Chapter 2

Diffusion and mixing in granular drainage

2.1 Introduction

Hopper drainage is perhaps one of the simplest examples of a dense granular flow, since it requires no external forcing and happens by gravity only. Over the past forty years, it has been extensively studied, and a number of continuum models [77, 90, 95, 108, 94, 92] have been proposed. However, these models primarily concentrated on the steady state mean flow. This chapter documents some early work that was carried out to study other aspects of granular drainage, particularly diffusion and mixing, which has relevance in a number of industrial situations.

2.2 Experimental motivation

The experiment shown in figure 1-1 made use of two different colors of particles. While the movement of the different colors allows us to visualize the mean flow, it also tells us something about diffusion. In frames (a) and (b), there appears to be very little mixing, but by the third frame, some of the colors appear to have blended together. Motivated by this work, Choi *et al.* carried out experiments to specifically address this issue [28]. They carried out experiments of glass beads of diameter $d = 3$ mm in

a quasi two dimensional hopper of size 20 cm $\times$ 2.5 cm $\times$ 2.5 cm, with aluminum side walls, and Plexiglas walls at the front and back. Drainage took place from a slit in the center of the container base, the width of which could be altered to change the overall flow rate. During the flow, the particles at the glass wall were imaged using a high-speed digital camera, capable of taking images at 1000 frames per second.

The steady state mean flow in the experiment is characterized by a parabolic converging region of flow above the orifice, that connects with roughly uniform flow higher up. The measurements of particle diffusion were made in the region of uniform flow, by imaging a $17d \times 87d$ region centered at a height of $150d$ above the slit. Three of their key conclusions were:

- Particle diffusion is extremely small. Péclet[1] numbers on the order of 300 were observed, meaning that a particle would have to fall twice the height of the silo before diffusing by a single particle diameter.

- The neighbors of a particle strongly persist: a particle at the top of the observation window will typically retain 90% of its neighbors by the bottom.

- Over a large range of different flow rates, the amount of particle diffusion collapsed if plotted as a function of distance dropped, and not time.

These results are very surprising: in gases and liquids, diffusion is facilitated by thermal fluctuations, and the amount of diffusion is proportional to the time. In these experiments, time appears to not be the most significant factor, and altering the flow rate gives the same results, but with time rescaled. It suggests that granular diffusion is distinctly athermal, and that the packing geometry of the particles is the most significant factor.

[1]The Péclet number characterises the ratio of advection to diffusion [105], and is typically defined as $\mathrm{Pe} = LV/D$, where V is the velocity, D is the diffusion constant, and L is a characteristic length scale. In Choi's experiments, this is defined as $\mathrm{Pe} = 2d\Delta h/\langle \Delta x^2 \rangle$, where Δh is the distance dropped, and Δx is the horizontal displacement.

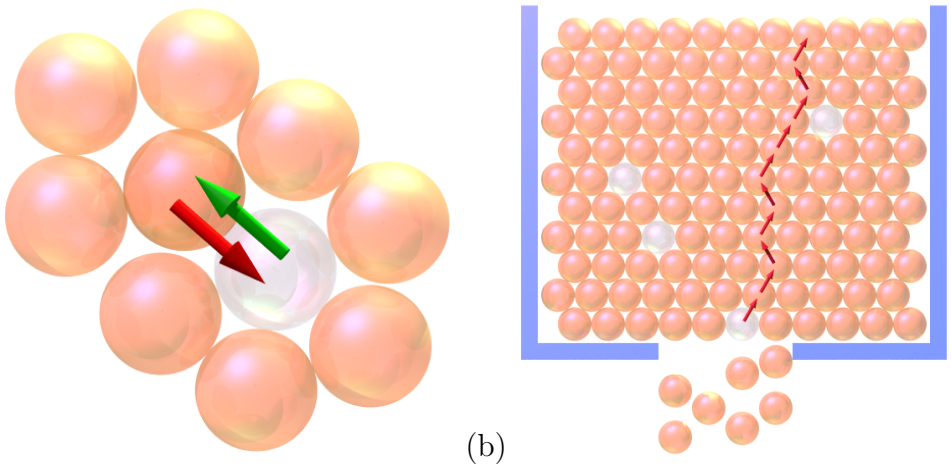

(a) (b)

Figure 2-1: (a) The void model microscopic mechanism: particles movie by exchanging position with voids. (b) The void model for granular flow: particles are imagined as lying in a two dimensional lattice, and move in response to voids diffusing upwards from the container orifice.

2.3 The void model

To explain these results requires a microscopic model for particle motion, and perhaps the only candidate in the literature is the concept of a void, shown in figure 2-1(a). The void mechanism has a long history, and was proposed by Eyring for viscous flow [42] but it has re-appeared in many other situations such as the glass transition [29], shear flow in metallic glasses [128], and compaction in vibrated granular materials [23].

The idea was first postulated in the context of granular materials by Litwiniszyn [77], and was subsequently studied by several other authors [90, 95, 137, 25] (with some related work making use of cellular automata models [15, 16]). Particles are imagined as being on a two-dimensional hexagonal lattice, as shown in figure 2-1(b). The flow is created by introducing voids at the orifice which propagate up through the packing according to a random walk, moving up-left or up-right with equal probability. To find the mean flow in this case, we introduce an infinite square lattice of sites labeled (M, N) in the (x, z) plane, with a horizontal and vertical spacings of d and h respectively. We consider a single void propagating upwards through the lattice, and let $V_N(M)$ be the probability that it is at horizontal site M when it

is in the Nth vertical layer. The random walk description allows us to write down a recurrence relation of the form

$$V_N(M) = \frac{V_{N-1}(M-1) + V_{N-1}(M+1)}{2}.$$

If we now introduce a continuum analogue η such that $V_N(M) = \eta(dM, hN)$ then we see that

$$\eta(x,z) = \frac{\eta(x-d, z-h) + \eta(x+d, z-h)}{2}$$

and by taking the limit $h, d \to 0$, in a way that $2bh = d^2$ where b is a constant, we arrive at a partial differential equation

$$\frac{\partial \eta}{\partial z} = b\frac{\partial^2 \eta}{\partial x^2}.$$

This is a diffusion equation, but with the time replaced by the vertical coordinate z. For the case of a point orifice, which would be represented by an initial condition of the form $\eta(x,0) = \delta(x)$, we would obtain the result

$$\eta(x,z) = \frac{1}{\sqrt{4\pi bz}} e^{x^2/4bz}. \tag{2.1}$$

The density of voids at a point is proportional to the vertical velocity, and hence the void model predicts Gaussian velocity profiles, with a width proportional to $\sqrt{z}$. This has been verified in several experimental studies [84, 83, 114, 27], and appears in good agreement with figure 1-1. The result depends on a single parameter b which is derived from geometrical considerations alone, and it controls the width of the velocity profile.

As described by Nedderman and Tüzün [95], this prediction for the velocity field $\mathbf{v}$ can also be derived entirely from a continuum standpoint, from two reasonable physical postulates:

- The material is incompressible, so that $\nabla \cdot \mathbf{v} = 0$.
- Particles drift towards regions of higher vertical velocity, so that $v_x = -b\frac{\partial v_z}{\partial x}$.

