

Physical test of a particle simulation model in a sheared granular system

Chris H. Rycroft*

*Department of Mathematics, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA and
Department of Mathematics, University of California, Berkeley, CA 94720, USA*

Ashish V. Orpe[†]

*Department of Physics, Clark University, Worcester, MA 01610, USA and
Chemical Engineering Division, National Chemical Laboratory, Pune 411008, India*

Arshad Kudrolli[‡]

Department of Physics, Clark University, Worcester, MA 01610, USA

We report a detailed comparison of a slow gravity-driven sheared granular flow with a discrete-element simulation performed in the same geometry. In the experiments, grains flow inside a silo with a rectangular cross-section, and are sheared by a rough boundary on one side and smooth boundaries on the other sides. Individual grain position and motion are measured using a particle index-matching imaging technique where a fluorescent dye is added to the interstitial liquid which has the same refractive index as the glass beads. The simulations use a Cundall–Strack contact model between the grains, using contact parameters that have been used in many other previous studies, and ignore the hydrodynamic effects of the interstitial liquid. Computations are performed to understand the effect of particle coefficient of friction, elasticity, contact model, and polydispersity on mean flow properties. We then perform a detailed comparison of the particle fluctuation properties as measured by the displacement probability distribution function and the mean square displacement. All in all, our study suggests a high level of quantitative agreement between the simulations and experiments.

I. INTRODUCTION

Granular flow is important in a number of industrial and geophysical phenomena but a well-established description of its properties does not exist. Although typical granular systems are composed of many particles, a hydrodynamic theory and a length scale other than the particle diameter has not been consistently identified, prompting a large amount of research interest in the past two decades and a variety of theoretical advances [1–6]. In addition to the development of new experimental techniques, computational models have been an invaluable tool in understanding granular flow, because they can provide complete information about a particle system, without the experimental difficulties in measurement.

However, simulating granular flows is difficult due to the problems of correctly capturing the physics of strongly-interacting hard particles. In one approach, particles are treated as being perfectly hard, and interact with each other via inelastic collisions, characterized by a coefficient of restitution. Often, simulations of this type are carried out with an event-driven approach [7, 8], advancing the simulation clock forward until the next time of a particle collision. These simulations are particularly well-suited for the study of fluidized granular materials at moderate density, although they can be more problematic for denser systems with static regions of particles, where “inelastic collapse” [9, 10] can take place and the time between successive particle collisions can tend to zero. However, some techniques have been developed to

avoid this issue, such as by making the collisions become perfectly elastic if they occur within some small characteristic timescale [11].

Another simulation approach is to model the particles to be slightly soft, and able to overlap slightly and interact via a contact force model. The contact force models employed typically include a stiff normal repulsion, and potentially other forces such as friction and viscous damping terms. They can be carried out using fixed-timestep integration via the discrete-element method (DEM) [12]. This approach can handle static packings and long-lasting particle contacts, but the stiffness of the equations requires a very small timestep to integrate accurately, making it computationally intensive, often requiring parallel simulation for problems of a practical size. Consequently, some approximations and simplifications are made in the contact models. It has been the authors’ experience that this can often lead to skepticism in the presentation of simulation results, as to what extent the contact model captures the underlying physics.

In this paper, we attempt to address these issues by directly comparing a three-dimensional (3D) DEM simulation to a gravity driven granular flow, sheared by a rough boundary wall inside a silo. Boundary generated shear is fairly common in granular flows and gives rise to a localized shear region, which is excellent for testing flow in the bulk, as well as boundary conditions. This system, presented in detail in Subsec. II A, was used recently to measure particle velocity auto-correlations [13] with an index-matching technique developed for measuring the motion of individual particles inside dense flows [14–16]. In that study, it was assumed that the flow and strain rate were small enough that the drag and lubrication forces caused by the presence of the interstitial liquid on the particle motion was negligible compared with the gravitational force acting on the particles. Thus, a comparison of the

*Electronic address: chr@math.berkeley.edu

[†]Electronic address: av.orpe@ncl.res.in

[‡]Electronic address: akudrolli@clarku.edu

simulations with the experiments is not only a physical test of the numerical model but also a test of this hypothesis.

For the DEM simulation, we make use of a modified version of a contact model originally employed by Cundall and Strack [17] for the simulation of cohesionless particulates, featuring a normal elastic interaction, viscoelastic terms, history-dependent tangential forces and a Coulomb friction criterion. This model is simulated by the Large-scale Atomic/Molecular Massively Parallel Simulator (LAMMPS) developed at Sandia National Laboratories [18]. This code is well-developed and supports a wide variety of particle force models, handling very large systems with good parallel efficiency. For granular materials, it has been used to study avalanching flow [19–21], static granular packings [22–24], granular drainage [25–27], granular shear and its relation to constitutive modeling [28–31], and granular segregation [32].

The Cundall–Strack contact model that we employ has several free parameters that determine the relative strength of the different forces, and it is worth noting that many of the studies listed above have used the same values for these. The original choice of these parameters can be traced to Silbert *et al.* [19] in the context of avalanching flow, making use of 2D and 3D studies with up to 24,000 particles, and determining appropriate sizes for elastic and viscoelastic forces for modeling the flowing granular regime. It was also later considered by Landry *et al.* [22] in examining static particle assemblies. Perhaps the largest approximation of these parameters is in the normal spring constant, which is four to five orders of magnitude smaller than what would be realistic for everyday hard materials. Since DEM simulations of this type are usually compared to experimental and theoretical results for dense granular flows of hard materials such as rock or glass, it would be advantageous to have a higher spring constant. However, due to computational limitations it could not be set higher, but it was found that the value used was large enough to capture the physics of the avalanching flow without too many detrimental elastic effects.

While testing the efficacy of any DEM simulation is interesting, we believe it is particularly useful to carry out a detailed analysis of the Cundall–Strack contact model with the commonly used choice of parameters, since the results are of direct relevance to the studies listed above. We note that Ref. [19] concentrates on macroscopic features of flow averaged over many particles, while Ref. [22] considers static, microscopic features. However, there is also much interest in quantities that are both microscopic and dynamic, such as velocity autocorrelations which impact calculations using kinetic theory approach [33–35], and we ask to what extent the contact model can capture such correlations. We also note that the original choice of parameters was made using 24,000 particles, but advances in computer technology mean that simulations of ten to twenty times the number of particles are now considered [25], and we examine whether the contact model remains valid for larger system sizes. By making direct, quantitative comparisons at the microscopic level to the index-matched experiment, we are able to address in detail if the simulation reproduces the essential physics of granular flow.

Our paper proceeds as follows: in Section II we describe the geometry of boundary-generated shear, and discuss in detail the methods used in the index-matching experiment, and the DEM simulation. In Section III we consider a variety of basic comparisons of the DEM simulation to experiment, examining the role of friction, polydispersity, and total flow rate. During this study we found evidence of short-timescale waves of velocity in simulation, and the presence of these is discussed in detail in Section IV. We then consider a detailed quantitative analysis of diffusion and velocity autocorrelations in Section V.

II. METHODS

A. Index-matching experiment

A schematic diagram of the silo with a rectangular cross-section inside which gravity driven granular flow occurs is shown in Fig. 1. The silo is filled with soda lime glass beads with average diameter $d = 1.0$ mm, and density $\rho_{\text{glass}} = 2.5 \times 10^3 \text{ kg m}^{-3}$. The beads exhibit a small amount of polydispersity, with diameters over the range $d \pm 0.1$ mm, with a majority ($\sim 80\%$) over the range $d \pm 0.05$ mm. The sides of the silo are composed of optically smooth transparent glass plates. A layer of the glass beads is glued to one of the sides of the silo in order to shear the flow relative to the other surfaces. The interstitial space between the grains is filled with a liquid with the same refractive index as the glass beads. The liquid has density $\rho_{\text{fluid}} = 1.0 \times 10^3 \text{ kg m}^{-3}$, and viscosity $\nu = 2.2 \times 10^{-2} \text{ kg m}^{-1} \text{ s}$. Side chambers (which are not shown in the schematic diagram) allow the interstitial fluids to redistribute as grains drain from the orifice. A dye added to the liquid is illuminated by a light sheet of thickness less than $0.1d$ and imaged from an orthogonal direction using a 512×480 pixel resolution CCD camera, where 20 pixels corresponds to one d . The particles in an image appear dark against a bright background with a flat intensity profile across each particle. We then make use of convolution procedure [36] to convert each image into a 2D map consisting of bright sharp peaks of intensity corresponding to particle centers which are then obtained using a centroid algorithm [37]. This procedure yields particle position in every image to within a twentieth of a particle diameter. Because of small variations of refractive index within the glass beads and defects, the accuracy with which we can determine the position of the particle diminishes with optical length within the index-matched sample. Therefore, we restrict our data acquisition to a window which is within $30d$ from a side wall. A sequence of images is recorded in the region of interest at a frame rate of 60 Hz and particle trajectories are obtained by comparing the particle centers in consecutive images. The particle trajectory data is further analyzed to obtain mean and fluctuating properties of the flow.

