

JAMMING IN THE GAS-FLUIDIZED BED

Adam R. Abate

A DISSERTATION

in

Physics and Astronomy

Presented to the Faculties of the University of Pennsylvania in Partial
Fulfillment of the Requirements for the Degree of Doctor of Philosophy

2006

Douglas J. Durian
Supervisor of Dissertation

Randall D. Kamien
Graduate Group Chairperson

COPYRIGHT

Adam R. Abate

2006

Acknowledgements

The most difficult challenge for a young physicist entering a field as diverse as soft condensed matter physics is to grasp what is known, what is unknown, and how ask questions, while still maintaining the stubbornness and ignorance that is essential to scientific discovery. In such an unfamiliar landscape, young scientists rely on the wisdom and skilful guidance of their elders, and I have been fortunate to be surrounded by the very wisest and most skilled scientists.

My advisor, Doug Durian, has undoubtedly had the greatest single impact over my professional development. It was clear to me early on that he has mastered many of the tools necessary to succeed in science, and I feel extremely fortunate to have the opportunity to learn from him. Practically speaking, I am grateful for his thoughtfulness and his ability to inject inspiration when an experiment inevitably grinds to a halt. I am also grateful for his availability and know that it has come at no small cost to his time. It has only been possible because his students are clearly a priority to him, and as one of his students I have benefited greatly.

There are many others over the years who have impacted me in ways I will never be able to adequately convey, for what person can convey adequately the subtle instants that change a life? I gratefully acknowledge Jerry Gollub, Ike Silvera, Donovan Hall, Arjun Yodh, Andrea Liu, Joe Rudnick, Paul Dixon, Jay Kikkawa, Rajesh Ojha, Klebert Feitosa, Pam Korda, Hiroaki Katsuragi, Paulo Arratia, Lynn Daniels, Kersten Nordstrom, Ted Brzinski, Glenn Kinen, Parag Shah, Mr. Adams, Mr. Walpole, Ms. Spina, Richard Feynman, Bertrand Russell, and Jess Porter.

Finally, my parents provided me with the upbringing and sense of self-worth that has given me personal strength and has led me to the present. My brothers taught me early on what it means to rely on other individuals and have them in turn rely on you. My family has no doubt had the greatest impact over my life and whatever success I might one day achieve I fully dedicate to them.

ABSTRACT

JAMMING IN THE GAS-FLUIDIZED BED

Adam R. Abate

Douglas J. Durian

This work focuses on the study of nonequilibrium systems driven close to jamming. Macroscopic particles agitated by turbulent air are observed with a high-speed camera. Driven this way, one or two ping-pong balls were found to perfectly obey statistical mechanics [64]. To probe the range of validity of the statistical mechanical description we systematically perturb the experiment searching for a governing parameter. Next, we considered dense packings of steel balls fluidized the same way and ask the question how the system's proximity to the jamming transition effects its structure and dynamics. Focusing our attention on systems very close to jamming, we observe spatially heterogeneous dynamics resembling dynamical structures seen in simulations of supercooled liquids and colloidal glasses. We measure the growth of these heterogeneities as a function of proximity to the jamming transition and find that dynamical time and length scales diverge in a way consistent with Mode-Coupling Theory and the Vogel-Tammann-Fulcher Equation, which adds support to the universality of the jamming concept. Motivated by the need to fully characterize spatially heterogeneous dynamics, we also develop novel topological methods based on time-varying Voronoi diagrams and correlations in the nearest neighbor network, which are both more accurate and more straightforward to measure than widely used four-point correlation functions. To investigate how the short time agitation of the balls by the air gives rise to long time system relaxation, we verify the consistency of three distinct effective temperatures that in thermal equilibrium are equal. We also study avalanches of flowing glass sand particles which offer a unique look at a system at different stages of jamming. Using a newly developed light scattering technique called Speckle-Visibility Spectroscopy we are able to probe the three-dimensional dynamics of the avalanche flows over ten orders of magnitude in time.

Contents

Acknowledgements	iii
Introduction	1
1 The Partition of Energy for Air-Fluidized Grains	4
1.1 Abstract	4
1.2 Introduction	4
1.3 Experimental Details	6
1.4 One Sphere in a Stadium	9
1.5 Five Spheres of Different Density	12
1.6 Two Spheres of Different Size	13
1.7 One Aspherical Grain	16
1.8 Conclusion	20
2 Approach to Jamming in an Air-Fluidized Granular Bed	21
2.1 Abstract	21
2.2 Introduction	21
2.3 Methods	23
2.4 General	25
2.5 Structure	28
2.5.1 Coordination Number	28
2.5.2 Circularity	28
2.5.3 Radial Distribution	31

2.5.4	Summary	34
2.6	Dynamics	34
2.6.1	Data	34
2.6.2	Short and Long-time Dynamics	38
2.6.3	Timescales	42
2.6.4	Jamming Phase Diagram	46
2.7	Conclusion	48
3	Direct Observation of Growing Dynamical Lengthscales for Driven Steel Beads	49
3.1	Introduction	49
3.2	Details	52
3.3	Strings	52
3.4	Four-Point Susceptibility	55
3.5	Divergent Dynamical Lengthscales	57
4	Topological Persistence Near Jamming	60
4.1	Abstract	60
4.2	Introduction	60
4.3	Background	62
4.4	Persistent Area	66
4.5	Persistent Bond	71
4.6	Applications	75
4.7	Approach to Jamming	78
4.8	Discussion	82
5	Three Effective Temperatures for Air-Fluidized Grains	84
5.1	Introduction	84
5.2	Methods	86
5.3	A Granular Temperature Thermometer	86
5.4	Einstein Temperature	90
5.5	Effective Temperature Comparison	94

6	Probing Avalanche Dynamics using Speckle-Visibility Spectroscopy	97
6.1	Abstract	97
6.2	Methods	98
6.2.1	Experimental Apparatus	98
6.2.2	Digital Video Imaging	98
6.2.3	Speckle-Visibility Spectroscopy	103
6.3	Avalanche Statistics	108
6.4	Anatomy of an Avalanche	110
6.5	Avalanche Dynamical Scales	113
7	Conclusions	115
A	Crystallized Steel Ball Bearings	118
B	Clustering in the Bed	123
C	Voronoi Shape Fluctuations	128
D	Sensitivity of the Four-Point Susceptibility on the Overlap Threshold Size	132
E	Energy Cascade in Driven Beads	135
F	Small Sand Avalanches	142
F.1	Small Sand Avalanche Statistics	144
F.2	Anatomy of a Small Sand Avalanche	145
G	Technical Appendix	151
G.1	Digital Video Imaging of Particles in the Gas-Fluidized bed.	151
G.1.1	Gas-fluidized Bed and Its Components	151
G.1.2	Image Analysis and Particle Tracking	155
G.2	Speckle-Visibility Spectroscopy of Avalanches	158
G.2.1	Avalanche Apparatus and Its Components	158
G.2.2	Avalanche Data Acquisition with LabVIEW	160
G.3	LabVIEW Miscellaneous	162

List of Tables

1.1 [aspherical] The density, length, width, height, mass, and moment of inertia of five aspherical grains. For computation of density and moment of inertia, the white, silver, and wood grains are approximated as cylinders, while the brown grain is approximated as a block. The dimer is composed of two polypropylene shells of thickness 0.14 cm.	17
---	----

List of Figures

1.1	[col-ero] (a) The measured blob area, and (b) the center-to-center separation, for two grains in collision. The grayed stripe indicates the time-interval the grains are in close enough proximity so that they are imaged as as a single blob (inset a), which is eroded producing two smaller separated circles (inset b).	8
1.2	[v-p-dis] Distributions for the components of (a) position and (b) velocity measured for a one-inch nylon sphere along the long and short axes of a stadium-shaped trap. To within measurement uncertainty, these distributions are all Gaussian (dashed curves).	10
1.3	[spr-cons] Shift in average position of a fluidized sphere with respect to the sine of the angle by which the entire apparatus is tilted. Data are shown by symbols, and prediction based on the equipartition assumption and the variance of the position and speed distributions are shown by shaded regions. The experimental conditions are the same as in Fig. 1.2.	11
1.4	[ene-dis] The energy distribution for a one-inch nylon sphere in a stadium-shaped trap, under the same conditions as in Figs. 1-2. The dashed curve shows the expectation based on the known effective temperature and the product of density of states times a Boltzmann factor. The insets show scatter plots of velocity and position components; successive points are separated by several seconds, longer than the decay time of the velocity autocorrelation function.	12
1.5	[diffdens] Speed distributions for grains of equal diameter but different density. In order of decreasing density, the sphere materials are ceramic (Al_2O_3), Teflon, Nylon, polypropylene, and wood. The inset shows a photograph of the system. . .	13

1.6	[ener] Average kinetic energy vs density, based on the speed distributions shown in Fig. 1.5 for five spheres of equal diameter. To within measurement uncertainty, the results are nearly constant, as shown by the solid horizontal line. The expectation based on single-grain behavior, Eq. (1.2), is shown by a dashed curve.	14
1.7	[TRK] (a) Kurtosis of the velocity distributions, (b) ratio of average kinetic energies, and (c) equal-time velocity cross-correlation for pairs of nylon spheres vs the ratio of their diameters. For the thermal analogy to hold, the kurtosis should equal 3, the kinetic energy ratio should equal 1, and the equal-time velocity cross-correlation should vanish. The ratio of average kinetic energies predicted by single-sphere behavior, Eq. (1.2), is shown by a dashed curve.	15
1.8	[AllShape] (a) Photographs of five aspherical grains; (b) the kurtosis of the speed distributions, and (c) the average kinetic energies, for the two translational and the one rotational degrees of freedom when these grains are individually fluidized. The airspeeds are 910 cm/s for the three leftmost grains, and 750 cm/s for the two rightmost grains.	18
1.9	[Silver] (a) The kurtosis of the speed distributions, and (b) the average kinetic energies, for the two translational and the one rotational degrees of freedom for the silver grain in Fig. 1.8, as a function of the fluidizing airspeed.	19
2.1	[VT](a) Example configurations for area fractions $\phi = 48.7\%$ and $\phi = 80.9\%$, with the big beads colored darker than the small beads. (b-c) Voronoi tessellations for these configurations with cells shaded darker for increasing coordination number and circularity, respectively.	26
2.2	[PZ] (a) Contour plot of the coordination number distribution; red is for large probability density and blue is for small. The dashed white line indicates $\phi = 0.74$. (b) Coordination number distributions for three area fractions, as labeled.	27

2.3	[PC] (a) Contour plot of the non-circularity shape factor distributions for the Voronoi tessellation polygons. Beyond about 74% (dashed white line) a well formed second peak develops and the distribution doesn't change much. (b) Shape factor distributions for a sequence of area fractions; the thick green curve is for $\phi = 74.4\%$. The labels 7, 6, 5, and 4 show the minimum shape factors for polygons with that number of sides.	29
2.4	[PCz] Non-circularity shape factor distributions for the Voronoi tessellation polygons for (a) $\phi = 48.7\%$ and (b) $\phi = 80.9\%$. Thick black curves are the actual distributions, and the thin colored curves are the contributions from cells with different numbers of sides, as labeled.	30
2.5	[GrC] The radial distribution function computed between (a,e) all beads; (b,f) big beads; (c,g) small and big beads; and (d,h) small beads. The top row shows contour plots, where blue is for large g and red is for small; the dashed white line represents $\phi = 0.74$. The bottom row shows data curves for different area fractions; the thick green curves are for $\phi = 0.744$. The grayed region represents the distance excluded by hard-core contact, for r less than the sum d_{ij} of radii.	32
2.6	[Strs] The area fraction dependence of (a) the occurrence probability $P(Z)$ of Voronoi cells with Z sides, and peak and valley values of (b) the shape factor distribution and (c) the pair distribution function for small beads.	33
2.7	[msd] (a-b) Mean-squared displacement, and (c-d) kurtosis of displacement probability distribution, all as a function of delay time. The left-hand plots are for the big beads and the right-hand plots are for the small beads. Area fraction color codes for the all plots are labeled in (c-d); the thick green curve is for $\phi = 74.4\%$. Note that the mean-squared displacement saturates at the square of the sample cell radius. The squares of bead diameters are $d_b^2 = 0.76 \text{ cm}^2$ and $d_s^2 = 0.40 \text{ cm}^2$. The square of the position resolution is $(0.0016 \text{ cm})^2 = 3 \times 10^{-6} \text{ cm}^2$	35
2.8	[Df] Density of vibrational states of frequency f , for various packing fractions; the thick green curve is for $\phi = 74.4\%$	37

2.9	[KKr] (a-b) Kinetic energy vs area fraction for big and small beads, respectively. The time-average for each bead is shown by a small dot; the ensemble average over all beads is shown by +. (c) The kinetic energy ratio of small-to-big beads. The line $(d_s/d_b)^3$ shows where beads move with the same mean-squared velocity. Insets show the same quantities on a logarithmic scale.	39
2.10	[D] (a-b) Diffusion coefficient vs area fraction for big and small beads, respectively. The time-average for each bead is shown by a small dot; the ensemble average over all beads is shown by +. (c) The diffusion coefficient ratio of small-to-big beads. Insets show the same quantities on a logarithmic scale.	41
2.11	[Inflects] (a) Logarithmic derivative λ of the mean-squared displacement vs delay time, for big beads at area fraction $\phi = 80\%$. Four special time scales can be defined from such data, as depicted: τ_b where $\lambda = 1.5$, τ_{\min} where λ is minimum, and τ_c and τ_r where λ is halfway between 1 and its minimum. (b) Density of state, for big beads at $\phi = 80\%$, marked with the four special timescales. (c) Contour plot of the logarithmic derivative for the big beads, where color indicates the value of λ , as a function of both area fraction and delay time. Red is slope two and blue is slope zero; caged dynamics are when the mean-squared displacement has a slope less than one, which is aqua-blue. The four special times defined by the behavior of λ , as well as a fifth time τ^* at which the displacement kurtosis is maximal, are superposed on the contour plot; symbol key is given in the legend.	43
2.12	[PointJ] The zero-stress plane of the jamming phase diagram showing our trajectories in phase space. The primary trajectory, which corresponds to sequences shown in all previous plots, is given by the diagonal dashed purple curve and open circles that intersects the “pre-jammed” boundary at 74% and that ends at Point J, $\{\phi \approx 0.83, T = 0\}$. The grayed region is forbidden for impenetrable beads. The solid circles, triangles, and square denote measurements of the phase boundary for large hollow polypropylene beads, solid steel beads, and tiny solid steel beads, as shown in the images. Each boundary point corresponds to the trajectory indicated by the dashed going through it.	47

3.1 [Dynamics773] Left, from top to bottom: mean square displacement $\langle r^2(t) \rangle$, average mobile bead cluster size $S_c(t)$, weight average string length $L_s(t)$, self contribution to the four point susceptibility $\chi_4^{ss}(t)$. As reported for liquids approaching their glass transition, the maximum in $S_c(t)$ and $L_s(t)$ occur in the late β - early α -relaxation regime and is maximum in the α -relaxation regime. Right: illustration depicting bead mobility over a time interval $t = 12$ seconds, $t_{S_c}^*$ or $t_{L_s}^*$, when mobile cluster size $S_c(t)$ and string length $L_s(t)$ are at a maximum. Top: beads are colored based on their relative mobility over the time window, as determined by the absolute displacement from their initial positions showing the spatially heterogeneous nature of the bead dynamics. Bottom: Close up, with vectors depicting the correlated motion of beads. Each vector connects the initial position of the bead to its new position at the end of the time window. 51

3.2 [Strings] Characterization of SHD for the most mobile beads in the system. (a) Depiction of a mobile bead cluster at $\phi = 0.773$ and $t = t_{S_c}^*$. Here, the beads with black arrows belong to a mobile bead cluster. Each arrow represents a vector drawn from the initial position of the bead to its new position $t_{S_c}^*$ seconds later. (b) Number average mobile bead cluster size $S_c(t)$ in units of number of beads as a function of ϕ . From bottom to top: $\phi = 0.597, 0.633, 0.647, 0.669, 0.693, 0.722, 0.742, 0.762, 0.773$. (c) All curves from (b) together on the same plot, where t is scaled by $t_{S_c}^*$ and $S_c(t)$ is scaled by $S_c(t_{S_c}^*)$. (d) Cluster size distribution for $\phi = 0.773$ at $t_{S_c}^*$ fit by a power law. (e) Depiction of a string of mobile beads at $\phi = 0.773$ at $t = 12$ seconds. Dark grey beads represent bead positions at $t = 0$ and light grey beads (mostly covered by dark grey beads) represent the new bead positions $t_{S_c}^*$ seconds later. Each black arrow represents a vector drawn from the initial position of the bead to its new position. (f) Weight average string length $L_s(t)$ in units of number of beads as a function of time. From bottom to top: $\phi = 0.597, 0.633, 0.647, 0.669, 0.693, 0.722, 0.742, 0.762, 0.773$. (g) All curves from (f) together on the same plot, where t is scaled by $t_{L_s}^*$ and $L_s(t)$ is scaled by $L_s(t_{L_s}^*)$. (h) String length distribution for $\phi = 0.773$ at $t_{L_s}^*$ fit by an exponential. 53

3.3	[X4] Characterization of the correlations among the most immobile (caged) beads in the system. Top: Average value of the overlap parameter $q_s(t)$ as a function of time window for different ϕ . From left to right: $\phi = 0.597, 0.633, 0.647, 0.669, 0.693, 0.722, 0.742, 0.762, 0.773$. Bottom: as a function of time window for different ϕ . In order of increasing peak height: $\phi = 0.597, 0.633, 0.647, 0.669, 0.693, 0.722, 0.742, 0.762, 0.773$. Inset: Log-log plot of as a function of time window for different ϕ . In order of increasing peak height: $\phi = 0.597, 0.633, 0.647, 0.669, 0.693, 0.722, 0.742, 0.762, 0.773$	56
3.4	[Scaling] Dependence of dynamic timescales and length scales on area fraction. Left: Best fits for the MCT equation with fixed parameter $\phi_c = 0.78$ and variable parameter x . From top to bottom: (a) Fit for timescale $t_{S_c}^*$. (b) Fit for peak height $S_c(t_{L_s}^*)$. (c) Fit for timescale $t_{L_s}^*$. (d) Fit for peak height $L_s(t_{L_s}^*)$. (e) Fit for timescale $t_{\chi_4}^*$. (f) Fit for peak height $\chi_4(t_{\chi_4}^*)$. Right: Best fits for the MCT equation with fixed parameter $\phi_0 = 0.83$ and variable parameter E . From top to bottom: (g) Fit for timescale $t_{S_c}^*$. (h) Fit for peak height $S_c(t_{S_c}^*)$. (i) Fit for timescale $t_{L_s}^*$. (j) Fit for peak height $L_s(t_{L_s}^*)$. (k) Fit for timescale $t_{\chi_4}^*$. (l) Fit for peak height $\chi_4(t_{\chi_4}^*)$. 59	59
4.1	[VDDemo] The Voronoi tessellation for six different delay-time snapshots. Each cell is permanently colored according to its location at $\tau = 0$. After a long time the beads move far from their initial positions and the cell colors mix.	62
4.2	[QTemp] (a) The average overlap order parameter and (b) its scaled variance, which is the self-part of the four-point susceptibility, for different area fractions as labelled. For high area fraction the configuration remains static for longer durations and a dynamical length scale associated with the peak of χ_4 diverges.	64
4.3	[PADemo] Persistent area images of the Voronoi tessellation for the start time $t = 0$ as a function of delay time, τ . Regions that are black have persisted within the same Voronoi cell for the time interval indicated. $A(t, \tau)$ is the total fraction of black area and is a measurement of the overlap between the configurations $\tau = 0$ and τ . Spatial heterogeneity due to stringlike cooperative motion is particularly evident at $\tau = 26$ and causes $A(t, \tau)$ to fluctuate with t	67

- 4.4 [Amv] (a) Measured persistent area $A(t, \tau)$ plotted separately for different start times t (colored curves), for $\phi = 0.792$. Its time average $\langle A(t, \tau) \rangle = A(\tau)$ (thick green curve with circles) is a measurement of the average configurational overlap for two snapshots separated by a time τ . (b) Its scaled temporal variance $\chi_A = N \cdot \sigma^2[A(t, \tau)]$ averaged over t , is a measurement of spatially heterogeneous dynamics due to strings. 68
- 4.5 [ATemp] (a) Average persistent area $A(\tau)$ for different area fractions showing two step relaxation due to intermediate time caging. (b) Its scaled temporal variance χ_A is analogous to χ_4 and measures spatially heterogeneous dynamics. In supercooled liquids and atomic and polymeric glassformers χ_4 peaks in the α -relaxation regime, about a decade after the mean-squared displacement breaks out of the subdiffusive regime, which is when spatial heterogeneity is a maximum and is exactly what we observe in χ_A 70
- 4.6 [PBDemo] Persistent bond images for $t = 0$ as a function of delay time τ . At $\tau = 0$ the persistent bond image is exactly the nearest neighbor network. For $\tau > 0$ beads change neighbors and bonds are broken. The persistent bond $B(t, \tau)$, is the fraction of nearest neighbor bonds that remain over the interval τ . Cooperatively rearranging beads cause shear zones and curled islets of broken bonds to open up. While the persistent area measures cooperative motion in the cage, the persistent bond is sensitive to large rearrangements that cause the breakup of the nearest neighbor network. 72
- 4.7 [BTemp] (a) The average persistent bond $B(\tau)$ for different area fractions showing the average rate at which the bonded network breaks apart and correlated beads move apart. (b) Its scaled temporal variance χ_B shows that spatial heterogeneity in bond breakup increases with area fraction and peaks about a decade after χ_A , when $A(\tau) = 0$ 74

- 4.8 [ZCiIvPaPb] Rows show snapshots in increments of $\tau_A^*/2$. First column: Voronoi cells at $\tau = 0$ colored according to the number of sides so that darker has more sides. Second column: Voronoi cells at $\tau = 0$ colored according to circularity so that darker is more circular. Third column: Instant velocity vectors at $\tau = \tau_A^*$ colored so that faster beads have large white vectors and slower beads have small red vectors. Fourth column: Persistent area images at $\tau = \tau_A^*$. Fifth column: Persistent bond images at $\tau = \tau_B^*$. There seems to be little relationship between structural defects and strings. Regions of highly cooperative motion in the persistent area are regions of high bond breakup much later in the persistent bond. Spatially heterogeneous dynamics evident in the average velocity vector images cause $A(t, \tau)$ and $B(t, \tau)$ fluctuate with t 76
- 4.9 [TabqLabq] (a) Characteristic peak times of the susceptibilities χ_A (blue circles), χ_B (red diamonds), and χ_4 (green triangles) as a function of area fraction. (b) Characteristic lengths computed from the ratio of the susceptibility to the order parameter, eq. 4.12. Both the characteristic times and lengths diverge on approach to the jamming transition density at $\phi = 0.83$ 79
- 4.10 [TimeLeng] Characteristic time and length scales compared with MCT scaling (left hand side) and the VFT equation (right hand side). Divergence is consistent with both power law or exponential form which is also true in supercooled liquids. Best fits to power law divergence for all topological measurements yield MCT critical density at $\phi = 0.79 \pm 0.01$ which is consistent with previous findings. Best fits to exponential divergence yield VTF critical density at $\phi = 0.83 \pm 0.01$ which is close to point J and also consistent with previous findings. 80
- 5.1 [PrvE] (a) Distribution of displacements of the thermometer's center-of-mass mechanical equilibrium. Inset shows a diagram of the thermometer with angular displacement θ . (b) Distribution of thermometer velocities. Inset shows a correlation plot between position and velocity. (c) Distribution of energies. Inset shows the probability distributions of the gravitational potential energy and kinetic energy binned simultaneously and fit to an exponential. 88

5.2	[Spectra] (a) Power spectra of the x- and y-position displacement from mechanical equilibrium of the thermometer. (b) Power spectra of the x- and y- velocity of the thermometer. The spectra show peaks at the natural frequency of the oscillator. . . .	89
5.3	[MSDVV] (a) The mean-squared displacement and (b) velocity autocorrelation function for motion perpendicular to the oscillating force. The diffusion coefficient is measured from the MSD after the mean collision time (when the VAF is first zero) and before saturation at the system size. It is nearly constant over the range 1 to 100 s.	91
5.4	[Fv] Scatter plot of the oscillating external force as a function of velocity response parallel to the force. The mobility is measured by computing the slope of the best linear fit to the scatter plot.	92
5.5	[Mobs] Mobility of test particles of different mass as a function of oscillating force frequency. Except for the very heavy steel ball, the data all collapse and are constant for frequencies less than 10 Hz.	93
5.6	[TbTtTe] (a) Comparison of the three effective temperatures as a function of area fraction and (b) as a function of the driving airspeed. The 5th area fraction in (a) is the same driving-density phase space point as the 5th airspeed in (b) and roughly marks when the Einstein temperature deviates from the other two temperatures. . . .	95
6.1	[setup] A picture and diagram of the experimental apparatus as seen from the side.	98
6.2	[BigTop] (a) An image from above of the top layer of sand taken with a high speed camera at 1000 fps. To obtain velocity profiles, the particles can either be tracked or the images can be analyzed using PIV. (b) The flow profile across the cell is nearly parabolic parallel to the flow and constant perpendicular to the flow. The fluctuations of the surface layer obtained by particle tracking are consistent with SVS measurements. (c) The flow profile down the cell is constant both parallel and perpendicular to the flow. The fluctuations are also nearly constant. For $x > 2.0$ cm the sand is close to the end of the channel and begins to free fall.	99

6.3 [SmallSand] (a) The flow speed in the x-direction as a function of y, which is the channel width. The profile is more plug-like than for the bigger spheres. Inset is an image taken from above of flowing glass spheres 100 μm in diameter. (b) The flow speed in the x-direction as a function of z, depth. The flow profile decreases non-linearly as a function of z and drops to zero abruptly after about 20 layers. Inset shows an image of the flowing spheres taken from the side. 100

6.4 [TwoAvalanches] (a) A blow-up of the speckle pattern as a function of time for an avalanche turning on, indicated as the gray region in the figures below. When the sand is static the speckle pattern is constant so that there are constant streaks as a function of time. As an event turns on the speckle pattern begins to change. The variance ratio (b) is a measurement at each instant of the ratio of the visibility of the speckle pattern for the exposure times of T and 2T. When the speckle pattern is constant (streaks) its visibility is independent of exposure duration, and thus the variance ratio is one. When the sand flows, the speckle pattern changes quickly and a short exposure is necessary for a sharp speckle pattern so that the variance ratio is less than one. Using theory for back-scattered light the variance ratio can be transformed numerically into the fluctuation speed of the sand particles (b) which is the granular temperature. 104

6.5 [Statistics] The top row shows correlation plots between time and velocity observables. (a) Correlation between the time-on in time-off durations. (b) Correlation between the maximum fluctuation speed in the time-off duration (c) Correlation between the time-on duration and the maximum fluctuation speed. The bottom row shows histograms for time and velocity observables. (e) Histogram of the duration of avalanche events. (f) Histogram for the duration between avalanche events. (g) Histogram of the maximum fluctuation speed obtained by the avalanche. 107

6.6 [AverageAva] Mean and standard deviation of $\delta v(t)$ timetraces for 1500 avalanche events. Inset is mean and standard deviation with $\delta v(t)$ normalized by δv_{max} and t normalized by t_{on} 110

6.7	[AverageOn] A blow up of turning on region of the average $\delta v(t)$ with its contour in the background. A wavefront is at the head of each avalanche. After the wavefront passes the sand flows smoothly and gradually slows down until finally turning off. .	111
6.8	[AverageOff] The standard deviation normalized by the mean for 1500 separate avalanche events. Inset shows the same quantity calculated for the normalized avalanche shape.	112
6.9	[PSFull] The power spectrum of $\delta v(t)$ for the full 35 hour data set. At very high frequencies the power spectrum saturates to a constant demarcating our measurement accuracy. Frequencies from 0.1 to 100 Hz are timescales of the sand fluctuation speed during avalanche flows and show $1/f^2$ noise. There is a peak at the average avalanche frequency and $1/f$ noise for lower frequencies having to do with spectrum many avalanche cycles.	113
A.1	[CrystalMSD] The images show the system from above for low density and high-density snapshots. At low densities the system is liquid-like and at high densities it is crystalline. In the figure below we plot the mean square displacement for the four densities studied. At low densities the mean-squared displacement is diffusive at long times and high densities when the system crystallizes, the the mean-squared displacement saturates at long times. Because of turbulent fluctuations, all densities are ballistic at short times.	119
A.2	[CrystalDOS] Density of phonon states computed from the Fourier transform of the velocity autocorrelation function, averaged over all particles. At low densities the system is liquid-like and the density of states approaches a finite value at zero frequency, so that there are translational modes. At high density when the system is crystalline, the density of states goes like $1/f$ as $f \rightarrow 0$, which is consistent with that of an atomic crystal in two dimensions.	120