Thus

$$\begin{aligned} 0 &= \nabla \cdot \mathbf{v} \\ &= \frac{\partial v_x}{\partial x} + \frac{\partial v_z}{\partial z} \\ &= \frac{\partial}{\partial x}\left(-b\frac{\partial v_z}{\partial x}\right) + \frac{\partial v_z}{\partial z} \end{aligned}$$

from which we see that

$$\frac{\partial v_z}{\partial z} = b\frac{\partial^2 v_z}{\partial x^2}$$

which is equivalent to the void model continuum prediction.

2.4 Diffusion in the void model

The void model was originally proposed to explain the mean flow in granular drainage, but the microscopic mechanism proposed in figure 2-1 also tells us how particles will move, and allows us to make predictions about the amount of mixing. In this section, we provide a mathematical description of this subject, comparing with experimental results of Choi [28]. Choi's third result about diffusion being proportional to distance dropped is immediately satisfied, since mixing occurs only in response to particle flow, and has no explicit time scale, but it remains to be seen whether the other two results remain true.

Continuing the analysis of the previous section, we introduce a quantity $P_N(M)$ to be the probability that a particle is at horizontal location M when it is in the Nth layer. We assume that the voids do not interact, and pass through the material independently of one another. A particle at (M, N) can move to either $(M-1, N-1)$ or $(M+1, N-1)$ depending on the direction the first void to arrive at (M, N) comes from. Thus the ratio of the probability of moving left to moving right is equal to the ratio of the probability of a void being at $(M-1, N-1)$ to being at $(M+1, N-1)$,

and hence

$$\begin{aligned}
\mathbb{P}((M,N) \to (M-1,N-1)) &= \frac{V_{N-1}(M-1)}{V_{N-1}(M-1)+V_{N-1}(M+1)} \\
&= \frac{V_{N-1}(M-1)}{2V_N(M)} \\
\mathbb{P}((M,N) \to (M+1,N-1)) &= \frac{V_{N-1}(M+1)}{V_{N-1}(M-1)+V_{N-1}(M+1)} \\
&= \frac{V_{N-1}(M+1)}{2V_N(M)}.
\end{aligned}$$

Thus we can write down a recurrence relation of the form

$$P_{N-1}(M) = P_N(M-1)\frac{V_{N-1}(M)}{2V_N(M-1)} + P_N(M+1)\frac{V_{N-1}(M)}{2V_N(M+1)}.$$

Introducing a continuum analogue ρ such that $P_N(M) = \rho(dM, hN)$ gives

$$2\rho(x, z-h) = \eta(x, z-h)\left(\frac{\rho(x-d,z)}{\eta(x-d,z)} + \frac{\rho(x+d,z)}{\eta(x+d,z)}\right).$$

The fastest way to find a continuum equation for ρ is to introduce the quantity $\sigma = \rho/\eta$. We see that

$$2\sigma(x, z-h) = \sigma(x-d,z) + \sigma(x+d,z)$$

and thus in the continuum limit $d, h \to 0$ with $2bh = d^2$, we obtain

$$-\frac{\partial \sigma}{\partial z} = b\frac{\partial^2 \sigma}{\partial x^2}.$$

Substituting back for σ gives

$$\begin{aligned}
\frac{\rho_z}{\eta} - \frac{\rho\eta_z}{\eta^2} &= -b\frac{\partial}{\partial x}\left(\frac{\rho_x}{\eta} - \frac{\eta_x\rho}{\eta^2})\right) \\
&= -b\left(\frac{\rho_{xx}}{\eta} - \frac{2\eta_x\rho_x}{\eta^2} - \frac{\eta_{xx}\rho}{\eta^2} + 2\frac{\eta_x^2\rho}{\eta^3}\right).
\end{aligned}$$

Using the continuum equation $\eta_z = b\eta_{xx}$ gives

$$\begin{aligned}\rho_z &= -b\rho_{xx} + 2b\left(\frac{\eta_x\rho_x}{\eta} + \frac{\eta_{xx}\rho}{\eta} - \frac{\eta_x^2\rho}{\eta^2}\right) \\ &= 2b\frac{\partial}{\partial x}\left(\frac{\rho\eta_x}{\eta}\right) - b\rho_{xx}.\end{aligned}$$

Thus the particle probability density follows an advection-diffusion equation. The amount of advection is proportional to the logarithmic derivative of η, meaning that particles preferentially drift towards regions of faster flow, which appears reasonable. However, this equation also predicts the diffusion of a particle is controlled by the same constant, b, as the diffusion of the voids. This appears to be at odds with Choi's experimental results, since we would expect the particles to diffuse on a much smaller length scale than the voids.

This can be seen more clearly by considering the explicit case of drainage from a point orifice, where η takes the form of equation 2.1. In that case, the advection is given by $\eta_x/\eta = -x/2bz$ and the above PDE becomes

$$\rho_z = -2\frac{\partial}{\partial x}\left(\frac{\rho x}{2z}\right) - b\rho_{xx}.$$

It is natural to consider the initial condition $\rho(x, z_0) = \delta(x - x_0)$, corresponding to a particle initially located at (x_0, z_0). In section A.1 of appendix A, two methods are provided to show that the exact solution of this equation is

$$\rho(x,z) = \frac{1}{\sqrt{4\pi bz\left(1-\frac{z}{z_0}\right)}}\exp\left(\frac{-(x-\frac{x_0 z}{z_0})^2}{4bz\left(1-\frac{z}{z_0}\right)}\right). \tag{2.2}$$

Thus the solution is always a gaussian, with variance $bz(1 - z/z_0)$ and mean x_0z/z_0. A typical solution is shown in figure 2-2. The mean drifts linearly towards the orifice. The variance is initially zero when $z = z_0$, corresponding to the delta function initial condition. It then increases to a maximum value of $bz_0^2/4$ when $z = z_0/2$. As $z \to 0$, the variance begins to decrease, since the particle must exit from the point orifice.