The resulting maximum mean flow velocity of the grains inside the silo in the region of interest is observed to be approximately $0.6d \text{ s}^{-1}$ [13]. At these measured velocities, the Reynolds number $Re \sim v_x d / \nu$ is about 10^{-2} , and the ratio of the viscous drag of the grain to the gravitational force is




FIG. 1: (Color online) Schematic diagram of the experimental apparatus with a typical image of particles observed in a thin slice in the shear plane obtained with a fluorescent particle index-matching technique. The co-ordinate system used is also indicated.

estimated to be less than 10^{-2} . The rough boundary is observed to shear the flow more than the smooth boundary. The resulting strain rate $\dot{\gamma}$ has a maximum of about 0.25 s^{-1} . This maximum strain rate is comparable in magnitude to that in experiments with a plate dragged on a granular bed [15], where the drag friction experienced by the plate was measured to be unchanged when the index-matching interstitial liquid was added. Based on these observations, we anticipate that the interstitial fluid does not affect the motion of the particles and neglect its presence in our simulation model.

B. Discrete-Element simulation

In the simulation, we define lengths in terms of a particle diameter d , and we define a natural mass unit m . Particles have unit density, and thus the mass of a particle is $m_p = 4\pi(0.5)^3 m/3 = 0.524m$. Gravity g acts in the negative x direction. Simulation results can then be expressed in terms of a natural time unit $\tau = \sqrt{d/g}$. If a particle and its neighbor are separated by \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

$$\mathbf{F}_n = f(\delta/d) \left(k_n \delta \mathbf{n} - \frac{\gamma_n \mathbf{v}_n}{2} \right) \quad (1)$$

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

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: $f(\zeta) = \sqrt{\zeta}$ for Hertzian particle contacts and $f(\zeta) = 1$ for Hookean 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 is upheld.

Much of the original choice of the parameters was carried out by Silbert *et al.* [19], and the reader should refer here for a detailed discussion. The normal damping term is set to $\gamma_n = 50\sqrt{g/d}$, and then tangential damping is set to zero for Hookean contacts, and equal to γ_n for Hertzian contacts. To approximate the Poisson ratio of real materials the tangential elastic constant is set to $k_t = \frac{2}{7}k_n$. The largest assumption of the model is the choice of the normal elastic constant, which is set to $k_n = 2 \times 10^5 mg/d$ and is significantly lower than what would be realistic for typical hard materials such as glass, where $k_n = O(10^{10} mg/d)$ would be more reasonable [19, 22, 25]. However, such a constant would be prohibitively computationally expensive, since the timestep required must have the form $\delta t \propto k_n^{-1/2}$. Both Silbert *et al.* [19] and Landry *et al.* [22] discuss that the chosen value of k_n is a reasonable compromise, which is small enough to feasibly simulate, but large enough to avoid the system exhibiting excessive elastic effects. Here, we make use of $\delta t = 10^{-4}\tau$.

The simulations were primarily carried out on the MiniMe64 test cluster at the Lawrence Berkeley Laboratory, featuring 19 dual-core Intel Xeon nodes with a fiber-optic Myrinet interconnection. For a typical simulation considered here featuring 147,000 particles, the MiniMe64 cluster computed one million timesteps in $4\frac{1}{2}$ hours using 24 processors. Additional simulations were carried out on a Mac Pro with two dual-core Intel Xeon processors, where one million timesteps would take 24 hours. The simulation produces text files of all particle positions at fixed intervals, and these were subsequently post-processed to analyze many different aspects of the flow.

The initial packings of particles were created by randomly pouring approximately 147,000 particles at a rate of $201\tau^{-1}$ from a height of $z = 205d$. The particles are introduced up to $t = 740\tau$, after which the system is run until $t = 3000\tau$ to allow them to come to rest. Six different particle models were considered, the details of which are shown in Table I. The main analysis was carried out with model B and unless otherwise stated, the results presented refer to this. The remaining models were used to analyze specific physical effects.

To create a rough wall analogous to that in the experiment, all particles whose centers satisfy $y < 1d$ are frozen in place,




FIG. 2: (Color online) A typical snapshot of discrete-element simulation, showing only those particles in the test region $80d < x < 100d$, $|z - 30d| < 10d$. The white (yellow) particles on the left are frozen in place to form the rough wall at $0 < y < 1d$. The light blue and dark blue particles are identical in physical characteristics, and initially form alternating layers of width $5d$. This snapshot was taken at $t = 124\tau$, after the bulk of the packing has dropped by approximately $5d$. The particles next to the rough wall undergo pronounced shear and form a boundary layer of slower flow.

Model	Contact type	Diameter range	k_n (mg/d)	μ
A	Hookean	d	2×10^5	0.2
B	Hookean	d	2×10^5	0.3
C	Hookean	d	2×10^5	0.4
D	Hookean	d	2×10^6	0.3
E	Hertzian	d	2×10^5	0.3
F	Hookean	$0.95d$ to $1.05d$	2×10^5	0.3

TABLE I: Detailed information about the six particle packings that were created and analyzed in this study.

so that their translational and angular velocities are kept at zero throughout the simulation. In the pouring process, the particles next to the walls are highly ordered, so it is worth noting that the vast majority of frozen particles lie close to $y = 0.5d$, giving a surface very similar to the glued particle layer used in experiment.

The drainage process is initiated by creating a $6d$ -wide slit in the center of the container base. To create as realistic a match to the experimental geometry as possible, the slit is modeled as a physical body filling the region $-5d < x < 0$, $|z - 30d| > 3d$, and all possible contacts with this body are considered, including the side walls of the orifice at $|z - 30d| = 3d$, and with the orifice edges at $|z - 30d| = 3d$, $x = 0$. In the simulation, particles which fall below $x < -10d$ are removed and no longer considered.

III. COMPARISON OF DEM TO EXPERIMENT

A. The effect of friction

To begin, drainage simulations were carried out using the particle models A, B, and C to investigate the effect of friction. For these simulation runs, snapshots of all particle positions were output at intervals of $\Delta t = 2\tau$. After an initial transient period of acceleration lasting until $t = 300\tau$, the top surface of the particle packing in these three runs descends at roughly constant velocity. Over a long time window ($300\tau < t < 2000\tau$) the top surface closely follows the linear relationship $x_{\text{top}} = 163d - 0.0584t$ (d/τ). Thus, to avoid any effects of the free surface, all data analyses were carried out over the time interval $300\tau \leq t \leq 900\tau$. At the end of the time window at $t = 900\tau$, the free surface is at $110d$, giving a ten particle buffer zone to the spatial test region.

For the three runs with different friction parameters, velocity profiles in the test region were computed across the y direction. The velocities are calculated based on the displacements of particles between successive simulation snapshots (effectively averaging them on a scale of 2τ) and stored in uniformly spaced bins and averaged. Figure 3(a) shows the three computed velocity profiles, when normalized and compared to experiment. In the shearing region near the rough wall, the profiles are almost identical, and closely match that seen in experiment. The weak dependence on friction is unsurprising, as previous studies have shown that much larger ranges for μ can have very little effect on macroscopic flow features in the bulk [28], since geometrical packing constraints play the dominant role; seeing a larger effect may require more significant changes in the material, such as using rough or angular particles.

However, it is surprising that the three runs exhibit different flow profiles at the smooth wall. For $\mu = 0.2$, there is no obvious boundary layer of slower velocities, but for $\mu = 0.3$ it becomes apparent, and for $\mu = 0.4$ it becomes more pronounced. While only a limited amount of data is available near the wall, it appears that the $\mu = 0.3$ curve most closely matches the experimental profile. It is worth noting that this value of friction is somewhat larger than the typical values for glass particles. In experiments with a plate dragged on a granular bed [15] – carried out using the same type of particles as used here – the effective friction was found to be in the range of 0.15 to 0.2 near a smooth boundary. However, it may not be possible to make a direct comparison to this work, since the confining pressure was different, and the boundary was held at a fixed for rather than fixed space condition. The results suggest potentially subtle effects in particle dynamics near the smooth wall, as to whether it is energetically more favorable for particles to fall as a plug and slip next to wall, or for particles next to the wall to roll and move slightly slower than those in the bulk. This behavior is difficult to precisely match between experiment and simulation.

Regardless of the details near the smooth wall, it is clear that in the shearing region of interest, the friction parameter has little effect on the velocity profile, and all three curves closely match the experiment. In the subsequent sections, we

concentrate on $\mu = 0.3$ only, as using $\mu = 0.2, 0.4$ had little effect on the presented results.

B. Layer positions and velocities

Much of the subsequent analysis is carried out in the shearing region near the rough wall, and in this section the positions and velocities of the grains are examined, using the chosen value of $\mu = 0.3$ in simulation. To illustrate and determine the positions of the layers, the particle number density was computed, and is shown in Fig. 4 for experiment and simulation, near the rough wall and near the smooth wall. In general, good agreement can be seen between the flowing states in simulation and in experiment. The peaks in the number density correspond to the layers of particles, and they are in very similar locations in experiment and simulation. To precisely determine the layer positions in simulation, Gaussians were locally fitted to the peaks in the number density. During flow, the first four layers are located at $(y_1, y_2, y_3, y_4) = (1.41d, 2.30d, 3.18d, 4.03d)$ corresponding to an average separation of $0.87d$.

