105, 2001.

- [12] P. Radziszewski, “Exploring total media wear,” *Minerals Engineering*, vol. 15, pp. 1073–1087, 2002.
- [13] P. W. Cleary, M. Sinnott, and R. Morrison, “Prediction of slurry transport in sag mills using sph fluid flow in a dynamic dem based porous media,” *Minerals Engineering*, vol. 19, pp. 1517–1527, 2006.
- [14] P. A. Cundall and D. L. Strack, “A discrete numerical model for granular assemblies,” *Geotechnique*, vol. 29, pp. 47–65, 1979.
- [15] H. Hertz, “Ueber der berührung fester elastischer korper,” *Journal fur die reine und angewandte*, vol. 92, pp. 156–171, 1881.
- [16] K. Johnson, *Contact Mechanics*. Cambridge University Press, 1985.
- [17] R. Y. Yang, R. P. Zou, and A. B. Yu, “Computer simulation of the packing of fine particles,” *Phys. Rev. E*, vol. 62, 3, pp. 3900–3908, 2000.
- [18] DEM Solutions, *EDEM 2.3 User Guide*, 2010.
- [19] N. Maw, J. Barber, and J. Fawcett, “The oblique impact of elastic spheres,” *Wear*, vol. 38, pp. 101–114, 1976.
- [20] R. Horn, J. Israelachvili, and F. Pribac, “Measurement of the deformation and adhesion of solids in contact,” *Journal of Colloid and Interface Science*, vol. 115, 2, pp. 480–492, 1987.
- [21] P. Claesson, T. Ederth, V. Bergeron, and M. Rutland, “Techniques for measuring surface forces,” *Advances in Colloid and Interface Science*, vol. 67, pp. 119–183, 1996.
- [22] H. A. Makse, N. Gland, D. Johnson, and L. Schwartz, “Granular packings: Nonlinear elasticity, sound propagation and collective relaxation dynamics,” *Phys. Rev. E*, vol. 70, 6, p. 061302, 2004.
- [23] J. Ottino and D. Khakhar, “Mixing and segregation of granular materials,” *Annual Review of Fluid Mechanics*, vol. 32, pp. 55–91, 2000.
- [24] R. D. Mindlin, “Compliance of elastic bodies in contact,” *Journal of Applied Mechanics*, vol. 16, pp. 256–270, 1949.
- [25] R. Mindlin and H. Deresiewicz, “Elastic spheres in contact under varying oblique forces,” *Journal of Applied Mechanics*, vol. 20, pp. 327–344, 1953.
- [26] A. Di Renzo and F. P. Di Maio, “Comparison of contact-force models for the simulation of collisions in dem-based granular flow codes,” *Chemical Engineering Science*, vol. 59, pp. 525–541, 2004.

- [27] C. Thornton and W. Randall, *Micromechanics of Granular Materials*, ch. Applications of Theoretical Contact Mechanics to Solid Particle System Simulations, pp. 133–142. Elsevier, Amsterdam, 1988.
- [28] N. Brilliantov and T. Poschel, “Rolling friction of a viscous sphere on a hard plane,” *Eurphysics Letters*, vol. 42, pp. 511–516, 1988.
- [29] N. V. Brilliantov, F. Spahn, J. M. Hertzsch, and T. Poschel, “Model for collisions in granular gases,” *Phys. Rev. E*, vol. 53, pp. 5382–5392, 1996.
- [30] P. W. Cleary, “Recent advances in dem modelling of tumbling mills,” *Minerals Engineering*, vol. 14, 10, pp. 1295–1319, 2001.
- [31] L. Vu-Quoc, X. Zhang, and O. R. Walton, “A 3-d discrete-element method for dry granular flows of ellipsoidal particles,” *Computer Methods in Applied Mechanics and Engineering*, vol. 187, pp. 483–528, 2000.
- [32] B. K. Mishra, “A review of computer simulation of tumbling mills by the discrete element method: Part i-contact mechanics,” *International Journal of Mineral Processing*, vol. 71, pp. 73–79, 2003.
- [33] Y. Tsuji, T. Kawaguchi, and T. Tanaka, “Discrete particle simulation of two-dimensional fluidized bed,” *Powder Technology*, vol. 77, pp. 79–87, 1993.
- [34] D. Zhang and W. Whitten, “Step size control for efficient discrete element simulation,” *Minerals Engineering*, vol. 14, pp. 1341–1346, 2001.
- [35] P. W. Cleary, R. Morrison, and S. Morrel, “Comparison of dem and experiment for a scale model sag mill,” *International Journal of Mineral Processing*, vol. 68, pp. 129–165, 2003.
- [36] A. Datta, B. K. Mishra, and R. K. Rajamani, “Analysis of power draw in ball mills by the discrete element method,” *Canadian Metallurgical Quarterly*, vol. 38, 2, pp. 133–140, 1999.
- [37] N. Taberlet, P. Richard, and J. Hinch, “S shape of a granular pile in a rotating drum,” *Phys. Rev. E*, vol. 73, p. 050301(R), 2006.
- [38] I. Govender, N. Mangesana, M. A. N., and J.-P. Franzidis, “Measurement of shear rates in a laboratory tumbling mill,” *Minerals Engineering*, vol. 24, pp. 225–229, 2011.
- [39] G. K. Batchelor, *An Introduction to Fluid Dynamics*. Cambridge University Press, 1967.
- [40] J. H. Ferziger and M. Peric, *Computational Methods for Fluid Dynamics*. Springer, 1997.
- [41] E. C. Bingham, *Fluidity and Plasticity*. McGraw-Hill (New York), 1922.