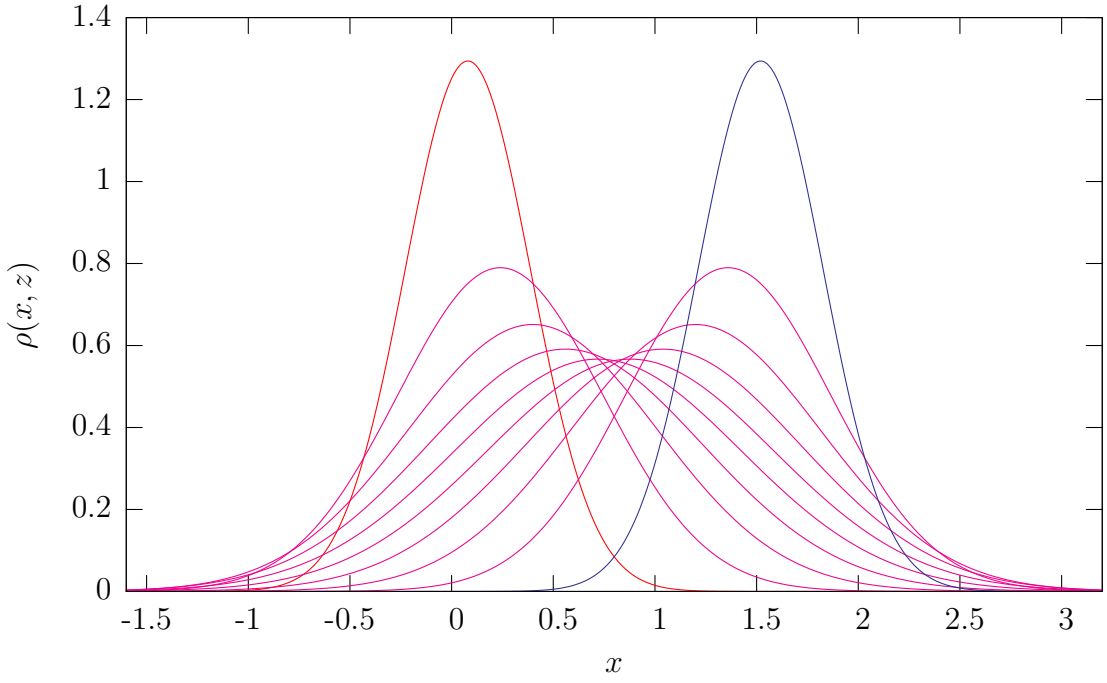

Figure 2-2: A typical solution to equation 2.2 using parameters $x_0 = 1.6$, $z_0 = 2$, and $b = 1$. Cross-sections are plotted for $z = 0.1$ (shown in blue), $z = 0.3, 0.5, \ldots, 1.7$ (shown in magenta), and $z = 1.9$ (shown in red).

The crucial result of figure 2-2 is that the particle's probability density will attain a width comparable of that of the mean flow, which is completely at odds with the experimental results that predict a particle's diffusion should be much more localized. In the experimental geometry of figure 1-1, the void model would predict that the colored bands would rapidly mix in the flowing region, when in fact they remain coherent even after a large amount of flow has taken place. We are therefore led to the conclusion that while the void model predicts a reasonable mean flow, it is based on a flawed model of the microscopic physics.

2.5 The spot model

Although the void model microscopic mechanism appears unrealistic, the spreading gaussian velocity profiles seen in experiment suggest that the notion of a diffusing quantity may still be worthwhile. It is therefore natural to ask if one could formulate an alternative model that would still predict the diffusing velocity profiles, but would have a better microscopic basis. Reconsidering the void mechanism, it is clear that one of its main problems is that particles are allowed to move too independently of one another: after a single void has passed, a particle will lose two of its original neighbors, which seems at odds with the persistent particle cages seen in experiment. In reality, a single particle in a granular packing is strongly constrained by its neighbors, and if it is going to move, it must do so *cooperatively*.

Based on these ideas, Bazant formulated the spot model [17] shown in figure 2-3(a). In this model, which can be applied in both two or three dimensions, particles are no longer constrained to lie on a lattice. They move in response to spots, shown by the blue circle, which correspond to a small amount of excess free volume spread across several particle diameters. When the spot is displaced by an amount, it causes a correlated motion in the opposite direction of all particles within range. In the simplest model, the displacement of each particle $\Delta\mathbf{x}_p$ is a fraction of the spot's displacement $\Delta\mathbf{x}_s$, so that $\Delta\mathbf{x}_p = -w\Delta\mathbf{x}_s$. The diameter of the spot D corresponds to a length scale characteristic of local particle rearrangement. Experimental evidence

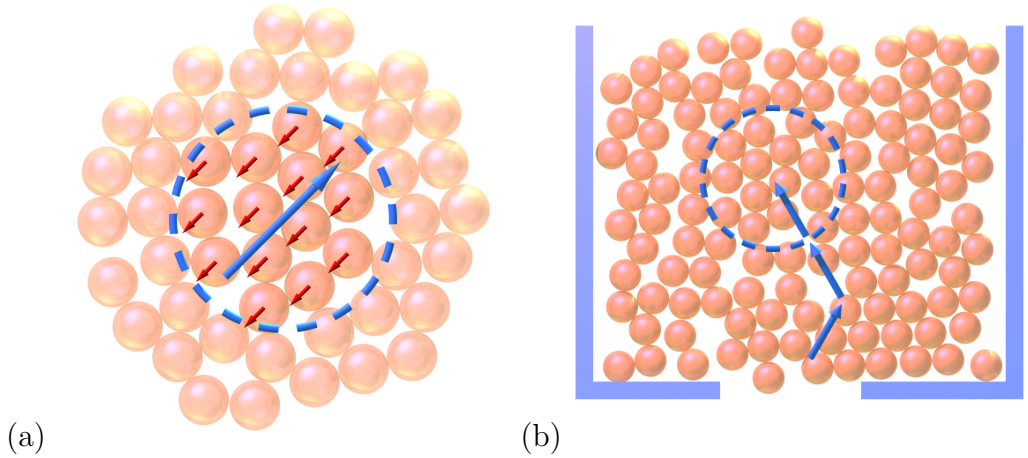

(a) (b)

Figure 2-3: The blue circle shown in (a) represents a spot, which corresponds to a small amount of free volume spread across spread across several particle diameters. When the spot is displaced by the blue arrow, it causes a small, correlated motion in the opposite direction of all particles within range, as shown by the red arrows. The size of the particle displacement is typically on the order of a hundredth of a particle diameter, but is shown as much larger here for clarity. Granular drainage is modeled by imagining a container full of off-lattice particles, and then injecting spots at the orifice, which propagate upwards according to a random walk.

suggests that a scale of three to five particle diameters is important in dense granular flow, and we typically take D in this range. In granular drainage, particles will only flow from an orifice of diameter $\sim 3d - 5d$, since for smaller orifices, particles will tend to jam. Granular flows often exhibit boundary layers of slower velocity with a length scale in this region.

In a typical three-dimensional monodisperse granular material made up of spherical particles, the packing fraction is approximately $\phi = 60\%$. For $D = 5d$, the number of particles that are influenced by a single spot is therefore

$$N = \frac{\frac{4}{3}\pi\left(\frac{5d}{2}\right)^3 \times 60\%}{\frac{4}{3}\pi d^3} \approx 75.$$

When the spot moves, it causes N particles to move by a fraction w in the opposite direction, and thus it makes sense to define the volume of the spot as

$$V_s = NwV_p$$

where V_p is the particle volume. When spots enter a region, they cause a drop in the packing fraction, and comparing this drop with experimental data on flowing packings allows us to gain an order-of-magnitude estimate on the volume carried by a spot. In the DEM simulations presented later in this thesis, and in work carried out by other authors [134, 67, 98, 99, 100] it is typical to expect the packing fraction to decrease by 5% from a static to a flowing state. Since spots could overlap, a reasonable estimate is to attribute $\Delta\phi/\phi \approx 1\%$ as an upper bound on the drop attributed with a single spot. In his paper [17], Bazant argued that

$$w = \frac{\Delta\phi}{\phi^2}.$$

However, the current author believes that

$$w = \frac{\Delta\phi}{\phi}$$

is a more appropriate definition, and for a full comparison between these two different perspectives, the reader should refer to section A.2 of appendix A. However, since ϕ is an order one quantity, both approaches would give an order of magnitude estimate for w in the range 10^{-2} to 10^{-3}. Correspondingly, we expect V_s/V_p to have an order of magnitude between 1 and 0.1.