A number density plot for the static initial packing is also shown in Fig. 4. Since this is based upon the single initial particle snapshot, rather than a time-average, the computed curve is noisier, and has to be computed using a larger bin size. Local Gaussian fitting gives the first four layers at $1.35d, 2.18d, 3.00d, 3.80d$, corresponding to an average separation of $0.82d$, suggesting that during flow, the particle layers expand slightly, allowing more space for particles to move past one another. Near the smooth wall, the static and flowing number density plots are very similar.

Subsequent computations within layers were carried out over the ranges $|y - y_i| < 0.3d$, where the width was chosen to match the experimental tolerance in the image processing. In experiment, the location y_i of the laser sheet was manually aligned so that the most grains were visible within each layer. The mean velocity component in the z direction and in the x direction in Layers 1 and 4 are also shown in Fig. 3(b,c). Very good agreement is observed over the spatial range over which experimental measurements could be made with the index-matching technique.

For subsequent analysis of the fluctuation properties, it was also important to determine the background velocity within each layer. To do this, average velocities were computed for each frame. In simulation, this was done using the method discussed in Subsec. III A, based upon successive snapshots between frames. Figure 5(a) shows the computed velocity as a function of time. Over the range $0 \leq t < 100\tau$, the speeds in the four layers begin to increase, as the particles in the test region begin to move, in response to the orifice being opened. After $t = 100\tau$, an approximate steady state in velocity is reached, although until $t = 300\tau$, the speeds in the four layers appear to decrease slightly. Measuring over the standard time window of steady flow, $300\tau < t < 900\tau$, the overall mean velocities were determined to be $-0.0372d/\tau, -0.0617d/\tau, -0.0718d/\tau, -0.0756d/\tau$ respectively for the four layers.




FIG. 3: (Color online) (a) The vertical component of the velocity v_x as a function of the y coordinate in the region $80d < x < 100d$ for different values of particle friction μ . The shearing region next to the rough wall at $0 < y < 1d$ is almost unchanged for the three different friction values, while the profile near the smooth wall at $y = 13.5d$ is slightly affected. (b) The z component of velocity within layer centered at y_1 and y_4 , as a function of the z coordinate. (c) x velocities as a function of the z coordinate.




FIG. 4: (Color online) Plots of the number density n_ϕ (the number of particle centers per unit volume) in the shearing region, as a function of the y coordinate. Plots during the flowing regime for the monodisperse and polydisperse runs are shown, using data from all snapshots in the range $300\tau < t < 900\tau$ and a bin size of $0.02d$ in the y direction. For the polydisperse plot, the contribution to the number density from each particle is weighted proportional to its volume. An additional plot for the static monodisperse packing is also shown using a bin size of $0.08d$. These plots include the frozen particles, that form the large peak at $y = 0.5d$.

While Fig. 5(a) shows a well-defined average velocity in each layer after $t > 300\tau$, there are surprisingly large variations on the order of 20% from one frame to the next. Furthermore, these variations appear to be strongly correlated between the layers. To illustrate this clearly, the vertical velocities computed from snapshots 0.01τ apart are shown in Fig. 5(b) for a small time interval. The “noise” is clearly correlated between layers and upon closer examination was found to be due to complex waves of velocity on a short timescale. Because, this has a significant effect on many of the subsequent simulation measurements, it is discussed in detail in Sec. IV.




FIG. 5: (Color online) (a) Background velocities on a scale of 2τ . (b) Background velocities on a scale of 0.01τ . The three vertical gray lines correspond to the times of the particle snapshots in Figure 10. (c) Variations in x and y background velocities computed on a timescale of 0.01τ for two typical layers.

The background velocities in the test region within each layer also show noticeable spatial variations both in experiment and simulation. Figure 3(b) shows the z velocities as a function of z across the test window for layers 1 and 4, for experiment and simulation, scaled according to the maximum downwards velocity in each layer. There is clear gradient across the window on the order of $\pm 5\%$, and good agree-

ment between experiment and simulation. Figure 3(c) shows the x velocities across the test window for each layer, this time scaled by the maximum velocity in the fourth layer to highlight the relative velocities between the layers. Again, we see good agreement between experiment and simulation, with slightly faster flow in the center of the test window than at the edges. These results are typical of drainage in wide rectangular hoppers [38] which tend to exhibit velocity profiles that spread with increased height near the orifice [39–41], although a precise comparison cannot be made to this previous work, as the rough wall used in the current study may have a strong effect on the flow. It is notable that the curves in Figs. 3(b) and 3(c) exhibit a significant amount of noise, even with a relatively coarse bin size of $0.5d$ using data from 301 snapshots. This may be due to the precise structure in the rough wall surface, introducing small fixed variations in velocity that are not removed by more time-averaging. This would be consistent with x velocities in Fig. 3(c) becoming more uniform for the higher-numbered layers, since those are further away from the rough surface and are less affected by its anisotropies.

C. Polydispersity

As previously noted, the glass particles used in the experiment exhibit a slight amount of polydispersity, with a diameter range of $\pm 10\%$ with the majority in $\pm 5\%$. While this is a relatively small range, it has been widely reported that flows in polydisperse particle packings can often exhibit fundamentally different behavior than monodisperse packings, as size variations decrease the tendency of particles to form regular crystalline arrangements. Such effects are particularly strong in two dimensional studies [42], and this plays a much weaker role in three-dimensional situations, where there is more geometrical freedom. For example, Tsai and Gollub [36] showed that crystallization in 3D monodisperse packings would only occur after many hours of shearing. To check the effect of particle size, a drainage simulation was carried out using model F using particle diameters uniformly distributed over $0.95d$ to $1.05d$.

Figure 4 shows a comparison of the number density as a function of the y coordinate, for the monodisperse model B, and the polydisperse model F. The curves appear almost identical, which is surprising, as it might be expected that polydispersity would smear out the peaks in the number density. The vertical velocity profiles, shown in Fig. 6 are also identical, almost to the level of statistical noise. Further measurements, not presented here, also showed near-perfect agreement. Since polydispersity introduces an additional numerical complication, and gave near-identical results, we therefore decided to concentrate on the monodisperse results.

D. Total flow rate

To directly relate the timescales in the simulation to the experimental results, the simulation time unit τ can be computed




FIG. 6: (Color online) Velocity profile comparison for several different drainage simulations using with friction coefficient 0.3 (using particle model B for the monodisperse case and particle model F for the polydisperse case). Plots for monodisperse and polydisperse runs are almost identical. A simulation with slower flow (created by inhibiting particle outflow at the orifice) is also almost identical once rescaled by the bulk velocity.

in terms of physical particle diameter and the effective gravity of the system that takes into account the upthrust on the particles due to the relative densities of the particles and the fluid. The effective gravity is

$$g_{\text{eff}} = g \left(1 - \frac{\rho_{\text{fluid}}}{\rho_{\text{glass}}} \right) = 5.9 \text{ m s}^{-2}$$

and thus $\tau = \sqrt{d/g_{\text{eff}}} = 0.013 \text{ s}$. Using this, the downwards velocity in the bulk of the packing, using data from figure 3(a) over the range $6d < y < 10d$, corresponds to $5.9 \times 10^{-1} \text{ cm/s}$.

In reality, the bulk downwards velocity in experiment in the region of interest is approximately $6 \times 10^{-2} \text{ cm/s}$, which significantly differs from the simulation result. As argued in Subsec. II A, the Stokes drag of the beads moving through the interstitial liquid is small compared to the gravitational force acting on the beads because the particles move slowly. However, in the orifice and in the converging flow region near the orifice, the grain flow velocity is about ten times faster than in the regions well above orifice where the flow is spread over wider cross-sectional area. Therefore, near the orifice the fluid forces may not be negligible, thus choking the flow and reducing the overall drainage rate, which in turn sets the bulk velocity in the test region.