A.3	[High] Three representations of the system for $\phi = 0.838$. The first shows a raw image of the system taken from above, showing that the system is clearly crystallized. The second image shows the Voronoi diagram for the configuration with each cell colored according to its number of sides. Six sided cells dominate the diagram, with a few 5-7 dislocations strewn throughout. The third image shows a snapshot of the average velocity vectors for a time interval 10 seconds. The velocity vectors tend to align along the crystal axis and are reminiscent of phonons.	121
B.1	[CircleImages] Representative snapshots of two different area fractions, as labelled. The big grains are colored light gray and small grains are colored dark gray and have a size that is proportional to the relative size of the grain. Clustering is evident at both area fractions and may be a consequence of negative non-additivity of the bead-air-bead interaction.	124
B.2	[StructureFactor] The structure factor averaged over time and space as a function of wave vector for three different area fractions. Each component of the structure factor is measured separately, as labelled. (a) Taking all beads, (b) big-big, (c) small-big, and (d) small-small contributions. The expectation based on previous negative non-additive systems is a small wave vector peak in the small-small component [36].	125
B.3	[EulerCharacter] The Euler characteristic measured for (a) the big beads and (b) the small beads, as a function of covering, or effective particle size. The curves have been displaced by a constant for clarity.	126
C.1	[Timetraces] The top figure is for a low density system, $\phi = 0.487$ and the bottom figure is for a high density system, $\phi = 0.792$. The light blue curve shows the percentage of cells that have six sides at given instant for each configuration. The purple curve shows the percentage of cells that have circularity within half a standard deviation of the average at each instant for each configurations. These topological averages fluctuate in a way that depends on the grain dynamics and structure.	129

C.2	[PSP6PC] The power spectra for the time traces $P_6(t)$ and $P_c(t)$. Previous simulations of atomic or colloidal systems undergoing a glass transition found that the power spectra is constant for a liquids and approached $1/f$ at low frequency for a crystal or glass.	131
D.1	[OOPParameterChoice] (a) The overlap order parameter for a single area fraction for different choices of the overlap threshold size. (b) The four-point susceptibility for different choices of overlap threshold size. (c) The size of the correlated motion for different choices of overlap threshold size. All measurements are sensitive to the overlap threshold size.	133
E.1	[MSDs] The mean-squared displacement for the two area fractions of interest. The blue curve is for low area fraction and the red curve is high area fraction. The low-density system saturates at long time due to finite system size, and is super diffusive between 10 and 100 seconds due to a convection current.	136
E.2	[S3rAF737] $\phi = 0.737$. The left-hand side shows the third moment of the structure function and the right hand side shows the average velocity vector fields for different time intervals as labelled. The time labels are in number of frames which are intervals of 1/120 of a second.	138
E.3	[S3rAF802] $\phi = 0.802$. The left-hand side shows the structure functions and right-hand side shows the average velocity vector fields exactly as in fig. E.2.	140
F.1	[SmallAvalanches] (a) Example time trace of $\delta v(t)$ for intermittent avalanche flow. Each spike corresponds to a single avalanche event. (b) Zoom of the grayed region in figure (a) showing the time trace for a single small sand avalanche.	143
F.2	[SmallStats] (a) Scatter plot of the time-on and time-off durations. (b) Scatter plot of the maximum fluctuation speed and the time-off duration (c) Scatter plot of the time-on duration and the maximum fluctuation speed. (d) Histogram of the duration of avalanche events. (e) Histogram of the duration between avalanche events. (f) Histogram of the maximum fluctuation speed obtained by the avalanche.	144

F.3	[SmallShape] Average dynamical avalanche shape of 450 separate small sand events. Similar to the big avalanches, there is a peak just after the turning on time followed by a gradual deceleration as the sand supply is exhausted. There is a high tail at long times $t > 3$ s, which is an artifact of averaging together events of different duration. Inset shows the average avalanche normalized shape, in which each event is normalized by its maximum fluctuation speed and its total time-on. Normalization removes the long time tail and uncovers the universal avalanche shape and the abruptness of jamming at turning-off.	146
F.4	[SmallOn] Zoom in of the average turning on of small sand avalanches. In the background is a contour plot of $\delta v(t)$ showing that there are three stages in how avalanches turn on.	148
F.5	[PSSmall] Power spectrum of the full $\delta v(t)$ time trace for the small sand avalanches. As with the big sand avalanches of Chapter 6, the small sand power spectrum is $1/f^2$ for timescales corresponding to the avalanche on-state, and has a gentle peak at the inverse average avalanche period. At low frequencies there are spurious oscillations which may be due to the on-state comprising a relatively small portion of the avalanche cycle so that $\delta v = 0$ most of the time.	149

Introduction

The study of driven-dissipative systems remains at the frontiers of scientific research today. Such systems are of vast importance to industry and technology and are found in natural forms at all scales in the universe. The theory that will one day describe all these systems has not yet been invented and many of the most basic questions remain unanswered. It has been my motivation to contribute in some small way to the understanding of such systems and provide information that might one day lead to a full theory. The majority of my work can be broadly separated into two types of experiments: grains fluidized by air in a two-dimensional bed and granular avalanches flowing in three-dimensions.

My first experiment originated as an extension of work done by my predecessor, Rajesh Ojha. Rajesh found that a single ping-pong ball fluidized by air in two dimensions perfectly obeyed statistical mechanics, as though it were a Brownian particle trapped in a harmonic potential. But what exactly validated a statistical mechanical description of this driven-dissipative system? To shed light on this question we went back to the original experiment and systematically changed one property at a time in an attempt to uncover a governing parameter. In the end, we discovered that statistical mechanics broke down gradually when grains of different size were fluidized together. The degree to which the system approximated a statistical mechanical system depended on the degree to which the grains differed in size, which provided us with a parameter that gradually caused the breakdown of the statistical mechanical description. This is the focus of Chapter 1.

We then turned our attention to disordered collections of many balls. The idea was to understand the behavior of granular systems driven at various proximities to the jamming transition. In fact, we found that at low densities and high driving the system behaved like a simple liquid both structurally and dynamically, and at high densities and low driving it behaved like a glass. By

keeping track of hallmark changes in the structure and dynamics for different driving and density trajectories, we were able to map out a jamming phase diagram. This is the focus of Chapter 2.

The analogy of granular systems close to jamming to atomic systems undergoing a glass transition extends beyond single particle quantities like the the mean-squared displacement and pair correlation functions measured in Chapter 2. For systems close to jamming, peculiar dynamical features become evident over a range of timescales. These dynamical “strings” resemble spatially heterogeneous dynamics that have been seen in simulations of supercooled liquids and dense colloidal suspensions near the colloidal glass transition. To explore the analogy further we study these strings as a function of proximity to the jamming transition, by changing the systems density and energy, and find that the dynamical time and length scales they represent diverge with proximity to the jamming point. In fact, this divergence was found to be consistent with Mode-Coupling Theory and the Vogel-Tammann-Fulcher equation, which further fuels the universality of the jamming concept. This is the focus of Chapters 3 and 4.

To understand how the short time agitation of the balls gives rise to long time relaxation, and also the further analogy with atomic systems undergoing a glass transition, we study three different effective temperatures that in thermal equilibrium are all equal. We measure the average kinetic energy of the balls, often called the granular temperature, and an Einstein temperature based on the ratio of diffusion to mobility. We also measure a novel effective temperature using a “granular thermometer” which is a weighted sphere oscillator. By measuring the average potential energy stored in the oscillator and its rotational and translational velocity we are able to measure yet another effective temperature having to do with grain collisions. With these three effective temperatures we are able to witness how the approach to jamming effects the validity of key players in statistical mechanics, such as ergodicity, equipartition of energy, and the fluctuation dissipation theorem, on which all three effective temperatures rely in order to be equal in thermal equilibrium. This is the focus of Chapter 5.

In the final chapter we present a light scattering study of avalanche flows. Even though sand-piles are some of the most common objects in the natural world, the fundamental physics that governs them is poorly understood, and as a result many basic questions, like what the average fluctuation speed of the sand particles is and how frequently avalanches occur, are yet to be answered. By using a new light scattering technique called Speckle-Visibility Spectroscopy, we are

able to observe the dynamics of 1500 independent avalanches down to the angstrom scale and over ten orders of magnitude in time. This light scattering technique and the enormous dynamical range it avails us is essential because as we will show, avalanche dynamics span an enormous range of time and length scales. With this data we are able to present for the first time the average dynamical shape of an avalanche, along with its frequency statistics and characteristic features. This is the focus of Chapter 6.

In the appendix are descriptions of short experiments and additional scientific analysis. The final appendix has technical details that will be useful to researchers attempting to reproduce these results or take them a step further.

Chapter 1

The Partition of Energy for Air-Fluidized Grains

1.1 Abstract

The dynamics of one and two identical spheres rolling in a nearly-levitating upflow of air obey the Langevin Equation and the Fluctuation-Dissipation Relation [Ojha *et al.* Nature **427**, 521 (2004) and Phys. Rev. E **71**, 016313 (2005)]. To probe the range of validity of this statistical mechanical description, we perturb the original experiments in four ways. First, we break the circular symmetry of the confining potential by using a stadium-shaped trap, and find that the velocity distributions remain circularly symmetric. Second, we fluidize multiple spheres of different density, and find that all have the same effective temperature. Third, we fluidize two spheres of different size, and find that the thermal analogy progressively fails according to the size ratio. Fourth, we fluidize individual grains of aspherical shape, and find that the applicability of statistical mechanics depends on whether or not the grain chatters along its length, in the direction of airflow.

1.2 Introduction

There is a growing list of driven, far-from-equilibrium systems where the dynamics of microscopic fluctuations are characterized by an effective temperature. One of the earliest examples is the kinetic energy associated with velocity fluctuations in a sheared granular material [5]. More recent

examples in granular physics include dilute grains driven within a horizontal plane [39, 63, 64, 69, 70], as well as flowing granular liquids [48, 51, 56, 58, 76, 82] and vertically-vibrated granular gasses [6, 16, 27, 65]. Wider ranging examples include chaotic fluids [26, 37], spin glasses [15, 35], glasses [9, 34], colloids [7, 77], and foams [67], which are all far away from equilibrium. In some of these cases [6, 16, 64, 67, 70, 82], the behavior is in perfect analogy with that expected for a system in thermal equilibrium, which we denote the thermal analogy. In other cases, the thermal analogy is more limited and does not hold in detail.

An outstanding question is how to predict whether or not the thermal analogy holds. What do the systems in Ref. [6, 16, 63, 64, 67, 70, 82] have in common, and how do they differ from other driven systems? Here we seek insight by systematically perturbing one case for which the analogy unarguably holds in all detail, in hopes that it may be progressively upset. In particular, we focus on a small number of grains fluidized in a nearly-levitating upflow of air. While grains thus never leave the plane, they can nevertheless be driven randomly within the plane by the random shedding of turbulent wakes at a rate set by the Strouhal number [3, 84]. The Reynold's number based on sphere size is of order 10^4 . Under these conditions, a single sphere confined to roll within a circular cell behaves exactly like a Brownian particle in a two-dimensional harmonic trap [64]. Specifically, the dynamics obey a Langevin equation where the random force autocorrelation is proportional to the viscous drag memory kernel and the effective temperature according to the Fluctuation-Dissipation Relation. For a variety of conditions, the root-mean-squared displacement of the sphere from the center of the trap and the mean-squared speed of the sphere, respectively, are given by [63]

$$\sqrt{\langle r^2 \rangle} = (0.20 \pm 0.02) R_{cell}, \quad (1.1)$$

$$\langle v^2 \rangle = 0.7 \left(\frac{\rho_{air}}{\rho_e} \right) \frac{u^3}{\sqrt{gD}}. \quad (1.2)$$

Here $\rho_e = m_e / [(4/3)\pi(D/2)^3]$; $m_e = m + I/(D/2)^2$ is the effective inertial mass of the sphere; m , I , and D are respectively the mass, moment of inertia and diameter of the sphere; ρ_{air} and u are respectively the density and flow speed of the air; $g = 980 \text{ cm/s}^2$ is gravitational acceleration; and R_{cell} is the radius of the sample cell. Physically, Eq. (1.2) can be understood by balancing energy input, via collision between the sphere and a sphere-sized volume of air, with energy dissipation via viscous drag. Geometrically, Eq. (1.1) can be understood by a picture of the repulsion between

the cell wall and the turbulent wake, which expands as it moves downstream.

The detailed thermal analogy for the behavior of one and two nearly-levitated gas-fluidized spheres was completely unexpected. In this paper, we seek insight via systematic perturbation of the original experiment. To begin, we first describe the experimental apparatus and analysis procedures used throughout. In the next four sections, we describe the perturbations and results, with one perturbation per section. We shall demonstrate that the thermal analogy is very robust with respect to some of these perturbations. We also shall demonstrate a control parameter by which the thermal analogy may be progressively upset.

1.3 Experimental Details

Our methods for fluidizing grains and tracking their positions are similar to those of Refs. [63,64], but with some embellishments that we describe in detail here. As before, the heart of the apparatus is a rectangular windbox, $1.5 \times 1.5 \times 4 \text{ ft}^3$, standing upright. A circular sieve with mesh-size $300 \mu\text{m}$ sits in a twelve-inch circular hole on the top. All systems except in the final section on non-spherical grains, have spherical grains that roll without slipping. The rotational motion is therefore coupled to the translational motion in such a way that the observed kinetic energy has both a translational and rotational part, together defined as m_e . A digital CCD video camera [Pulnix 6710, 8 bits deep, 120 Hz frame rate], and a ring of six 100 W incandescent lights, are located approximately three feet directly above the sieve, mounted to a scaffolding which in turn is mounted to the windbox. A blower is connected at the base of the windbox to provide air flow perpendicular to the sieve. A hotwire anemometer measures the flow rate, and verifies its uniformity. Previously, a perforated metal sheet was fixed in the middle of the windbox to break up large scale structures in the airflow. To ensure even more uniformity, we now use two perforated metal sheets with a one-inch thick foam air-filter sandwiched in between.

The control of the camera, and all image processing, is accomplished within LabVIEW. In all runs, images are harvested at 120 frames per second and written to hard disk for post processing. To minimize the size of the data set, and hence to optimize the maximum possible run length, we first threshold the images to binary so that each grain appears as a white blob on a black background. The illumination and thresholding level are adjusted so that each blob corresponds closely to the

entire projected area of the grain. Successive binary images are encoded as a lossless format AVI movie [Microsoft RLE]. Previously, we used a custom encoding scheme that is optimal only for a very small number of spherical grains. The AVI format requires more disk space, but is also more flexible for large numbers of grains.

For post-processing we also use LabVIEW. If the grains are far enough apart as to be distinct white blobs that are completely surrounded by black, then we use LabVIEW's "IMAQ Particle Analysis" algorithm to locate the center of brightness of each blob to sub-pixel accuracy. However, when two grains collide and their blobs touch, then this algorithm identifies only the center of the combined two-grain blob. When the total number of identified blobs falls below the known number of grains, we must modify our tracking procedure. The widely-used technique of Ref. [14] cannot be invoked, because it requires the grain separation be large compared to grain size. Instead, we apply an erosion algorithm, in which a square mask consisting of ones (white) and zeros (black) is run over the binary image. The output at each pixel is one (white) if *all* the image pixels under the white region of the mask are white; otherwise, the output is zero (black). For spherical grains, we choose a mask that is about $2/3$ the grain size, that is white throughout the largest inscribed circle, and that is black outside this region. This construction preserves the circular shape of the blobs, while eroding their size. It also separates blobs that are in contact, and optimizes their circularity after separation. After applying such an erosion, we then invoke the same centroid-finding algorithm as before. These procedures are demonstrated in Fig. 1.1, which shows two grains before, during, and after collision.

There are two more steps. First, the grain coordinates measured in each frame must be identified with the correct corresponding grains in the previous frame. This is aided by the fast frame rate of our camera, which is such that the maximum displacement in one frame is much less than the grain size. Finally, position vs time data is fitted to a third order polynomial within a window of ± 5 points in order both to smooth and to differentiate to second order. Gaussian weighting that is nearly zero at the edges is used to ensure continuity of derivatives. The rms deviation of the raw data from the polynomial fit is 0.001 cm, which we take as an estimate of position accuracy. This and the frame rate give an estimate of speed accuracy as 0.1 cm/s. Indeed these numbers correspond to a visual inspection of the level of noise in time traces.

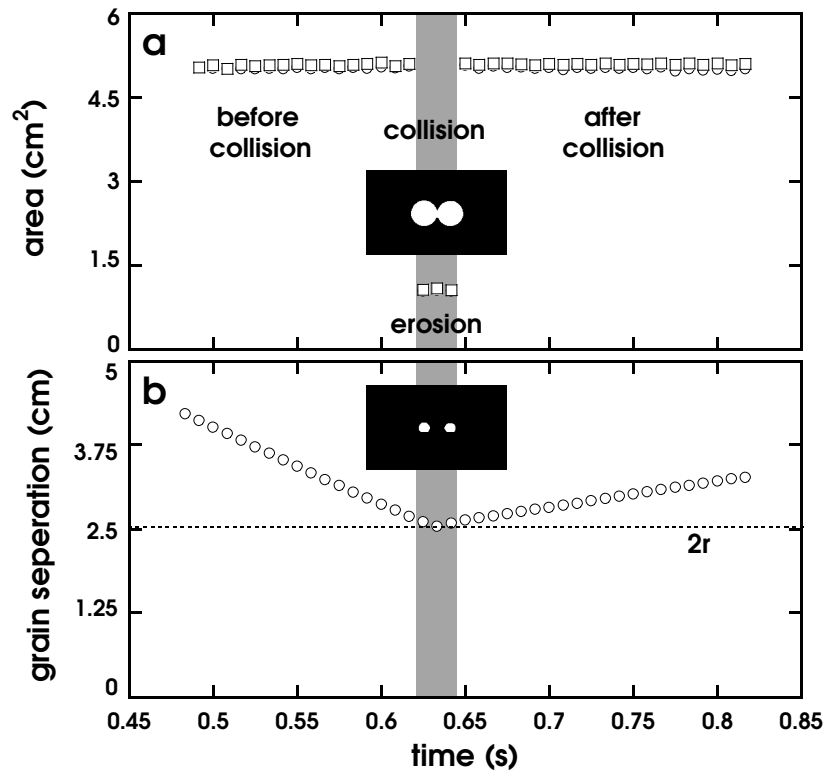


Figure 1.1: [col-ero] (a) The measured blob area, and (b) the center-to-center separation, for two grains in collision. The grayed stripe indicates the time-interval the grains are in close enough proximity so that they are imaged as as a single blob (inset a), which is eroded producing two smaller separated circles (inset b).

1.4 One Sphere in a Stadium

Our first perturbation is motivated by the very form of Eq. (1.1), which says that the rms position of the sphere is set by the radius of the sieve. So instead of using a circular sample cell, we now construct a stadium-shaped sample cell by placing appropriate wooden inserts into contact with the sieve both above and below. A binary image of our stadium, with a $D = 1$ inch nylon sphere, is shown in the inset of Fig. 1.2. Certainly the elongated boundary will affect the confining ball-wall potential, with the sphere expected to move farther along the long axis. Due to loss of symmetry, the rms position and speed of the grain could now both be different along the long and short axes, which would be a direct violation of the thermal analogy. To investigate, we fluidize the nylon ball with an upflow of air at speed $u = 750$ cm/s, and we track its position with the methods described above.

Results for the position and speed distributions along the two axes of the stadium are displayed in Fig. 1.2. All four have the same Gaussian shape, characteristic of a Brownian particle in a harmonic trap, as seen before. Though it couldn't have been expected, the speed distributions remain identical along the two axes. However, the sphere now has wider excursions along the longer axis. The observed rms displacements are $\sqrt{\langle x_l^2 \rangle} = 0.83$ cm and $\sqrt{\langle x_s^2 \rangle} = 0.56$ cm for long and short axes, respectively. The ratio of these displacements is 1.48, which is very close to the ratio of long to short dimensions of the stadium, $(22.8 \text{ cm})/(15.2 \text{ cm})=1.50$, in accord with the scaling of Eq. (1.1).

If the thermal analogy holds for the translational, in addition to the kinetic, degrees of freedom then the spring constants along the two axes can be deduced from the equipartition of energy:

$$T = m_e \langle v_l^2 \rangle = m_e \langle v_s^2 \rangle = k_l \langle x_l^2 \rangle = k_s \langle x_s^2 \rangle, \quad (1.3)$$

where T is the effective temperature measured in units of energy. To test this relation, we compare with an auxiliary mechanical measurement of the spring constants. As in Ref. [64], we tip the entire apparatus by a small angle θ away from horizontal and measure the shift $\langle \Delta x \rangle$ in the average position of the sphere down the plane. The new average position is where the spring force balances the force of gravity acting within the plane:

$$k \langle \Delta x \rangle = mg \sin \theta. \quad (1.4)$$

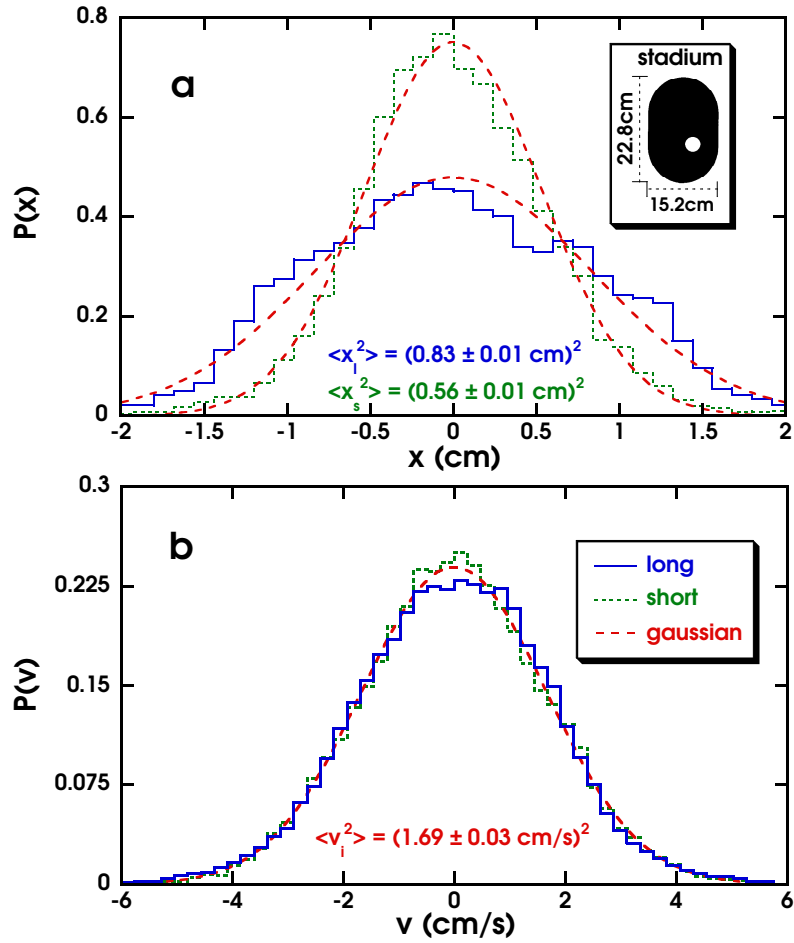


Figure 1.2: [v-p-dis] Distributions for the components of (a) position and (b) velocity measured for a one-inch nylon sphere along the long and short axes of a stadium-shaped trap. To within measurement uncertainty, these distributions are all Gaussian (dashed curves).




Figure 1.3: [spr-cons] Shift in average position of a fluidized sphere with respect to the sine of the angle by which the entire apparatus is tilted. Data are shown by symbols, and prediction based on the equipartition assumption and the variance of the position and speed distributions are shown by shaded regions. The experimental conditions are the same as in Fig. 1.2.

This is done for orientations of the stadium with the long axis both parallel and perpendicular to the tilting direction. The results for the shift in average position are plotted as symbols vs the sine of the tilt angle in Fig. 1.3. The expectations based on Eq. (1.3) and the position and speed statistics are also plotted in Fig. 1.3, now as a shaded region that reflects measurement uncertainty in the rms displacements and speeds. Indeed, the two results agree well.

To drive home the validity of the thermal analogy for a nearly-levitated sphere in a stadium-shaped cell, we now compute the total mechanical energy E as the sum of kinetic and potential energies at each instant in time. We then compute the distribution of total energy sampled over the entire run. The data, shown in the main plot of Fig. 1.4, agree nicely with the expectation for a thermal system, $P(E) = (E/T^2) \exp(-E/T)$, which is given by the density of states times the Boltzmann factor with no adjustable parameters. The insets show no correlation in phase-space scatter plots of speed vs position.

While we might have hoped to tune the validity of the thermal analogy by the value of the aspect ratio of the sample boundary, apparently it is robust with respect to this perturbation and we




Figure 1.4: [ene-dis] The energy distribution for a one-inch nylon sphere in a stadium-shaped trap, under the same conditions as in Figs. 1-2. The dashed curve shows the expectation based on the known effective temperature and the product of density of states times a Boltzmann factor. The insets show scatter plots of velocity and position components; successive points are separated by several seconds, longer than the decay time of the velocity autocorrelation function.

must look elsewhere.

1.5 Five Spheres of Different Density

Our second perturbation is motivated by the form of Eq. (1.2), which specifies the mean-squared speed of individually fluidized spheres as a function of the air and sphere properties. Note that Eq. (1.2) gives the scaling of the mean kinetic energy with sphere density as $K \sim 1/\rho_e$, due to the way energy is injected by turbulent wakes. If this relation also holds when multiple spheres of different density are in the sample cell, then the spheres would have different temperatures. To test this possibility we now simultaneously fluidize five solid spheres of the same diameter, $D = 2.54$ cm, but of different density. The materials and effective densities of the spheres are as follows: wood 0.95 g/cc; polypropylene 1.29 g/cc; nylon 1.57 g/cc; teflon 3.02 g/cc; Al_2O_3 ceramic 5.33 g/cc. We note that the wooden ball is slightly aspherical, and its diameter is about 0.5% smaller than the others. The air speed is $u = 600$ cm/s, the trap is circular, and the sieve is perpendicular to gravity.




Figure 1.5: [diffdens] Speed distributions for grains of equal diameter but different density. In order of decreasing density, the sphere materials are ceramic (Al_2O_3), Teflon, Nylon, polypropylene, and wood. The inset shows a photograph of the system.

The normalized speed distributions are displayed in Fig. 1.5 for the five spheres. All have the same Gaussian form as for a thermal particle in two dimensions, $P(v) = (2v/\langle v^2 \rangle) \exp(-v^2/\langle v^2 \rangle)$. But evidently the lighter spheres move faster, on average, than the denser spheres. The mean kinetic energy for each sphere is plotted in Fig. 1.6 vs effective density. There is a slight upward trend, but to within measurement uncertainty the mean kinetic energy is almost constant. The prediction of Eq. (1.2), $K \sim 1/\rho_e$, is completely wrong for multi-ball systems. The spheres exchange energy, perhaps through interaction of their wakes as well as through direct collisions, and thereby come to the same temperature. Again, the thermal analogy is robust with respect to perturbation.

1.6 Two Spheres of Different Size

Our next perturbation is also motivated by the form of Eq. (1.2), which implies that the mean kinetic energy of an individually fluidized sphere scales with sphere diameter as $K \sim D^{5/2}$. If this holds when multiple spheres of different diameter are simultaneously fluidized, then the spheres would have different temperatures. To test this possibility, we fluidize pairs of nylon spheres of different diameter. The airspeed is $u = 770$ cm/s and the trap is circular. By varying the choice of




Figure 1.6: [ener] Average kinetic energy vs density, based on the speed distributions shown in Fig. 1.5 for five spheres of equal diameter. To within measurement uncertainty, the results are nearly constant, as shown by the solid horizontal line. The expectation based on single-grain behavior, Eq. (1.2), is shown by a dashed curve.

spheres, we have examined the behavior for seven diameter ratios ranging from about 0.5 to 5.

The speed distributions are always nearly Gaussian. This is quantified in Fig. 1.7a, which shows the kurtosis $\langle v^4 \rangle / \langle v \rangle^4$ of the speed distribution for each sphere as a function of diameter ratio. The values are close to 3, the Gaussian expectation, except for two cases. This in and of itself is a violation of the thermal analogy. However, it is not so drastic that the mean kinetic energy, and hence the effective temperature, become ill-defined.

The ratio of mean kinetic energies of the two balls, which equals the ratio of their effective temperatures, is plotted vs diameter ratio in Fig. 1.7b. The kinetic energies are nearly equal for diameter ratios of less than two. But for increasing diameter ratio, the larger sphere becomes progressively hotter than the smaller sphere. Evidently, the diameter ratio is a control parameter that can be varied to systematically break the thermal analogy. This breakdown appears to be quite gentle, though. The temperature ratio is not as great as expected by Eq. (1.2), which again we find to be incorrect for multi-ball systems. Also, the leading behavior is not linear, but rather quadratic in the diameter ratio. This may be amenable to theoretical modelling.