The granular drainage experiment can be modeled by introducing spots at the orifice which propagate upwards according to a random walk, as shown in figure 2-3(b). At each stage, we assume that a spot makes a fixed step Δz_s upwards, and makes a random step $\Delta\mathbf{x}_s$ in the remaining horizontal dimensions. The spot diffusion length is defined by

$$b_s = \frac{\mathrm{Var}(\Delta x_s)}{2d_h\,|\Delta z_s|}$$

where d_h is the number of horizontal dimensions. The particle diffusion length is given by

$$b_p = \frac{\mathrm{Var}(\Delta x_p)}{2d_h\,|\Delta z_p|} = \frac{w^2\,\mathrm{Var}(\Delta x_s)}{2d_h w\,|\Delta z_p|} = w b_s$$

and thus we see that the spot model predicts that particle diffusion occurs on a scale

approximately two to three orders of magnitude smaller than the spot diffusion. The spot model therefore appears to qualitatively explain all the experimental features seen by the Choi experiments. Like the void model, it predicts parabolic velocity profiles, and that diffusion will be driven by geometry, but it also successfully captures the strongly correlated motion and small diffusion of the particles. The spot model remains simple enough for mathematical analysis, and as described by Bazant [17], makes a number of predictions about particle motion in granular flow. Perhaps most importantly, it suggests that particles in a granular flow should exhibit spatial velocity correlations. Consider the spatial velocity correlation tensor

$$C^{\alpha\beta}(\mathbf{r}_1, \mathbf{r}_2) = \frac{\left\langle v^\alpha(\mathbf{r}_1) v^\beta(\mathbf{r}_2) \right\rangle}{\sqrt{\left\langle v^\alpha(\mathbf{r}_1) \right\rangle \left\langle v^\beta(\mathbf{r}_2) \right\rangle}}$$

where $v^\alpha(\mathbf{r})$ is the α component of instantaneous velocity at position $\mathbf{r}$, and the angular brackets refer to time averages. We can simplify this by considering a single horizontal velocity component u, and exploiting translational and rotational symmetry to write

$$C(\mathbf{r}) = C(r) = \frac{\langle u(\mathbf{0}) u(\mathbf{r}) \rangle}{\sqrt{\langle u(\mathbf{0}) \rangle \langle u(\mathbf{r}) \rangle}}.$$

In the spot model, the velocities of two nearby particles will be correlated if they fall within the displacement of the same spot. If we assume that spot displacements occur randomly, then $C(r)$ will be given by the probability that a spot influencing a particle at $\mathbf{0}$ influences a particle a distance r away as well. The correlations in two dimensions are given by

$$C(r) = \begin{cases} 1 - \frac{r\sqrt{1-(r/2R)}}{\pi R} - \frac{2}{\pi} \sin^{-1}\left(\frac{r}{2R}\right) & \text{for } r < 2R \\ 0 & \text{for } r \geq 2R \end{cases}$$

and in three dimensions by

$$C(r) = \begin{cases} \left(1 + \frac{r}{4R}\right)\left(1 - \frac{r}{2R}\right)^2 & \text{for } r < 2R \\ 0 & \text{for } r \geq 2R. \end{cases}$$

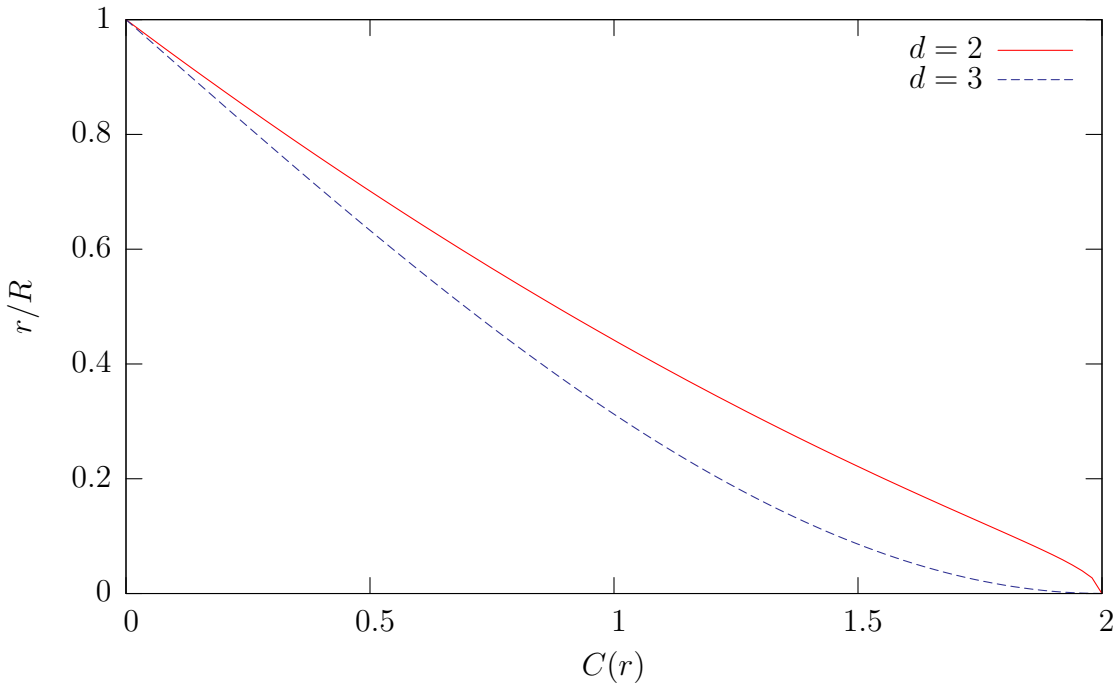

Figure 2-4: Plots of the theoretical velocity correlations for the uniform spot potential, in two and three dimensions.

These are plotted in figure 2-4 – if the spot model for particle motion exists, then we would expect to see correlations that have a decay length similar to the spot size. Choi searched for these correlations in the tall silo experimental geometry, and the results are shown in figure 2-5 – we do indeed see correlations decaying on an intermediate scale, although local ordering effects are visible. Since the measurement is taken at the wall, many particles locally arranged in a hexagonal crystalline formation. It would be expected to see higher correlations at values of $1d$, $\sqrt{3}d$, and $2d$ since particles which are locally crystallized will have very strong geometrical constraints with their neighbors.