- [42] C. W. Hirt and B. D. Nichols, "Volume of fluid (vof) method for the dynamics of free boundaries," *Journal of Computational Physics*, vol. 39, pp. 201–2225, 1979.
- [43] D. Gueyffier, J. Li, A. Nadim, R. Scardovelli, and S. Zaleski, "Volume-of-fluid interface tracking with smoothed surface stress methods for three-dimensional flows," *Journal of Computational Physics*, vol. 152, pp. 423–456, 1999.
- [44] J. E. Pilliod Jr. and E. G. Puckett, "Second-order accurate volume-of-fluid algorithms for tracking material interfaces," *Journal of Computational Physics*, vol. 199, pp. 465–502, 2004.
- [45] M. S. Brennan, "Multiphase cfd simulations of dense medium and classifying hydrocyclones," in *Third International Conference on CFD in the Minerals and Process Industries*, 2003.
- [46] J. A. Delgadillo and R. K. Rajamani, "A comparative study of three turbulence-closure models for the hydrocyclone problem," *International Journal of Mineral Processing*, vol. 77, pp. 217–230, 2005.
- [47] *FLUENT 12.0 Theory Guide*, May 2008.
- [48] C. G. Speziale, "Analytical methods for the development of reynolds stress closures in turbulence," tech. rep., Institute for Computer Applications in Science and Engineering, 1990.
- [49] J. O. Hinze, *Turbulence*. McGraw-Hill Book Company, Inc., 1959.
- [50] B. E. Launder and D. B. Spalding, "The numerical computation of turbulent flows," *Computer Methods in Applied Mechanics and Engineering*, vol. 3, no. 2, pp. 269 – 289, 1974.
- [51] B. V. Antohe and J. L. Lage, "A general two-equation macroscopic turbulence model for incompressible flow in porous media," *International Journal of Heat and Mass Transfer*, vol. 40, no. 13, pp. 3013–3024, 1997.
- [52] V. Yakhot and S. A. Orszag, "Renormalization-group analysis of turbulence," *Physical Review Letters*, vol. 57, 14, pp. 1722–1724, 1986.
- [53] T.-H. Shih, W. W. Liou, A. Shabbir, Z. Yang, and J. Zhu, "A new k- ϵ eddy viscosity model for high reynolds number turbulent flows," *Computers & Fluids*, vol. 24, 3, pp. 227–238, 1994.
- [54] F. C. Blake, "The resistance of packing to fluid flow," *Transactions of the American Institute of Chemical Engineers*, vol. 14, p. 415, 1922.
- [55] R. Beetstra, M. A. van der Hoef, and J. A. M. Kuipers, "Drag force of intermediate reynolds number flow past mono- and bidisperse arrays of spheres," *A.I.Ch.E. Journal*, vol. 53, 2, pp. 489–501, 2007.