Since a restricted flow at the orifice could potentially affect the simulation results in the test region, we carried out an additional drainage run, using particle model B, with a lower orifice outflow. An obvious method for reducing the outflow would be to make the exit slit thinner, but this is not feasible, since the granular packings tend to jam and flow intermittently for orifice sizes smaller than $5d$. To circumvent this problem, we kept the slit the same size, but restricted the velocities of particles in the orifice, so that those in the range $-5d < x < -d$ would have their velocities enforced to $\mathbf{v} = (-0.20d/\tau, 0, 0)$. This simple process works effec-

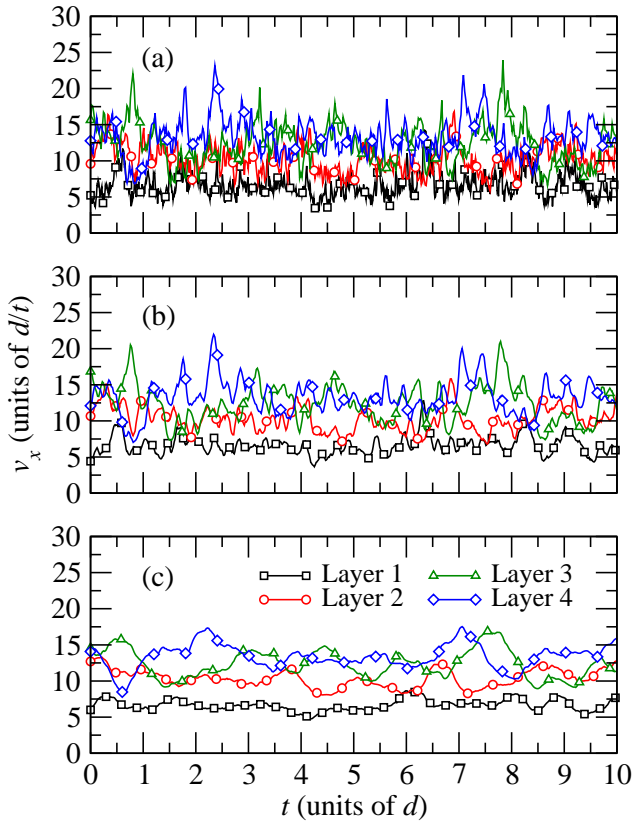


FIG. 7: (Color online) The background velocity observed in the experiments on timescales corresponding to a mean drop of grains over a distance of (a) $0.01d$, (b) $0.1d$, and (c) $0.5d$.

tively, and creates a smooth outflow, with a bulk downwards velocity in the test region of $0.0164d/\tau$, which is a precisely a factor of 4.73 from the corresponding unrestricted case. If the velocities are rescaled by this factor, then the velocity profile in the test region in the y direction closely matches with the unrestricted case as shown in Fig. 6. This points strongly to rate independence, that scaling the total flow rate results in a rescaling of the time variable, but does not create large changes in the particle dynamics. Such rate independence has been noted in other studies in flows inside silos [16, 27, 43], where particle diffusion was shown to be a function of distance dropped, as opposed to total flow rate. It provides further justification for making comparisons between experiment and simulation, by scaling out the overall flow rate. The presence of rate-dependent effects in granular materials has been reported [29, 44], but this may only become important for larger values of strain rate.

IV. TEMPORAL FLUCTUATION OF THE PARTICLE VELOCITIES

Figure 5(a) showed the presence of large variations in background velocity between successive frames, that was correlated between layers. However, the timescale of 2τ between frames is too coarse to properly resolve this behavior, so an

additional run was carried out using 10,000 snapshots at intervals of 0.01τ , over the range $500\tau < t < 600\tau$. Figure 5(b) shows part of the computed background velocities on this scale, showing unsteady oscillations in velocity, that are strongly correlated between layers. On this timescale, the differences in velocity become even more pronounced, with velocities in the fourth layer oscillating over a large range $-0.15d/\tau < v_x < -0.01d/\tau$. While the minima of the oscillations take different values of v_x , the maxima take values that are close together.

Figure 5(c) shows plots of typical y and z velocity averages on the 0.01τ timescale over the same time interval, which display a very different structure to the x variations. The scale of this plot is much smaller, and the structure mainly appears to be due to random statistical variation. There is no obvious wave-like behavior, and no strong correlation between layers, although it is noteworthy that the biggest velocities are approximately correlated with the largest velocity waves in Fig. 5(b).

Because these waves occur on short timescales, it is possible that the integration timestep may play a role. However, simulations making use of $\delta t = 5 \times 10^{-5}\tau$ and $\delta t = 2.5 \times 10^{-5}\tau$ show no appreciable difference in the wave structure. To more carefully quantify the differences between the velocities in each coordinate a Discrete Fourier Transform (DFT) was carried out using the data from the fast run. If the computed velocities in a layer are denoted by v_n for $n = 0, \dots, N-1$ where $N = 10,000$, then the DFT is computed as

$$V_k = \sum_{n=0}^{N-1} v_n e^{-\frac{2\pi i}{N} kn}.$$

Figure 9(a) shows the magnitude of the first four hundred Fourier modes for velocities in each of the three coordinates. For this plot, we made use of velocities in the third layer, although data from all layers shows the same picture. For $n \geq 150$, the modes for all three curves are similar in structure, suggesting that the short-timescale statistical noise is similar in all three directions. However, for the x velocity data, there is a pronounced region of modes from $n = 10$ to $n = 50$, corresponding to wavelengths in the range 2τ to 10τ , which is in agreement with the scale of the oscillations seen in Fig. 5(b).

In the experimental data, there are also variations in mean velocity in the x direction, and a selection of these is shown in Fig. 7 for three different timescales corresponding to a mean particle drop of $0.01d$, $0.1d$, and $0.5d$ between frames. The plots also show variations, but these appear to be more consistent with random noise, and show no evidence of unsteady oscillations and no correlation between layers. The absence of these oscillations in experiment suggest that they are most likely attributed to the approximations made in the simulation contact model, such as the normal spring interaction being smaller than realistic values. However the oscillations that are seen are not pure elastic modes due to the particle contact model. The natural frequency associated with the normal




FIG. 8: (Color online) (a) Variations in x velocities computed on a timescale of 0.01τ for a drainage simulation with a Hertzian contact model (thin lines), and a contact model using a normal spring coefficient of $k_n = 2 \times 10^6 mg/d$, a factor of ten bigger than usual (heavy lines). (b) Background velocities in the $k_n = 2 \times 10^5 mg/d$ simulation in layer 3, computed in ten strips of width $10d$ in the x direction, on a timescale of 0.01τ . The lowest and highest strips are highlighted, and the intermediate strips show a steady progression from noisy random walk behavior to wave-like behavior.

spring interaction is

$$t = \frac{2\pi}{\sqrt{k_n/m_p - \gamma_n^2/16m_p^2}} = 0.0324\tau,$$

which is significantly smaller than the scale observed here. Waves on the elastic timescale can be seen in simulation, but only by looking at even shorter snapshot intervals. The behavior of Fig. 5(b) happens on an intermediate timescale, larger than the particle interaction timescale, but smaller than the timescale of the macroscopic flow features.

To investigate the importance of the details of the contact model, two more short runs with 10,000 snapshots were carried out, using particle model D with $k_n = 2 \times 10^6 mg/d$, and particle model E with Hertzian contact forces. Although the extra factor of $\sqrt{\delta/d}$ in the Hertzian contact model precludes the assignment of a single natural frequency to the particle interaction, we expect that elastic oscillations will happen on

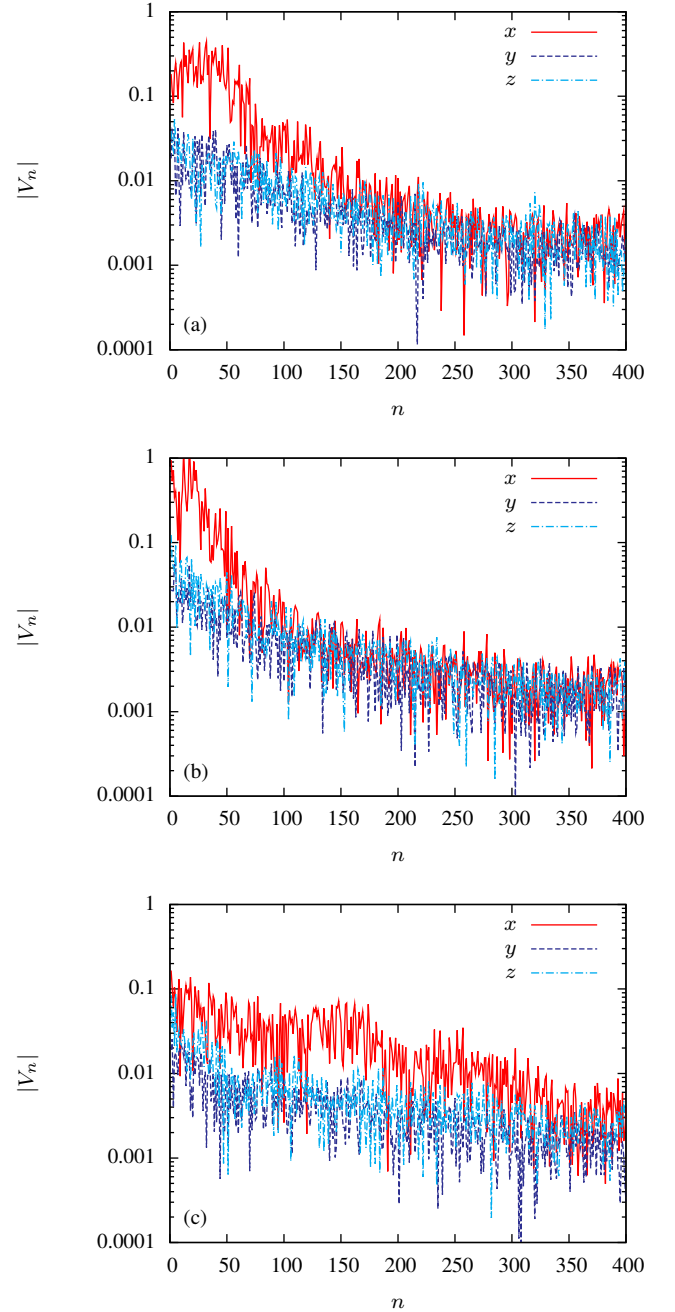


FIG. 9: (Color online) The magnitudes of the first four hundred Discrete Fourier Transform modes, for velocities in all three directions, computed in the third layer, for (a) the original simulation, (b) the Hertzian simulation, and (c) the high spring constant simulation with $k_n = 2 \times 10^6 mg/d$.