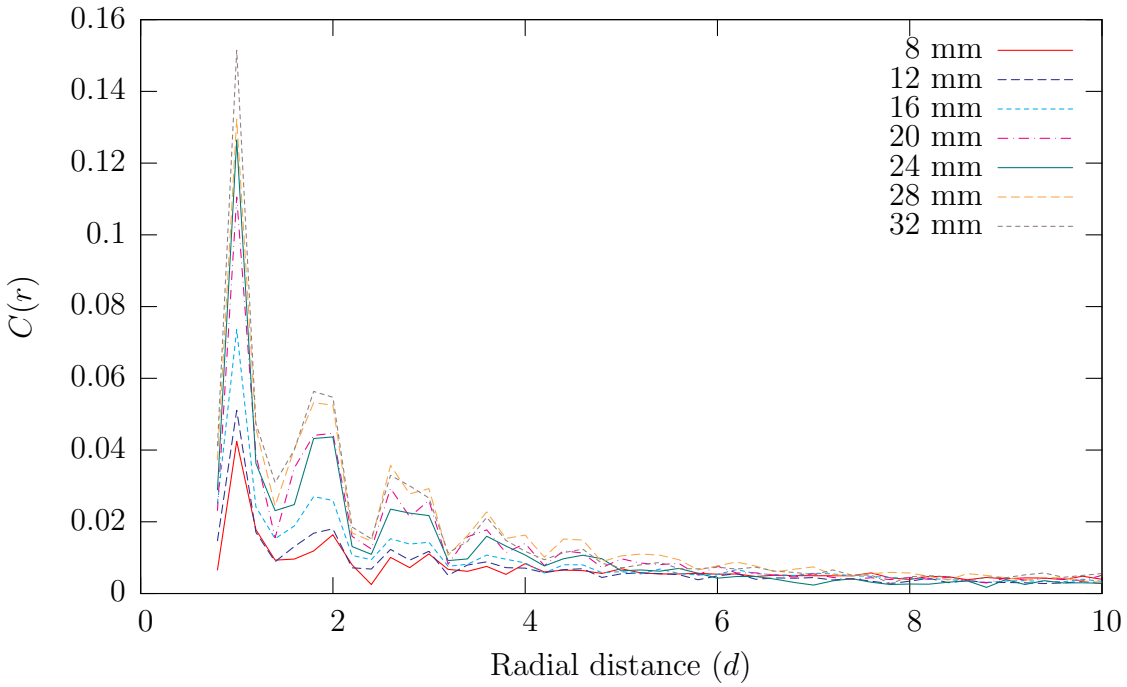

Figure 2-5: Velocity correlations in the x direction for Choi's granular drainage experiment.

2.6 Discrete Element Simulation and comparison to experiment

While the experiments of Choi provided important insight into granular diffusion, they have one unavoidable drawback: measurements can only be made of the particles at the glass wall. It is hoped that the low friction coefficient of the glass causes particles near the wall to have similar velocities to those in the bulk, but the presence of the wall will have a significant effect on the packing geometry. It is a well-known fact that packing properties of monodisperse spherical particles are very different in two and three dimensions, with particles in two dimensions exhibiting a much larger tendency to form hexagonal crystalline arrangements. It is therefore natural to ask whether the velocity correlations of figure 2-5 are also present in the bulk.

Carrying out large-scale particle-based simulations would therefore be particularly advantageous, since they can provide us with complete three-dimensional information

about packing arrangements. However, simulating granular materials is a difficult task that has only become computationally practical in the last five to ten years. The most obvious problem is the large number of particles that must be considered, even for a fairly small system. The simulation of particles made of hard materials is also very difficult. One approach is to model the particles as inelastic hard spheres, which interact according to a coefficient of restitution. Simulations of this type are event-driven, with the code calculating the next time until a collision occurs between any two particles in the system. Simulation studies have used this method to look at dense granular flows [48, 46, 47], but they have only been able to consider a fully flowing packing, since some particles in static regions undergo "inelastic collapse" where they will want to undergo an infinite number of collisions in over a finite time.

An alternative model is to treat particles as interacting via a spring with a very high spring constant, and then simulate the resulting positions using a traditional fixed timestep – this is referred to as a Discrete Element Method (DEM). The high spring constant means that the equations are stiff, and require a very small timestep, but the approach has become computationally feasible in the past five years, particularly since the interactions between particles are short range, which allows the code to be efficiently parallelized, to take advantage of the recent and growing trend toward multicore and parallel computing. In 2001 and 2002 Grest, Landry, Plimpton, and Silbert at Sandia National Laboratories implemented a parallel DEM code called GranFlow, written in Fortran 90. This has since been incorporated into the Large-scale Atomic/Molecular Massively Parallel Simulator (LAMMPS) [1] which is written in C++. GranFlow is used in the work presented here and in chapters 3 and 5, while LAMMPS was used in the later work of chapters 6 and 7. The code has been very successful and has been used by many groups to study jammed granular packings [124, 125, 71], vibrated granular systems [129], and granular flows [123, 126, 24, 112, 113, 26].

The DEM simulations presented in this thesis use of a common set of conventions. In all cases, we make use of a three co-ordinate system (x, y, z), and make measurements in terms of the particle diameter d. The simulations act under gravity

g, pointing in the negative z direction. From this, a natural simulation time unit $\tau = \sqrt{d/g}$ is defined. Given a particle diameter d in terms of a physical length, we can make comparisons to experiment by computing τ in terms of seconds. For example, to make the correspondence to Choi's simulation data, where $d = 3$ mm, we have

$$\tau = \sqrt{\frac{3 \text{ mm}}{9.81 \text{ ms}^{-2}}} = 0.0174 \text{ s}.$$

The mass of the particles is referred to by m, although most of the time, it is not necessary to consider this explicitly, since it is only important in relation to the details of the interaction model.

The contact model is based on that developed by Cundall and Strack [32] to model cohesionless particulates. If a particle and its neighbor are separated by a distance $\mathbf{r}$, and they are in compression, so that $\delta = d - |\mathbf{r}| > 0$, then they experience a force $\mathbf{F} = \mathbf{F}_n + \mathbf{F}_t$, where the normal and tangential components are given by

$$\mathbf{F}_n = f(\delta/d)\left(k_n \delta \mathbf{n} - \frac{\gamma_n \mathbf{v}_n}{2}\right) \tag{2.3}$$

$$\mathbf{F}_t = f(\delta/d)\left(-k_t \Delta \mathbf{s}_t - \frac{\gamma_t \mathbf{v}_t}{2}\right). \tag{2.4}$$

Here, $\mathbf{n} = \mathbf{r}/|\mathbf{r}|$. $\mathbf{v}_n$ and $\mathbf{v}_t$ are the normal and tangential components of the relative surface velocity, and $k_{n,t}$ and $\gamma_{n,t}$ are the elastic and viscoelastic constants, respectively. Two different force models are considered: $f(z) = \sqrt{z}$ for Hertzian particle contacts and $f(z) = 1$ for Hookean particle contacts. $\Delta \mathbf{s}_t$ is the elastic tangential displacement between spheres, obtained by integrating tangential relative velocities during elastic deformation for the lifetime of the contact. If $|\mathbf{F}_t| > \mu |\mathbf{F}_n|$, so that a local Coulomb yield criterion is exceeded, then $\mathbf{F}_t$ is rescaled so that it has magnitude $\mu |\mathbf{F}_n|$, and $\Delta \mathbf{s}_t$ is modified so that equation 2.4 is upheld.

Particle-wall interactions are treated identically, but the particle-wall friction coefficient is set independently. For the current simulations we set $k_t = \frac{2}{7} k_n$, and choose $k_n = 2 \times 10^5 mg/d$. While this is significantly less than would be realistic for glass spheres, where we expect $k_n \sim 10^{10} mg/d$, such a spring constant would be prohibitively computationally expensive, as the time step must have the form $\delta t \propto k_n^{-1/2}$

for collisions to be modeled effectively. Previous simulations have shown that increasing k_n does not significantly alter physical results [71]. We make use of a timestep $\delta t = 10^{-4}\tau$ The normal damping term is set to $\gamma_n = 50\sqrt{g/d}$, and the tangential damping γ_t is set to zero for Hookean contacts, and equal to γ_n Hertzian contacts.