- [56] S. Ergun, "Fluid flow through packed columns," *Chemical Engineering Progress*, vol. 48, pp. 89–94, 1952.
- [57] R. Fand, B. Y. K. Kim, A. C. C. Lam, and R. T. Phan, "Resistance to the flow of fluids through simple and complex porous media whose matrices are composed of randomly packed spheres.," *Journal of fluids engineering*, vol. 109, pp. 268–274, 1987.
- [58] N. G. Deen, M. Van Sint Annaland, M. Van der Hoef, and J. A. M. Kuipers, "Review of discrete particle modeling of fluidized beds," *Chemical Engineering Science*, vol. 62, pp. 28–44, 2007.
- [59] W. Sobieski, "Momentum exchange in solid-fluid system modeling with the eulerian multi-phase model," *Drying Technology*, vol. 27, pp. 653–671, 2009.
- [60] R. J. Hill, D. L. Koch, and A. J. C. Ladd, "The first effects of fluid inertia on flows in ordered and random arrays of spheres," *J. Fluid Mech.*, vol. 448, pp. 213–241, 2001.
- [61] R. J. Hill, D. L. Koch, and A. J. C. Ladd, "Moderate-reynolds-number flows in ordered and random arrays of spheres," *J. Fluid Mech.*, vol. 448, pp. 243–278, 2001.
- [62] A. Z. Zinchenko, "Algorithm for random close packing of spheres with periodic boundary conditions," *Journal of Computational Physics*, vol. 114, pp. 298–307, 1994.
- [63] G. A. Bokkers, M. van Sint Annaland, and J. A. M. Kuipers, "Mixing and segregation in a bidisperse gas-solid fluidised bed: a numerical and experimental study," *Powder Technology*, vol. 140, pp. 176–186, 2004.
- [64] S. Benyahia, M. Syamlal, and T. J. O'Brien, "Extension of hill-koch-ladd drag correlation over all ranges of reynolds number and solids volume fraction," *Powder Technology*, vol. 162, pp. 166–174, 2006.
- [65] M. Van der Hoef, R. Beetstra, and J. A. M. Kuipers, "Lattice-boltzmann simulations of low-reynolds-number flow past mono- and bidisperse arrays of spheres results for the permeability and drag force," *Journal of Fluid Mechanics*, vol. 528, pp. 233–254, 2005.
- [66] A. V. Potapov, M. L. Hunt, and C. S. Campbell, "Liquid-solid flows using smoothed particle hydrodynamics and the discrete element method," *Powder Technology*, vol. 116, pp. 204–213, 2001.
- [67] P. Condori, M. A. N., I. Govender, and M. S. Powell, "A mechanistic approach to modelling slurry transport in ag/sag mills: Transport through the charge," *Proceedings of XXIV International Mineral Processing Congress*, vol. 2, pp. 384–391, 2008.

- [68] B. C. Blakey and D. F. James, "Characterizing the rheology of laterite slurries," *International Journal of Mineral Processing*, vol. 70, pp. 23–29, 2003.
- [69] B. Naylor, "Binary space partitioning trees as an alternative representation of polytopes," *Computer-Aided Design*, vol. 22, pp. 250–252, 1990.
- [70] S. Owen, "A survey of unstructured mesh generation technology," in *Proceedings of the 7th International Meshing Roundtable*, 1998.
- [71] D. Parker, M. R. Hawkesworth, C. J. Broadbent, P. Fowles, T. D. Fryer, and P. McNeil, "Industrial positron-based imaging: Principles and applications," *Nuclear Instruments and Methods A*, vol. 348, pp. 583–592, 1994.
- [72] D. Parker, A. E. Dijkstra, T. W. Martin, and J. P. K. Seville, "Positron emission particle tracking studies of spherical particle motion in rotating drums," *Chemical Engineering Science*, vol. 52 (13), pp. 2011–2022, 1997.
- [73] D. Szasz, "Boltzmann's ergodic hypothesis, a conjecture for centuries?," in *Lecture given at the international symposium in honour of Boltzmann's 150th birthday*, 1994.
- [74] M. R. Hawkesworth, M. O'Dwyer, J. Walker, P. Fowles, J. Heritage, P. Stewart, R. Witcomb, J. Bateman, J. Connolly, and R. Stephenson, "A positron camera for industrial application," *Nuclear Instruments and Methods A*, vol. 253, pp. 145–157, 1986.
- [75] P. Rosin and E. Rammler, "The laws governing the fineness of powdered coal," *Journal of the Institute of Fuel*, vol. 7, pp. 29–36, 1933.
- [76] S. V. Patankar and D. B. Spalding, "A calculation procedure for heat, mass and momentum transfer in three-dimensional parabolic flows," *International Journal of Heat and Mass Transfer*, vol. 15, pp. 1787–1806, 1972.
- [77] J. P. Van Doormaal and G. D. Raithby, "Enhancements of the simple method for predicting incompressible fluid flows," *Numerical Heat Transfer*, vol. 7, pp. 147–163, 1984.
- [78] C. Song, P. Wang, and H. A. Makse, "A phase diagram for jammed matter," *Nature*, vol. 453 (7195), pp. 629–632, 2008.
- [79] J. A. Canny, "A computational approach to edge detection," *IEEE Trans. Pattern Analysis and Machine Intelligence*, vol. 8 (6), pp. 679–698, 1986.
- [80] O. Ozcan, B. B. Gonul, A. N. Bulutcu, and H. Manav, "Correlations between the shear strength of mineral filter cakes and particle size and surface tension," *Colloids and Surfaces A: Physicochemical and Engineering Aspects*, vol. 187–188, pp. 405–413, 2001.