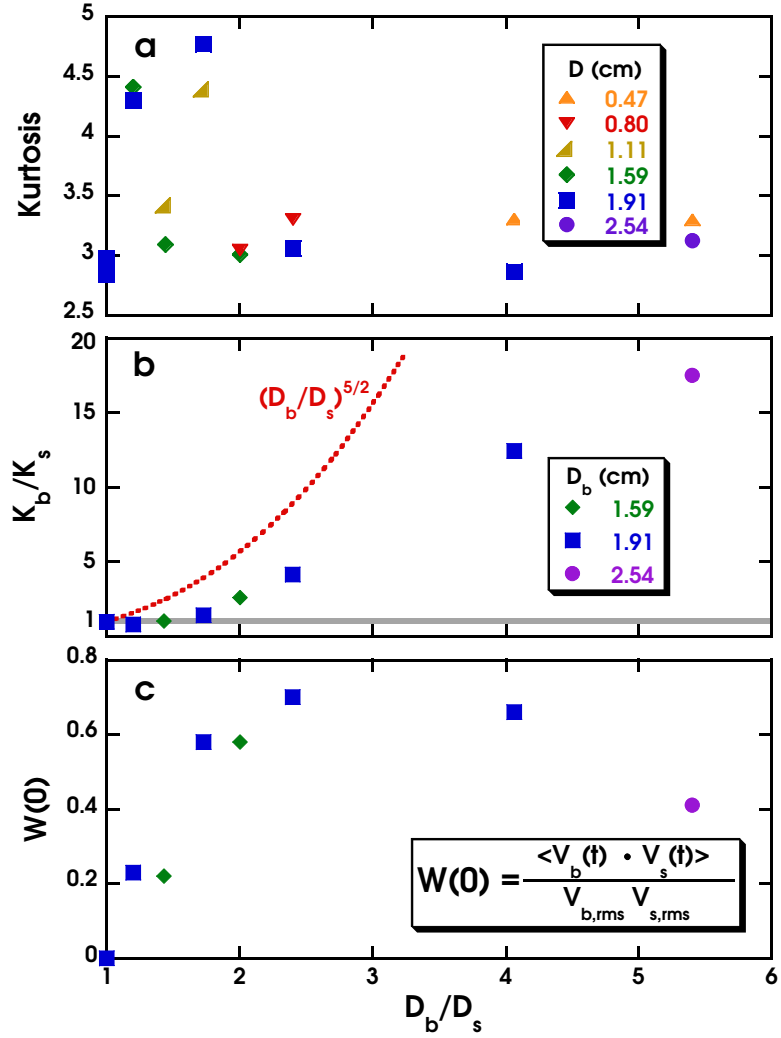


Figure 1.7: [TRK] (a) Kurtosis of the velocity distributions, (b) ratio of average kinetic energies, and (c) equal-time velocity cross-correlation for pairs of nylon spheres vs the ratio of their diameters. For the thermal analogy to hold, the kurtosis should equal 3, the kinetic energy ratio should equal 1, and the equal-time velocity cross-correlation should vanish. The ratio of average kinetic energies predicted by single-sphere behavior, Eq. (1.2), is shown by a dashed curve.

Before closing this section, we now consider the physical origin of the breakdown of the thermal analogy vs diameter ratio. The reason, actually, is immediately obvious when viewing the system directly. The two spheres usually repel one another through interaction of their wakes, as discussed in Ref. [63]. However, if they approach close enough, then they come into lasting contact and the upflow of air exerts a net total force on the pair causing them to accelerate straight across the cell until reaching the boundary. The direction of motion is such that the large sphere appears to chase the small sphere out of its territory. We speculate that the loss of symmetry of the two-ball pair causes the vortices to be shed preferentially along the line of centers, resulting in a net force.

The prevalence of this “chasing” phenomenon may be quantified by the equal-time velocity cross-correlation, $\langle v_b(t)\dot{v}_s(t) \rangle$, where the subscripts denote “big” and “small” spheres as before. For the thermal analogy to hold, this quantity must vanish because all kinetic degrees of freedom must be independently populated. During a chase, however, the two velocities are equal and hence perfectly correlated. Data for the equal-time velocity cross-correlation, made dimensionless by the rms speeds of the two balls, are plotted vs diameter ratio in Fig. 1.7c. By contrast with the effective temperature ratio, this rises abruptly from zero for diameter ratios greater than one. Also by contrast, it reaches a peak for a diameter ratio of about 3 and then decreases. If the size disparity is too small, then the loss of symmetry is not enough to cause much chasing. If the size disparity is too great, then the large grain slowly rolls without regard for the small ball, which quickly flits about and is repelled as though from a stationary object. We believe the thermal analogy is recovered in this limit, but with the two balls being differentially “heated” by the upflow of air.

1.7 One Aspherical Grain

In the above sections, and also in Refs. [63, 64], the shape of the grain is spherical. This is special because it permits the grain to roll freely in all directions without sliding. It is also special because it permits vortices to be shed equally in all directions. To explore for qualitatively new phenomena, and to seek another means of violating the thermal analogy, we now perturb the grain shape. The objects we fluidize are listed in Table 1.1 and pictured in Fig. 1.8: various pharmaceutical pills, a cylindrical wooden rod, and a dimer consisting of two connected hollow plastic spheres. When individually fluidized by an upflow of air, the pictured grains all translate and rotate seemingly at

Name	ρ (g/cc)	L (cm)	W (cm)	H (cm)	m (g)	I (g-cm ²)
white	0.685	2.12	0.848	0.848	0.82	0.344
silver	0.951	1.52	0.586	0.586	0.39	0.083
brown	0.937	1.94	0.966	0.888	1.56	0.611
wood	0.671	4.75	0.540	0.540	0.73	1.386
dimer	0.256	5.08	2.540	2.540	4.40	11.34

Table 1.1: [aspherical] The density, length, width, height, mass, and moment of inertia of five aspherical grains. For computation of density and moment of inertia, the white, silver, and wood grains are approximated as cylinders, while the brown grain is approximated as a block. The dimer is composed of two polypropylene shells of thickness 0.14 cm.

random. A few of the grains are axisymmetric, like the dimer; however, they exhibit virtually no rotation about the axis of continuous symmetry. While the spheres in previous sections roll without slipping, here the aspherical grains must slide in order to translate or rotate.

We may characterize the motion of these grains in terms of the time dependence of their center-of-mass position and their angular orientation. The former is deduced as per the spherical grains from the center of brightness. The latter is deduced from second moments of the spatial brightness distribution. Then we differentiate to measure both translational and rotational speed distributions. None of the grains is chiral by design, since that would lead to steady whirling in one direction. Nonetheless, some whirling can occur due to imperfections in shape. Therefore we measure both the average angular speed $\langle\omega\rangle$ as well as fluctuations $\delta\omega$ about this average.

A summary of the results for all grains, individually fluidized, is shown in Fig. 1.8. The kurtosis of the translational and rotational speed distributions is shown in the top plot. The results appear statistically greater than 3, the gaussian result, except for the translational velocity components of the dimer. The average kinetic energies are shown in the bottom plot. They too exhibit a violation of the thermal analogy since the translational kinetic energy is greater than the rotational kinetic energy. At this airspeed, the energy associated with whirling, $I\langle\omega\rangle^2/2$, is at least ten times smaller than the energy of angular speed fluctuations, $I\langle\delta\omega^2\rangle/2$. Therefore shape imperfections are not responsible for the breakdown of the thermal analogy.

To systematically explore the range of behavior for aspherical grains, we now vary the airflow for just one shape. We choose the silver pill, for which the thermal analogy works best in Fig. 1.8. Results for the average energy in each of the three kinetic degrees of freedom, $m\langle v_x^2\rangle/2$, $m\langle v_y^2\rangle/2$,

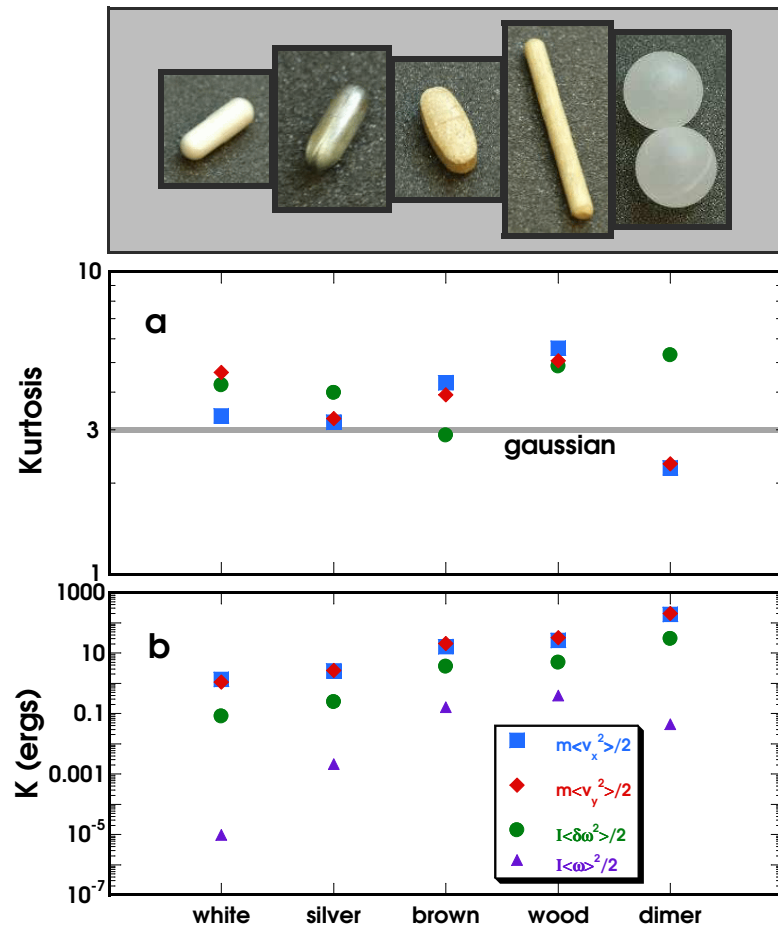


Figure 1.8: [AllShape] (a) Photographs of five aspherical grains; (b) the kurtosis of the speed distributions, and (c) the average kinetic energies, for the two translational and the one rotational degrees of freedom when these grains are individually fluidized. The airspeeds are 910 cm/s for the three leftmost grains, and 750 cm/s for the two rightmost grains.

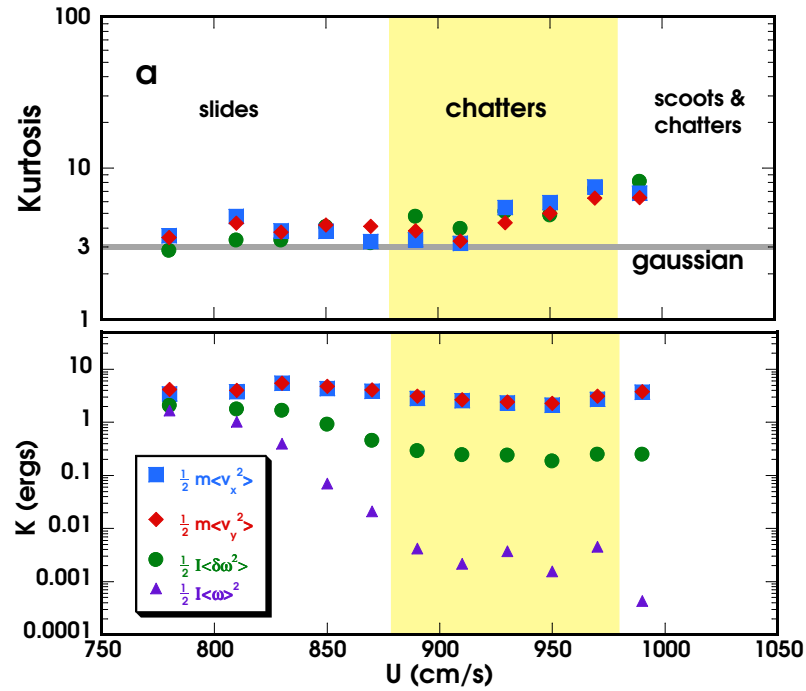


Figure 1.9: [Silver] (a) The kurtosis of the speed distributions, and (b) the average kinetic energies, for the two translational and the one rotational degrees of freedom for the silver grain in Fig. 1.8, as a function of the fluidizing airspeed.

and $I\langle\delta\omega^2\rangle/2$, as well as the whirling energy $I\langle\omega\rangle^2/2$, are shown in Fig. 1.9 along with the kurtosis of the distributions. Counter to intuition, and also counter to Eq. (1.2), the translational kinetic energy is nearly constant while the rotational kinetic energy actually decreases with increasing airspeed. As airspeed decreases, the kurtosis values decrease towards three and both the whirling and rotational fluctuation energies approach the translational kinetic energies. Except for the whirling, the motion is more nearly thermal at lower airspeeds.

The sequence of behavior in Fig. 1.9 correlates with the motion of the grain *perpendicular* to the sieve, which cannot be captured by our usual video methods. At low airspeeds, the grain is in physical contact with the sieve; translational motion thus requires sliding. At intermediate airspeeds, the center of mass is raised somewhat and the grain chatters back and forth along its length. This chattering becomes more prevalent as the airspeed increases. At the highest airspeeds, the chattering motion continues but with the important difference that occasionally the grain scoots

rapidly across the cell. This is somewhat reminiscent of the intermittent chasing observed for two spheres of different size, and it too ruins the thermal analogy. Perpendicular motion is important for the other aspherical grains, as well. At the given airspeeds in Fig. 1.8, the white and silver grains both chatter steadily. The brown grain, wooden rod, and dimer all slide without chattering, like the silver grain at low airspeeds. To fully characterize and understand the behavior of aspherical grains, it would be necessary to measure their out-of-plane motion.

1.8 Conclusion

In summary we have explored four systematic perturbations to an experiment on nearly-levitated spheres that was previously [63, 64] discovered to behave in perfect analogy to a thermal system. Here we find that the statistical mechanical description is robust with respect to variation of both the shape of the sample cell and with respect to the densities of the spheres. This adds to the growing list of driven out-of-equilibrium systems for which an effective temperature may be defined and used in the usual statistical mechanical sense. However, we also find that the spheres must have the same diameter or else the thermal analogy progressively breaks down as the size disparity increases. Furthermore, the analogy is well-controlled only for spherical grains. It can work for pill-shaped objects, but depends on out-of-plane motion that has not yet been well characterized. We hope that the smooth, gradual breakdown as a function of diameter ratio will stimulate theoretical work. This could lead to a better general understanding of when the concepts and tools of statistical mechanics can be invoked for driven far-from-equilibrium systems.

Chapter 2

Approach to Jamming in an Air-Fluidized Granular Bed

2.1 Abstract

Quasi-2D bidisperse amorphous systems of steel beads are fluidized by a uniform upflow of air, so that the beads roll on a horizontal plane. The short-time ballistic motion of the beads is stochastic, with non-Gaussian speed distributions and with different average kinetic energies for the two species. The approach to jamming is studied as a function of increasing bead area fraction and also as a function of decreasing air speed. The structure of the system is measured in terms of both the Voronoi tessellation and the pair distribution function. The dynamics of the system is measured in terms of both displacement statistics and the density of vibrational states. These quantities all exhibit tell-tale features as the dynamics become more constrained closer to jamming. Though the system is driven and athermal, the behavior is remarkably reminiscent of that in dense colloidal suspensions and supercooled liquids.

2.2 Introduction

One of the grand challenges in physics today is to understand non-equilibrium systems, which evolve with time or remain in a steady state by injection of energy. The concept of jamming is helping to unify such seemingly disparate non-equilibrium systems as supercooled liquids and

dense collections of droplets, bubbles, colloidal particles, grains, and traffic [13,49,50]. In all cases, the individual units can be jammed – stuck essentially forever in a single packing configuration – by either lowering the temperature, increasing the density, or decreasing the driving. But what changes occur in the structure and dynamics that signal the approach to jamming? Which features are generic, and which depend on details of the system or details of the driving? *A priori* all non-equilibrium systems are different and all details should matter; therefore, the ultimate utility and universality of the jamming concept is not at all obvious.

Granular materials have a myriad of occurrences and applications, and are being widely studied as idealized non-thermal systems that can be unjammed by external forcing [22,40,59]. Injection of energy at a boundary, either by shaking or shearing, can induce structural rearrangements and cause the grains to explore different packing configurations; however, the microscopic grain-scale response is not usually homogeneous in space or time. This can result in fascinating phenomena such as pattern formation, compaction, shear banding, and avalanching, which have been explained by a growing set of theoretical models with disjoint underlying assumptions and ranges of applicability. To isolate and identify the universally generic features of jamming behavior it would be helpful to explore other driving mechanisms, where the energy injection is homogeneous in space and time. One approach is gravity-driven flow in a vertical hopper of constant cross-section, where flow speed is set by bottom opening. Diffusing-wave spectroscopy and video imaging, combined [56], reveal that the dynamics are ballistic at short times, diffusive at long times, and subdiffusive at intermediate times when the grains are ‘trapped’ in a cage of nearest neighbors. Such a sequence of dynamics is familiar from thermal systems of glassy liquids and dense colloids [50]. Another approach is high-frequency vertical vibration of a horizontal granular monolayer. For dilute grains, this is found to give Gaussian velocity statistics, in analogy with thermal systems [6]. For dense monodisperse grains, this is found to give melting and crystallization behavior also in analogy with thermal systems [66,75].

In this paper, we explore the universality of the jamming concept by experimental study of disordered granular monolayers. To achieve homogeneous energy injection we employ a novel approach in which grain motion is excited by a uniform upflow of air. For dilute grains, the shedding of turbulent wakes was found earlier to cause stochastic motion that could be described by a Langevin equation respecting the Fluctuation-Dissipation Theorem, in analogy with thermal

systems [1, 63, 64]. Now we extend this approach to dense collections of grains. To prevent crystallized domains, and hence to enforce homogeneous disorder, we use a bidisperse mixture of two grain sizes. The extent of grain motion is gradually suppressed, and the jamming transition is thus approached, by both raising the packing fraction and by decreasing the grain speeds in such a way as to approach Point-J [60, 61] in the jamming phase diagram. For a given state of the system, we thoroughly characterize both structure and dynamics using a broad set of statistical measures familiar from study of molecular liquids and colloidal suspensions. In addition to such usual quantities as coordination number, pair-distribution function, mean-squared displacement, and density of states, we also use more novel tools such as a shape factor for the Voronoi cells and the kurtosis of the displacement distribution. We shall show that the microscopic behavior is not in perfect analogy with thermal systems. Rather, the two grain species have different average kinetic energies, and their speed distributions are not Gaussian. Nevertheless, the systematic change in behavior on approach to jamming is found to be in good analogy with thermal systems such as supercooled liquids and dense colloidal suspensions. Our findings support the universality of the jamming concept, and give insight as to which aspects of granular behavior are generic and which are due to details of energy injection.

2.3 Methods

The primary granular system under investigation consists of a 1:1 bidisperse mixture of chrome-coated steel ball-bearings with diameters of $d_b = 11/32$ inch = 0.873 cm and $d_s = 1/4$ inch = 0.635 cm; the diameter ratio is 1.375; the masses are 2.72 g and 1.05 g, respectively. These beads roll on a circular horizontal sieve, which is 6.97 inches in diameter and has a 100 μm mesh size. The packing fraction, equal to the fraction of projected area occupied by the entire collection of beads, is varied across the range $0.487 < \phi < 0.826$ by taking the total number of beads across the range $262 < N < 444$.

The motion of the beads is excited by a vertical upflow of air through the mesh at fixed superficial speed 950 ± 10 cm/s. This is the volume per time of air flow, divided by sieve area; the air speed between the beads is greater according to the value of ϕ . The air speed is large enough to

drive stochastic bead motion by turbulence ($Re \approx 10^4$), but is small enough that the beads maintain contact with the sieve and roll without slipping. The uniformity of the airflow is achieved by mounting the sieve atop a $1.5 \times 1.5 \times 4 \text{ ft}^3$ windbox, and is monitored by hot-wire anemometer.

The system of beads is illuminated by six 100 W incandescent bulbs, arrayed in a 1 ft diameter ring located 3 ft above the sieve. Specularly reflected light from the very top of each bead is imaged by a digital CCD camera, Pulnix 6710, placed at the center of the illumination ring. The sensing element consists of a 644×484 array of $10 \times 10 \mu\text{m}^2$ square pixels, 8 bits deep. Images are captured at a frame rate of 120 Hz, converted to binary, and streamed to hard-disk as AVI movies using the lossless Microsoft RLE codec. Run durations are 20 minutes. The threshold level for binary conversion is chosen so that each bead appears as a small blob about 9 and 18 pixels in area for the small and large beads, respectively. Note that the spot size is smaller than the bead size, which aids in species identification and ensures that colliding beads appear well-separated. The minimum resolvable bead displacement, below which there is a fixed pattern of illuminated pixels within a blob, is about 0.1 pixels.

The AVI movies are post-processed using custom LabVIEW routines, as follows, to deduce bead locations and speeds. For each frame, each bead is first identified as a contiguous blob of bright pixels. Bead locations are then deduced from the average position of the associated illuminated pixels. Individual beads are then tracked uniquely vs time, knowing that the displacement between successive frames is always less than a bead diameter. Finally, positions are refined and velocities are deduced by fitting position vs time data to a cubic polynomial. The fitting window is ± 5 points, defined by Gaussian weighting that nearly vanishes at the edges; this choice of weighting helps ensure continuity of the derivatives. The rms deviation of the raw data from the polynomial fits is 0.0035 cm, which corresponds to 0.085 pixels – close to the minimum resolvable bead displacement. The accuracy of the fitted position is smaller according to the number of points in the fitting window: about $0.0035 \text{ cm}/\sqrt{5} = 0.0016 \text{ cm}$. This and the frame rate give an estimate of speed accuracy as 0.2 cm/s. An alternative analysis approach is to run the time-trace data through a low pass filter using Fourier methods. We find very similar results to the weighted polynomial fits for a range of cut-off frequencies, as long as the rms difference between raw and filtered data is between about 0.1 and 0.3 pixels. While there is no significant difference in the plots presented below, we slightly prefer the polynomial fit approach based on qualitative inspection – it

is much slower but appears better at hitting the peaks without giving spurious oscillations.

2.4 General

For orientation, Fig. 2.1a shows representative configurations for two area fractions, $\phi = 48.7\%$ and $\phi = 80.9\%$. To mimic an actual photograph, a disk with the same diameter as the bead has been centered over each bead's position, with a darker shade for the big beads. Note that the configurations are disordered, and that the two bead species are distributed evenly across the system. With time, due to the upflow of air, the beads move about and explore different structural configurations. Over the duration of a twenty minute run, at our lowest packing fractions, each bead has time to sample the entire cell several times over. The beads never crystalize or segregate according to size. Thus, the system appears to be both stationary and ergodic. However, the beads tend to idle for a while if they come in contact with the boundary of the sieve. Therefore, care is taken to prevent the contamination of bulk behavior by boundary beads.

In the next sections we quantify first structure and then dynamics, and how they both change with increasing packing fraction on approach to jamming. A two-dimensional random close packing of bidisperse hard disks can occupy a range of area fractions less than about 84%, depending on diameter ratio and system size [61, 87]. We find the random close packing of our system of bidisperse beads to be at an area fraction of $\phi_c = 0.83$. If we add more beads, in attempt to exceed this value, then there is not enough room for all beads to lie in contact with the sieve – some are held up into the third dimension by enduring contact with beads in the plane. So we expect the jamming transition to be at or below $\phi_c = 0.83$ depending on the strength and range of bead-bead interactions. Earlier, we found that the upflow of air creates a repulsion between two isolated beads that can extend to many bead diameters [63]. Nevertheless, we shall show here that our system remains unjammed all the way up to ϕ_c and that it develops several tell-tale signatures on approach jamming.




Figure 2.1: [VT](a) Example configurations for area fractions $\phi = 48.7\%$ and $\phi = 80.9\%$, with the big beads colored darker than the small beads. (b-c) Voronoi tessellations for these configurations with cells shaded darker for increasing coordination number and circularity, respectively.




Figure 2.2: [PZ] (a) Contour plot of the coordination number distribution; red is for large probability density and blue is for small. The dashed white line indicates $\phi = 0.74$. (b) Coordination number distributions for three area fractions, as labeled.

2.5 Structure

2.5.1 Coordination Number

Perhaps the simplest structural quantity is the coordination number Z , equal to the number of nearest neighbors for each bead. This can be most conveniently measured by constructing a Voronoi tessellation, which is dual to the position representation, and by counting the number of sides of each polygonal Voronoi cell. Examples are shown in Fig. 2.1b for the same configurations shown in Fig. 2.1a. Here the Voronoi cells are shaded darker for greater numbers of sides. The coordination number ranges between 3 and 9, but by far the most common numbers are 5, 6, and 7 irrespective of area fraction. It seems that 5- and 7-sided cells appear together, and that 6-sided cells sometimes appear in small compact clusters.

The distribution $P(Z)$ of coordination numbers, and trends vs area fraction, are displayed in Fig. 2.2. The results are obtained by averaging over all times and over all beads away from the boundary. The bottom plot shows actual distributions for three area fractions, two of which are the same as in Fig. 2.1. The main effect of increasing the area fraction is to increase the fraction of 7-sided cells at the expense of all others, until essentially only 5-6-7 sided cells remain. This trend is more clearly displayed in a contour plot, Fig. 2.2a, where the value of $P(Z)$ is indicated by color as a function of both Z and area fraction ϕ . For area fractions above about 74%, indicated by a dashed white line, the 4-, 8-, and 9-sided cells have essentially disappeared. This effect is rather subtle, owing to the discrete nature of the coordination number. Six-sided cells are always the most plentiful; their abundance gently peaks near 74%.

2.5.2 Circularity

A more dramatic measure of structural change upon approach to jamming is found by considering the *shapes* of the Voronoi cells. One choice for a dimensionless measure of deviation from circularity is

$$\zeta = P^2/(4\pi A), \quad (2.1)$$

where P is the cell perimeter and A is the cell area. This quantity was recently used to study crystallization of two-dimensional systems, both in simulation [57] and experiment [75]. By construction ζ equals one for a perfect circle, and is higher for more rough or oblong shapes; for a




Figure 2.3: [PC] (a) Contour plot of the non-circularity shape factor distributions for the Voronoi tessellation polygons. Beyond about 74% (dashed white line) a well formed second peak develops and the distribution doesn't change much. (b) Shape factor distributions for a sequence of area fractions; the thick green curve is for $\phi = 74.4\%$. The labels 7, 6, 5, and 4 show the minimum shape factors for polygons with that number of sides.

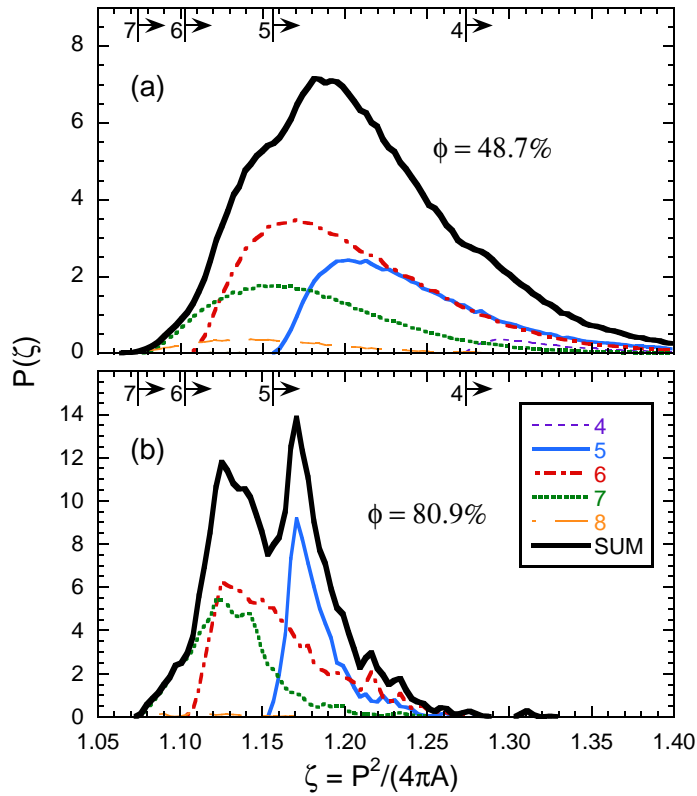


Figure 2.4: [PCz] Non-circularity shape factor distributions for the Voronoi tessellation polygons for (a) $\phi = 48.7\%$ and (b) $\phi = 80.9\%$. Thick black curves are the actual distributions, and the thin colored curves are the contributions from cells with different numbers of sides, as labeled.

regular Z -sided polygon it is

$$\zeta_Z = (Z/\pi) \tan(\pi/Z), \quad (2.2)$$

which sets a lower bound for other Z -sided polygons. As an example in Fig. 2.1c the Voronoi cells are shaded darker for more circular shapes, i.e. for those with smaller non-circularity shape factors. Since ζ_Z decreases with increasing Z , it may be expected that the shape factor is related to coordination number. The advantage is that ζ is a continuous variable while Z is discrete.