Unless stated otherwise, the Hertzian contact model was employed for the simulations presented in this thesis. The Hertzian model is motivated by the fact that when spheres come into contact, the amount of volume by which they overlap grows faster than linearly, and thus the force should grow faster than linearly too. While it may be more physically desirable to use this model, it was found here and by others that it can sometimes lead to simulation artefacts, such as elastic "breathing modes" propagating through the simulations. Therefore, in cases where precise information about microscopic motion was needed, the Hookean contact model was used.

The simulations were carried out on MIT's Applied Mathematics Computational Laboratory (AMCL), a Beowulf cluster consisting of 16 nodes each with a dual processor. During the simulations, snapshots of all particle positions were saved every 2τ, corresponding to 20,000 integration timesteps. The LAMMPS code is written in C++ and can be run on any number of processors, by decomposing the computational domain into a rectangular grid of subdomains of equal size. Interactions between particles in neighboring domains are handled using message passing. For problems where the particles are split evenly between the subdomains, the LAMMPS code scales very well, and doubling the number of processors can frequently result in almost a doubling of speed.

2.7 Comparison to experiment: velocity profiles

One of the first tasks was to validate the simulations against experiment. A simulation was run where the parameters were chosen as closely as possible to the experiments of Choi. Based on a particle diameter of $d = 3$ mm, a quasi-two dimensional container with walls at $x = \pm 33\frac{1}{3}d$, $y = \pm 4\frac{1}{6}d$ and $z = 0d$ was created. The interparticle friction coefficient and the front and back wall friction coefficients were set to be 0.2

to approximately model a glass/glass contact. The side and base walls were modeled using a friction coefficient of 0.3, to correspond with a glass/metal contact.

Throughout this thesis, initial packings were created by pouring the particles into a container. To pour the particles, an insertion region over a range $z_{\text{low}} < z < z_{\text{high}}$ is created. When the simulation starts, the insertion region is filled with particles up to a packing fraction of 10%, by making random insertion events. When a particle is inserted at a location (x, y, z), its horizontal velocity components are set to zero, and its vertical velocity is set to $v_z = \sqrt{2g(z_{\text{high}} - z)}$ corresponding to the speed it would have attained had it fallen from rest at z_{high}. The x and y coordinates of the insertion events are chosen uniformly. The z coordinates are chosen so that v_z is uniformly distributed. The simulation continues, and when the particles at the top of the insertion region have dropped to z_{low}, another insertion event is carried out, and another batch of particles is created. This process creates the effect of a steady stream of particles being introduced at a constant rate from $z = z_{\text{high}}$. Once the required number of particles has been introduced, the batch insertion events are terminated and simulation continues. Typically, the system is simulated for a long time after the last insertion event to ensure that the particles have come to rest.

For the current simulation, 174,249 particles were poured using $z_{\text{low}} = 333\frac{1}{3}d$ and $z_{\text{high}} = 366\frac{2}{3}d$, to fill the container to a height of $264d$. The drainage process was then carried out by opening a slit of width $5\frac{1}{3}d = 16$ mm in the center of the container base. Figure 2-9(a) shows a snapshot of the system after a certain amount of flow has taken place.

Figure 2-6 shows a comparison of the vertical velocity profiles between the simulation and the experiment for two different cross sections, where the simulation data has been converted into physical units. The match is extremely good, both in terms of the overall shape, and the total flow rate.

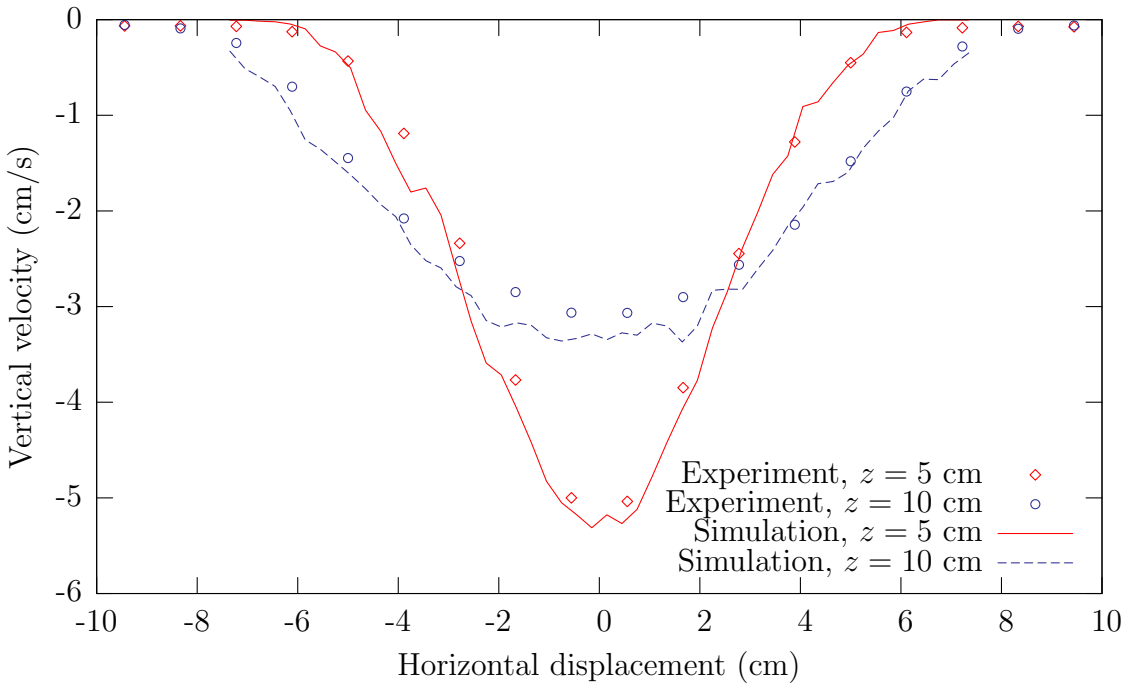

Figure 2-6: Comparison of horizontal profiles of downwards velocity between DEM simulation and Choi's granular drainage experiment.

2.8 Velocity correlations

Velocity correlations were also compared to experiment. However, since searching for velocity correlations requires a large amount of data over which to gain a statistical average, it was chosen to carry out these measurements in a smaller simulation, which could be run for longer. Since the velocity correlations are a local effect, they should not be dependent on the precise geometry used. Walls were placed at $x = \pm 25d$, $y = \pm 4d$, and $z = 0$, and approximately 55,000 particles were poured from a height of $170d$ to fill the container up to a height of $110d$. For this simulation, $\mu = \mu_w = 0.5$, which is a commonly-used figure when running the DEM code. To create the flow, a circular hole of radius $4d$ centered on $x = y = 0$ was opened. The simulation was treated as periodic in the z direction, so that particles falling out at $z = 0$ would be reintroduced at $z = 181d$, creating a continuous flow. For making measurements of velocity correlations, there were two regions of interest: the particles near the front wall, whose centers satisfy $y > 3.2d$, and the particles in the central slice, whose