- [81] B. An and D. D. Tannant, "Discrete element method contact model for dynamic simulation of inelastic rock impact," *Computers and Geosciences*, vol. 33, pp. 513–521, 2007.
- [82] M. K. Abd El-Rahman, B. K. Mishra, and R. K. Rajamani, "Industrial tumbling mill power prediction using the discrete element method," *Minerals Engineering*, vol. 10, pp. 1321–1328, 2001.
- [83] R. I. Al-Raoush and C. S. Willson, "A pore-scale investigation of a multiphase porous media system," *Journal of Contaminant Hydrology*, vol. 77, pp. 67–89, 2005.
- [84] R. Beetstra, M. A. van der Hoef, and J. A. M. Kuipers, "Numerical study of segregation using a new drag force correlation for polydisperse systems derived from lattice-boltzmann simulations," *Chemical Engineering Science*, vol. 62, pp. 246–255, 2007.
- [85] C. E. Brennen, *Fundamentals of Multiphase Flows*. Cambridge University Press, 2005.
- [86] A. Di Renzo and F. P. Di Maio, "Homogeneous and bubbling fluidization regimes in dem-cfd simulations: Hydrodynamic stability of gas and liquid fluidized beds," *Chemical Engineering Science*, vol. 62, pp. 116–130, 2007.
- [87] Y. T. Feng, K. Han, and D. R. J. Owen, "Coupled lattice boltzmann method and discrete element modelling of particle transport in turbulent fluid flows: Computational issues," *International Journal for Numerical Methods in Engineering*, vol. 72, pp. 1111–1134, 2007.
- [88] F. Fichot, P. Meekunnasombat, J. Belloni, F. Duval, A. Garcia, and M. Quintard, "Two-phase flows in porous media: Prediction of pressure drops using a diffuse interface mathematical description," *Nuclear Engineering and Design*, vol. 237, pp. 1887–1898, 2007.
- [89] K. Han, Y. Feng, and D. R. J. Owen, "Coupled lattice boltzmann and discrete element modelling of fluid-particle interaction problems," *Computers and Structures*, vol. 85, pp. 1080–1088, 2007.
- [90] H. Herrman and S. Luding, "Modeling granular media on the computer," *Continuum Mechanics and Thermodynamics*, vol. 10, pp. 189–231, 1998.
- [91] W. E. Horst and E. J. Freeh, "Mathematical modelling of a continuous comminution process," *Transactions of the American Institute of Minerals Engineers*, vol. 252, pp. 160–167, 1972.
- [92] D. L. Koch and R. J. Hill, "Inertial effects in suspension and porous-media flows," *Annu. Rev. Fluid Mech.*, vol. 1, pp. 619–647, 2001.

- [93] L. Labous, A. Rosato, and R. Dave, “Measurements of collisional properties of spheres using high-speed video analysis,” *Phys. Rev. E*, vol. 56, 5, pp. 5715–5725, 1997.
- [94] J. Link, L. Cupers, N. G. Deen, and J. A. M. Kuipers, “Flow regimes in a spout-fluid bed: A combined experimental and simulation study,” *Chemical Engineering Science*, vol. 60, pp. 3425–3442, 2005.
- [95] B. K. Mishra, “A review of computer simulation of tumbling mills by the discrete element method part ii - practical applications,” *International Journal of Mineral Processing*, vol. 71, pp. 95–112, 2003.
- [96] E. G. Nezami, Y. M. Hashash, D. Zhao, and J. Ghaboussi, “A fast contact detection algorithm for 3-d discrete element method,” *Computers and Geotechnics*, vol. 31, pp. 575–587, 2004.
- [97] F. N. Shih and T. J. Napier-Munn, “A model for slurry rheology,” *International Journal of Mineral Processing*, vol. 47, pp. 103–123, 1996.
- [98] E. Simsek, B. Brosch, S. Wirtz, V. Scherer, and F. Krüll, “Numerical simulation of grate firing systems using a coupled cfd/discrete element method (dem),” *Powder Technology*, vol. 193, pp. 266–273, 2009.
- [99] D. E. Stewart and J. C. Trinkle, “An implicit time-stepping scheme for rigid body dynamics with inelastic collisions and coulomb friction,” *International Journal for Numerical Methods in Engineering*, vol. 39, pp. 2673–2691, 1996.
- [100] T. Tsuji, K. Yabumoto, and T. Tanaka, “Spontaneous structures in three-dimensional bubbling gas-fluidized bed by parallel dem-cfd coupling simulation,” *Powder Technology*, vol. 184, pp. 132–140, 2008.
- [101] O. Ubbink, *Numerical prediction of two fluid systems with sharp interfaces*. PhD thesis, University of London, 1997.
- [102] B. H. Xu and A. B. Yu, “Numerical simulation of the gas-solid flow in a fluidized bed by combining discrete particle method with computational fluid dynamics,” *Chemical Engineering Science*, vol. 16, pp. 2785–2809, 1997.
- [103] W. Zhong, Y. Xiong, Z. Yuan, and M. Zhang, “Dem simulation of gas-solid flow behaviors in spout-fluid bed,” *Chemical Engineering Science*, vol. 61, pp. 1571–1584, 2006.
- [104] A. Z. Zinchenko and R. H. Davis, “Algorithm for direct numerical simulation of emulsion flow through a granular material,” *Journal of Computational Physics*, vol. 227, pp. 7841–7888, 2008.