a longer timescale, as the factor will be always smaller than 1. Figure 8 shows plots of the background velocity in the test region for the four layers for these two simulations. For the higher spring constant, the waves are smaller in magnitude and happen on a faster timescale, while for the Hertzian contacts, the waves are larger and slower, to a level where occasionally (such as at $t = 523.5\tau$) the mean velocity points

upwards. Discrete Fourier Transforms of these two runs confirm this: for the Hertzian simulation (Fig. 9(b)) the modes are larger and shifted to the left, while for the high spring constant (Fig. 9(c)) the modes are smaller and cover a wider range of frequencies. While our data suggests the Hertzian contacts are more problematic, it is important to stress that this may not be a verdict on the relative merits of the two contact models. It may be a consequence of the specific method in which k_n is typically defined for Hertzian contacts, where it occurs along with a factor of $(\delta/d)^{1/2}$ which is always significantly smaller than 1, effectively decreasing the oscillation timescale of normal elastic interactions. Higher values of k_n may be appropriate when the Hertzian contact model is being used.

The correlation between the contact model and the velocity wave timescale strongly suggests that while the oscillations are not directly attributable to the normal spring constant, they are an indirect manifestation of it. Such behavior has been noted in previous discrete-element simulations, although the precise reason is unclear. Figure 10 shows a plot of snapshots in the four layers for three different times that correspond to the vertical gray lines in Figure 5(b) that happen before, during, and after a large oscillation. The wave affects all layers, although the differences are largest in the layers furthest from the wall. Clusters of faster and slower moving particles can be seen, sometimes across several layers. This complex behavior is perhaps indicative of periodic relaxation events, where built-up energy is released as particles are re-configured. The complex spatial structure of the waves makes them hard to deal with using a simple mean velocity subtraction. It is also undesirable that some particles end up moving upwards, as particle contacts may break and re-form, resetting the history-dependent terms in the contact model.

To examine the origin of the velocity waves, the background velocity in the third layer was computed in ten different horizontal strips of width $10d$ from $x = 0$ to $x = 100d$, over the central section of the packing $|z - 30d| < 10d$, and the results are shown in Figure 8(b). Near the orifice, over the range $0 < x < 10d$, the variations in velocity do not exhibit the wave-like behavior. However, in the higher strips, the waves of velocity become progressively more pronounced, suggesting a positive feedback mechanism; similar behavior has been noted in other larger 3D simulations in cones [45]. Since the waves take time to propagate upwards through the container, the curves in Figure 8(b) are also shifted rightwards from strip to strip. The fact that the waves become more pronounced in the higher parts of the container points to the overall system size as being a major factor, since a smaller simulation would not give the waves enough range to develop.

The presence of these oscillations must be carefully considered in the subsequent analysis. They are an undesirable feature when comparing with grains composed of hard materials like glass, and keeping the spring constant as high as possible may help, since the faster and smaller waves are easier to time-average. However, with the current numerical capabilities, it is impossible to eliminate them, so the best approach is to appreciate their scale and structure, and make sure that any computed statistics are influenced by them as little as possible.

V. ANALYSIS

A. Particle diffusion

In this section, we examine particle rearrangement and diffusion in experiment and for the simulation with contact model B. In the previous sections, the overall flow characteristics were shown to be in good agreement, but here we ask whether this agreement is also present at the microscopic scale.

A precise description of particle diffusion in slow, dense granular flow is difficult and there is relatively little precedent work in this area. There is no thermal equilibrium in the conventional sense, and particle rearrangement only occurs in response to outside forcing. At sub-particle length scales there is no Brownian motion, but rather particles move in response to collisions with their neighbors. Choi *et al.* [43] made experimental measurements of diffusion in rectangular silo flow, although this study was confined to making two-dimensional measurements imaged through a Plexiglas wall. By considering a range of flow rates, their results pointed to a rate-independence in granular mixing, allowing time to be scaled out, and the mean-squared particle displacements to be expressed in terms of the total deformation only. In three dimensions Campbell [46] has considered particle diffusion in a cell with Lees-Edwards boundary conditions, although for a more rapid regime. Here, although we can consider measurements in three dimensions, we have the additional complication of a layered structure in the y direction which our analysis must take account of.

We begin by computing the PDFs of particle displacements Δx in the three coordinate directions as a function of time, and several complications must be addressed. The first question that must be addressed is how to separate a “particle fluctuation” from the “background flow”. In Ref. [43], the background flow was computed from a spatially varying mean field $\bar{v}(\mathbf{x})$, but in reality there may also be large-scale temporal fluctuations in velocity, such as the elastic wave effects seen in the previous section, that would be better characterized as a mean flow as opposed to a particle fluctuation. Also, care must be taken in correctly defining the ensemble of particle trajectories over which the measurement is made. A computation of mean-squared displacements based on examining particles that remain within a test box of side length L , will bias against the more mobile particles that move outside the test box, and will never be able to measure displacements on scales larger than L . A solution to this is to continue to track those particles that move outside the test box into a new region, although this approach requires care, as the particles may behave differently in that region, and it may not be desirable for the displacement PDFs to incorporate that behavior.

Based on these concerns, our general approach to calculating displacement PDFs has made use of two regions. The first, R_s , represents the region of interest, for which the measurement is made. Any particle that is within R_s starts making a particle trajectory. The second region R_c , represents the space over which the particles are allowed to move. A particle continues making a trajectory while it remains within R_c , and in

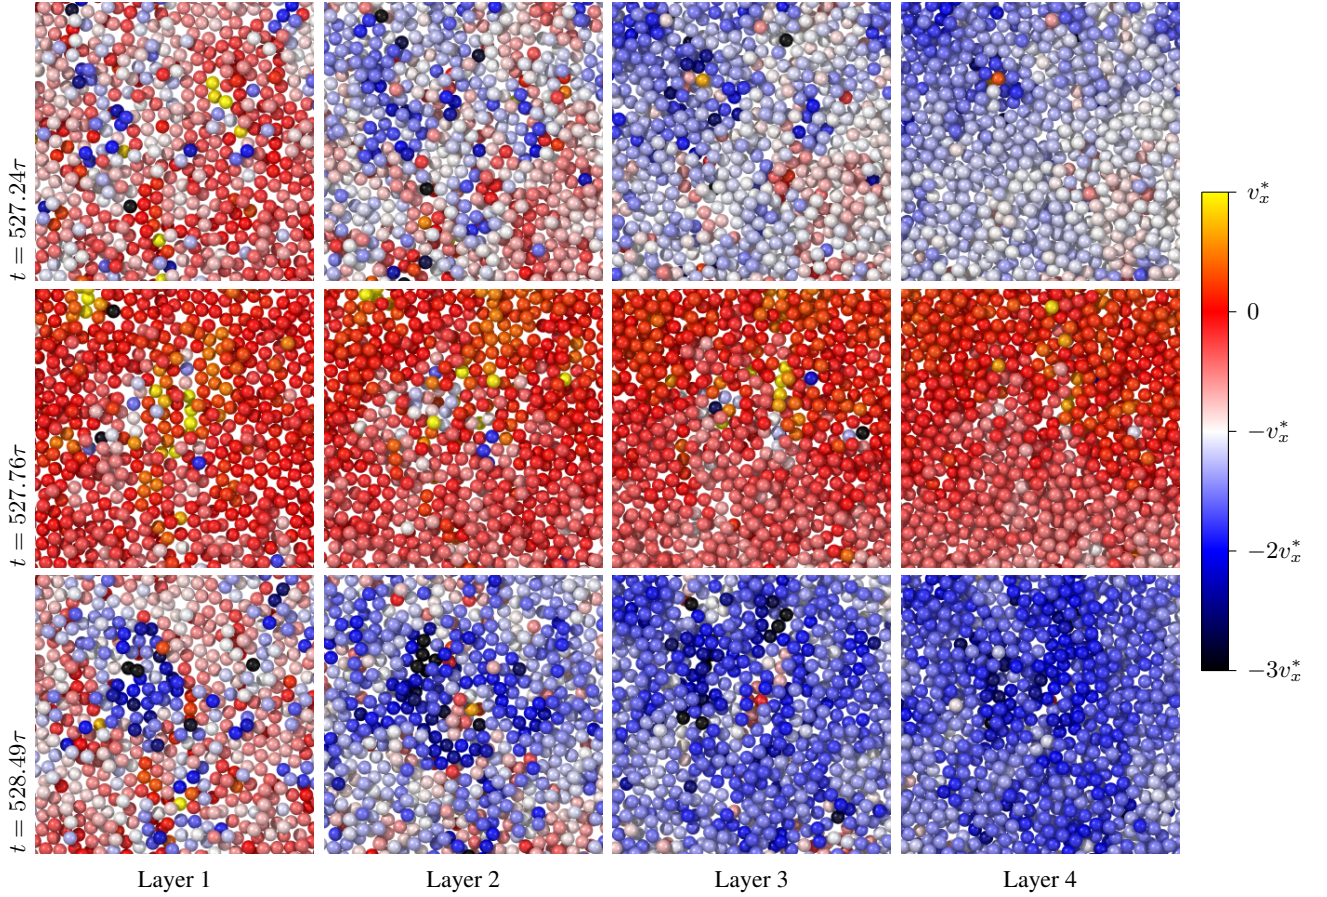


FIG. 10: (Color) Snapshots of particle positions in the simulation (using contact model B in table I) within the four layers for three different times, showing velocities computed based on particle displacements on a time interval of 0.01τ . The same color scheme is used in all four layers, and is expressed in terms of a typical vertical velocity $v_x^* = 0.03d/\tau$. The plots show the complexities of dealing with elastic waves that cause large variations in velocity over short timescales.

general it should be made as large as possible to minimize biases. It is restricted either by limitations in the available information, or because moving to a new region would result in fundamentally different particle behavior.