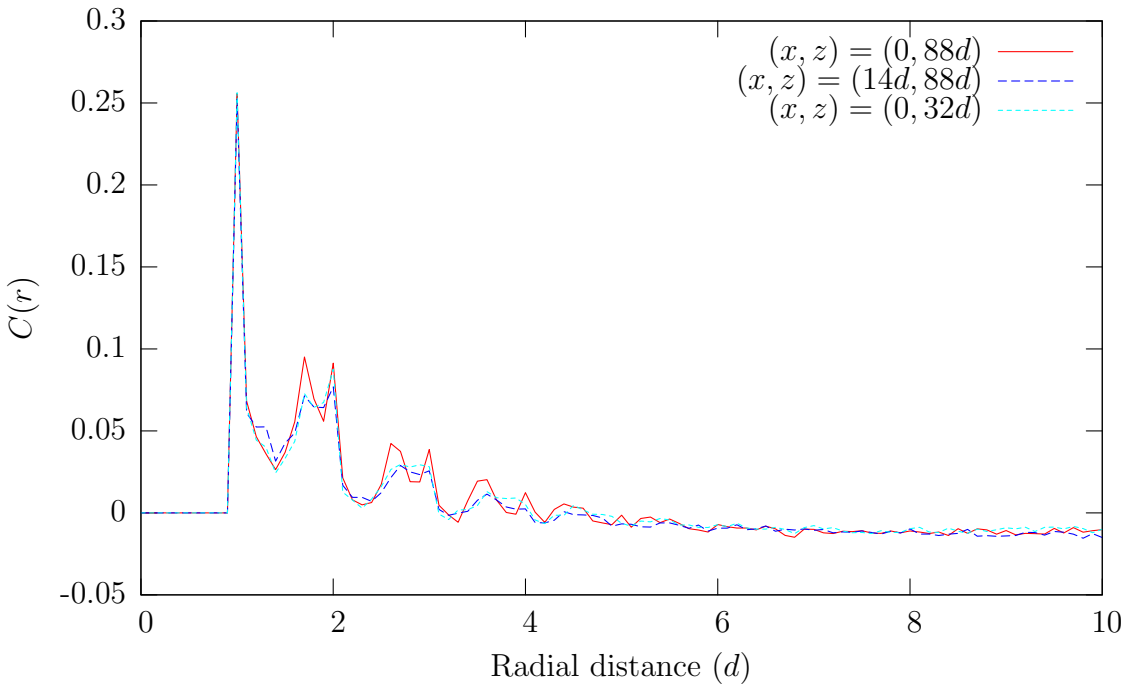

Figure 2-7: Velocity correlations at the boundary of a DEM simulation, computed in square test regions of side length $16d$ centered at three different locations. The functional form closely matches Choi's experimental data.

centers satisfy $|y| < 2d$.

The simulation was first run to calculate a mean background velocity, by averaging the instantaneous particle velocities in cubic grid of side length $1d$. Since the system is symmetric in the x and y directions, this can be exploited to improve the accuracy of the computed field by a factor of four. Once computed, the mean velocity at a specific point could be found by linearly interpolating from the eight neighboring cells.

A run was then carried out to calculate velocity correlations in three different test squares of side length $16d$, both at the boundary and in the bulk. One test square was centered on $(x, z) = (0, 88d)$, high in the central region of the container where the flow is approximately uniform. In addition, an off-axis test region at $(x, z) = (14d, 88d)$, plus a lower test region at $(x, z) = (0, 32d)$, were tested. We computing the correlation statistics, the previously computed mean velocity field is subtracted, so that the results capture the local variations.

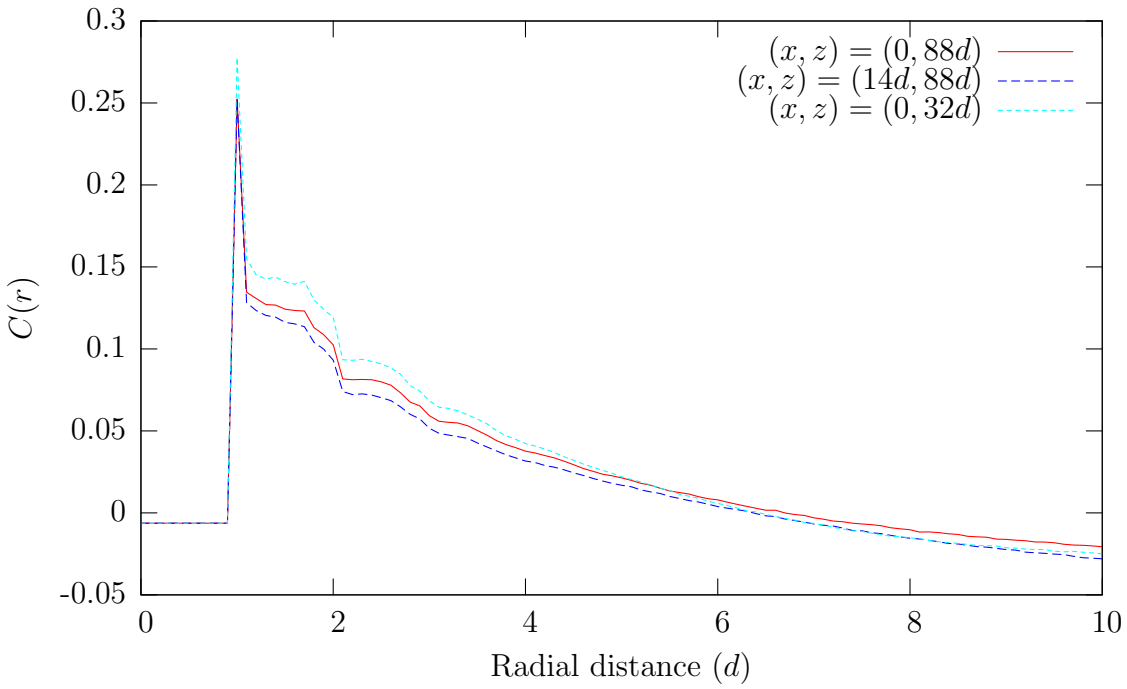

Figure 2-8: Velocity correlations at in the bulk slice of a DEM simulation, computed in square test regions of side length $16d$ centered at three different locations.

Figure 2-7 shows the computed velocity correlations in the boundary slice. In all three test regions, the functional form closely matches Choi's correlation data, and captures the same lattice effects. Along with the mean flow data presented in the previous section, this provides further evidence that the simulations are a faithful reproduction of experiment.

Figure 2-8 shows the velocity correlations in the same three test regions, but looking in the boundary slice. We see a decaying correlation, similar in magnitude and decay length to the result at the boundary. However, as expected, the local ordering effects are much weaker: there is a strong signature at $r \approx d$, corresponding to particles in contact, but for larger separations, the computed function is smooth. We can thus infer that the spot model for correlated particle motion appears to be a general phenomenon in dense granular flow, that holds both at the boundaries and in the bulk.