Shape factor distributions, $P(\zeta)$, and the way they change with increasing area fraction, are displayed in Fig. 2.3b. These are obtained by constructing Voronoi tessellations, and averaging over all times and over all beads. At low area fractions, $P(\zeta)$ exhibits a single broad peak. At higher area fractions, this peaks moves to lower ζ , i.e. to more circular domains, and eventually *bifurcates* into two sharper peaks. This trend can be seen, too, in the contour plot of Fig. 2.3a where

color indicates the value of $P(\zeta)$ as a function of both non-circularity ζ and area fraction ϕ . This plot shows that the double peak becomes essentially completely developed around $\phi = 0.74$. This is the same area fraction singled out by a subtle change in the coordination number distribution. Thus the shape factor and its distribution are useful for tracking the change in structure as a liquid-like system approaches a disordered jammed state.

The origin of the double peak in the shape factor distribution $P(\zeta)$ can be understood by considering the contributions $P_Z(\zeta)$ made by cells with different coordination numbers. These contributions are defined such that $P(\zeta) = \sum_{Z=3}^{\infty} P_Z(\zeta)$ and $P(Z) = \int_{\zeta_Z}^{\infty} P_Z(\zeta) d\zeta$; in particular, cells with coordination number Z contribute a sub-distribution $P_Z(\zeta)$ which must vanish for $\zeta < \zeta_Z$ according to Eq. (2.2) and which subtends an area equal to the coordination number probability. As an example, the shape factor distribution is shown along with the individual contributions in Fig. 2.4. At low area fractions, in the top plot, the broad peak in $P(\zeta)$ is seen to be composed primarily of overlapping broad contributions from 5-, 6-, and 7-sided cells. At high area fractions, in the bottom plot, the double peak in $P(\zeta)$ is seen to be caused by 5-sided cells for the right peak and by overlapping contributions of 6- and 7-sided cells for the left peak. At these higher area fractions, the individual contributions $P_Z(\zeta)$ are more narrow and rise more sharply from zero for $\zeta > \zeta_Z$. In other words, the Voronoi cells all become more circular at higher packing fractions. Due to disorder, there is a limit to the degree of circularity that cannot be exceeded and so the changes in the circularity distribution eventually saturate. For our system this happens around $\phi = 0.74$, which is well below random close packing at $\phi_c = 0.83$.

2.5.3 Radial Distribution

We now present one last measure of structure that is commonly used in amorphous systems: the radial or pair distribution function $g(r)$. This quantity relates to the probability of finding another bead at distance r away from a given bead. For large r , it is normalized to approach one – indicating that the system is homogeneous at long length scales. For small r in hard-sphere systems like ours, it vanishes for r less than the sum of bead radii. Since we have two species of beads, big and small, there are four different distributions to consider: between any two beads, between only big beads, between big and small beads, and between only small beads.

Data for all four radial distribution functions are collected in Fig. 2.5. As usual, these data

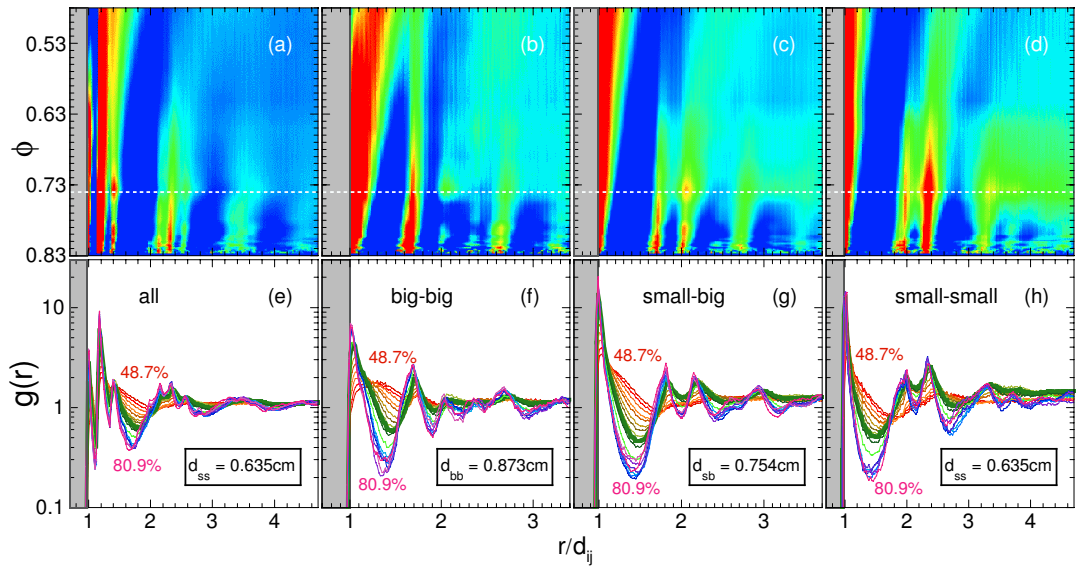


Figure 2.5: [GrC] The radial distribution function computed between (a,e) all beads; (b,f) big beads; (c,g) small and big beads; and (d,h) small beads. The top row shows contour plots, where blue is for large g and red is for small; the dashed white line represents $\phi = 0.74$. The bottom row shows data curves for different area fractions; the thick green curves are for $\phi = 0.744$. The gray region represents the distance excluded by hard-core contact, for r less than the sum d_{ij} of radii.

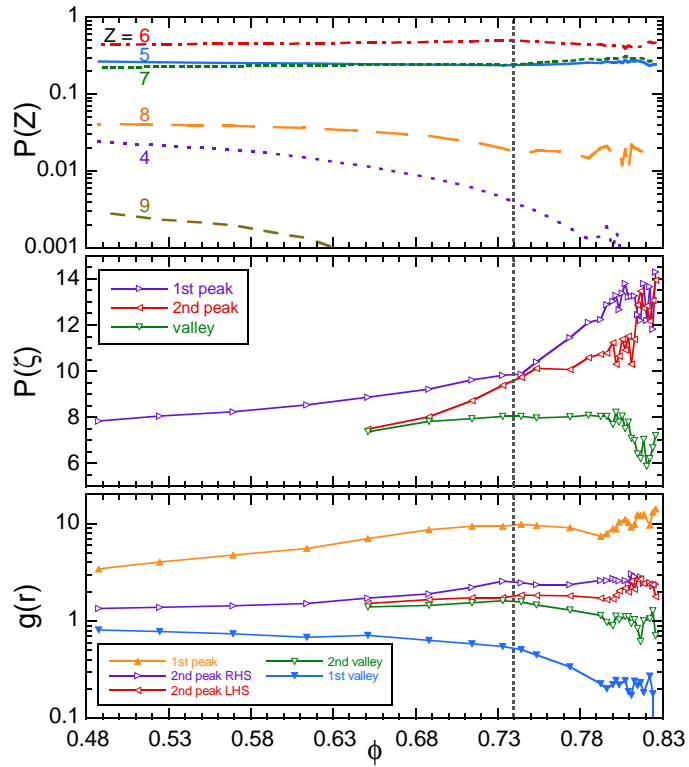


Figure 2.6: [Strs] The area fraction dependence of (a) the occurrence probability $P(Z)$ of Voronoi cells with Z sides, and peak and valley values of (b) the shape factor distribution and (c) the pair distribution function for small beads.

were obtained by averaging over all times and over all beads away from the boundary. The bottom row shows functions plotted vs r for different area fractions, while the top row shows contour plots where color indicates the value of $g(r)$ as a function both r and area fraction; radial distance is scaled by the sum of bead radii d_{ij} . All four radial distribution functions display a global peak at hard core contact, $r/d_{ij} = 1$. For increasing area fractions, these peaks become higher and more narrow, while oscillations develop that extend to larger r . Also, near $r/d_{ij} = 2$ there develops a second peak that grows and then bifurcates into two separate peaks. Both the growth of the peak at $r/d_{ij} = 1$, relative to the deepest minimum, and the splitting of the peak near $r/d_{ij} = 2$ have been taken as structural signatures of the glass transition [12, 79, 95]. For our system the split second peak becomes fully developed for area fractions greater than about $\phi = 0.74$, the same area fraction noted above with regards to changes in the Voronoi tessellations.

2.5.4 Summary

Well-defined features develop in several structural quantities as the area fraction increases. The most subtle is an increase in the number of 7-sided Voronoi cells at the expense of all other coordination numbers, and the disappearance of nearly all 4- and 9-sided cells. The most obvious are the splitting of peaks in the shape factor distribution and the radial distribution functions. These key quantities are extracted from our full data sets and are displayed vs area fraction ϕ in Fig. 2.6. The top plot shows the fraction $P(Z)$ of Voronoi cells with Z sides, for several coordination numbers; the middle plot shows peak and valley values of the probability density $P(\zeta)$ for Voronoi cells with shape factor ζ ; the bottom plot shows peak and valley values of the pair distribution function $g(r)$ for small beads. These three plots give a consistent picture that $\phi = 0.74$ is the characteristic area fraction for structural change.

2.6 Dynamics

2.6.1 Data

For the remainder of the paper we focus on bead motion, and how it changes in response to the structural changes found above an area fraction of 74%. The primary quantity we measure and analyze is the mean-squared displacement (MSD) that the beads experience over a time interval




Figure 2.7: [msd] (a-b) Mean-squared displacement, and (c-d) kurtosis of displacement probability distribution, all as a function of delay time. The left-hand plots are for the big beads and the right-hand plots are for the small beads. Area fraction color codes for the all plots are labeled in (c-d); the thick green curve is for $\phi = 74.4\%$. Note that the mean-squared displacement saturates at the square of the sample cell radius. The squares of bead diameters are $d_b^2 = 0.76 \text{ cm}^2$ and $d_s^2 = 0.40 \text{ cm}^2$. The square of the position resolution is $(0.0016 \text{ cm})^2 = 3 \times 10^{-6} \text{ cm}^2$.

τ : $\langle \Delta r^2(\tau) \rangle$ vs τ . The MSD is readily measured directly from time- and ensemble averages of the position vs time data; it can also be computed efficiently from position autocorrelation data using Fourier methods. Results are shown in Figs. 2.7a-b for the big and small beads, separately. At short times, the bead motion is ballistic and characterized by a mean-squared speed according to $\langle \Delta r^2(\tau) \rangle = \langle v^2 \rangle \tau^2$. Our frame rate is 120 Hz, corresponding to a shortest delay time of $\tau = 0.0083 \text{ s}$, which is fast enough that we are able to observe ballistic motion over about one decade in time for all area fractions. At long times, the bead motion is diffusive and characterized by a diffusion coefficient according to $\langle \Delta r^2(\tau) \rangle = 4D\tau$. Our run durations are 20 minutes, corresponding to a longest delay time of $\tau = 1200 \text{ s}$, which are long enough for the beads to explore the entire system several times at the lowest area fractions. The crossover from ballistic to diffusive regimes becomes progressively slower as the area fraction increases. Our run durations are long enough to fully capture the diffusive regime at all but the highest area fractions. Thus our full position vs time dataset, for all beads and area fractions, should suffice for a complete and systematic study of changes in dynamics as jamming is approached.

The MSD has long been used to characterize complex dynamics. In simple systems there is a single characteristic time scale given by the crossover from ballistic to diffusive regimes. In supercooled or glassy systems, the crossover is much more gradual and there are two characteristic time scales. The shortest, called the ‘ β ’ relaxation time, is given by the end of the ballistic regime. The longest, called the ‘ α ’ relaxation time, is given by the beginning of the diffusive regime. At greater degrees of supercooling in glass-forming liquids, and at greater packing fractions in colloidal suspensions, the α relaxation time increases and a corresponding plateau develops in the MSD. As seen in our MSD data of Figs. 2.7a-b, this same familiar sequence of changes occurs in our system as well. The greater the delay in the onset of diffusive motion, the more ‘supercooled’ is our system and the closer it is to being jammed – where each bead has a fixed set of neighbors that never changes.

Another similarity between the dynamics in our system and thermal systems can be seen by examining the kurtosis of the displacement distribution. For a given delay time, there is a distribution $P(\Delta x)$ of displacements. By symmetry the average displacement and other odd moments must all vanish. The second moment is the most important; it is the MSD already discussed. If the distribution is Gaussian, a.k.a. normal, then all other even moments can be deduced from the value of the MSD. For example the ‘kurtosis’ is the fourth moment scaled by the square of the MSD and with the Gaussian prediction subtracted; by construction it equals zero for a Gaussian distribution, and otherwise is a dimensionless measure of deviation from ‘normality’. The kurtosis of the displacement distribution has been used in computer simulation of liquids, both simple [72] and supercooled [38, 44, 85], as well as in scattering [88] and imaging experiments of dense colloidal suspensions [41, 42, 54, 91]. These works consider a quantity α_2 equal to 1/3 of the kurtosis, and find that the displacement distribution is Gaussian in ballistic and diffusive regimes, but becomes non-Gaussian with a peak in the kurtosis at intermediate times. We computed the kurtosis of the displacement distribution for our system, with results displayed vs delay time in Figs. 2.7c-d for many area fractions. As in thermal systems, our kurtosis results display a peak at intermediate times and becomes progressively more Gaussian at late times. Furthermore, the peak increases dramatically once the area fraction rises above about 74%, particularly for the large beads.

The kurtosis data in Figs. 2.7c-d do not vanish at short times, by contrast with thermal systems. Instead, we find that the kurtosis decreases to the left of the peak and saturates at a nonzero constant




Figure 2.8: $\langle \mathcal{D}(f) \rangle$ Density of vibrational states of frequency f , for various packing fractions; the thick green curve is for $\phi = 74.4\%$.

upon entering the ballistic regime. At these short times, the displacement distribution has the same shape as the velocity distribution since $\Delta x = v_x \tau$. Indeed our velocity distributions are non-Gaussian with the same kurtosis as the short-time displacement distributions. This reflects the non-thermal, far-from-equilibrium, driven nature of our air-fluidized system. Another difference from thermal behavior, as we'll show in the next subsection, is that the two bead species have different average kinetic energies – which is forbidden by equipartition for a thermal systems.

While the second and fourth moments of the displacement distributions capture many aspects of bead motion, another dynamical quantity has been considered recently [?, 61, 96, 97]: the density $\langle \mathcal{D}(f) \rangle$ of vibrational states of frequency f . At high frequencies, the behavior of $\langle \mathcal{D}(f) \rangle$ reflects the short time ballistic nature of bead motion. At low frequencies, the behavior of $\langle \mathcal{D}(f) \rangle$ reflects on slow collective relaxations. If the system is unjammed, there will be zero-frequency translational relaxation modes with relative abundance set by the value of $\langle \mathcal{D}(0) \rangle$. If the system is fully jammed, by contrast, there can be no zero-frequency modes and hence $\langle \mathcal{D}(f) \rangle$ must vanishes for decreasing f . Thus the form and the limit of $\langle \mathcal{D}(f) \rangle$ at low frequencies give a sensitive dynamical signature of the approach to jamming. This shall be our focus, while by contrast in Refs. [61, 78, 96, 97] the focus was on the behavior of low frequency modes above close packing on approach to *unj*jamming.

The density of states may be computed from the velocity time traces, $v_i(t)$, for all beads i [21]. The mass-weighted ensemble average of time-averaged velocity autocorrelations, $w(\tau) =$

$\sum m_i \langle \mathbf{v}_i(t) \cdot \mathbf{v}_i(t+\tau) \rangle / \sum m_i$, is the key intermediate quantity; its Fourier transform is $w(f)$ and has units of cm^2/s . The final step is to compute the modulus,

$$\langle \mathcal{D}(f) \rangle = \sqrt{w(f)w^*(f)}. \quad (2.3)$$

The angle-brackets in this notation are a reminder that this is an ensemble average of single-grain quantities. By construction, the integral of Eq. (2.3) over all frequencies is equal to the mass-weighted mean-squared speed of the beads. While Eq. (2.3) appears to be a purely mathematical manipulation, the identification of the right-hand side with the density of states requires that modes be populated according to the value of kT ; hence, for non-thermal systems like ours, the result is an effective density of states that only approximates the true density of states. Whatever the accuracy of this identification, both the expression for $\langle \mathcal{D}(f) \rangle$ in Eq. (2.3) and the mean-squared displacement may be computed from velocity autocorrelations, and thus do not embody different physics; rather, they give complementary ways of looking at the same phenomena and serve to emphasize different features.

The effective density of vibrational states for our system of air-fluidized beads is shown in Fig. 2.8, with separate curves for different area fractions. At low ϕ there are many low frequency modes and $\langle \mathcal{D}(f) \rangle$ is nearly constant until dropping off at high frequencies. At higher ϕ , the number of low frequency modes gradually decreases; $\langle \mathcal{D}(f) \rangle$ is still constant at low f , but it increases to a peak before dropping off at high f . Even at the highest area fractions, $\langle \mathcal{D}(f) \rangle$ is nonzero at the smallest frequencies observed, as given by the reciprocal of the run duration.

2.6.2 Short and Long-time Dynamics

The quantities that specify the ballistic and diffusive motion at short and long times, respectively, are the mean-squared speeds and the diffusion coefficients. These may be deduced from the mean-squared displacement data, separately for each bead. We now examine trends in these dynamical quantities as a function of area fraction.

We begin with the mean-squared speed. To highlight the contrast with a thermal system, we convert it to a mean kinetic energy and plot the results in Figs. 2.9a-b for the big and small beads, respectively. We show a small point for each individual bead, as well as a larger symbol for the ensemble average of these values. The average kinetic energies appear to decrease nearly




Figure 2.9: [KKr] (a-b) Kinetic energy vs area fraction for big and small beads, respectively. The time-average for each bead is shown by a small dot; the ensemble average over all beads is shown by +. (c) The kinetic energy ratio of small-to-big beads. The line $(d_s/d_b)^3$ shows where beads move with the same mean-squared velocity. Insets show the same quantities on a logarithmic scale.

linearly towards zero as the area fraction is raised towards random close packing. The scatter in the points is roughly constant, independent of ϕ , and reflects the statistical uncertainty in our velocity measurements. There is no evidence of nonergodicity or inhomogeneity in energy injection by our air-fluidization apparatus; namely, all the beads of a given species have the same average kinetic energy to within measurement uncertainty. Beyond about $\phi = 0.81$ the uncertainty becomes comparable to the mean, as set by our speed resolution, and we can no longer readily discern the average kinetic energies. This may be more evident in the insets, which show kinetic energies on a logarithmic scale.

Note that the big and small beads have different average kinetic energies. Equipartition is thus violated for our athermal, driven system. This is highlighted in Fig. 2.9c, which shows the kinetic energy ratio of small to large beads. This ratio is roughly constant at about 0.6 for the lowest area fractions. It decreases gradually for increasing ϕ , at first gradually and then more rapidly beyond about $\phi = 0.78$. There is no obvious feature at $\phi = 0.74$, where structural quantities changed noticeably. Interestingly, however, the kinetic energy ratio appears to head toward the cube of the bead diameter ratio on approach to random close packing. This means that the mean-squared speeds are approaching the same value, perhaps indicating that only extremely collective motion is possible very close to jamming. In order for one bead to move, the neighboring bead in the path of motion must move with the same speed. This seems a natural geometrical consequence of nearly close packing, but would be a violation of equipartition in a thermal system.

Turning now to late-time behavior, we display diffusion coefficients vs area fraction in Figs. 2.10a-b, for big and small beads respectively. As in Fig. 2.9, a small dot is shown for each individual bead and a larger symbol is shown for the ensemble average of these results. The average diffusion coefficients appear to decrease linearly with increasing area fraction, starting at the lowest ϕ . Here there is considerable scatter in the data, reflecting a level of uncertainty set by run duration. Linear fits of diffusion coefficient vs area fraction in this regime extrapolate to zero at the special value $\phi = 0.74$. However, just before reaching this area fraction, the data fall below the fit. As shown on a logarithmic scale in the insets, the diffusion coefficients are nonzero and continue to decrease with increasing ϕ . Beyond about $\phi = 0.81$ the system becomes non-ergodic over the duration of our measurements. Namely, some beads remain stuck in the same nearest neighbor configurations




Figure 2.10: [D] (a-b) Diffusion coefficient vs area fraction for big and small beads, respectively. The time-average for each bead is shown by a small dot; the ensemble average over all beads is shown by +. (c) The diffusion coefficient ratio of small-to-big beads. Insets show the same quantities on a logarithmic scale.

while others have broken out. Indeed, at these high area fractions the MSDs in Fig. 2.7 grow sub-linearly at the latest observed times. To obtain reliable diffusion coefficient data in this regime would require vastly longer run durations.

2.6.3 Timescales

Lastly we focus on characterizing several special times that serve to demarcate the short-time ballistic and the long-time diffusive regimes seen in the displacement statistics data. The time scale-dependent nature of the dynamics, which is obvious in the MSD or the density of states, is also obvious in real-time observations or in AVI movie data. For high area fractions, the bead configuration appears immutable at first; the beads collide repeatedly with an apparently-fixed set of neighbors that cage them in. Only with patience, after hundreds or thousands of collisions, can the beads be observed to break out of their cages, change neighbors, and begin to diffuse throughout the system.

To measure unambiguously the characteristic timescales, we consider the slope of the mean-squared displacement as seen on a log-log plot:

$$\lambda(\tau) = \frac{\partial \ln[\langle \Delta r^2(\tau) \rangle]}{\partial \ln \tau}. \quad (2.4)$$

An example for one of the higher area fractions is shown in Fig. 2.11 a. At the shortest delay times, when the motion is perfectly ballistic, this slope is 2; at the longest delay times, when the motion is perfectly diffusive, this slope is 1. For high area fractions, such as shown, there is a subdiffusive regime with $\lambda < 1$ at intermediate times; this is when a typical bead appears by eye to be stuck in a cage, rattling against a fixed set of neighbors. For low area fractions, not shown, there is no such ‘caging’ and λ decreases monotonically from 2 to 1.

Several natural time scales can now be defined with use of $\lambda(\tau)$ vs τ data. The shortest is the delay time τ_b at which the logarithmic slope falls to $\lambda(\tau_b) = 1.5$. This demarcates the ballistic regime, below which the bead velocity is essentially constant. At high area fractions, τ_b is a typical mean-free time between successive collisions; at low area fractions, τ_b it is also the time for crossover to diffusive motion. The other time scales that we define all refer to the subdiffusive, caging dynamics at high area fractions. The most obvious is the delay time τ_{\min} at which $\lambda(\tau_{\min})$ is minimum; this corresponds to an inflection in the MSD on a log-log plot. Below τ_{\min} most

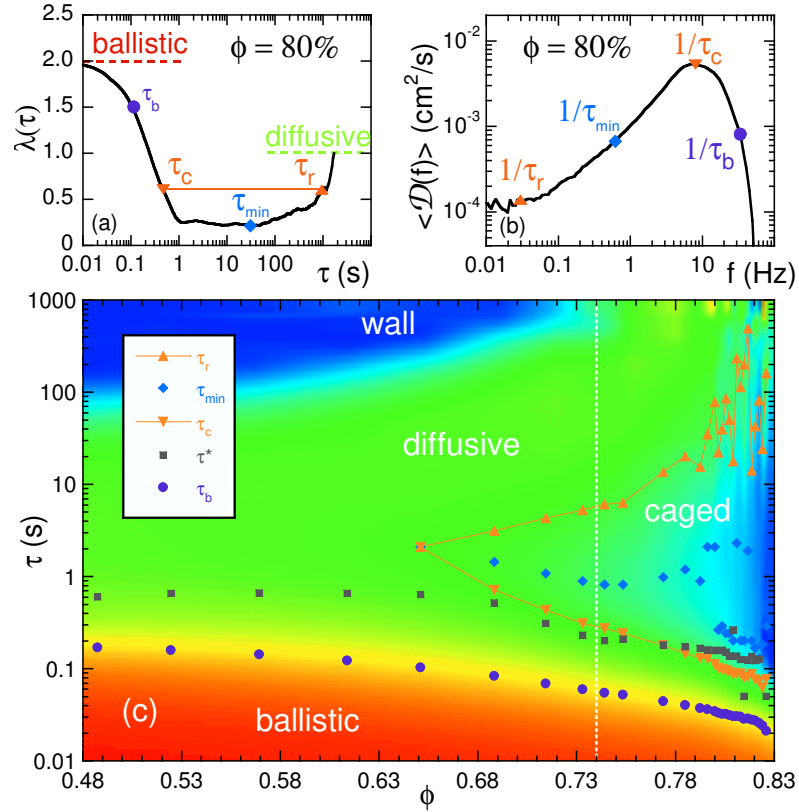


Figure 2.11: [Infects] (a) Logarithmic derivative λ of the mean-squared displacement vs delay time, for big beads at area fraction $\phi = 80\%$. Four special time scales can be defined from such data, as depicted: τ_b where $\lambda = 1.5$, τ_{\min} where λ is minimum, and τ_c and τ_r where λ is halfway between 1 and its minimum. (b) Density of state, for big beads at $\phi = 80\%$, marked with the four special timescales. (c) Contour plot of the logarithmic derivative for the big beads, where color indicates the value of λ , as a function of both area fraction and delay time. Red is slope two and blue is slope zero; caged dynamics are when the mean-squared displacement has a slope less than one, which is aqua-blue. The four special times defined by the behavior of λ , as well as a fifth time τ^* at which the displacement kurtosis is maximal, are superposed on the contour plot; symbol key is given in the legend.

beads remain within a fixed cage of neighbors. The last two special times specify the interval when the motion is subdiffusive, with a logarithmic slope falling in the range below 1 and above its minimum. The smaller is the delay time τ_c at which the logarithmic slope decreases half-way from 1 down towards its minimum: $\lambda(\tau_c) = [1 + \lambda(\tau_{\min})]/2$. This is the time at which the beads have explored enough of their immediate environment to ‘realize’ they are trapped at least temporarily within a fixed cage of neighbors; it is longer, and distinct from, the mean-free collision time. The longest special time scale is the delay time τ_r at which the logarithmic slope increases half-way from its minimum up towards 1: $\lambda(\tau_r) = [1 + \lambda(\tau_{\min})]/2$. This is the time beyond which the beads rearrange and break out of their cages; it demarcates the onset of fully diffusive motion.

Before continuing, we note that these four special times also correspond to features in the density of states. As shown in Fig. 2.11b, the density of states drops to zero precipitously for frequencies above $1/\tau_b$, the reciprocal of the ballistic- or collision mean-free time. The density of states reaches a peak at $1/\tau_c$, corresponding to the time that beads ‘realize’ they are stuck at least temporarily within a cage. For lower frequencies, the logarithmic derivative of $\langle \mathcal{D}(f) \rangle$ vs f has an inflection at roughly $1/\tau_{\min}$. And at the lowest frequencies, $\langle \mathcal{D}(f) \rangle$ approaches a constant value below about $1/\tau_r$. Thus we could have analyzed all data in terms of $\langle \mathcal{D}(f) \rangle$; however we prefer to work with the mean-squared displacement since it does not involve numerical differentiation.

Results for the special timescales are collected in Fig. 2.11 as a function of area fraction. Actual data points are superposed on top of a contour plot of the logarithmic derivative, where color [on-line] fans through the rainbow according to the value of λ : red for ballistic, green for diffusive, and blue for subdiffusive. With increasing ϕ , the ballistic time scale decreases steadily by a factor of nearly ten as close as we can approach random close packing. Note that with our definition of τ_b , these data points lie near the center of the yellow band demarcating the end of the red ballistic regime. Below $\phi = 0.65$ the motion is never subdiffusive, and τ_b is the only important time scale. Right at $\phi = 0.65$ the motion is only barely subdiffusive, for a brief moment, so that all three associated times scales nearly coincide. Above $\phi = 0.74$ the caging is sufficiently strong that the time scale for rearrangement τ_r is ten times longer than the time scale τ_c for ‘realization’ that there is caging. At progressively higher area fractions this separation in time scales grows ever stronger. On approach to jamming at random close packing, the cage realization time decreases towards a nonzero constant, while the cage break-up or rearrangement time τ_r increases rapidly towards our

run duration and appears to be diverging.

Note that the two time scales τ_c and τ_r capture the subdiffusive caging dynamics better than just the delay time τ_{\min} at which the logarithmic derivative is minimum and the motion is maximally subdiffusive. One reason is that τ_{\min} is difficult to locate for extremely subdiffusive motion, where there is a wide plateau in the MSD and hence where λ is nearly zero over a wide range of delay times. This difficulty can be seen in both Fig. 2.11a, where λ remains close to its minimum for over two decades in delay time, as well as in Fig. 2.11c above about $\phi = 0.8$, where τ_{\min} data jump randomly between about 2 s and about $3\tau_c$. The other reason is that τ_{\min} does not appear to be either a linear or geometric average of the cage ‘realization’ and breakup times. Rather, the shape of $\lambda(\tau)$ vs $\ln(\tau)$ is asymmetric, with a minimum closer to the cage ‘realization’ time. Thus it is useful to know the two times τ_c and τ_r that together specify the subdiffusive dip of $\lambda(\tau)$ below one, just as it is useful to know the full-width half-max of a spectrum of unspecified shape.