Once the collection of trajectories has been defined, the displacement PDFs at a time Δt are computed by evaluating all pairs of points $(\mathbf{x}_i, \mathbf{x}_f)$ on the trajectories which are separated by an amount Δt , such that the initial point \mathbf{x}_i lies within R_s . The initial point of this pair is advected according to the mean background velocity for an amount Δt to give a new point \mathbf{x}_i' . The distance between the advected point and the final point defines a vector $\Delta\mathbf{x} = \mathbf{x}_f - \mathbf{x}_i'$, from which the PDFs can be calculated. If $R_s = R_c$, then this method becomes equivalent to just evaluating all trajectories which lie wholly within the region of interest.

Our first analysis examines the displacement PDFs within each layer, in the x and z coordinates. In the experiment, where information about particles is strictly limited to the field of observation, we employed $R_s = R_c = \{80d < x < 100d, |y - y_i| < 0.3d, 22d < z < 38d\}$. The typical displacements $\Delta\mathbf{x}$ of a particle during its drop across the test box are significantly smaller than the test box size, so any introduced biases will be minimal.

As seen from Fig. 3(b,c), the velocities are reasonably uniform spatially within the test region selected for calculating the mean-squared displacements with slight variations near the border of the region. To examine this spatial variation in more detail we reduced the test region, thus eliminating the small variations in velocities, but found no appreciable changes in the measurements. Having confirmed the near homogeneity of the velocities in spatial directions, a background velocity was taken as a single mean downwards velocity per frame to account for the temporal variations. Plots of the PDFs of the displacements after $1d$ of mean drop in the x and z directions are shown in Figs. 11(b) and 11(c) respectively. As expected, the fluctuations increase for the layers closer to the wall, where the shear rate is higher. For a normal diffusive process, the curves would be expected to appear quadratic on a semi-log plot, but here we see a slower decay and larger tails corresponding to ballistic motion on a sub-particle length scale. Also visible in Fig. 11(c) is a small positive skewness, with more large steps in the positive x direction than in the negative x direction. This anisotropy is caused by gravity: when a gap opens up in the particle packing, it is possible that particles above will fall downwards to fill it, resulting in a large downwards displacement, even once the mean velocity




FIG. 11: (Color online) Plots of the PDFs of particle displacements Δx after $1d$ of drop within the test region in the three coordinate directions for experiment (symbols) and simulation (lines). Layers 1 to 4 are shown with black squares, red circles, green triangles, and blue diamonds respectively.

is subtracted.

In the simulation analysis, $R_s = \{80d < x < 100d, |y - y_i| < 0.3d, 20d < z < 40d\}$, while R_c is a slightly larger region, $\{70d < x < 100d, |y - y_i| < 0.3d, 15d < z < 45d\}$ to reduce any biases associated with disregarding trajectories crossing the boundary. We investigated several possibilities for the mean flow subtraction. If a homogeneous background velocity was employed in the simulation calculation, then it was found that the mean squared displacements in the x and z directions would scale according to Kt^α where $\alpha > 1$ (typically in the range $1.1 < \alpha < 1.3$). In the x and z directions, the average environment within each layer is homogeneous, and thus for long timescales, when particles have undergone several randomizing collisions, it should be expected that the displacements Δx and Δz should exhibit normal diffusive scaling, with their variance scaling like Kt . The higher power




FIG. 12: (Color online) Plots of the PDFs of particle displacements Δy and Δz after $8d$ of drop within the test region in experiment (symbols) and simulation (lines). Layers 1 to 4 are shown with black squares, red circles, green triangles, and blue diamonds respectively.

of Kt^α suggests that a systematic drift is also being measured, arising from the small inhomogeneities in background velocity seen in Fig. 3(b,c). To circumvent this, we employed a spatially and temporally varying background velocity. First, a spatial velocity field was calculated in $5d \times 3d$ boxes over R_c , using the standard time window $300\tau < t < 900\tau$. Second, an overall mean velocity was computed for each frame. The background velocity is taken to be the bilinear interpolation of the spatial velocity field, plus an additional overall temporal correction for each frame. Using this procedure gives the asymptotic behavior $\text{Var}(\Delta x) \sim Kt$, $\text{Var}(\Delta z) \sim Kt$ indicative of normal diffusion. The PDFs of the displacements after $1d$ of mean drop in the x and z directions are shown as lines in Figs. 11(a) and 11(c) respectively. In both directions and in all four layers, we see excellent quantitative agreement with experiment.

To examine particle displacements in the y direction in experiment, we make use of measurements in the xy plane, and consider $R_s = R_c = \{80d < x < 100d, |y - y_i| < 0.5d, |z - 30d| < 0.3d\}$. For the background flow, we make use of a single mean velocity computed at each frame. Unlike the x and z measurements, this procedure is more susceptible to biases, as the width of the test region in the y direction is comparable to the scale of particle displacements. The particles that take large steps and move between layers will be discounted. Also, Fig. 11(c) shows that a large number of particles will undergo z displacements that are comparable with the thickness of the viewing plane, meaning that many parti-

cles will be lost when they can no longer be tracked.

In simulation, where we have the freedom to track particles wherever they go, the regions were chosen to remove some of the above biases. We make use of $R_s = \{80d < x < 100d, |y - y_i| < 0.3d, 20d < z < 40d\}$, tracking particles across the entire test region rather than restricting to a single slice. We choose $R_c = \{70d < x < 100d, 15d < z < 45d\}$, continuing trajectories even if they pass from one layer to another. For the mean velocity subtraction, we make use of the same procedure as in the xz measurements, combining a spatial field with a frame-by-frame temporal correction. It is worth noting that this causes no problem for particles moving between layers, as for each pair of positions $(\mathbf{x}_i, \mathbf{x}_f)$ that is considered, the mean velocity is applied to \mathbf{x}_i which is always within R_s and hence inside the layer itself.

Figure 11(b) shows a comparison of the y displacement PDFs for experiment and simulation after $1d$ of drop. Despite handling particle trajectories differently, there is very good agreement between the curves, as the particle displacements are small enough that biases do not factor in heavily. However, the PDFs after $8d$ of average displacement in the flow, shown in Fig. 12(a) show large differences. While the central peaks are similar in size, additional peaks can be seen in the simulation data corresponding to those particles that have moved into neighboring layers. This complicated behavior makes it hard to assign a meaningful diffusion constant in this direction. On the length scales that can be observed, the PDFs do not appear to tend to anything resembling a Gaussian, and particle motion is a combination of stochastic behavior and layer confinement. It is also possible to carry out an analysis of x displacements using the xy plane measurements. For these cases, as shown in Fig. 12(b), larger amounts of particle mixing can be measured, as particles are also separated due to the velocity gradient between layers.

Figure 13 shows a logarithmic plot of the mean squared displacement $\text{Var}(\Delta\mathbf{x})$ as a function of mean distance dropped for the three coordinate directions. The x and z plots were based on measurements in the xz plane, and we see quite good quantitative agreement in all four layers. As noted by Choi *et al.*, a transition from superdiffusive behavior with slopes greater than one, to normal diffusion with slope close to one, can be seen at a distance dropped of approximately $1d$. A slight disagreement is seen in the z direction, particularly for layers 1 and 2 for very large displacements (*ie.* more than $4d$), wherein we see the experimental mean squared displacements to be slightly higher than in simulations. We believe that this bias may be introduced perhaps due to loss of particles between layers over long displacements since the tracking was done only within one layer. The effect is not seen for layers 3 and 4 wherein the particle hopping between the layers is almost negligible. These results probably state the lack of reliability in experiments to determine fluctuations over very long time scales and displacements, thus stressing the importance of simulation data over these scales.