Furthermore, we see that for the three different test regions, the structure of the correlations, and their overall correlation length, are remarkably similar. It suggests

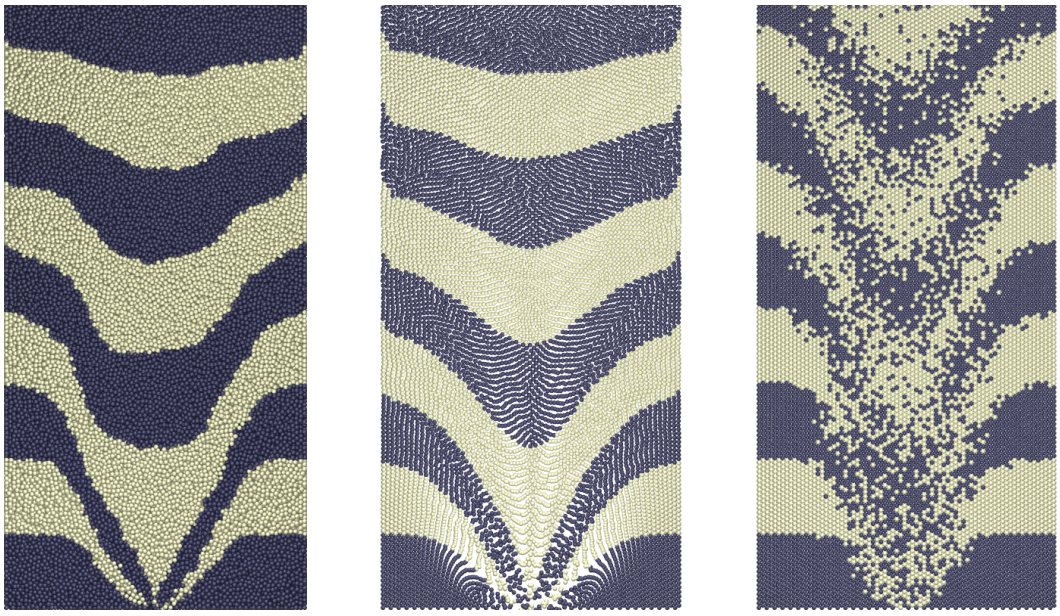

Figure 2-9: A snapshot of a DEM simulation of granular drainage in Choi's experimental geometry (left), compared to 2D simulations of the spot model (center) and void model (right), all shown at $t = 200\tau$. The total flow rates of the spot and void models were calibrated to match DEM. The spot size and diffusion rate were calibrated using the same process discussed in the following chapter.

that the process by which particles flow is similar throughout the container, and thus it may be a reasonable approximation to treat spots as having a fixed radius. The evidence also supports the later work in chapter 7 where we make use of a fixed-size granular element throughout a granular simulation.

2.9 Comparison of DEM, spot and void simulations

Figure 2-9 shows a snapshot from the DEM simulation of Choi's geometry, and compares it to simulations in two dimensions of the spot and void models. The simulation of the void model shows an unrealistic amount of diffusion, and the interfaces between the colored layers is rapidly smeared out, agreeing with the theoretical predictions of previous sections.

The spot model is a significantly better match, and the interfaces show an amount

of diffusion which matches that seen in the DEM. Although there are some discrepancies, the spot model tracks the overall flow profile reasonably well. However, this image shows one very significant drawback of the spot model as a simulation technique: looking near the orifice, we see that many particles have become overlapped. This is seen more clearly in the progression of close-up images shown in figure 2-10.

This should be expected, since in the spot model of microscopic particle motion shown in figure 2-3(a), there is nothing explicitly enforcing packing constraints. In this figure, it is clear that particles near the edge of the spot may end up overlapping with their neighbor by small amount, which may become larger and larger as more spots pass through. The above figures show that the problem is largest in areas undergoing a large amount of shear.

This represents a significant hurdle in using this model as a simulation technique. The basic model of particle motion is good for a mathematical analysis, and can make reasonable predictions about the statistics of motion of a single particle, but it cannot be used to generate reasonable predictions about the motions of all particles in the system. This problem is not present in the void model, since by enforcing particles to lie on a lattice they will never become overlapped.

Several modifications to the spot model microscopic motion were proposed to solve this problem. Making the spots carry out persistent random walks [55], adding a diffusive term to the particle motion, and adding a rotation were all tested. While some of these techniques mitigated the effect of the overlap, none of them were able to create fully valid packings. This is unsurprising, since it should be expected that in order to maintain a valid packing of particles, one must take into account the geometry of packing structure at a local level. In the following chapter, a solution to the density problem taking this approach will be described.

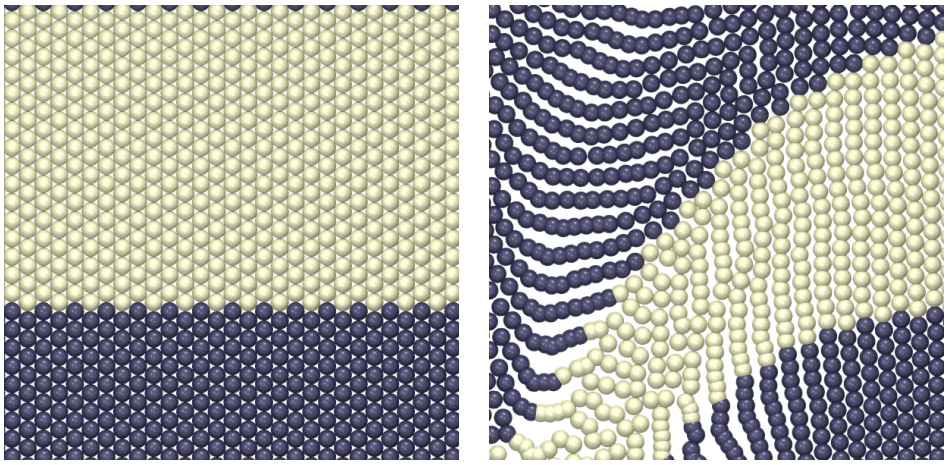

$t = 0$ s

$t = 1.74$ s

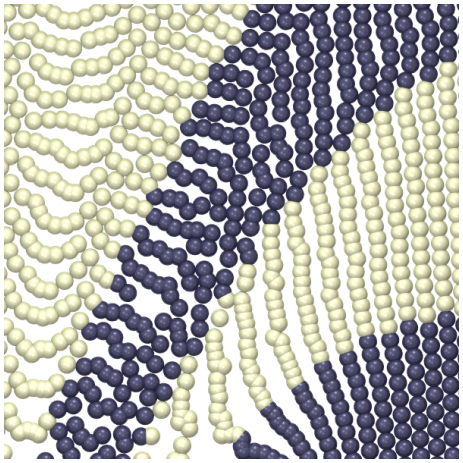

$t = 3.50$ s

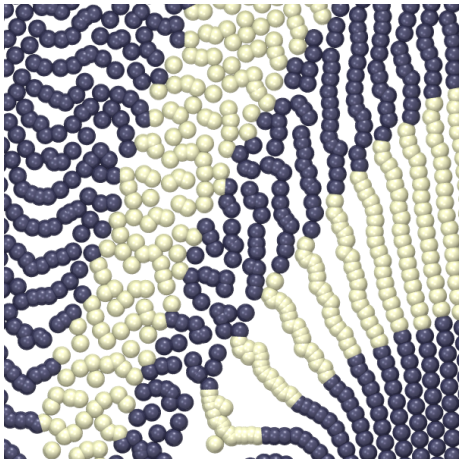

$t = 5.25$ s

Figure 2-10: Four close-up snapshots of the spot model simulation from figure 2-9, taken in the region $0 < x < 7.5$ cm, 2.5 cm $< z < 10$ cm. The overlapping particles highlight the "density problem" of the basic spot model.