Note too that τ_c and τ_r are distinct from the time τ^* at which the kurtosis is at maximum. Results for τ^* are extracted from our kurtosis data, and are displayed as open squares along with other characteristic times in Fig. 2.11. At low area fractions, even when there is no subdiffusive regime, the kurtosis exhibits a maximum at delay time τ^* that is several times the ballistic / collision time τ_b . With increasing area fraction, τ^* decreases. Once a subdiffusive regime appears, the value of τ^* is close to the time τ_c at which grains ‘realize’ they are stuck in a cage. Data in Fig. 2.11 for both τ_c and τ^* decrease with increasing area fraction, ϕ , on approach to jamming, while the time τ_r signalling the end of the subdiffusive regime appears to grow without bound. The decrease of τ^* with ϕ agrees with previous observations on colloids [42, 54, 91], but contrasts with statements that τ^* corresponds to the cage-breakout α -relaxation time beyond which the motion is diffusive. To emphasize, we find that the tail of the displacement distribution is largest relative to a Gaussian, and hence that the kurtosis is maximal, at the *beginning* of the subdiffusive regime. At that time most beads have been turned back by collision, but a few prolific beads move at ballistic speed roughly one diameter – which is long compared to the rms displacement. This observation is not an artifact of limited spatial or temporal resolution or of limited packing fraction range, which are all notably better than in previous experiments. It is also not an artifact of our analysis method; indeed, if we filter too strongly then the kurtosis peak shifts to later times.

2.6.4 Jamming Phase Diagram

Before closing we now summarize our structure and dynamics observations and place them in the context of the jamming phase diagram. By adding more beads to our system at fixed air speed, we control the approach to jamming two ways. First, most obviously, the area fraction increases towards random close packing at $\phi_c = 0.83$. Second, the average kinetic energy of the beads decreases – roughly linearly with ϕ according to the kinetic energy data shown in Fig. 2.9. Both aspects are captured by the trajectory on a jamming phase diagram plot of kinetic energy vs area fraction. To get a single dimensionless measure of kinetic energy, we divide the mass-weighted average kinetic energy of the beads by a characteristic gravitational energy given by average ball weight times diameter; the result is denoted as a scaled effective temperature $T/(mgd)$. This scaling reflects the natural energy for air-mediated interactions in a gas-fluidized system. The particular $T/(mgd)$ vs ϕ trajectory followed by our experiments is shown by open symbols in Fig. 2.12. It is a diagonal line that terminates at $\{T = 0, \phi = \phi_c\}$, which is the special point in the jamming phase diagram known as ‘Point- J ’ [60, 61]. On approach to Point- J our system remains unjammed, but develops tell-tale features in the Voronoi tessellations and in the mean-squared displacement indicating that jamming is near. In particular the structural changes saturate, and the ratio τ_r/τ_c of cage breakup to ‘realization’ times τ_r exceeds ten, for points along the trajectory closer to Point- J than $\{T/(mgd) = 0.007, \phi = 0.74\}$; we then say the system is ‘pre-jammed’. In this pre-jammed region, we did not observe any signs of aging; the structure and dynamics change from that of a simple liquid but do not appear to evolve with time.

An extended region of ‘pre-jammed’ behavior must exist near Point- J , and can be mapped out by changing the experimental conditions. The simplest variation is to examine vertical trajectories, at fixed ϕ for the same system of steel beads, where the effective temperature is changed by the speed of the upflowing air. We did this for several different area fractions, locating the special effective temperatures below which the system becomes pre-jammed. Furthermore, to better test the universality of our phase diagram and the choice of energy scaling, we examined two other 1:1 bidisperse mixtures of spheres. This includes several constant- ϕ trajectories for large hollow polypropylene spheres with diameters of 1-1/8 and 1-3/8 inches, and one constant air speed trajectory for very small steel spheres with diameters of 5/32 and 1/8 inches. The resulting trajectories and pre-jamming boundary points are shown in Fig. 2.12 as dashed lines and symbols, respectively.




Figure 2.12: [PointJ] The zero-stress plane of the jamming phase diagram showing our trajectories in phase space. The primary trajectory, which corresponds to sequences shown in all previous plots, is given by the diagonal dashed purple curve and open circles that intersects the “pre-jammed” boundary at 74% and that ends at Point J, $\{\phi \approx 0.83, T = 0\}$. The grayed region is forbidden for impenetrable beads. The solid circles, triangles, and square denote measurements of the phase boundary for large hollow polypropylene beads, solid steel beads, and tiny solid steel beads, as shown in the images. Each boundary point corresponds to the trajectory indicated by the dashed going through it.

In spite of the vast differences in bead systems, a remarkably consistent region of pre-jammed behavior appears in the temperature vs area fraction jamming phase diagram.

2.7 Conclusion

The quasi-2D system of air-fluidized beads studied in this paper is fundamentally different from an equilibrium system. Here all microscopic motion arises from external driving, and has nothing to do with thermodynamics and ambient temperature. Rather there is a constant input of energy from the fluidizing air, and this excites all motion. As a result of thus being far from equilibrium, the velocity distributions are not Gaussian. Furthermore, the average kinetic energy of the two species is not equal because of how they interact with the upflow of air.

In spite of these differences, our system exhibits hallmark features upon approach to jamming that are very similar to the behavior of thermal systems. In terms of structure, our system develops a split peak in the circular factor distribution and a split second peak in the pair distribution function. In terms of dynamics, our system develops a plateau in the mean-squared displacement at intermediate times, between ballistic and diffusive regimes, where the beads are essentially trapped in a cage of nearest neighbors. And also like a thermal system, the prominence of these features increases with both increasing packing fraction and with decreasing particle energy, in a way that can be summarized by a jamming phase diagram.

The significance of the above conclusions is to help reinforce the universality of the jamming concept beyond just different types of thermal systems, to a broader class of non-equilibrium systems as well. This suggests that the geometrical constraints of disordered packing plays the major role. Our system of air-fluidized beads may now serve as a readily-measured model in which to study further aspects of jamming that are not readily accessible in thermal systems. For example it should now be possible to characterize spatial heterogeneities and dynamical correlations in our system, expecting the results to shed light on all systems, thermal or not, that are similarly close to being jammed.

Chapter 3

Direct Observation of Growing Dynamical Lengthscales for Driven Steel Beads

This work was done as a collaboration with Sharon Glotzer and Aaron Keys at the University of Michigan, Ann Arbor and has been submitted for publication.

3.1 Introduction

Recent experiments have shown that the anomalous dynamical characteristics known as “spatially heterogenous dynamics” (SHD) previously observed in supercooled liquids and colloids near the glass transition are present in granular materials as well. This supports the long-standing hypothesis that a fundamental link exists between the glass transition and “jamming” in granular materials. Here, we perform a complete characterization of SHD in a system of air-driven granular beads as a function of increasing density and show that as with the glass transition, growing dynamical lengthscales accompanies the jamming transition. Moreover, the divergence of these dynamical lengthscales is accurately described by standard equations originally developed for glass forming liquids. Our results demonstrate strong quantitative parallels between jamming and the glass transition, and provide a significant step forward in the quest for a unified theory of “jammed” phenomena.

Supercooled liquids and dense colloidal suspensions exhibit anomalous yet ubiquitous dynamical characteristics known as spatially heterogeneous dynamics (SHD), which become increasingly pronounced upon cooling towards the glass transition [25, 42, 54, 91]. Key features of SHD include unusual correlations [8] in which particles move in one-dimensional paths (“strings [18]”) that aggregate into clusters [20], and dynamical correlations as measured by a dynamic four-point susceptibility [10, 19, 45]. Granular systems undergo a jamming transition that is in many ways similar to the glass transition in liquids, an observation that has elicited speculation that the two classes of matter may share many dynamical features on approach to a universal jammed state [49, 79]. Recently, two studies have confirmed that similarities between glassy and so-called “jammy” dynamics exist by studying sheared granular systems under compaction¹⁴ and constant density [55]. Ref. [17] has shown direct statistical evidence of SHD in granular materials by measuring the time-window dependence of χ_4 for a sheared system of granular beads at a single packing density.

Using diagnostics developed for thermal glass-forming liquids, we perform a complete characterization of SHD in an athermal, far-from-equilibrium system of air-driven steel spheres over a range of packing densities and temperatures where characteristic relaxation times change by several orders of magnitude. Thus, we are able to cover the dynamical range necessary to permit a meaningful comparison of the anomalous dynamics in quiescent granular and liquid systems for the first time. We find that our granular system exhibits many of the same dynamical characteristics as supercooled liquids, which, in addition to strong correlations among caged particles, includes clusters of mobile particles, string-like motion, growing timescales, and evidence of growing dynamical length scales with approach to the jamming transition. We show that the observed growth of time and length scales can be described using standard models developed for supercooled liquids, and that these models can be used to predict the jamming transition area fraction and a “pre-jamming” transition [2] that may be analogous to the mode coupling transition in supercooled liquids [33]. Thus, for the first time, we draw a quantitative connection between jamming and the glass transition.

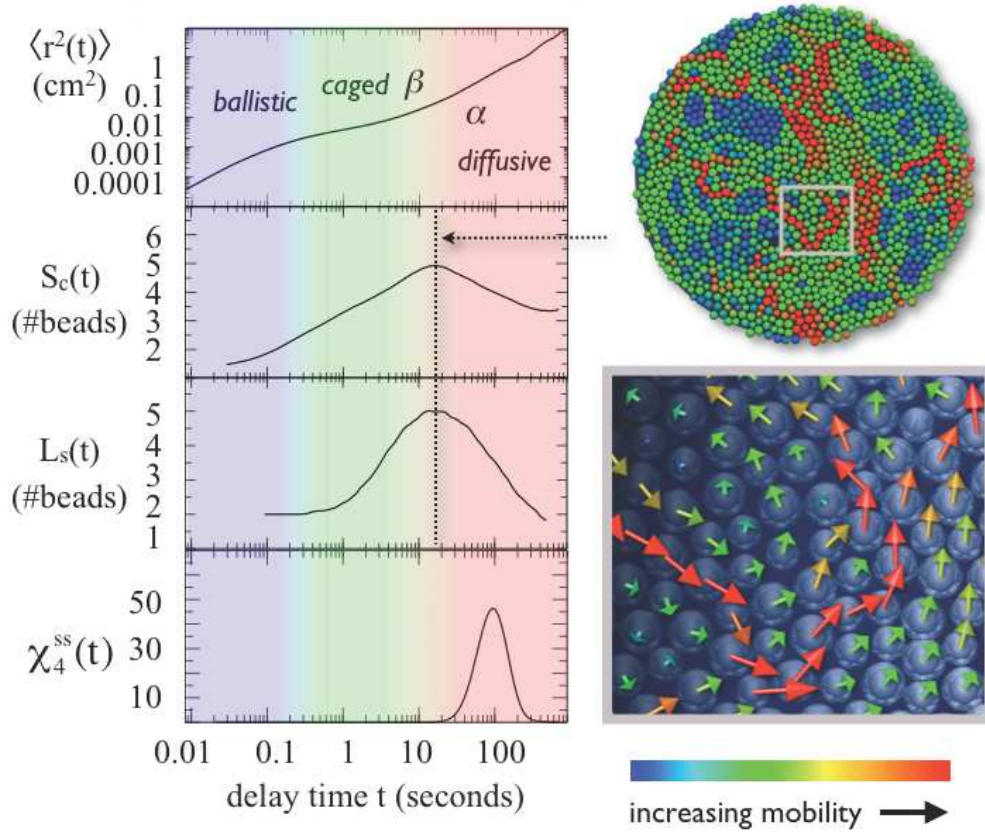


Figure 3.1: [Dynamics773] Left, from top to bottom: mean square displacement $\langle r^2(t) \rangle$, average mobile bead cluster size $S_c(t)$, weight average string length $L_s(t)$, self contribution to the four point susceptibility $\chi_4^{ss}(t)$. As reported for liquids approaching their glass transition, the maximum in $S_c(t)$ and $L_s(t)$ occur in the late β - early α -relaxation regime and is maximum in the α -relaxation regime. Right: illustration depicting bead mobility over a time interval $t = 12$ seconds, $t_{S_c}^*$ or $t_{L_s}^*$, when mobile cluster size $S_c(t)$ and string length $L_s(t)$ are at a maximum. Top: beads are colored based on their relative mobility over the time window, as determined by the absolute displacement from their initial positions showing the spatially heterogeneous nature of the bead dynamics. Bottom: Close up, with vectors depicting the correlated motion of beads. Each vector connects the initial position of the bead to its new position at the end of the time window.

3.2 Details

In this study, we consider a granular system consisting of a 1:1 bidisperse mixture of steel spheres (“beads”) of diameters $d_s = 0.318$ cm and $d_l = 0.397$ cm and masses $m_s = 0.130$ g and $m_l = 0.266$ g. The packing density (or area fraction) ranges from $\phi = 0.597$ to $\phi = 0.773$ and is controlled by varying the number of beads (from $N = 1470$ to $N = 1904$) for a fixed circular system of diameter 17.70 cm. Bead motion is restricted to rolling within a horizontal plane, and is excited by an upward flow of air at a fixed rate of 545 cm/s. Bead positions are identified by reflecting incandescent light from their chrome surface to a camera three feet above. Each experiment runs a total length of 20 minutes. By contrast with the molecules in a supercooled liquid, the beads are macroscopic objects driven by a continuous input of energy. Consequently the speed distributions are non-Maxwellian, and the average kinetic energies of the two bead species are unequal. Nevertheless the dynamics of the system mimic those of a simple liquid for low ϕ and show the first signs of caging around $\phi = 0.65$. Beyond a pre-jamming transition at $\phi_{pj} = 0.74$, the structure and dynamics show hallmark features of the approach to jamming, including diverging cage times with proximity to the measured jamming transition at $\phi_j = 0.83$ at which point all motion stops. Over the range of ϕ studied, characteristic relaxation times increase by 2-3 orders of magnitude.

3.3 Strings

We characterize SHD in this system by performing three measurements developed for supercooled liquids, all of which are plotted together with the mean squared displacement for $\phi = 0.773$ in fig. 3.1, and described in detail below. We begin by considering the formation of clusters of the most mobile beads in the system, whose net displacement from their original position greatly exceeds the average value as depicted by the red beads in fig. 3.1. Specifically, we label beads as “mobile” if their displacement ranks among the top 10% of all bead displacements after a given delay time t , which maximizes the distinction between mobile and immobile beads. Mobile beads that are nearest neighbors (i.e, their Voronoi cells share an edge) are said to belong to the same cluster, as depicted in fig. 3.2a.

Fig. 3.2b shows the temporal evolution of the number average mobile bead cluster size defined

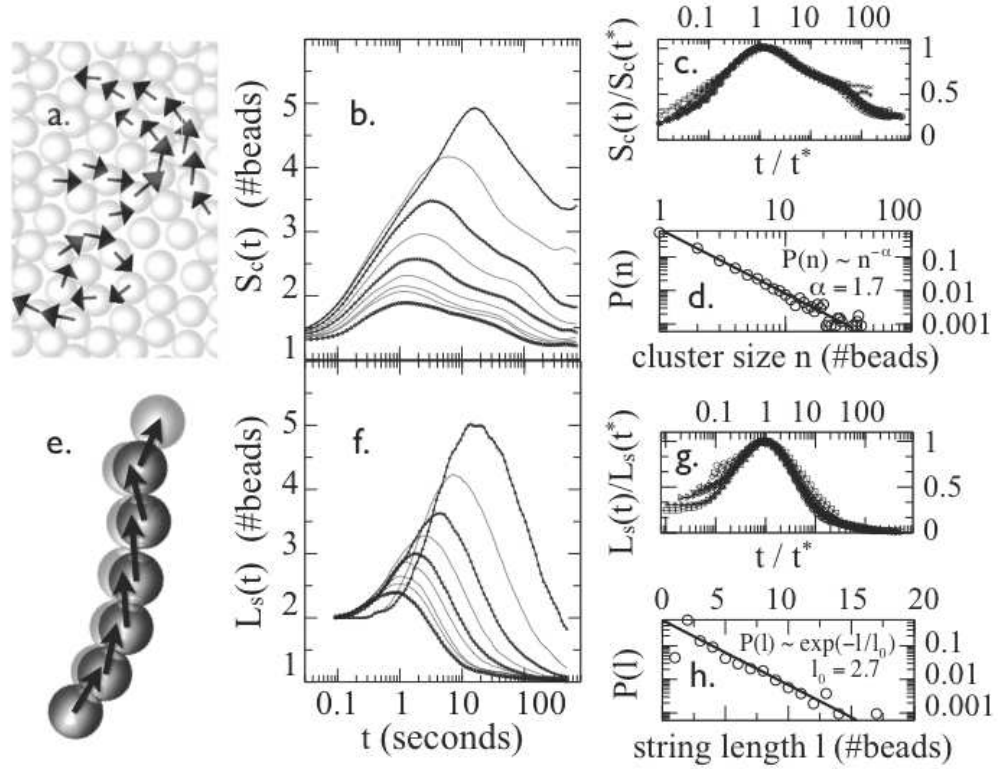


Figure 3.2: [Strings] Characterization of SHD for the most mobile beads in the system. (a) Depiction of a mobile bead cluster at $\phi = 0.773$ and $t = t_{S_c}^*$. Here, the beads with black arrows belong to a mobile bead cluster. Each arrow represents a vector drawn from the initial position of the bead to its new position $t_{S_c}^*$ seconds later. (b) Number average mobile bead cluster size $S_c(t)$ in units of number of beads as a function of ϕ . From bottom to top: $\phi = 0.597, 0.633, 0.647, 0.669, 0.693, 0.722, 0.742, 0.762, 0.773$. (c) All curves from (b) together on the same plot, where t is scaled by $t_{S_c}^*$ and $S_c(t)$ is scaled by $S_c(t_{S_c}^*)$. (d) Cluster size distribution for $\phi = 0.773$ at $t_{S_c}^*$ fit by a power law. (e) Depiction of a string of mobile beads at $\phi = 0.773$ at $t = 12$ seconds. Dark grey beads represent bead positions at $t = 0$ and light grey beads (mostly covered by dark grey beads) represent the new bead positions $t_{S_c}^*$ seconds later. Each black arrow represents a vector drawn from the initial position of the bead to its new position. (f) Weight average string length $L_s(t)$ in units of number of beads as a function of time. From bottom to top: $\phi = 0.597, 0.633, 0.647, 0.669, 0.693, 0.722, 0.742, 0.762, 0.773$. (g) All curves from (f) together on the same plot, where t is scaled by $t_{L_s}^*$ and $L_s(t)$ is scaled by $L_s(t_{L_s}^*)$. (h) String length distribution for $\phi = 0.773$ at $t_{L_s}^*$ fit by an exponential.

as

$$S_c(t) = \langle n \rangle = \sum_n nP(n) \quad (3.1)$$

where $P(n)$ is the probability of observing a cluster of size n . For short times (in the ballistic and plateau regimes), the most mobile beads are randomly distributed throughout the system and the cluster size is small. At an intermediate time $t_{S_c}^*$ the clusters achieve a maximum average size. As with glass-forming liquids [30,31,89] and colloids [91], $t_{S_c}^*$ coincides with the crossover from the plateau regime to the diffusive regime in the mean square displacement. As ϕ is increased, $S_c(t)$ shifts to longer times and larger cluster sizes, indicating - as in liquids - that the cooperative motion of larger groups of particles is required for beads to escape their cages as the jamming transition is approached from below.

Fig. 3.2c shows that data for all collapse onto a single master curve for intermediate to late times when t is scaled by $t_{S_c}^*$ and S_c is scaled by $S_c(t_{S_c}^*)$, which implies that $t_{S_c}^*$ is a characteristic time for this system. Fig. 3.2d shows a typical cluster size distribution $P(n)$ at $t_{S_c}^*$. The distribution approaches a power-law of the form $\lambda = 1.7$ with increasing ϕ . A power-law distribution is consistent with the presence of a percolation transition of mobile clusters within the region of ϕ studied. Analogous behavior has been observed for simulations of supercooled liquids near the mode-coupling temperature with similar values of the scaling exponent λ .

In glass-forming liquids particles within mobile-particle clusters have been shown to move in quasi-1d paths called “strings”, depicted in Fig. 3.2e. Clusters of strings arise naturally in dynamic facilitation theory [28] and the random first-order theory of glasses [83] and their shape reflects the fractal nature of dynamical motion in these systems. Strings are also a crucial ingredient in a recent theory of dynamics near the glass transition [47]. We test for the presence of strings in our system by considering mobile beads that “follow” one another. Specifically, if a bead replaces the initial position of another bead within a tolerance $c = 0.3d_s$ after a delay time t , then both beads are considered to engage in string-like motion. Neighboring beads that engage in string-like motion are said to belong to the same string.

We detect strings of mobile beads for all ϕ . Previously, such dynamical objects had only been measured statistically in numerical simulations of atomic [18, 30, 98] molecular [89], and polymer [4] glass formers, although they have been observed visually in colloidal systems [42, 54, 91] approaching the glass transition and in a granular system [17] near the jamming transition.

Fig. 3.2f shows the temporal evolution of the weight average string length, defined as l_s , where l is the number of beads in a string [4, 30]. The weight average corresponds to the average length of a string containing a bead chosen at random, provided that the bead belongs to any string. We observe that for short times, all strings have a length of two beads, indicating the presence of very short-ranged correlated motion. At longer times, we observe larger strings, thus demonstrating the increasing length scale over which bead motion is correlated. As for liquids, the evolution of the string length mirrors the evolution of the mobile cluster size, with maxima $t_{S_c}^*$ and $t_{L_s}^*$ at nearly identical positions. For delay times much longer than $t_{L_s}^*$ in the diffusive regime, beads follow beads that are not nearest neighbors and the average string size approaches one. Fig. 3.2g shows that data for all collapse onto a single master curve for intermediate and long times when t is scaled by $t_{L_s}^*$ and L_s is scaled by $L_s(t_{L_s}^*)$. The poor collapse at short times is due to the fact that all strings initially have the same length of two.

Fig. 3.2h shows that the strings roughly follow an exponential length distribution given by $P(l) \propto \exp(-l/l_0)$. This behavior is consistent with that reported in simulations of several different supercooled liquids [4, 18, 30]. We find that for a given ϕ the value of l_0 is equal to the mean string length at t^* plus a constant value of 1.3.

3.4 Four-Point Susceptibility

Thus far, we have considered the cooperative motion of the most mobile beads in the system. We next consider the susceptibility $\chi_4^{SS}(t)$, which is a measure of the fluctuations in the number of caged particles in the system, and indicates, through its relation to a suitably defined pair correlation function [46], the extent of correlations among caged particles. A large (small) value of $\chi_4^{SS}(t)$ indicates a high (low) degree of spatial correlation among caged particles, and is thus an alternative measure of the extent to which the dynamics in the system is spatially heterogeneous on a timescale t . Mathematically, $\chi_4^{SS}(t)$ is expressed as fluctuations in the average of the self overlap order parameter $q_s(t)$ as,

$$\chi_4 = N [\langle q_s(t)^2 \rangle - \langle q_s(t) \rangle^2], \quad (3.2)$$

where

$$q_s(t) = (1/N) \sum_{i=1}^N w (|r_i(t) - r_i(0)|) \quad (3.3)$$




Figure 3.3: [X4] Characterization of the correlations among the most immobile (caged) beads in the system. Top: Average value of the overlap parameter $q_s(t)$ as a function of time window for different ϕ . From left to right: $\phi = 0.597, 0.633, 0.647, 0.669, 0.693, 0.722, 0.742, 0.762, 0.773$. Bottom: as a function of time window for different ϕ . In order of increasing peak height: $\phi = 0.597, 0.633, 0.647, 0.669, 0.693, 0.722, 0.742, 0.762, 0.773$. Inset: Log-log plot of as a function of time window for different ϕ . In order of increasing peak height: $\phi = 0.597, 0.633, 0.647, 0.669, 0.693, 0.722, 0.742, 0.762, 0.773$.

with $w = 1$ if $|r_i(t) - r_i(0)| < a$ and 0 otherwise, and averages are taken over the starting time “0”. Here, N is the total number of beads, $r_i(t)$ is the position of bead i at time t and a is the overlap cutoff taken as $0.5d_s$ for this study. $\chi_4^{SS}(t)$ represents the self-contribution to the generalized four-point susceptibility $\chi_4(t)$ and is the dominant contribution.

Fig. 3.3 shows the time evolution of $\langle q_s(t) \rangle$ and $\chi_4^{SS}(t)$ for different ϕ . As beads begin to move from their initial position, the fraction of self-overlapping beads $\langle q_s(t) \rangle$ decreases from its initial value of unity towards zero. At a time $t_{\chi_4}^*$ the positions of the caged particles are most correlated and $\chi_4^{SS}(t)$, so that the fluctuations in $\langle q_s(t) \rangle$ are maximum. As with supercooled liquids, $t_{\chi_4}^*$ occurs in the α -relaxation regime [19, 32, 46, 86, 90] and both the peak value of $\chi_4^{SS}(t)$ and $t_{\chi_4}^*$ increase with increasing ϕ . A similar peak in $\chi_4(t)$ was recently reported at a single packing density for a sheared granular system of beads [17].

Thus we observe that the peak values $S_c(t_{S_c}^*)$, $L_s(t_{L_s}^*)$ and $\chi_4^{SS}(t_{\chi_4}^*)$ and the timescales $t_{S_c}^*$, $t_{L_s}^*$ and $t_{\chi_4}^*$ increase with increasing ϕ . It has been hypothesized that this phenomenon, which occurs in supercooled liquids as a function of temperature, is indicative of the system approaching a critical point where a growing length scale causes the timescale to diverge. Whereas for critical phenomena this growing length scale is static, for the glass and jamming transitions it is dynamic with no obvious change in static structure [2, 23, 52, 78]. We next show how the timescales t^* and the dynamical length scales (given roughly by the peak height at t^* in the case of S_c and L_s , and as a lower bound $\xi = \sqrt{\chi_4^{SS}(t^*)/\pi}$ for χ_4^{SS}) scale for different quantities as a function of ϕ . The functional form of these quantities has been predicted by simulation for glass formers as a function of temperature, although there are no such predictions for granular materials.

3.5 Divergent Dynamical Lengthscales

To draw further analogy to the glass transition we fit our measurements from the previous sections to the equations of mode-coupling theory (MCT) and the Vogel-Tammann-Fulcher (VTF) equation. Motivated by application of MCT to colloids, we fit the data to a power-law relationship of the form

$$f(\phi) \propto |\phi - \phi_c|^{-x} \quad (3.4)$$

where f is a characteristic timescale or peak height measured above. Here we let both ϕ_c and x float during the fit. For the VTF equation, we fit the data to the function

$$f(\phi) \propto \exp(E/|\phi - \phi_0|) \quad (3.5)$$

where we let E and ϕ_0 float during the fit. The parameters ϕ_c and ϕ_0 are critical points that are analogous to the mode coupling temperature and the ideal glass transition temperature in the thermal MCT and VTF equations, respectively.

We first obtain best fits to the MCT and VTF equations for each measurement independently. Remarkably, the predicted value of ϕ_0 from the VTF equation has an average of $\phi_0 = 0.83 \pm 0.02$, which is coincident with the value of $\phi_j = 0.83$ at the jamming transition measured for this system. That is, the VTF equation predicts a critical point at the measured jamming transition area fraction for the system. The predicted values of ϕ_c from the MCT equation are somewhat lower, with an average of $\phi_c = 0.78 \pm 0.02$, close to the pre-jamming transition at $\phi_{pj} = 0.74$. Fig. 3.4 shows fits to the MCT and VTF equation using the best estimate of the critical area fractions, given by the average values $\phi_0 = 0.78$ and $\phi_0 = 0.83$, respectively. We find that both expressions yield excellent fits to the data within the error bars for the range of ϕ considered. The fact that ϕ_c is lower than ϕ_0 is consistent with the behavior observed for supercooled liquids where the predicted mode coupling temperature is generally higher than the glass transition temperature. Interestingly, the pre-jamming transition boundary in granular media demarcates a region in which the dynamical behavior deviates substantially from that of a low density system; in supercooled liquids, the mode-coupling transition boundary coincides with a crossover to potential energy landscape-dominated dynamics [68, 81].

In summary, our results show that the SHD observed in a far-from-equilibrium system of fluidized granular beads is essentially indistinguishable from that observed in thermal systems like supercooled liquids and dense colloids, suggesting that (1) SHD arises solely from packing constraints; (2) thus, fluidized granular systems may serve as models for glass-forming liquids, just as colloids have in the past; and (3) the behavior of jammed systems may be understood using the theoretical tools developed for liquids. Our results open the door to future theoretical insight into the relationship between jamming and the glass transition and, perhaps ultimately, a unified theory of “jammed” phenomena.