- [105] A. Z. Zinchenko and R. H. Davis, “A multipole-accelerated algorithm for close interaction of slightly deformable drops,” *Journal of Computational Physics*, vol. 207, pp. 695–735, 2005.

Appendix A

Animations

Original copies of this thesis include a disk with a number of video files showing animations of one-way coupled simulations from Section 5.5. All of the videos show the axially-averaged slurry concentration profile, with some smoothing between voxels. There is some apparent strobing at the edge of the mill in some of the videos, which is a product of the sampling rate and the very edge of the mill being out of view.

Naming

All of the file names have the following format `30cm_[s]speed_31.25load_slurryc[c]_[model]_[averaging]`, where the parts have the following meanings:

- `30cm`: indicates that it is the 30 cm mill.
- `[s]`: an integer that gives the percentage of v_{crit} that the mill is operating at.
- `31.25load`: indicates that the charge fills 31.25% of the available volume.
- `[c]`: an integer that gives the slurry concentration percentage by mass.
- `[model]`: the name of the drag model employed. Either Ergun or Beetstra.
- `[averaging]`: if not present, indicates that the normal 0.05 s time-averaging was used. If “fullaverage”, indicates that full-run time-averaging was used.

For example, `30cm_75speed_31.25load_slurryc20_beetstra_fullaverage` indicates that the mill is operating at 75% v_{crit} , the slurry concentration is 20%, the Beetstra drag correlation was used, and that full-run time-averaging was used.

Name	Speed	Slurry	Drag model	Time-averaging
30cm_60speed_31.25load_slurryc20_ergun	60% v_{crit}	20% by mass	Ergun	0.05 s
30cm_60speed_31.25load_slurryc30_ergun	60% v_{crit}	30% by mass	Ergun	0.05 s
30cm_60speed_31.25load_slurryc30_beetstra	60% v_{crit}	30% by mass	Beetstra	0.05 s
30cm_60speed_31.25load_slurryc40_beetstra	60% v_{crit}	30% by mass	Beetstra	0.05 s
30cm_75speed_31.25load_slurryc20_ergun	75% v_{crit}	20% by mass	Ergun	0.05 s
30cm_75speed_31.25load_slurryc20_beetstra	75% v_{crit}	20% by mass	Beetstra	0.05 s
30cm_75speed_31.25load_slurryc30_ergun	75% v_{crit}	30% by mass	Ergun	0.05 s
30cm_75speed_31.25load_slurryc30_beetstra	75% v_{crit}	30% by mass	Beetstra	0.05 s
30cm_75speed_31.25load_slurryc40_ergun	75% v_{crit}	40% by mass	Ergun	0.05 s
30cm_75speed_31.25load_slurryc40_beetstra	75% v_{crit}	40% by mass	Beetstra	0.05 s

Table A.1: Table of videos containing data presented in Sub-section 5.5.1

Name	Speed	Slurry	Drag model	Time-averaging
30cm_75speed_31.25load_slurryc20_beetstra_fullaverage	75% v_{crit}	20% by mass	Beetstra	Full-run
30cm_75speed_31.25load_slurryc40_beetstra_fullaverage	75% v_{crit}	40% by mass	Beetstra	Full-run

Table A.2: Table of videos containing data presented in Sub-section 5.5.3

Full Listing

Table A.1 lists the videos containing data presented in Sub-section 5.5.1. Table A.2 lists the videos containing data presented in Sub-section 5.5.3.

File Format

The files are all in the .mp4 format, with standard H.264/MPEG-4 compression.