The y plot is based on measurements in the xy plane. We see good agreement for small distances but the curves begin to diverge for larger distances, due to the simulation methods counting those particles moving between layers. The plot for




FIG. 13: (Color online) Logarithmic plots of mean squared particle displacement versus mean distance dropped for the three coordinate directions, in experiment (symbols) and simulation (lines). Here, V refers to the mean velocity within each layer. Layers 1 to 4 are shown using black squares, red circles, green triangles, and blue diamonds respectively.

layer 1 exhibits fundamentally different behavior since particles can only jump to layers in one direction as opposed to both. The curves do not exhibit slopes close to one, confirming that a diffusion constant in the y direction cannot be meaningfully defined.

Figure 14 shows the diffusion constants in the x and z directions as a function of shear rate, computed from the mean squared displacement measurements. In the simulation, where the velocity profiles are known with high spatial resolution, the shear rate $\dot{\gamma}(y)$ is computed by finding the slope of the function $v_x(y)$ over the range $y \pm 1d$, using a local linear regression. The scale of $1d$ is used, since this corresponds to the distance at which a particle may directly come into contact with another. The shear rate within the layer is then calculated as the average of $\dot{\gamma}(y)$ across the layer $|y - y_i| < 0.3d$, weighted by the particle number density $n_\phi(y)$. In the experiments, where the velocity profiles are not known with as much spatial resolution, we fit an eighth order polynomial to the velocity data points shown in Fig. 3, and then differentiate the polynomial numerically at each point to obtain the local




FIG. 14: (Color online) Diffusion constants in the x and z directions as a function of shear rate in the four particle layers, for experiment (a) and simulation (b).

shear rate. In the figure, roughly linear growth with shear rate can be seen. However, with only four data points, and potentially different behavior for the first layer that is next to the wall, it is hard to say anything conclusive about scaling. Another striking feature of these measurements is that the diffusion constant in the flow x direction is greater than in the z direction. Such an anisotropy in the diffusion constants has been anticipated in sheared athermal suspensions [33]. Here, we see that such anisotropy persists in both our numerical and experimental granular systems, and points to the importance of particle geometry and local packing in determining local rearrangement and diffusion of particles rather than details of interaction between particles. Some differences in the fluctuations of the particles in the experiments and numerics can be observed. However, these differences appear to arise because particles are systematically lost in the experiments because of the limitation of tracking particles over long times, rather than due to physical differences.

B. Velocity autocorrelations

The calculation of the velocity autocorrelation function $\psi(t)$ for some time t is based upon finding a collection of velocity pairs $(\mathbf{v}_i, \mathbf{v}_f)$ that are separated by t and then computing the product-moment correlation coefficient. Defining the collection of velocity pairs is subject to the same biasing problems that were faced in the diffusion measurements. Because of this, we make use of the same definition of parti-

cle trajectories as previously, using a starting region R_s and a continuing region R_c . Since autocorrelations are based upon velocities, constructed from the difference of two positions, their calculation is more sensitive than the diffusion measurements.

Autocorrelations in the experiments have been reported previously [13]. To carry out autocorrelations in the x and z directions, we made use of $R_s = R_c = \{80d < x < 100d, |y - y_i| < 0.3d, 20d < z < 40d\}$ and took snapshots in intervals corresponding to exactly $0.01d$ of mean drop. Based on these, velocities were computed on a scale of $0.1d$ by looking at particle displacements ten frames apart.

The precise timescale on which velocities are computed could potentially have a significant effect on the autocorrelation function, so to obtain the best match possible, the simulation snapshots were recorded at the same intervals corresponding to a $0.01d$ drop, as in experiments. Initially, an autocorrelation was attempted using the standard contact model B, but the results were problematic. As shown by the dashed gray lines in Fig. 15, the correlations in the x direction exhibit chaotic oscillations at large times. This appears unphysical, since after a particle has fallen by several times its diameter and undergone many collisions with neighbors, it is velocity is unlikely to be correlated with its previous velocity. The problem seen in the graphs is due to the waves of velocity moving through the system that were discussed in Section IV. Since the waves are larger in the higher-numbered layers, the autocorrelation oscillations are more significant there.

Several procedures were tried to improve these results. A mean velocity subtraction per frame can mitigate the worst of the oscillations, but there is still a significant amount of noise. We therefore decided to switch to carrying out simulations using the particle model with $k_n = 2 \times 10^6 mg/d$, that were previously shown to have fewer velocity waves. To increase the amount of available data, we carried out a second drainage run by taking the static packing for model D, rotating it by 180° around the x axis, and making the rough wall by freezing the particles that were now in the range $0 < y < 1d$. This creates a second data set with a different particle configuration without the need to generate a completely new packing by pouring. Since the pouring process is the most time-consuming part of the simulations, particularly for this contact model where a smaller timestep is needed, it was best to avoid generating more. The velocity measurements from the two simulations were treated as a single ensemble of pairs $(\mathbf{v}_i, \mathbf{v}_f)$ that were used to compute autocorrelations.

We also employed a spatial background mean flow using a bilinear interpolation on a 4×5 grid. Without this subtraction, the plots look almost identical, except that the plots in the higher layers are shifted upwards by a small amount and do not tend to zero at large separations. For the x and z measurements within layers, we employed $R_s = R_c = \{80d < x < 100d, |y - y_i| < 0.3d, 20d < z < 40d\}$, and for the y measurements we used $R_c = \{80d < x < 100d, 20d < z < 40d\}$ to continue trajectories the move to other layers. In addition, a temporal mean velocity subtraction was applied in the x direction to remove some of the velocity waves. The resulting autocorrelation functions in the three directions are shown in




FIG. 15: (Color online) Autocorrelations in the three coordinate directions as a function of the mean distance dropped, in the simulation with $k_n = 2 \times 10^6 mg/d$, based on the average of two drainage runs. A spatial mean velocity subtraction is applied in all three directions. For the x direction, an additional frame-by-frame correction is subtracted to mitigate the effect of the intermediate-timescale velocity waves. In this direction two additional lines are shown for the lower spring constant calculation without a frame-by-frame correction, highlighting the problem with the velocity waves, which create large chaotic oscillations in the autocorrelation function, particularly for the higher-numbered layers. Here, V refers to the mean velocity within each layer.

Fig. 15. The plots are a significant improvement over the results with the lower spring constant. Although some noise is visible, the curves decay to zero for large separations. Negative correlations are visible for separations of around $0.5d$ in all three directions. This phenomenon has been widely reported in autocorrelation measurements for interacting hard spheres and is due to “back-scattering” [47–49], whereby a particle moving in one direction will be likely to undergo a collision with a neighbor, and thus on average be moving in an opposite direction after a certain interval. The negative autocorrelation is particularly strong in the y direction, which we attribute to layer confinement, as particles moving out of a layer are particularly likely to come into contact with a particle in the adjacent layer. Such an effect was not found in the experimental autocorrelation measurements [13], as the particles were not tracked between layers. For the first layer, there is evidence of a small peak in the autocorrelation function at around $1d$, although more testing is necessary to determine if this feature is robust.

Figure 16 shows a comparison between the experimental and simulation, plotted on a logarithmic scale. In the x and z directions we see good agreement, particularly for layers 1 and 2, where the simulation measurements are less affected by velocity waves. Further, it appears that the initial decay is closer to an exponential decay, as opposed to a long time decay tail, such as $t^{-3/2}$ decay that has been observed for elastic spheres at equilibrium [48] and recently for unsheared inelastic spheres [16]. While our results demonstrate the computation of autocorrelations within the discrete-element method, it is a significantly larger computational challenge than many of the other measurements considered in this study. Gaining detailed, precise information about the decay would require smaller timesteps and larger ensembles, both of which increase the amount of computation needed.

VI. CONCLUSION

For a wide variety of flow features that have been considered, our results have shown a high degree of quantitative agreement between index-matched experiment and DEM simulation with the Cundall–Strack granular contact model with typical values of the contact parameters. Despite the two completely different procedures, we have been able to show close matches between macroscopic flow features (*eg.* velocity profiles), as well as microscopic particle properties (*eg.* number density profiles, particle diffusion, and velocity autocorrelations). Our results provide validation that both techniques can be reliably used to study granular flows. While both procedures have potential shortcomings, such as interstitial fluid effects in experiment at high flow rates, or approximations in the simulation contact model, the fundamental physics of granular flow and particle rearrangement appears largely similar.

This successful matching can be partially attributed to the fact that the contact models employed in discrete-element simulation are a close reproduction of the contact physics of the index-matched flow. However, our results are also indicative that many key features of slow, dense granular flow may




FIG. 16: (Color online) A comparison of autocorrelations between experiment (symbols) and simulation (lines) in the three coordinate directions as a function of mean distance dropped. Here, V refers to the mean velocity within each layer. Layers 1 to 4 are shown with black squares, red circles, green triangles, and blue diamonds respectively. The solid heavy black lines are exponential fits.

exhibit some degree of universality across a wide variety of situations. Despite small differences near the smooth wall, our results showed that the velocity profile in the shearing region was largely similar for friction values over the range from 0.2 to 0.4. This is consistent with previous work [28], where the same result was shown for a larger range of $0.1 < \mu < 0.9$. A small amount of polydispersity, while a critical issue in two dimensional packings, appeared to have a minimal effect on the velocity profiles and packing structure. Also, while there is clear evidence of rate-dependent effects [29, 44] at faster flow rates, our results suggest that in the slow, dense regime, the total flow rate can be scaled out of the measurements, making it much easier to quantitatively compare to experimental results.