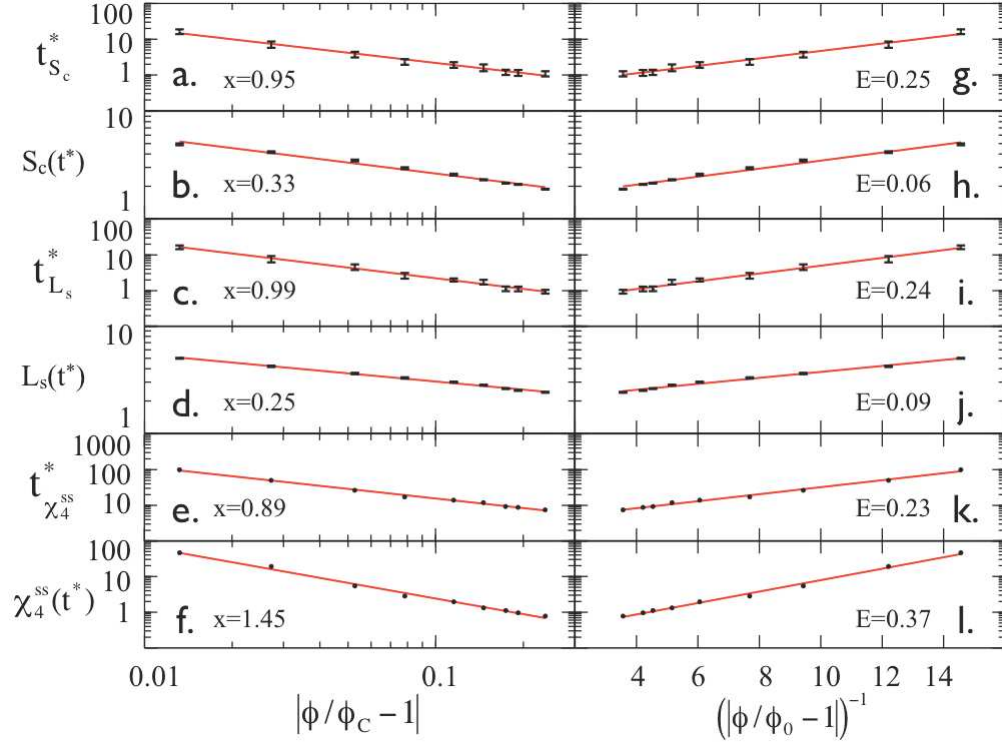


Figure 3.4: [Scaling] Dependence of dynamic timescales and length scales on area fraction. Left: Best fits for the MCT equation with fixed parameter $\phi_c = 0.78$ and variable parameter x . From top to bottom: (a) Fit for timescale $t_{S_c}^*$. (b) Fit for peak height $S_c(t_{S_c}^*)$. (c) Fit for timescale $t_{L_s}^*$. (d) Fit for peak height $L_s(t_{L_s}^*)$. (e) Fit for timescale $t_{\chi_4}^*$. (f) Fit for peak height $\chi_4(t_{\chi_4}^*)$. Right: Best fits for the MCT equation with fixed parameter $\phi_0 = 0.83$ and variable parameter E . From top to bottom: (g) Fit for timescale $t_{S_c}^*$. (h) Fit for peak height $S_c(t_{S_c}^*)$. (i) Fit for timescale $t_{L_s}^*$. (j) Fit for peak height $L_s(t_{L_s}^*)$. (k) Fit for timescale $t_{\chi_4}^*$. (l) Fit for peak height $\chi_4(t_{\chi_4}^*)$.

Chapter 4

Topological Persistence Near Jamming

4.1 Abstract

We introduce topological methods for quantifying spatially heterogeneous dynamics in a driven system of steel ball-bearings close to jamming. The persistent area and susceptibility χ_A , are based on the time-varying Voronoi tessellation and capture two-step relaxation and correlated motion. The persistent bond and susceptibility χ_B , are based on the time-varying Delaunay triangulation and capture heterogeneities on much longer timescales that have to do with cage rearrangements and the breakup of the nearest neighbor network. While the timescales of the heterogeneities differ by a decade, their spatial size and physical makeup are the same so that long-time breakup of the nearest neighbor network is initiated by short-time correlations of caged particles. The scaling behavior of the characteristic times and lengths of the heterogeneities is also seen to mimic the dynamics of atomic and polymer glasses adding a quantitative perspective on the approach to jamming.

4.2 Introduction

Non-equilibrium systems remain one of the most studied and least understood branches of physics today. Despite enormous effort and ever increasing pressure from fields like biology, many of the most basic questions about non-equilibrium systems remain unanswered. Recently, a radical idea was put forward that proposed a deep and fundamental relationship between jamming and the glass

transition [49] which has stimulated a cascade of activity and breathed fresh life into two very old fields of study. If the glass transition is connected to a wider and much more general class of non-equilibrium systems, then one may well be able to understand aspects of the glass transition by studying granular flows, the dispersal of a crowd from a music hall, and even traffic jams on the freeways of Los Angeles. Conversely, one might better understand traffic, crowds, and the movement of grain by extending knowledge harvested in the study of glass.

There is already evidence supporting a deep relationship between jamming and the glass transition. Atomic and polymeric glassformers [28–30,83,89,98], supercooled liquids [8,18,24,31,46,80], and even dense colloidal suspensions [20,42,54,91–93] have been verified to share structural and dynamical features. Structurally, they are all composed of dense packings of particles in a disordered array. Dynamically, they all exhibit anomalously slow relaxation, and small changes in temperature and density can result in enormous changes in dynamics. What’s more, the dynamics are also heterogeneous in time and space and are characterized by a diverging dynamical length scale that grows with proximity to the glass transition and is accompanied by no obvious changes in structure [10,19,32]. Understanding it is critical since it may be a general feature of all systems that are close to jamming, whether composed of atoms, polymers, colloids, grains, people or cars.

The discussion has therefore turned to quantitatively characterizing heterogeneous dynamics in a wide variety of systems close to jamming. Mode-Coupling Theory (MCT) and the Vogel-Tammann-Fulcher (VTF) equation have both been applied to the glass transitions of atoms, polymers, and colloids, and each capture certain aspects of a growing dynamical length scale [10], predicting scaling behavior and extrapolating to critical points that are close to the measured glass transition temperature [29]. Non-equilibrium granular systems provide a substantial step forward because they are insensitive to thermal fluctuations. By removing temperature from the mix the universality of spatially heterogeneous dynamics (SHD) can be tested, shedding light on which features are general characteristics of the jamming phase and which depend on the particulars of the system. To this end, stable packings of grains are being studied in detail to uncover the relationship between local and global structure [53]. The propagation of forces through the granular network [62,78] has helped uncover how force chains composed of local contacts in a disordered array can span the entire system and may act as the backbone of a dynamical length scale.




Figure 4.1: [VDDemo] The Voronoi tessellation for six different delay-time snapshots. Each cell is permanently colored according to its location at $\tau = 0$. After a long time the beads move far from their initial positions and the cell colors mix.

In this paper we study both the structure and dynamics of a granular system of steel ball bearings driven close to jamming. By supplying a constant injection of energy with air-turbulence this granular systems can be driven in a steady state close to jamming, offering a unique look at the roles of temperature, bead interactions, driving, and dissipation [1, 63, 64]. One enormous advantage of this system is that it can be easily and directly visualized over a wide range of time, length, and density scales [2, 43]. We have exploited these characteristics and developed novel topological methods that allow us to quantify heterogeneous dynamics with the range and accuracy necessary for reasonable comparison with glass theory. These methods are described in detail in the sections to follow and are based on dynamical correlation functions of the system’s geometrical topology. Fits to theory from extracted dynamical time and length scales provide scaling exponents that signal the onset of jamming and extrapolate to critical density points consistent with the system’s measured jamming transition at $\phi_j = 0.83$, adding to the triad a thoroughly different sort of system, driven far from equilibrium, where neither temperature nor statistical mechanics plays a role, that nonetheless behaves in near perfect analogy with atomic, polymeric, and colloidal systems undergoing a glass transition.

4.3 Background

The system under investigation is unchanged from previous studies [2, 43]. A bi-disperse monolayer of beads, diameters $5/32$ and $1/4$ inches in equal numbers, are confined in a circular flat sieve 6.97 inches in diameter and with mesh size 150 microns. Energy is injected by blowing air up through the sieve at a fixed rate high enough to generate turbulence ($Re \approx 10^4$) but low enough so

that the beads maintain contact with the sieve. The flow rates are between 560-700 cm/s depending on area fraction. Three feet above the sieve are six incandescent lights aligned in a ring one foot in diameter. At the center of the ring is a 120 Hz Pulnix 6710 video camera with a 644x484 array of 8BIT square pixels and a zoom lens. Light reflects specularly off the centers of the beads and is imaged by the camera. Spots are tracked so that every bead's position and size is known for the entire experiment.

We begin the discussion with a visual depiction of our system. Fig. 4.1 shows the Voronoi diagram of area fraction 79.2% at different times. Each Voronoi cell belongs to a bead that is at its approximate center. The particular shape of the Voronoi cell depends on the particular arrangement of each bead and its nearest neighbors. Beads with five nearest neighbors have five sided Voronoi cells, beads with seven neighbors have seven sided Voronoi cells and so forth. When the Voronoi cells for all the beads are drawn together, they fill (or tessellate) space completely. We have colored the Voronoi cells according to their order from the top of the system at $\tau = 0$ so that the cell of the bead at the very top is colored white, the cell of the bead at the very bottom is colored black, and those in between are colored a grayscale value. The cell color of each bead is constant over the entire experiment. Fig. 4.1, $\tau = 0$ shows the Voronoi diagram at the beginning of the experiment so that there is a continuous gradient in color. As the beads jostle about the system explores new configurations and the colors begin to mix. By the end, the majority of beads are in a different location and have gone through many generations of nearest neighbors. Voronoi diagrams for intermediate times hint of spatially heterogeneous dynamics. For example, at $\tau = 10$ and 100 s, beads near one another appear to have moved as a unit and the Voronoi diagram is broken into regions of similar color.

A standard way to characterize spatially heterogeneous dynamics is a four-point correlation function known as the overlap order parameter,

$$Q(t, \tau) = \frac{1}{N^2} \sum_{i=1}^N \sum_{j=1}^N w(|r_i(t) - r_j(t + \tau)|) \quad (4.1)$$

with,

$$w(|r_1 - r_2|) = \begin{cases} 1 & : |r_1 - r_2| < a_0 \\ 0 & : \text{otherwise} \end{cases} \quad (4.2)$$

where a_0 is the overlap threshold size and is a parameter that must be specified. The overlap order




Figure 4.2: [QTemp] (a) The average overlap order parameter and (b) its scaled variance, which is the self-part of the four-point susceptibility, for different area fractions as labelled. For high area fraction the configuration remains static for longer durations and a dynamical length scale associated with the peak of χ_4 diverges.

parameter can be broken up into self contributions ($i = j$), in which the particle overlaps with itself, and distinct contributions ($i \neq j$), in which the particle overlaps with a different particle. Considering only self contributions, $Q(t, \tau)$ is the fraction of particles that have remained within a distance a_0 of their position at t over the interval τ . The choice of a_0 greatly influences the time and length scales captured by $Q(t, \tau)$ and must be chosen with care. At $\tau = 0$, $Q(t, 0) = 1$ since all particles are at their initial location. As particles move around $Q(t, \tau)$ decays to zero in the time that it takes for all the beads to displace from their initial positions a distance a_0 at least once. Repeating the measurement of $Q(t, \tau)$ for different times and averaging together we obtain the time-averaged overlap order parameter, $Q(\tau)$. $Q(\tau)$ relates how long, on average, it takes for the system's configuration to change. Fig. 4.2a shows $Q(\tau)$ measured for different area fractions for $a_0 = \langle R \rangle$, the average bead radius. For low area fractions the beads move about easily and the system samples new configurations quickly so that $Q(\tau)$ decays after about a second. For high area fractions the beads are cramped and the configuration remains jammed into place for a long time. The higher the area fraction, the more cramped beads are and the longer the configuration is locked-in, representative of a diverging relaxation time on approach to jamming. Indeed, fig. 4.2a shows that $Q(\tau)$ takes 100 times longer to decay at $\phi = 0.797$ than at $\phi = 0.626$.

The average overlap order parameter therefore captures a slowdown in dynamics with increasing area fraction. Motivated by previous studies [30, 43] we introduce the normalized variance of the overlap order parameter χ_4 , which is the self-part of the four-point susceptibility and is sensitive to spatially heterogeneous dynamics,

$$\chi_4 = N[\langle Q(\tau)^2 \rangle - \langle Q(\tau) \rangle^2]. \quad (4.3)$$

Clusters of fast and slow particles have been seen in simulations of supercooled-liquids [30] and dense colloidal suspensions [91]. Recently, evidence for a diverging dynamical length scale associated with clusters of fast and slow particles was observed in glass-forming liquids [10]. Only recently has χ_4 been measured in a granular system [17] and even more recently has it been measured as a function of area fraction to determine the scaling of spatially heterogeneous dynamics on approach to jamming [43]. Fig. 4.2b shows χ_4 for different area fraction for our system. At low area fractions when the configuration is free to evolve without too much difficulty and the overlap order parameter decays quickly, there is a modest peak in χ_4 at about the time that $Q(\tau)$ decays to zero, which shows that even at low area fractions there is a small amount of cooperative motion.

At higher area fraction bead dynamics slow considerably and cooperative motion increases. As a result, the peak of χ_4 occurs at later time and increases in scale, coming to a maximum around the time that $Q(\tau)$ decays to zero.

$Q(\tau)$ and χ_4 therefore provide a picture of spatially heterogeneous dynamics in this driven non-equilibrium system. But this picture is unsatisfactory for a few reasons. First, the overlap order parameter requires the specification of the overlap threshold size a_0 , which must be chosen carefully to simultaneously optimize both the average overlap order parameter and the four-point susceptibility. There is no rigorous or clear-cut way to do this. Second, at short times and at lengths smaller than a_0 , $Q(\tau) = 1$ even though the beads are not perfectly still. At long times when particles have moved just beyond a_0 , $Q(\tau) = 0$ forcing $\chi_4 = 0$ even if there is long time dynamical heterogeneity. Hence, even with the ideal choice of a_0 , the dynamical range that is observable is limited to times and lengths that are close to a_0 . Third, the discrete nature of the overlap function w , influences the magnitude of the fluctuations in the overlap order parameter $Q(t, \tau)$, and effects the overall scale of χ_4 . In the next section we introduce novel methods that overcome these difficulties and provide a more accurate and detailed picture of spatially heterogeneous dynamics.

4.4 Persistent Area

The persistent area follows naturally from the time-varying Voronoi images shown in fig. 4.1. As beads move about the system samples new configurations and the Voronoi diagrams change in kind. The persistent area is the amount of area that persists within the same Voronoi cell over the interval τ . Fig. 4.3 shows a few persistent area images for $\phi = 0.792$ for start time $t = 0$ s and delay times, τ . Area persisting within the same Voronoi cell is colored black and all other area is colored white. Initially, the entire system is black because all of the Voronoi cells trivially overlap with themselves. As time moves forward the configuration begins to change and the amount of area persisting within the same Voronoi cell decreases. Eventually, all the area is turned white since no area persists within the same Voronoi cell over the course of the experiment. The persistent area is an alternate measurement of configurational overlap that is continuous and does not require the specification input parameters. Defining $A(t, \tau)$ as the fraction of area that remains within the

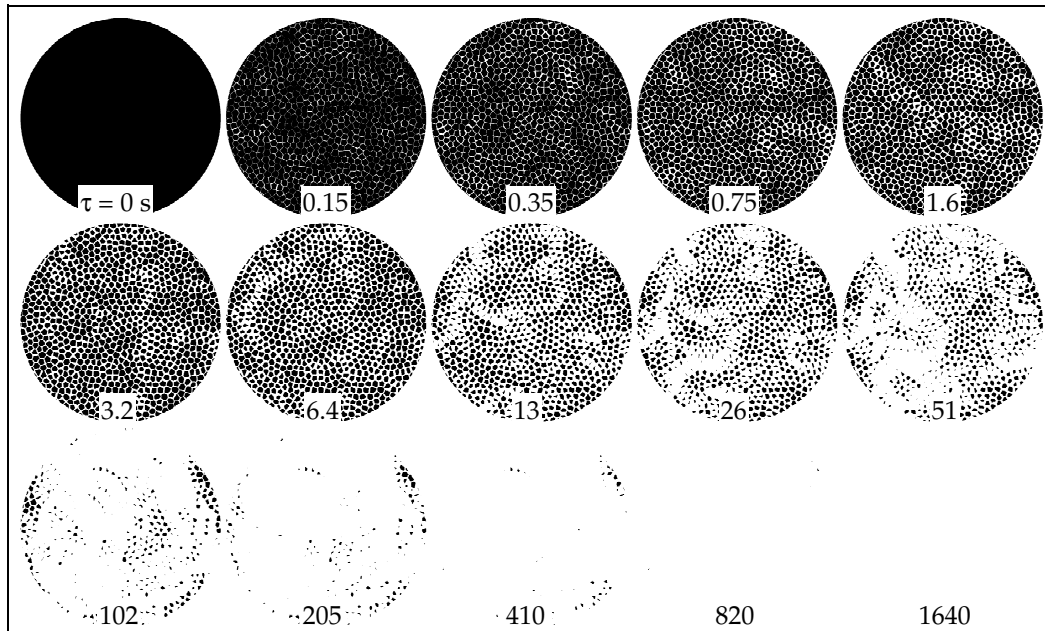


Figure 4.3: [PADemo] Persistent area images of the Voronoi tessellation for the start time $t = 0$ as a function of delay time, τ . Regions that are black have persisted within the same Voronoi cell for the time interval indicated. $A(t, \tau)$ is the total fraction of black area and is a measurement of the overlap between the configurations $\tau = 0$ and τ . Spatial heterogeneity due to stringlike cooperative motion is particularly evident at $\tau = 26$ and causes $A(t, \tau)$ to fluctuate with t .




Figure 4.4: [Amv] (a) Measured persistent area $A(t, \tau)$ plotted separately for different start times t (colored curves), for $\phi = 0.792$. Its time average $\langle A(t, \tau) \rangle = A(\tau)$ (thick green curve with circles) is a measurement of the average configurational overlap for two snapshots separated by a time τ . (b) Its scaled temporal variance $\chi_A = N \cdot \sigma^2[A(t, \tau)]$ averaged over t , is a measurement of spatially heterogeneous dynamics due to strings.

same Voronoi cell over an interval starting at t and ending at $t + \tau$,

$$A(t, \tau) = \frac{1}{A_s} \int_{x'} \int_{y'} A(x', y', t, \tau) dx' dy' \quad (4.4)$$

where,

$$A(x', y', t, \tau) = \begin{cases} 1 & : \quad \forall(x', y') \in C_i(t, 0) \cap C_i(t, \tau) \\ 0 & : \quad \text{otherwise} \end{cases} \quad (4.5)$$

with A_s the area of the region of interest, and $C_i(t, \tau)$ the x-y domain enclosed by the i-th Voronoi cell at time t with itself at time $t + \tau$. The images in fig. 4.3 graphically depict a single measurement of the persistent area over the time interval $t = 0$ to 3000 s. Following the procedure for the overlap order parameter we measure $A(t, \tau)$ for different start times, colored curves in fig. 4.4a, and average over time to obtain the average persistent area $A(\tau)$, which is analogous to the average overlap order parameter and relates the average fractional overlap for two configurations separated by time τ . Fig. 4.5a shows the average persistent area for different area fractions. Direct comparison with the overlap order parameter Fig. 4.2a, shows that the two are grossly consistent, but that much of the detail evident in the persistent area is missing in the overlap order parameter. At short times the persistent area partially decays due to caged motion while overlap order parameter is exactly one. At intermediate times the persistent area has two-step relaxation due to caging while the overlap order parameter does not. At longer times when the beads move on the order of the cage the slopes of the persistent area are slightly different while for the overlap order parameter they are all the same. The root cause of all of these problems is that the overlap order parameter is discreet and sensitive primarily to the length scale a_0 . As a consequence, much of the rich detail at different time and length scales is absent.

The persistent area images also offer a striking look at cooperative motion. In fig. 4.3, $\tau = 26$ s, stringlike swirls of cooperatively rearranging beads have swept out large domains of white while other regions where motion is cramped remain mostly black. The time trace $A(t, \tau)$ over this interval is unique to a degree that depends on the particular spatial heterogeneities that are present. Fig. 4.4a bears this out. While on average the persistent area for different initial configurations decays in a regular way, there is a systematic increase in the fluctuations that culminates at 26 s when spatially heterogeneous dynamics are evident in the persistent area images. Motivated by the overlap order parameter and four-point susceptibility we define χ_A , the size-scaled variance of the




Figure 4.5: [ATemp] (a) Average persistent area $A(\tau)$ for different area fractions showing two step relaxation due to intermediate time caging. (b) Its scaled temporal variance χ_A is analogous to χ_4 and measures spatially heterogeneous dynamics. In supercooled liquids and atomic and polymeric glassformers χ_4 peaks in the α -relaxation regime, about a decade after the mean-squared displacement breaks out of the subdiffusive regime, which is when spatial heterogeneity is a maximum and is exactly what we observe in χ_A .

of the persistent area,

$$\chi_A = N[\langle A(t, \tau)^2 \rangle - \langle A(t, \tau) \rangle^2] \quad (4.6)$$

plotted for $\phi = 0.792$, fig. 4.4b. χ_A peaks at 35 s because temporal fluctuations (fig. 4.4a) due to spatial heterogeneities (fig. 4.3) are a maximum at that delay time. In fig. 4.5b we plot χ_A for the different area fractions. For low area fractions χ_A has a modest peak at short delay that is pushed higher and to later times at high area fraction. This has a simple physical interpretation. For low area fractions, bead motion is relatively free so that most beads can move about without the cooperation of their nearest neighbors. At high area fraction, the beads are packed together so that for any one to move its neighbors must permit it by either moving out of the way or moving with it. As a result, the higher the area fraction, the larger the number of beads involved in cooperative motion, and hence the higher the peak of χ_A and the longer it takes to occur.

4.5 Persistent Bond

The persistent bond offers a complimentary look at grain dynamics on very long timescales. It is based on the nearest neighbor bonded network which is dual to the Voronoi diagram. For a given configuration the nearest neighbor bonded network is generated by performing a Delaunay triangulation to connect nearest neighbors with a bond, fig. 4.6, $\tau = 0$. The persistent bond goes one step further. It is the time-correlation function of the nearest neighbor bonded network. As beads move apart bonds break and the network dissolves. A useful characteristic of the persistent bond is that it is insensitive to translations that leave the nearest neighbor network unchanged. In systems with spatially heterogeneous dynamics it is important to distinguish between the coordinated translation of local network regions from the rearrangement of the network itself. Fig. 4.6 shows a series of persistent bond images for $\phi = 0.792$. Comparison with the persistent area images in fig. 4.3 shows that the two appear to be related but are nonetheless different. As the beads move bonds are broken as they seek out new nearest neighbors. This process is slow at first because beads tend to remain localized within a cage. On longer timescales they remain connected to the same nearest neighbors because, as the persistent area shows, they are moving in a highly correlated way. The nearest neighbor network remains mostly intact until about 90 s when correlated motion finally leads to bond breakup and completes only after correlated beads move apart.




Figure 4.6: [PBDemo] Persistent bond images for $t = 0$ as a function of delay time τ . At $\tau = 0$ the persistent bond image is exactly the nearest neighbor network. For $\tau > 0$ beads change neighbors and bonds are broken. The persistent bond $B(t, \tau)$, is the fraction of nearest neighbor bonds that remain over the interval τ . Cooperatively rearranging beads cause shear zones and curled islets of broken bonds to open up. While the persistent area measures cooperative motion in the cage, the persistent bond is sensitive to large rearrangements that cause the breakup of the nearest neighbor network.

Fig. 4.6 graphically depicts a single measurement of the persistent bond beginning at $t = 0$ s and ending at $t = 3000$ s. Using the pictures as a guide we define $B(t, \tau)$ as the fraction of bonds that remain connected over the interval t to $t + \tau$,

$$B(t, \tau) = \frac{1}{2} \sum_{i=1}^N \sum_{j=1}^N \frac{B_{ij}(t, \tau)}{B_{ij}(t, 0)} \quad (4.7)$$

where,

$$B_{ij} = \begin{cases} 1 & : \quad i \text{ and } j \text{ nearest neighbors} \\ 0 & : \quad \text{otherwise} \end{cases} \quad (4.8)$$

The persistent bond begins at one and decays to zero over the time takes for all the beads to exchange their nearest neighbors at least once. By measuring the persistent bond for independent initial configurations we obtain many independent persistent bond measurements in analogy with fig. 4.4a for the persistent area. Averaging over initial configurations at different times we obtain the average persistent bond $B(\tau)$, which relates how long on average it takes for the nearest neighbor network to break apart. Fig. 4.7a shows the average persistent bond measured for different area fractions. At low area fractions the persistent bond decays quickly since beads can move away from each other with ease. At high area fraction the persistent bond decays much more slowly for two reasons. First, motion is cramped and beads tend to stay put and so their nearest neighbors remain the same. Second, because of correlated motion neighboring beads tend to move together. Cooperatively rearranging beads preserve some of their nearest neighbor bonds longer than they would if they were moving independently. As a consequence, bond breakup only completes on timescales over which cooperatively rearranging beads move apart.

Turning back to fig. 4.6, bond breakup occurs in stringy domains and curled islets that border regions of high connectivity. Cooperatively rearranging beads can therefore cause spatial heterogeneities in the bonded network where shear zones of broken bonds separate highly mobile cooperatively rearranging beads from more static ones. To quantify the degree of spatial heterogeneity in bond breakup we compute the time-average fluctuations of the persistent bond time traces. In analogy to χ_A , we define χ_B as the size-scaled variance of the persistent bond,

$$\chi_B = N[\langle B(t, \tau)^2 \rangle - \langle B(t, \tau) \rangle^2] \quad (4.9)$$

In fig. 4.7b we plot χ_B for different area fractions. The shapes of χ_B resemble those of χ_A with the important difference that they peak nearly a decade later in time. Interestingly, despite this wide




Figure 4.7: [BTemp] (a) The average persistent bond $B(\tau)$ for different area fractions showing the average rate at which the bonded network breaks apart and correlated beads move apart. (b) Its scaled temporal variance χ_B shows that spatial heterogeneity in bond breakup increases with area fraction and peaks about a decade after χ_A , when $A(\tau) = 0$.

separation of timescales, the heights of the peaks in χ_B are comparable to the heights of the peaks in χ_A , implying that the average size of the heterogeneities in the persistent area and the persistent bond are comparable: correlated motion initiated at short time and registered first in the persistent area eventually leads to bond breakup of comparable size much later.

4.6 Applications

To tie everything together we construct fig. 4.8 for area fraction $\phi = 0.792$. The first two columns relate to structure and the last three columns relate to dynamics. Rows are spaced in time increments of $\tau_A^*/2$ which based on the persistent area is the time over which the configuration just begins to change and is close to the “cage-breakout” time (crossover from subdiffusive to diffusive dynamics) in the mean-squared displacement. On this timescale strings persist so that there is continuity from one row of images to the next.

Even though no static structural length scale has been directly related to dynamical heterogeneities, structure is important, and many believe that it must be the root cause of spatially heterogeneous dynamics. To more completely characterize this system and to make evident relationships that may exist between structure and dynamics, we construct the Voronoi diagram for $\tau = 0$ colored according to the number of sides of each cell in fig. 4.8. Cells with many sides are colored darker than cells with few. The vast majority of the light colored cells are pentagons and the very dark cells are heptagons. The moderate gray colored cells are hexagons which have the ideal number of sides for packing in two dimensions and therefore correspond to the “most ordered” regions. Five and seven sided cells tend exist in pairs (five-seven dislocations) and link up in connected strings. Four sided cells occur rarely, about once every four images, and eight or greater sided cells are extremely rare for this particular size ratio and area fraction. Previous experiments [2] showed that the side distribution narrows at high area fraction so that most cells are five, six, or seven sided, which is what we observe. Scanning down and stepping forward through time, the Voronoi diagrams change considerably in detail but have the same overall structure. Regions that are six sided in one row may have a five-seven dislocation string slicing through them in the next so that there is little continuity in changes in coordination number structure on this timescale.

One problem with analyzing structure in terms of the number of sides of the Voronoi cells




Figure 4.8: [ZCivPaPb] Rows show snapshots in increments of $\tau_A^*/2$. First column: Voronoi cells at $\tau = 0$ colored according to the number of sides so that darker has more sides. Second column: Voronoi cells at $\tau = 0$ colored according to circularity so that darker is more circular. Third column: Instant velocity vectors at $\tau = \tau_A^*$ colored so that faster beads have large white vectors and slower beads have small red vectors. Fourth column: Persistent area images at $\tau = \tau_A^*$. Fifth column: Persistent bond images at $\tau = \tau_B^*$. There seems to be little relationship between structural defects and strings. Regions of highly cooperative motion in the persistent area are regions of high bond breakup much later in the persistent bond. Spatially heterogeneous dynamics evident in the average velocity vector images cause $A(t, \tau)$ and $B(t, \tau)$ fluctuate with t .

is that it is a discreet observable and insensitive to small structural variations. The circularity factor of Voronoi cells has recently been proposed as a far more sensitive alternative that is a continuous variable [57]. The second column shows the Voronoi diagram for $\tau = 0$ colored according circularity. More circular cells are colored darker than less circular and tend to have more sides. Regions of high circularity tend to be clumped together comprised primarily of six sided cells. Like the side-colored images, the circularity-colored images change with little continuity on this timescale. Evidently, structural changes registered this way change on timescales faster than the lifetime of strings. Viewed from this perspective, there seems to be little relationship between local structure and the emergence of strings. Still, both measurements confirm that the system is disordered, that less ordered regions tend to occur in linked strings of five-seven dislocations, and that more ordered regions of hexagons tend to occur in clumps.