Despite the successes, our results do highlight several potential areas of concern. In experiment, the inability to build a complete three-dimensional map of the particles means that a number of properties of the flow cannot be quantified. When

presenting diffusion and autocorrelation measurements, the importance of choosing trajectories was discussed, but in experiment the approach was limited by the lack of information when particles moved outside of the laser sheet.

In simulation, our study has highlighted several possible areas of difficulty. As discussed in the introduction, much of the initial choice of the contact model parameters was carried out by examining macroscopic flow properties, and microscopic packing structure, and in general our results have shown excellent agreement in these areas. However, our results suggest that for examining microscopic dynamical features, such as autocorrelations, using a stiffer spring constant may be required to achieve a reasonable match with realistic flows.

The presence of velocity waves as described in Sec. IV also presents a large cause for concern. Our results suggest that the overriding factor in the generation of these waves is the total system size, since they become progressively larger with height. Again, we note that the original choice of contact parameters, that occurred eight years ago when less computational power was available, made use of much smaller system sizes featuring 24,000 particles, meaning that the packings were small enough that these effects may not play a significant role. The waves are undesirable for several reasons. They occur on an intermediate timescale much larger than the natural contact frequency, potentially interfering with a variety of measurements. There is also potential for particle contacts to successively break and re-form during the passage of a velocity wave, which may have a significant effect on the history-dependent terms of the contact model.

Furthermore, the waves appear to have no analog in the experimental data. But it should be also noted that our study does not provide enough evidence to show that these waves are unphysical in all situations: it may be that particles composed of a softer material such as acrylic glass, where a normal spring constant that is closer to that used in the simulation, would show waves of this type, and we believe this could be an interesting direction for further study. However, in previous studies, DEM simulations have been compared with the body of theoretical and experimental results using hard materials, and when comparing rapid features of flow our results suggest this should be done with caution.

Our results indicate that increasing the normal contact stiffness by a factor of ten may be a useful remedy. While this does not remove the waves completely, it does make them smaller and more rapid, allowing for them to be more easily removed by time-averaging. Since the simulation makes use of a second order scheme, this requires a three-fold increase in computational cost, which would be reasonable in many situations.

Acknowledgments

This work was supported by the Director, Office of Science, Computational and Technology Research, U.S. Department of Energy under Contract Nos. DE-AC02-05CH11231 and the National Science Foundation under grants DMS-0410110 and DMS-070590, and at Clark University under Grant No.

-
- [1] H. M. Jaeger and S. R. Nagel, *Science* **255**, 1523 (1992).
 - [2] H. M. Jaeger, S. R. Nagel, and R. P. Behringer, *Rev. Mod. Phys.* **68**, 1259 (1996).
 - [3] T. Halsey and A. Mehta, eds., *Challenges in Granular Physics* (World Scientific, 2002).
 - [4] P. G. de Gennes, *Rev. Mod. Phys.* **71**, S374 (1999).
 - [5] J. T. Jenkins and S. B. Savage, *J. Fluid Mech.* **130**, 187 (1983).
 - [6] I. S. Aranson and L. S. Tsimring, *Rev. Mod. Phys.* **78**, 641 (2006).
 - [7] A. Ferguson, B. Fisher, and B. Chakraborty, *Europhysics Letters* **66**, 277 (2004).
 - [8] M. P. Allen and D. J. Tidesley, eds., *Computer Simulations of Liquids* (Oxford University Press, New York, 1989).
 - [9] S. McNamara and W. R. Young, *Phys. Fluids A* **4**, 496 (1992).
 - [10] M. Alam and C. M. Hrenya, *Phys. Rev. E* **63**, 061308 (2001).
 - [11] S. Luding and S. McNamara, *Granular Matter* **1**, 113 (1998).
 - [12] P. A. Cundall, in *Proc. Symp. Int. Soc. for Rock Mechanics, Nancy, France* (1971), pp. 129–136.
 - [13] A. Orpe, V. Kumaran, A. Reddy, and A. Kudrolli, *Europhys. Lett.* **84**, 64003 (2008).
 - [14] J.-C. Tsai, G. A. Voth, and J. P. Gollub, *Phys. Rev. Lett.* **91**, 064301 (2003).
 - [15] S. Siavoshi, A. V. Orpe, and A. Kudrolli, *Phys. Rev. E* **73**, 010301(R) (2006).
 - [16] A. Orpe and A. Kudrolli, *Phys. Rev. Lett.* **98**, 238001 (2007).
 - [17] P. A. Cundall and O. D. L. Strack, *Geotechnique* **29**, 47 (1979).
 - [18] <http://lammmps.sandia.gov/>.
 - [19] L. E. Silbert, D. Ertaş, G. S. Grest, T. C. Halsey, D. Levine, and S. J. Plimpton, *Phys. Rev. E* **64**, 051302 (2001).
 - [20] L. E. Silbert, J. W. Landry, and G. S. Grest, *Phys. Fluids* **15**, 1 (2003).
 - [21] R. C. Brewster, J. W. Landry, G. S. Grest, and A. J. Levine, *Phys. Rev. E* **72**, 061301 (2005).
 - [22] J. W. Landry, G. S. Grest, L. E. Silbert, and S. J. Plimpton, *Phys. Rev. E* **67**, 041303 (2003).
 - [23] L. E. Silbert, D. Ertaş, G. S. Grest, T. C. Halsey, and D. Levine, *Phys. Rev. E* **65**, 031304 (2002).
 - [24] L. E. Silbert, G. S. Grest, and J. W. Landry, *Phys. Rev. E* **66**, 061303 (2002).
 - [25] C. H. Rycroft, M. Z. Bazant, G. S. Grest, and J. W. Landry, *Phys. Rev. E* **73**, 051306 (2006).
 - [26] H. P. Zhu and A. B. Yu, *J. Phys. D* **37**, 1497 (2004).
 - [27] C. H. Rycroft, G. S. Grest, J. W. Landry, and M. Z. Bazant, *Phys. Rev. E* **74**, 021306 (2006).
 - [28] K. Kamrin, C. H. Rycroft, and M. Z. Bazant, *Modelling Simul. Mater. Sci. Eng.* **15**, S449 (2007).
 - [29] C. H. Rycroft, K. Kamrin, and M. Z. Bazant, *J. Mech. Phys. Solids* **57**, 828 (2009).
 - [30] X. Cheng, J. B. Lechman, A. Fernandez-Barbero, G. S. Grest, H. M. Jaeger, G. S. Karczmar, M. E. Mobius, and S. R. Nagel, *Phys. Rev. Lett.* **96**, 038001 (2006).
 - [31] M. Depken, J. B. Lechman, M. van Hecke, W. van Saarloos, and G. S. Grest, *Europhys. Lett.* **78**, 58001 (2007).
 - [32] J. Sun, F. Battaglia, and S. Subramaniam, *Phys. Rev. E* **74**, 061307 (2006).
 - [33] D. R. Foss and J. F. Brady, *J. Fluid Mech.* **401**, 243 (2006).
 - [34] V. Kumaran, *Phys. Rev. Lett.* **96**, 258002 (2006).
 - [35] M. Otsuki and H. Hayakawa, *cond-mat arXiv:0711.1421* (2007).
 - [36] J.-C. Tsai and J. P. Gollub, *Phys. Rev. E* **70**, 031303 (2004).
 - [37] J. C. Crocker and D. G. Grier, *Journal of Colloid and Interface Science* **179**, 298 (1996).
 - [38] A. Medina, J. A. Cordova, E. Luna, and C. Trevino, *Physics Letters A* **220**, 111 (1998).
 - [39] J. Mullins, *J. Appl. Phys.* **43**, 665 (1972).
 - [40] R. M. Nedderman and U. Tüzün, *Powder Technology* **22**, 243 (1979).
 - [41] H. Caram and D. C. Hong, *Phys. Rev. Lett.* **67**, 828 (1991).
 - [42] P. M. Reis, R. A. Ingale, and M. D. Shattuck, *Phys. Rev. Lett.* **96**, 258001 (2006).
 - [43] J. Choi, A. Kudrolli, R. R. Rosales, and M. Z. Bazant, *Phys. Rev. Lett.* **92**, 174301 (2004).
 - [44] P. Jop, Y. Forterre, and O. Pouliquen, *Nature* **441**, 727 (2006).
 - [45] C. H. Rycroft, K. Kamrin, and G. S. Grest (2009), (in preparation).
 - [46] C. S. Campbell, *J. Fluid Mech.* **348**, 85 (1997).
 - [47] J. H. Dymond and B. J. Alder, *J. Chem. Phys.* **45**, 2061 (1966).
 - [48] B. J. Alder and T. E. Wainwright, *Phys. Rev. Lett.* **18**, 988 (1967).
 - [49] R. D. Wildman, J.-P. Hansen, and D. J. Parker, *Physics of Fluids* **14**, 232 (2002).