The lack of static structural signatures of dynamical heterogeneities has been well-documented in glassformers, colloids, and granular systems. Yet, dynamical heterogeneities do exist and are obvious when viewed directly. The third column bears this out, showing the average velocity vector field for the interval t to $t + \tau_A^*$. Large white vectors correspond to fast moving beads and small reddish vectors to slow moving beads. Not only are certain beads moving much faster than others but the fast moving beads are lined up in quasi-one-dimensional paths. Such strings are only evident on a narrow range of timescales. At short times motion is ballistic and the vectors of nearby beads are uncorrelated [2]. At long times the motion is diffusive and the beads behave independent of one another. But at intermediate times the vectors are strongly correlated into strings and swirls, as shown in the images, and the time and length scales of the strings both increase with proximity to the jamming transition. Our motivation has been to quantitatively analyze these strings and to uncover whether their scaling behavior is consistent with similar structures observed in glasses and colloids approaching a critical density point.

The persistent area provides us with the quantitative perspective that is necessary for suitable comparison with theory. The timescale chosen to construct the velocity vector field τ_A^* , is the characteristic time of the peak of χ_A for this area fraction. Comparing the average velocity vector field images with the persistent area images in the fourth column, large velocity vectors of stringlike swirls correspond to large whited out domains in the persistent area images, and regions with small vectors where beads are localized are mostly black in the persistent area images. The persistent

area relates the ratio of white-to-black area directly to the system's dynamics across the time interval. By measuring the quantity of white-to-black we measure the dynamics. If the dynamics are heterogeneous in space, so will be the distribution of white and black area, which is clearly visible in the persistent area images on this timescale. A first order quantification of the spatial heterogeneity of colored area is the scaled variance of the color values over the entire image, which is precisely χ_A .

After the beads have moved beyond their initial Voronoi cell, the entire system is whited out and the persistent area is no longer able to detect heterogeneities. Any cooperative motion that exists over such large length scales will not be visible. But because the nearest neighbor network dissolves only after correlated beads move apart, the persistent bond is sensitive to rearrangements of beads, and is not limited to a particular time or length scale, and in this way is complimentary to the persistent area. Inspection of persistent bond images in column five show that bond breakup tends to occur in regions where strings were active. The intermediate time correlated motion of strings captured by the persistent area eventually leads to correlated bond breakup: beads moving in strings at intermediate times eventually burst through their cages together and rearrange as a unit.

4.7 Approach to Jamming

The persistent area and persistent bond highlight characteristic time and length scales associated with spatially heterogeneous dynamics in this system. The magnitudes of the susceptibilities, χ_A and χ_B , are proportional to the sizes of the heterogeneities. Both sets of curves, fig. 4.5b and 4.7b, have peaks at the times corresponding to maximums in spatially heterogeneous dynamics. In the case of χ_A the peak corresponds to the time at which correlated motion is a maximum. In the case of χ_B it is when correlated bond breaking due to rearrangements is a maximum. In fig. 4.9a we plot the times of the peaks of τ_A^* , τ_B^* , and τ_4^* vs. area fraction. Notice that τ_A^* and τ_4^* are nearly equal for all area fractions. This is because for the overlap parameter chosen, $a_0 = \langle R \rangle$, χ_A and χ_4 are essentially the same except for overall scale. χ_B on the other hand peaks nearly a decade later because this is the amount of time it takes for correlated motion in the cage to lead to rearrangements.

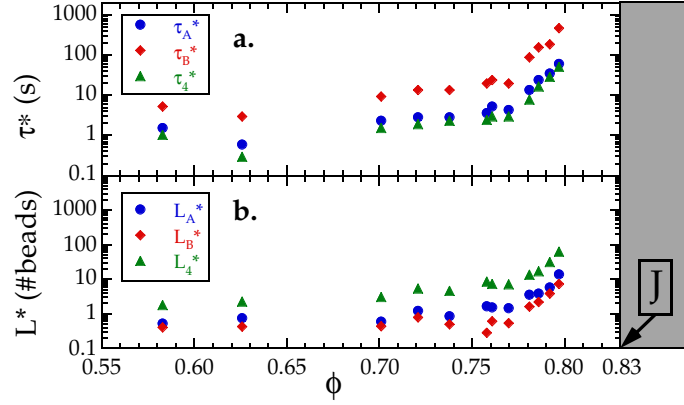


Figure 4.9: [TabqLabq] (a) Characteristic peak times of the susceptibilities χ_A (blue circles), χ_B (red diamonds), and χ_4 (green triangles) as a function of area fraction. (b) Characteristic lengths computed from the ratio of the susceptibility to the order parameter, eq. 4.12. Both the characteristic times and lengths diverge on approach to the jamming transition density at $\phi = 0.83$.

The peaks of these susceptibilities therefore highlight characteristic timescales due to correlated motion. The heights of the susceptibilities at the peaks are related to characteristic length-scales of the correlated motion. These lengths can be derived statistically. Defining S to be the number of strings that exist at any given instant and assuming it to be Gaussian distributed,

$$S = \langle S \rangle \pm \sqrt{\langle S \rangle}. \quad (4.10)$$

Defining n_s to be the average number of beads per string, N the total number of beads, and Q the total fraction of beads in strings, then $S = NQ/n_s$ and,

$$Q = \langle Q \rangle \pm \sqrt{\left\langle \frac{n_s Q}{N} \right\rangle} \quad (4.11)$$

so that,

$$n_s = \frac{N\delta Q^2}{\langle Q \rangle} = \frac{\chi_C}{C(\tau)} \quad (4.12)$$

which provides the average size of the strings in terms of the general order parameter C and associated susceptibility χ_C . Using eq. 4.12 we compute the sizes of the heterogeneities and plot them in fig. 4.9b. Notice that L_A and L_B are of a similar size. Evidently, the characteristic lengths of strings measured by χ_A at τ_A^* lead to heterogeneities in bond breakup measured by χ_B at τ_B^* about a decade later.




Figure 4.10: [TimeLeng] Characteristic time and length scales compared with MCT scaling (left hand side) and the VFT equation (right hand side). Divergence is consistent with both power law or exponential form which is also true in supercooled liquids. Best fits to power law divergence for all topological measurements yield MCT critical density at $\phi = 0.79 \pm 0.01$ which is consistent with previous findings. Best fits to exponential divergence yield VTF critical density at $\phi = 0.83 \pm 0.01$ which is close to point J and also consistent with previous findings.

Taken as a whole, the behavior of the characteristic time and length scales of the persistent area and persistent bond point in a single direction. As area fraction goes up, relaxation times diverge because the number of beads that must rearrange cooperatively spans an ever increasing fraction of the system and eventually diverges. This is reminiscent of the divergence of the relaxation time and viscosity of atomic and polymeric systems undergoing a glass transition. Armed with our data we can compare the form of our divergence with what is known about the glass transition.

Mode-coupling theory predicts the functional form of the divergence of the glass transition to be power-law,

$$\tau^* \propto |\phi - \phi_c|^{-\lambda}. \quad (4.13)$$

The task is to apply the MCT form to our data and judge the quality of the fit and the meaning of

the fit parameters. The parameter ϕ_c is the mode-coupling density and must be obtained empirically from a best fit to all of our data. For our system $\phi_c = 0.79 \pm 0.01$. The plots on the left side of fig. 4.10 show best fits holding ϕ_c constant. Interestingly, 0.79 is close to the “pre-jamming” transition point previously measured for a similar system at $\phi_{pj} = 0.74$ [2]. ϕ_{pj} roughly separates when the structure and dynamics go from resembling a simple-liquid to that of a glass, and its precise value depends on both the energy and density of the system. In glassformers and supercooled liquids the mode-coupling density coincides with a crossover to potential energy landscape dominated dynamics [81].

An alternate functional form that fits the glass transition is the VTF equation, an exponential divergence,

$$\tau^* \propto \exp\left(\frac{E}{|\phi - \phi_0|}\right). \quad (4.14)$$

Again, the critical density ϕ_0 is obtained empirically by best fit of eq. 4.14 to all the data. For this system we find that $\phi_0 = 0.83 \pm 0.01$ which is consistent with the previously reported jamming transition point at $\phi_j = 0.825$ [2]. Holding ϕ_0 constant and fitting eq. 4.14 to our data, fig. 4.10 shows that there is excellent agreement with the VTF equation.

It is well known in glassformers that both MCT and the VTF equation fit the data, and that $\phi_c < \phi_0$, which is analogous to $\phi_{pj} < \phi_j$, and is what we observe in this system. Just as the dynamics of glassy systems are influenced below the transition by the critical point ϕ_0 , so are the dynamics of this non-equilibrium system influenced by ϕ_j for $\phi < 0.83$. As low as $\phi = 0.74$ there are hallmark features of jamming: the pair-correlation function and the circularity factor distribution both exhibit peak splitting and the mean-squared displacement develops a subdiffusive plateau [2]. At long enough times the system behaves dynamically like a liquid but at shorter times there is some inherent rigidity. Analogously, ϕ_c demarcates the point beyond which the dynamics are dominated by the potential energy landscape so that the system has to pay the cost of climbing potential energy barriers in order to relax. At $\phi_j = 0.83$ this system jams regardless of the magnitude of the driving force [2]. Analogously, ϕ_0 is the point at which the relaxation time and viscosity of a glass are essentially infinite.

4.8 Discussion

We have introduced methods based on the Voronoi diagram and Delaunay triangulation that are sensitive to spatially heterogeneous dynamics and provide the precision and rigor necessary for comparison with mode-coupling theory and the Vogel-Tammann-Fulcher equation. To date the most common ways to characterize SHD have been the overlap order parameter and related four-point susceptibility. Both require the specification of a particle overlap threshold ahead of time that limits the dynamical range observable. The measurements we introduce are based on dual transformations of the particle configurations. As such, no displacement information is lost in the calculation and no input parameters need be specified.

The persistent area is a measurement of the structural correlation between two configurations. Comparison with the overlap order parameter shows that the two are grossly consistent, but that the persistent area shows far finer detail and two-step relaxation. The scaled fluctuations in the overlap order parameter are defined as χ_4 and are representative of spatially heterogeneous dynamics. Analogously, we define χ_A as the scaled fluctuations in the persistent area. χ_A contains much of the same information of χ_4 but with much greater resolution at short and intermediate time and without the need to specify input parameters.

The persistent bond and related susceptibility χ_B , are long-time compliments to the persistent area and χ_A . The persistent bond is sensitive to the exchange of neighbors and captures spatially heterogeneous dynamics that occur on much longer timescales due to the breakup of cooperatively rearranging regions. To date, such very long time evolution of the nearest neighbor network has been largely ignored. The persistent bond offers a way to distinguish between small translations of local network regions and the breakup of the global network, which are two distinct and essential aspects of spatially heterogeneous dynamics.

The characteristic times of the persistent area and persistent bond differ by nearly a decade in time. Nonetheless, the characteristic lengths are nearly the same so that breakup of the bonded network at long times is initiated by correlated motion much earlier. Scaling of the characteristic times and lengths with area fraction show critical divergences over a wide enough range for comparison with MCT and the VTF equation. While neither theory offers a complete description of the glass transition, both lead to scaling exponents that signal the onset jamming and highlight

critical densities near 83%, and do so in a manner that perfectly mimics atomic and polymer systems undergoing a glass transition. The take home message is twofold: (1) The glass transition and the jamming transition for this non-equilibrium system are quantitatively related by eqs. 4.13 and 4.14. (2) Spatially heterogeneous dynamics may be a generic feature of systems close to jamming, whether they be composed of atoms, polymers, colloids, or even steel ball bearings.

Finally, we remark that the methods we introduce naturally apply to other systems at higher dimension. In the case of foam, the persistent area and χ_A can be measured simply by summing up images so that the often impossible task of bubble identification is avoided. If particle locations are known then one can reconstruct the three-dimensional Voronoi tessellation and Delaunay triangulation leading directly to the persistent volume and to a three-dimensional version of the persistent bond. Both would be powerful tools for quantifying spatially heterogeneous dynamics with fine detail over a wide range of timescales for a variety of three dimensional systems.

Chapter 5

Three Effective Temperatures for Air-Fluidized Grains

5.1 Introduction

For decades geologists, geophysicists, and granular physicists have sought a theoretical framework to describe the granular state matter. This unique state of matter spans an enormous range of length scales, from nearly microscopic powders to weblike galaxy clusters, but as yet lacks adequate scientific description. Similarly, the glass transition has persisted to the current day as one of the greatest unsolved mysteries in all of science, and remains an enormous gap in our understanding of atomic thermal systems. Recently, the unifying concept of jamming has helped point out the direction of progress in both of these fields, by proposing that the two are in fact related, that granular systems close to jamming are fundamentally similar to atomic systems near the glass transition. The ultimate challenge for physicists in both cases is to develop a non-equilibrium statistical mechanics that can describe atomic systems that are in equilibrium and granular systems that are driven far from equilibrium, and in which jamming is a critical phenomena transparent in the theory. Such a lofty goal will require an enormous amount of effort and it is difficult in the present to even know where to begin and what observables will be adequate to characterize the similar behavior of so many very different systems.

Many believe that the first order of business must be to discover an effective temperature that can play the role in non-equilibrium systems that the thermal temperature plays in equilibrium

systems. Unfortunately, the wide variety of non-equilibrium systems precludes from the very outset that such an effective temperature will be simple, since it must be abstract enough to be apply to many disparate systems. In such unfamiliar territory, experiments must be done to determine which effective temperatures are appropriate. In this experiment, we explore the roles of effective temperatures in a driven dissipative granular system that in the past has acted as a testbed for the extension of equilibrium statistical mechanics to non-equilibrium systems [1, 2, 43, 63, 64]. By comparing different effective temperatures we keep track of the breakdown of key players of equilibrium statistical mechanics, such as the fluctuation dissipation theorem, ergodicity, and equipartition of energy, and learn much about which temperatures are suitable and bear heavily on the behavior of this non-equilibrium driven dissipative system.

We consider three possible candidates for effective temperatures. First, we measure the energy stored in a harmonic oscillator submerged in the system, which is sensitive to the grain-grain interaction energy and provides a effective temperature. Second, we measure the Einstein temperature based on a fluctuation dissipation theorem between diffusion and mobility of a driven test particle. Third, the simplest, we measure the average kinetic energy of the grains, or “granular temperature.” The first and third effective temperatures are ensemble averages over local quantities, and rely on ergodicity and equipartition to be consistent in the bulk. The last effective temperature utilizes the time average of a single driven test particle and relies on the Fluctuation Dissipation Theorem in order to be consistent. In thermal equilibrium all of these effective temperatures are equal to each other and to thermal temperature, demonstrating both the generality and beauty of equilibrium statistical mechanics.

But as this is not an equilibrium system we have no guarantee that such agreement exists here. We continuously inject and dissipate anergy from the macroscopic grains so that the system attains a steady-state that in certain instances and with respect to certain variables does resemble an equilibrium system and can be described by equilibrium statistical mechanics. In these cases the system is ergotic, energy is equipartitioned, and the fluctuation and response of the system is consistent with a fluctuation-dissipation theorem. In the present experiment we study the system over a wide range of densities and driving energies in which such tenants of statistical mechanics are not assured, allowing us to study consistency of the effective temperatures and the gradual breakdown of the thermal analogy.

5.2 Methods

The system under investigation is a 50-50 bidisperse mixture plastic hollow polypropylene balls, diameters 1 and 1-1/8 in. The balls are confined to a circular system and sit atop a perforated wire mesh that is flat and has a uniform hole spacing. Air is blown up through the wire mesh using a carpet dryer connected through a Variac autotransformer allowing us to control the airflow from 0 to 1000 cm/s with a resolution of 10 cm/s, which we verify with an anemometer. We also verify that the airflow is uniform across the sieve. Based on the sphere size the Reynolds number is 10^4 which is well in the turbulent regime. As a result, the spheres on the mesh shed wakes in random directions that interact with the walls and the wakes of other balls. These wakes both agitate the balls and cause them to repel one another.

We observe the balls from above using a 120 Hz Pulnix 6710 camera and stream AVI movies to disk using Microsoft's RLE codec. With this setup we can stream movies continuously for many hours at 120 Hz, which provides us with the dynamical range necessary. After movies are saved we process them to track all particles for the whole observation time using our own LabVIEW tracking routines. We have also written tracking routines to measure the orientation of a weighted sphere and also the tilt angle of the apparatus. For the weighted sphere we track a black dot at its north pole that displaces with respect to the center when the thermometer is tilted. To track of the tilt angle of the apparatus we image in a shiny thumbtack that remains static in the lab frame. Thus, when the apparatus tilts the static thumbtack appears to move. In both cases we use elementary geometry to convert the locations of these dots in the images to orientations of both the thermometer and the apparatus at each instant and while simultaneously tracking the balls. To tilt the apparatus we use an electric motor controlled by a computer which can oscillate with a frequency from 0.01 Hz to 100 Hz. The motor is connected to a swivelling platform on which sits the apparatus that allows it to tilt from 0.001 to 1 deg.

5.3 A Granular Temperature Thermometer

The first effective temperature that we present comes from a granular thermometer. The granular thermometer is a sphere identical to the other spheres but with a mass of dried Elmer's Glue at its bottom. By controlling the amount of added weight, we control the moment of inertia and location

of the thermometer's center of mass. By tracking the location and orientation of the oscillator at each instant we can measure its instantaneous energy and average over time to obtain an effective temperature. The potential energy comes from the tilt of the oscillator away from its equilibrium position, storing energy in the gravitational field just like a pendulum. The kinetic energy comes from the translational and rotational speed of the oscillator as it is driven by the air and collides with other balls. On average and in equilibrium, the energy of each degree of freedom is set by the temperature and they are therefore equal, but there is no guarantee in this out of equilibrium system that such an equality will exist.

In fig. 5.1 we show distributions of the quantities that are necessary in order to measure the potential and kinetic energy of the granular thermometer – i.e. its orientation and location as a function of time. Fig. 5.1a shows the distribution of thermometer tilt displacements and fig. 5.1b shows the distribution of thermometer velocities. The inset of fig. 5.1a shows a diagram of the thermometer displaced from its stable upright position. The power spectra in fig. 5.2 show that there is a strong signal at the natural frequency of the oscillator. When the thermometer is perfectly upright its potential energy is equal to zero. When tilted the restoring force is $F = -mgR_{cm} \sin \theta$ so that on average the energy is $\langle U \rangle = mgR_{cm}^2 \langle \sin^2 \theta \rangle$. For an oscillator in equilibrium these displacements must be sampled from a Gaussian distribution so that the mean-square average characterizes the average potential energy stored in the oscillator. In fig. 5.1b we plot the distribution of displacements from equilibrium along with a Gaussian with a mean-square equal to the means-square of the data. On average, the thermometer is displaced by about 0.6 rad with an average tilt energy of about 14 ergs. In fig. 5.1b we plot the velocity distribution with a Gaussian with a mean-square equal to the mean-square velocity of the oscillator. The velocity distribution has high tails compared to a Gaussian, which is a common feature of many driven-dissipative granular systems. Nevertheless, the distribution is approximately Gaussian and the inset plot shows that the position and velocity are uncoupled and populated at random. Each of these distributions corresponds to two degrees of freedom so that there are four possible effective temperatures. To determine whether these degrees of freedom are populated in accordance with statistical mechanics we plot the distribution of the total energy of the thermometer, which is the potential energy plus the kinetic energy each instant fig. 5.1c. For a thermal particle at equilibrium the energy must be sampled from a Boltzmann distribution. We plot a Boltzmann distribution on top of the data with an average energy equal to the

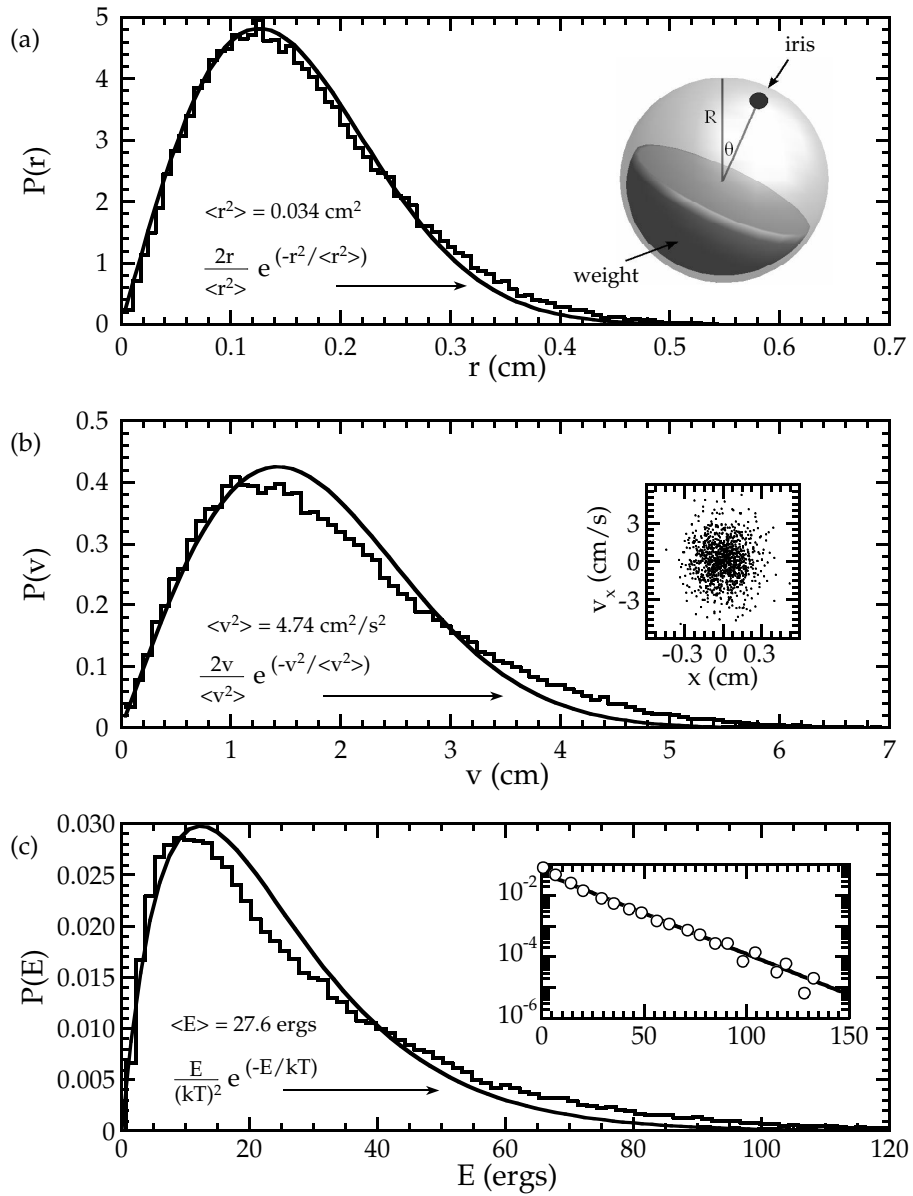


Figure 5.1: [PrvE] (a) Distribution of displacements of the thermometer's center-of-mass mechanical equilibrium. Inset shows a diagram of the thermometer with angular displacement θ . (b) Distribution of thermometer velocities. Inset shows a correlation plot between position and velocity. (c) Distribution of energies. Inset shows the probability distributions of the gravitational potential energy and kinetic energy binned simultaneously and fit to an exponential.




Figure 5.2: [Spectra] (a) Power spectra of the x- and y-position displacement from mechanical equilibrium of the thermometer. (b) Power spectra of the x- and y- velocity of the thermometer. The spectra show peaks at the natural frequency of the oscillator.

average energy of the thermometer. Plotting the potential energy and kinetic energy distributions separately in the inset, they are statistically identical, and approximately exponential, suggesting that energy is equipartitioned and that the effective temperatures from each of these degrees of freedom are equal.

We do not require that the thermometer perfectly obey statistical mechanics in order to obtain an effective temperature from it. Still, we are impressed that the position and velocity distributions are approximately Gaussian, the energy distributed according to a Boltzmann factor and partitioned equally among the oscillator's degrees of freedom. These features may serve to impart greater thermodynamic meaning to our granular thermometer effective temperature.

5.4 Einstein Temperature

The next effective temperature we consider is the Einstein temperature, defined in terms of the ratio of fluctuation to response. One can define any number of Einstein temperatures related to different fluctuation and response functions. We consider the simple case of diffusion to mobility along the lines of a classic Einstein relation. Specifically, we submerge a heavy test particle in the bath of light hollow plastic balls and subject the entire system to an oscillating gravitational force by tilting. We can control both the frequency and amplitude of the tilt and also the mass of the test particle, allowing us to explore a wide range of force amplitudes and frequencies. If the test particle is lighter than the bath particles it would feel a buoyant force opposing the direction of gravity. If it is heavier than the bath particles then gravity wins out. Either way, parallel to the applied force the test particle responds with a velocity, and perpendicular to the applied force performs a random walk. By measure perpendicular diffusion and parallel mobility we measure the Einstein temperature.

One condition that we must verify is that the perpendicular motion is indeed diffusive. In fig. 5.3a we plot the mean-squared displacement of the test particle perpendicular to the force for a low-density liquid-like system, with density $\phi = 0.569$ and driving airspeed $U = 375$ cm/s. At short times, the mean-squared displacement is ballistic and at long times it is diffusive. At very long times it saturates because the test particle has diffused the confining walls. Below we plot the velocity autocorrelation function which decays to zero after one second, which is the mean




Figure 5.3: [MSDVV] (a) The mean-squared displacement and (b) velocity autocorrelation function for motion perpendicular to the oscillating force. The diffusion coefficient is measured from the MSD after the mean collision time (when the VAF is first zero) and before saturation at the system size. It is nearly constant over the range 1 to 100 s.




Figure 5.4: $[Fv]$ Scatter plot of the oscillating external force as a function of velocity response parallel to the force. The mobility is measured by computing the slope of the best linear fit to the scatter plot.

collision time. Taken together, the perpendicular motion of the test particle is indeed diffusive for three orders of magnitude in time, from about 1 to 100 s. In this region the diffusion coefficient is nearly constant. Therefore, we measure it as $D = \text{MAX}[\langle \Delta x(\tau)^2 \rangle / \tau]$ and verify that it is insensitive to the particular time chosen as long as it is within the diffusive regime.

According to the mean-squared displacement and velocity autocorrelation function the system behaves dynamically like a simple liquid. The next step in our analysis is to measure the viscosity of this strange simple liquid, which we accomplish by considering motion parallel to the applied force. If the response of the medium is indeed linear and the drag on the test particle is indeed viscous, then the velocity should be linear in the force on average. In fig. 5.4 we show a scatter plot of the instantaneous force acting on the particle as a function of its instantaneous velocity. The cloud of force data has a clear upward trend as a function of velocity, so that positive forces applied to the test particle result in positive velocities and negative forces to negative velocities. The wide scatter in the data is a consequence of applying a small force to the test particle – it can still jostle about due to collisions with balls but on average it moves in the direction of the applied force. We




Figure 5.5: [Mobs] Mobility of test particles of different mass as a function of oscillating force frequency. Except for the very heavy steel ball, the data all collapse and are constant for frequencies less than 10 Hz.

fit to a line and extract the slope which is the inverse mobility.

For this particular density and driving rate, and for this particular test particle, we are able to simultaneously measure the diffusion coefficient and the mobility. However, in addition to linearity in the force amplitude, we must verify linearity in the oscillating force frequency, as well as independence to the test particle parameters. To verify linear frequency response we perform a frequency sweep holding the system density and driving airspeed constant. To verify the mobility is independent of test particle parameters and force amplitudes, we repeat the measurement with different test particles composed of materials with different mass and friction coefficients. The legend in fig. 5.5 shows the test particle parameters and associated tilt amplitudes. The test particle materials used in descending molecular weight are Steel (St), Teflon (Tf), Acrylic (Ac) and Wood (Wd). All are solid spheres with different surface roughness and friction coefficient appropriate to the material type. Consulting the data we see the mobility is indeed independent of the applied force amplitude and frequency. We also see that it is independent of the friction coefficient and weight of the test particle as long as it isn't too heavy. The mobility measured with the steel

ball clearly deviates from the others which we believe is due to the fact that it feels much greater gravitational and frictional forces and plows through the system without regard to the bath particles. In any case, as long as we choose one of the other test particles and measure the mobility for frequencies less than 10 Hz, the response will be linear in force frequency and amplitude for this simple but strange liquid.

5.5 Effective Temperature Comparison

To this point we have introduced two distinct effective temperatures. We have introduced the thermometer temperature, which is defined in terms of the average energy stored in a weighted sphere, and we have introduced the Einstein temperature, defined in terms of the ratio of diffusion to mobility of a test particle dragged through the system. The final effective temperature is the granular temperature of the bath particles, which we define as their area-weighted kinetic energy. Armed with these three effective temperatures change the systems density and driving to see how they respond on the approach to jamming.

There is more than one way to approach jamming for our driven-dissipative granular system. We can either increase density or decrease driving rate. Fig. 5.6a shows the result of holding the driving rate constant while increasing the density. Fig. 5.6b shows the result of holding density constant while decreasing the driving. In both cases we have measured each of the three effective temperatures. For the thermometer temperature we have measured it for two different thermometers with different moments of inertia and mass, to ensure that the results are independent of thermometer composition. For the Einstein temperature we choose a wooden ball at a frequency of 1 Hz and amplitude of 0.009 deg.

Holding airspeed constant and increasing density we see that the system follows a diagonal trajectory on this temperature-density phase space as it approaches random close packing. On this trajectory all three effective temperatures agree until approximately $\phi = 0.55$. Beyond $\phi = 0.55$ the granular temperature and the thermometer temperature agree but the Einstein temperature begins to systematically deviate from the others. By $\phi = 0.70$ it is an order of magnitude smaller than the other temperatures. Holding density constant and decreasing the driving airspeed fig. 5.6b, we see a similar trend since this is another route to jamming.

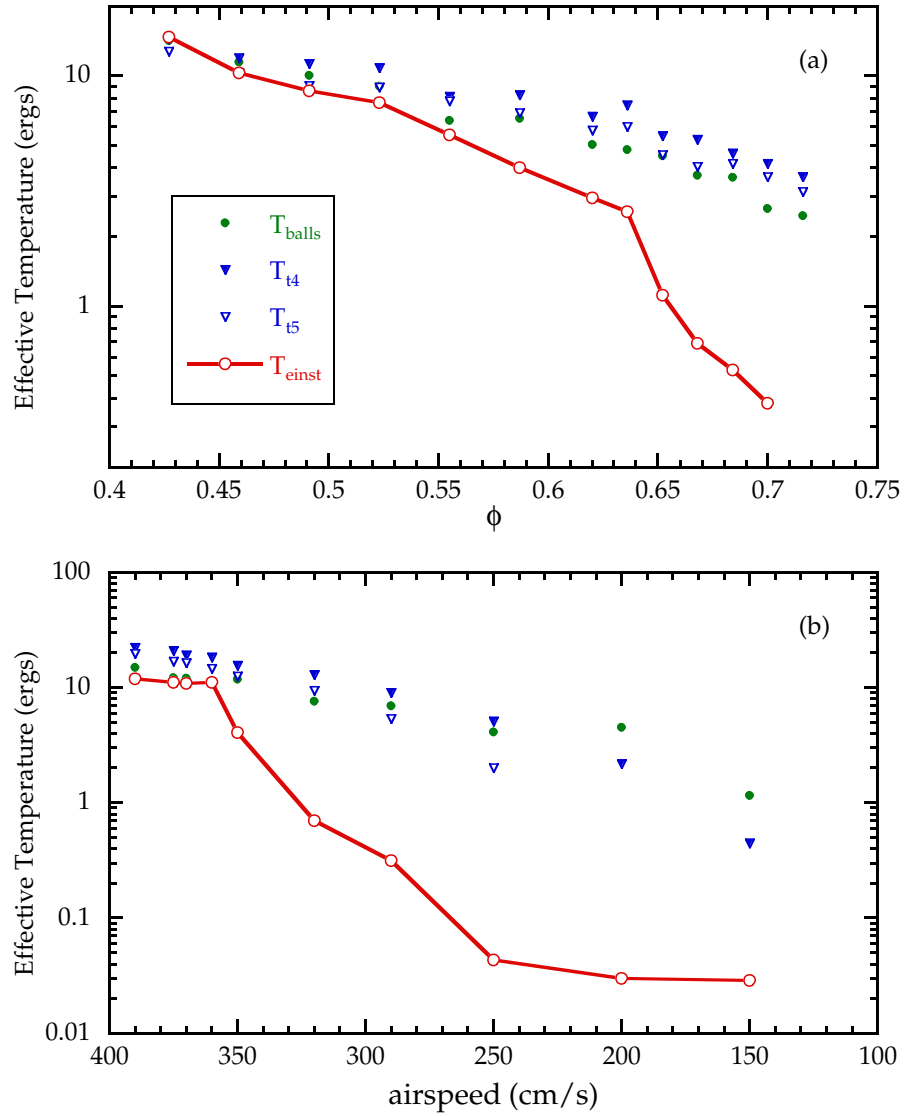


Figure 5.6: [TbTtTe] (a) Comparison of the three effective temperatures as a function of area fraction and (b) as a function of the driving airspeed. The 5th area fraction in (a) is the same driving-density phase space point as the 5th airspeed in (b) and roughly marks when the Einstein temperature deviates from the other two temperatures.

At high driving, the system behaves like a simple liquid and there is good agreement between all effective temperatures. At low driving or high density the granular temperature and thermometer temperatures agree but the Einstein temperature drops precipitously. Taken together, as jamming is approached by increasing density or decreasing driving, we see similar deviations of the effective temperatures. The Einstein temperature which is related to the relaxation of the system through rearrangements and the sampling of phase space, is far smaller than the other two effective temperatures. At short times the balls can move in their local neighborhood but on longer times the system cannot rearrange globally. Thus, there is a disconnect between the short time ballistics of the particles and longtime relaxation of the system, and a straightforward application of the Einstein relation predicts a temperature that is much “cooler” than the granular temperature or thermometer temperature.

Chapter 6

Probing Avalanche Dynamics using Speckle-Visibility Spectroscopy

6.1 Abstract

We apply a new light scattering technique called Speckle-Visibility Spectroscopy to the study of avalanches. By directly relating the rate of change of the scattered speckle pattern to the fluctuation dynamics of the flowing sand particles, we attain a precision of 0.1 mm/s. Running for 35 hours at 58 kHz, we simultaneously observe the microscopic short-time fluctuations of the sand particles and the long time behavior of thousands of avalanche events, and thus report avalanche frequency statistics and average shape. Interestingly, while all avalanches turn on in 0.3 s and in a similar way, there is a wide variation in how avalanches turn off. The fluctuation speed reaches a maximum just after the avalanche begins, it remains constant for a while, and then decays to zero. Power spectra of the full data set show that as avalanches slow the dynamics are self-similar ($\propto 1/f^2$) and the normalized variance of different events diverge at the turning off time.




Figure 6.1: [setup] A picture and diagram of the experimental apparatus as seen from the side.

6.2 Methods

6.2.1 Experimental Apparatus

The sand of our avalanches is $300 \mu\text{m}$ glass spheres with polydispersity of about 10%. The spheres are clear shiny balls that can be easily imaged with far-field lighting. The apparatus is built of two parallel non-static plexiglass walls with a square aluminum bar sandwiched at the bottom and one at the top, Fig. 6.1. The non-static coating is necessary to prevent sand from sticking to the walls. The aluminum bar at the bottom is the “floor” of the apparatus, on top of which the sand pile sits. The aluminum bar above obstructs the sand flowing onto the pile, slowing it down so that it does not impart much momentum initiating an avalanche. Sand is stored in a large metal funnel above the apparatus and slowly trickles through a sand valve with a knife edge, which allows us to control the flow rate to 0.01 g/s . Sand flows between the plates and a pile builds until the maximum angle of stability is reached. As more sand is added an avalanche flows until excess sand is exhausted, the angle of repose is reached, and the sand comes to rest, allowing the process to repeat.

6.2.2 Digital Video Imaging

The big hurdle to directly observing granular avalanches is that even if the grains are perfectly clear, as they are in this experiment, entering photons are multiple-scattered so the pile appears opaque. One way around this is to submerge the grains in an index matched fluid, but this severely alters the system. Another method is to track particles at the surface. We image the surface of

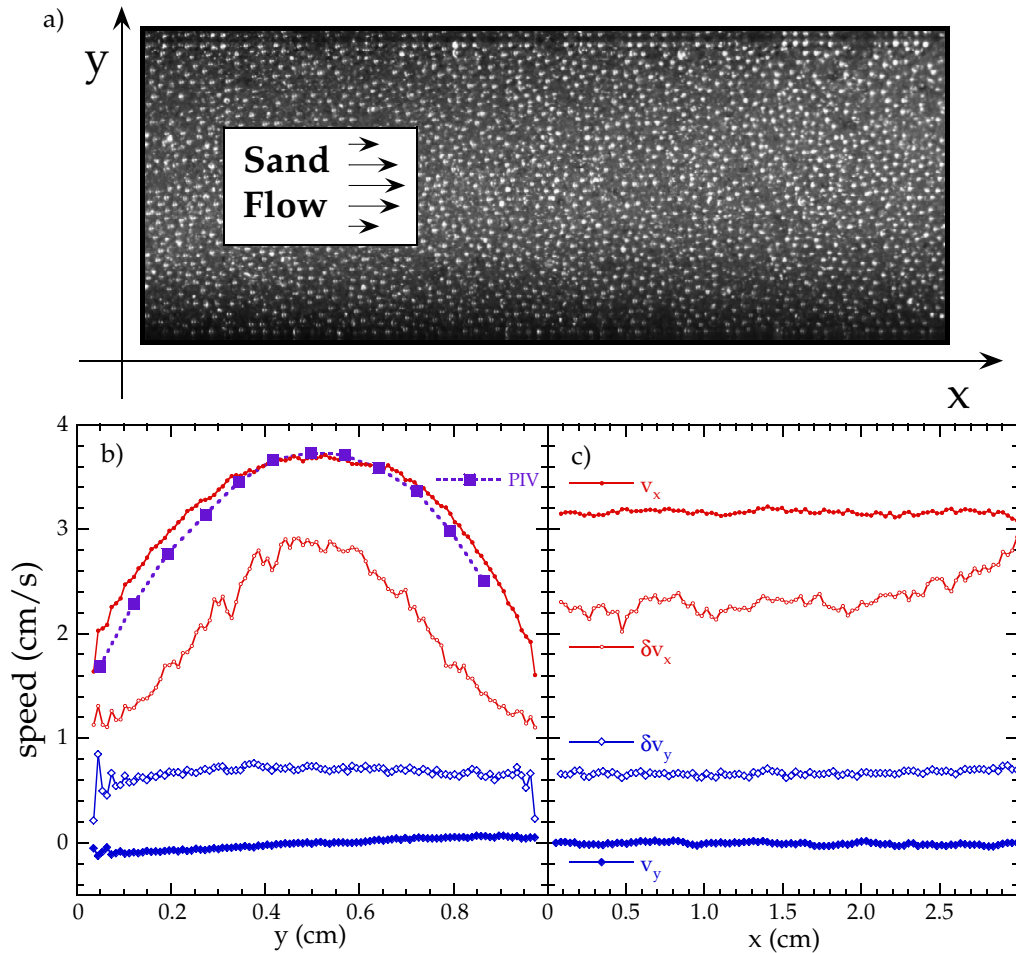


Figure 6.2: [BigTop] (a) An image from above of the top layer of sand taken with a high speed camera at 1000 fps. To obtain velocity profiles, the particles can either be tracked or the images can be analyzed using PIV. (b) The flow profile across the cell is nearly parabolic parallel to the flow and constant perpendicular to the flow. The fluctuations of the surface layer obtained by particle tracking are consistent with SVS measurements. (c) The flow profile down the cell is constant both parallel and perpendicular to the flow. The fluctuations are also nearly constant. For $x > 2.0$ cm the sand is close to the end of the channel and begins to free fall.




Figure 6.3: [SmallSand] (a) The flow speed in the x -direction as a function of y , which is the channel width. The profile is more plug-like than for the bigger spheres. Inset is an image taken from above of flowing glass spheres $100 \mu\text{m}$ in diameter. (b) The flow speed in the x -direction as a function of z , depth. The flow profile decreases non-linearly as a function of z and drops to zero abruptly after about 20 layers. Inset shows an image of the flowing spheres taken from the side.

continuously flowing sand from multiple angles with a high-speed CCD Phantom camera, storing 2 s movies at 10000 fps. We use standard particle tracking routines available in both IDL and LabVIEW to locate particles within the 1 cm x 3 cm field of view. For a continuous flow rate of 2.5 g/s we typically track 10000 individual particles for 0.1 s each.

An example image from one movie is shown in fig. 6.2a. Gravity points into the image at an angle of about 30 degrees, and has a projection in the x - y plane parallel to the x -axis. Thus, the net flow is along the x -axis. By accumulating statistics from many such images we can measure the flow profile of the sand along different directions as a function of both x and y . In fig. 6.2b we show the flow profiles as a function of y . $v_x(y)$ is greater in the center of the channel than at the edges. We verify these particle tracking results using a fourier implemented particle image velocimetry (FPIV) that we have developed. In our method a grid is overlaid on the image. In each grid box is a section of the image. Each section is transformed using a two-dimensional fast-fourier transform routine. Each image transform is then multiplied with its transform at later time and back-transformed. The result for each image section is a two-dimensional peaked function where the displacement of the peak from the origin is the distance that the the two original images must be shifted in order to obtain maximum overlap – this is the two-dimensional equivalent of computing the cross-correlation between two nearly identical functions displaced by a constant. This method is fast and works well provided that (1) the images have not changed much over the time interval (i.e. the functions are essentially identical aside from a uniform displacement) and (2) the displacement is small compared to the size of the image in the grid box so that edge effects do not affect the location of the cross-correlation peak.

In general, our FPIV method is far easier to implement and more reliable than particle tracking because specific image features do no need to be identified. We have measured the flow profile in the $v_x(y)$ direction using both particle tracking and FPIV. Both show that the profile is parabolic across the channel. Evidently, grains near the wall feel a friction force that slows them down. In addition to the average flow velocity, we also measure the fluctuations in the x -direction as a function of y , $\delta v_x(y)$. Previous work showed that the fluctuations scale with the derivative of the average flow [48]. We see in fig 6.2b that for $v_x(y) \propto y^2$, we have $\delta v_x(y) \propto y$. Such velocity fluctuations in a granular systems are often referred to as a the “granular temperature” which in our case is not constant across the channel. Another profile of interest is $v_y(y)$. On average it is

zero. However, fig. 6.2b, shows that for $y < 0.5$, $v_y < 0$ and for $y > 0.5$, $v_y > 0$. This small upward trend suggests that sand particles flowing down the channel tend to migrate towards the walls, which is visible to the naked eye in the movies and implies that the flow is probably not constrained entirely to two-dimensions. Nevertheless, the velocity fluctuations δv_y are constant across the channel and are also less than δv_x .

Fig. 6.2c shows the flow profile as a function of y , which is distance down the channel. Not surprisingly, the average velocity of the flowing sand $v_x(y)$ is essentially constant down the channel for continuous flows, with a slight increase in $\delta v_x(y)$ as the channel edge is reached. At $y = 4$ cm the channel ends and the sand falls over the edge into a bin. This “sand fall” looks similar to a water fall and according to the increase in $\delta v_x(y)$ may influence the flow dynamics 2 cm upstream. The perpendicular component of the flow is constant and zero even though the fluctuations are non-zero. Comparison with all fluctuations show that the granular temperature is spatially dependent even when the sand is flowing in a steady-state. In addition, the granular temperature is different along different directions – it is larger in the direction of larger flow.

To investigate whether the flow profile depends on the size of the sand particles with respect to the width of the channel, we also measure it for smaller sand particles, 100 μm glass spheres, in the same geometry with the same channel width. Like with the big grains, we supply sand at a flow rate large enough so that the sand flows down the channel continuously without avalanching. In fig. 6.3 we show just the primary profiles of interest, which are obtained using the same camera and particle tracking routines we use for the big sand particles. Fig. 6.3a shows the flow profile as a function of channel width $v_x(y)$. The first thing to notice is that the flow is faster with the smaller grains. Indeed, all the dynamics of the smaller grains are faster – they flow faster, they fluctuate faster, and in the intermittent regime, avalanches are shorter and more frequent. Interestingly, the flow profile from above is more plug-like than for the larger grains. Indeed, there is a shear-band near either edge that is about 20 grains thick. The grains in the center are moving at a more constant velocity and nearly twice as fast as those near the edges, which is the same that we see from above for the larger grains. But the top profile only relates one projection of the total flow. Fig. 6.3b shows the flow profile as a function of depth $v_x(z)$, which is the projection perpendicular to the top projection, as indicated by the inset images. The flow speed in the x -direction as seen from the side decreases as a function of depth. In fact, the speed scales exponentially with depth, as is

shown in the fit in the inset. Thus, the side profile shows an instantaneous snapshot of the flow at different stages of jamming as a function of depth. In the topmost layer the sand flows rapidly without hindrance. Down below it gradually decreases until reaching a jamming depth at about $z = -0.22$ cm. To verify whether the flow profile within is similar to what we see in the surface layers, we can estimate the total sand flow by assuming that profiles we observe on the surface persist in the bulk. Computing this extrapolation and comparing with time averaged measurements of the flow rate using a scale and stopwatch, we find agreement to our profile estimate so that flow in bulk is probably not too different from what is visible at the edges.

By using particle tracking and FPIV on the surface layers of continuously flowing sand we are able to characterize three-dimensional aspects of continuous granular flow down the channel. Strictly speaking, the flow profile we observe is only applicable to continuously flowing sand, and additional video experiments will be necessary to characterize the intermittent flow of avalanches. Nevertheless, we do uncover that the velocity profile is not constant across the channel or as a function of depth. We also learn that the shape of the profile depends on the size of the sand particles with respect to the channel width, and is more plug-like for smaller grains. We speculate that that with even smaller grains or a wider channel, one would observe an even more plug-like profile, with a constant region in the center and shear boundaries at either edge. In our system the existence of a non-constant profile across the channel could be symptomatic of boundary effects that may influence the long-time behavior and periodicity of avalanches.

6.2.3 Speckle-Visibility Spectroscopy

Particle tracking and FPIV on the surface layer provide a partial visual depiction of granular flow down the channel, and estimations based on the surface profiles show that the flow in the bulk is similar to that on the surface. However, video imaging has major drawbacks. The frame rate necessary to track particles is so high that we are limited to observation durations of only a few seconds at a time, which is insufficient for observing an avalanche. Moreover, the smallest observable displacements are set by the size of the particles themselves and most of the small sand fluctuation will be missed. Even then, we can only see the surface layers. A complimentary observation method that overcomes many of the difficulties and limitations of digital imaging is a recently developed light scattering technique called Speckle-Visibility Spectroscopy (SVS).




Figure 6.4: [TwoAvalanches] (a) A blow-up of the speckle pattern as a function of time for an avalanche turning on, indicated as the gray region in the figures below. When the sand is static the speckle pattern is constant so that there are constant streaks as a function of time. As an event turns on the speckle pattern begins to change. The variance ratio (b) is a measurement at each instant of the ratio of the visibility of the speckle pattern for the exposure times of T and $2T$. When the speckle pattern is constant (streaks) its visibility is independent of exposure duration, and thus the variance ratio is one. When the sand flows, the speckle pattern changes quickly and a short exposure is necessary for a sharp speckle pattern so that the variance ratio is less than one. Using theory for back-scattered light the variance ratio can be transformed numerically into the fluctuation speed of the sand particles (b) which is the granular temperature.

SVS works by measuring the visibility of a speckle pattern on a CCD camera as a function of exposure duration. We shine a solid-state, diode-pumped, frequency-doubled Nd:Vanadate laser at 532 nm at a fixed output of 4 W. The laser is passed through a lens that expands the beam to 1.3 cm in diameter. The diameter of the beam spot on the pile is defined as twice the radial distance from the center to the point that the laser intensity drops to $1/e^2$, its peak intensity. The channel width is 1 cm in diameter. Thus, the beam spot fills the channel. The beam penetrates to a depth of about one centimeter so that we sample a cubic-centimeter volume. Photons from the coherent source enter the pile and are multiple-scattered. Upon exiting the spectrum is broadened and creates an interference (speckle) pattern on the CCD, which we observe at 58 kHz using a 2×1024 8-bit Basler linescan camera. The particular interference pattern imaged at the CCD depends on the location of the scattering sand particles with respect to the laser and camera. Thus, if the sand particles move, the speckle pattern changes. Using theory for backscattered light [71] we can directly relate the rate of change of the speckle pattern to the motion of the sand particles. Because the speckle pattern is an interference pattern, this method is sensitive to motion on the scale of the wavelength of the scattered light; we do not measure the average flow of the sand but rather the velocity fluctuations and thus the granular temperature.

We supply sand to the pile at a rate 0.1 g/s, which is too slow to drive a continuous flow, and so we are in the intermittent (avalanching) flow regime. In fig. 6.4a we show a pixel vs. time image of the speckle pattern as an avalanche turns on. When the pile is static the speckle pattern is constant and the pixels form streaks as a function of time. When an avalanche begins the sand particles fluctuate and the speckle pattern changes. As in Ref. [71] we define the visibility of the speckle pattern as the normalized variance,

$$V_2(T) = \frac{1}{\beta} [\langle I^2 \rangle_T / \langle I \rangle^2 - 1] \quad (6.1)$$

$$= \int_0^T 2(1 - t/T) [g_1(t)]^2 dt / T. \quad (6.2)$$

where T is the exposure duration, $\langle I^2 \rangle_T$ is the intensity moment across the CCD for the exposure duration T , and $g_1(t)$ is the electric field autocorrelation function. The visibility of the speckle pattern on the CCD depends both on the exposure duration and the rate of change of the speckle pattern. If the CCD is exposed for a short duration compared to the rate of decay of $g_1(t)$, one observes a sharp speckle pattern. If it is exposed for a long duration compared to the rate of decay

of $g_1(t)$, one observes an averaged out speckle pattern. We save only the first and second intensity moments of the two exposure durations T and $2T$, which allows us to stream data continuously at a frame rate of 58 kHz, where $T = 17.24 \mu\text{s}$, while still saving enough information to measure the variance ratio of the two exposure durations. We then transform the variance ratio into the sand fluctuation speed via the formula,

$$\frac{V_2(2T)}{V_2(T)} = \frac{e^{-4x} - 1 + 4x}{4(e^{-2x} - 1 + 2x)}. \quad (6.3)$$

To measure δv eqn. 6.3 must be inverted for x , which we do numerically. For back-scattered light $x \approx 4\pi T \delta v / \lambda$, with $\lambda = 532 \text{ nm}$ for our laser, which gives us the sand fluctuation speed.

Fig. 6.2b shows the variance ratio time trace of the visibility of the speckle pattern for the exposure durations T and $2T$. In this 80 s snippet of data we observe two avalanches. Before the first avalanche the sand is static, the speckle pattern is constant, and the variance ratio for the two exposure durations is one. After the avalanche turns on the ratio drops below one to a value that depends on the fluctuation speed of the sand. Fig. 6.2c shows $\delta v(t)$ for the corresponding variance ratio time trace above. When the visibility of the speckle pattern is independent of exposure duration the sand particles are static. Thus, when the variance ratio is one, δv is zero. As the grains begin to fluctuate the $V(2T)/V(T) < 1$ and $\delta v > 0$. The avalanche turn-on time is set by the amount of time it takes the head of the avalanche to traverse the beam spot. For a beam spot size of 1.3 cm and a typical flow speed of 3.1 cm/s (fig. 6.2), we estimate the turn on time to be about 0.4 s.

Using this method we are not only able to observe the turn on dynamics but all the dynamics of thousands of avalanches. As discussed, our method is sensitive to the velocity fluctuations of the sand particles and we therefore measure the instantaneous granular temperature. Because we stream highly compact data continuously we can observe the system for tens of hours depending on the patience of the experimenter and the size of the hard drive. Typically, 1 hour of data consumes about 4 GB of space. Thus, given a data capture rate of 58 kHz and a modest observation time of 100 hrs, we attain ten orders of magnitude in dynamical range.




Figure 6.5: [Statistics] The top row shows correlation plots between time and velocity observables. (a) Correlation between the time-on in time-off durations. (b) Correlation between the maximum fluctuation speed in the time-off duration (c) Correlation between the time-on duration and the maximum fluctuation speed. The bottom row shows histograms for time and velocity observables. (e) Histogram of the duration of avalanche events. (f) Histogram for the duration between avalanche events. (g) Histogram of the maximum fluctuation speed obtained by the avalanche.

6.3 Avalanche Statistics

The enormous advantage of SVS for this system is that we can simultaneously observe the microscopic fluctuations of the sand particles and also the long time behavior of thousands of avalanches. This allows us to directly see how the microscopic motion of the sand relates to avalanche frequency and properties, such as the relationship between the composition of the pile and the types of avalanches that occur.

To begin, we compile statistical data on three observables, shown in fig. 6.5. Figs. 6.5(d-f) are histograms and (a-c) are correlation plots between the variables. Fig. 6.5a is a histogram of the duration of each avalanche, which is the time when motion is first observed to the time the pile comes to rest. Previous studies [48] showed that the average time of an avalanche is independent of the rate at which sand is added to the pile provided that it is well below the transition from intermittent (avalanching) to continuous. In such a “quasi-static” regime, the flow rate only effects the time between events. Therefore, we choose a flow rate of $Q = 0.07$ g/s which is well below the transition point at $Q_c \approx 0.4$ g/s, but still high enough that we observe avalanches often. For this choice avalanches occur periodically, with duration distributed asymmetrically about an average lifetime of 16 s, and having a particularly high tail due to uncommonly long-lived avalanches. Indeed, while avalanches tend to occur roughly periodically, they are not perfectly periodic, and there is a preference for longer lived avalanches which exhaust much more of the sand supply than usual. We speculate that the source of this wide variability among avalanche lifetimes is the wide range of stable packing densities and angles of the granular pile. Small differences in granular configuration can be the difference between whether an avalanche forms or not and will ultimately govern how much sand is available and how long an avalanche will live.

The amount of time that the pile is static depends on how much sand must be replenished to initiate another avalanche. The amount of sand needed depends on the amount of sand exhausted by the previous avalanche. Because there is a wide variation in avalanche duration there is a wide variation in the amount of sand exhausted, and therefore also in the time between events. This is shown in fig. 6.5e. The average time between events is about 50 s. Yet, the distribution is neither narrow nor symmetric. Like the distribution of times on, the distribution of times off has a high tail for particularly long times between events. In some instances the time between events is hundreds of seconds, many times the average, so that the pile was either uncommonly low to begin with, or

built to an uncommonly high but stable angle. Another peculiar feature is that the distribution has a small second hump at smaller times. This hump may be the result of there being two characteristic stable angles just before avalanches flow, a steep angle initiating a big avalanche that exhausts much sand, and a smaller angle that initiates small avalanche that exhausts less sand. If this is the case, it is interesting that the duration of events is approximately gaussian with a single peak.

The time of events and the time between events relate to the long time behavior of the avalanche cycle. To connect this behavior to the microscopic dynamics of the sand particles we compile a histogram of the maximum fluctuation speed attained by the sand particles in each avalanche, fig. 6.5f. This speed always occurs just after the avalanche turns on. Unlike the distributions of times on and off, this distribution is symmetric about the average value of $\delta v = 5.5$ mm/s. However, like the others, it has fat tails and is peaky (“leptokurtic”), so that it is non-gaussian. These tails may be yet another manifestation of the wide variability of avalanche size since it is likely that big avalanches and small avalanches start with different fluctuation speeds. To investigate this possibility further and directly relate the microscopic behavior of the sand particles to the macroscopic behavior of the avalanches they make up, we construct correlation plots between the different observables, fig. 6.5(a-c). Fig. 6.5a shows the correlation between time on and off. The double peak in the $H(t_{off})$ results in two blobs in the correlation plot. However, the two blobs do not have the same average suggesting that there is a correlation between the duration of events and the time between events. In particular, avalanches that are longer lived are preceded and followed by longer times off. Avalanches that are short lived occur with less time between them. Given that the flow rate is approximately constant, both observations support that the size of the avalanche is determined by the amount of sand that is available to be exhausted, and that this amount varies greatly. Yet more evidence is available in the correlation between the maximum fluctuation speed of an avalanche and time off, fig. 6.5b. While there is no correlation among the averages (i.e. δv is constant), the variability of δv does seem to increase the shorter the time between events: as events get closer together, the dynamics become less predictable. Fig. 6.5c shows a plot of the average time on versus the maximum fluctuation speed. Because each of these distributions had a single peak, we see only a single blob. However, there is a strong correlation between these two observables – longer lived avalanches tend to start faster. This again supports the notion that there is a variation in avalanche size, because if there were not, then one would expect that fast avalanches would end quicker than




Figure 6.6: [AverageAva] Mean and standard deviation of $\delta v(t)$ timetraces for 1500 avalanche events. Inset is mean and standard deviation with $\delta v(t)$ normalized by δv_{max} and t normalized by t_{on} .

slow avalanches, which is in direct opposition to what we see here. A more likely possibility is that there are both small and big avalanches and that the big avalanches start off with an uncommonly large amount of sand, and therefore the pile is at a steep angle just before the avalanche starts. The steep angle gives rise to both a fast starting and long-lived avalanche.

6.4 Anatomy of an Avalanche

Amazingly many of the most basic questions about avalanches remain unanswered. For example how similar are different avalanches, and how long do they last? Do all avalanches start and end the same way, and if so, is there a universal avalanche shape? These questions and others can be partially answered by observing the average behavior of many avalanches. In fig. 6.6 we show the average “shape” of an avalanche in δv vs. t space. To generate this figure we have observed thousands of events and aligned them according to their start time, which is defined as the moment

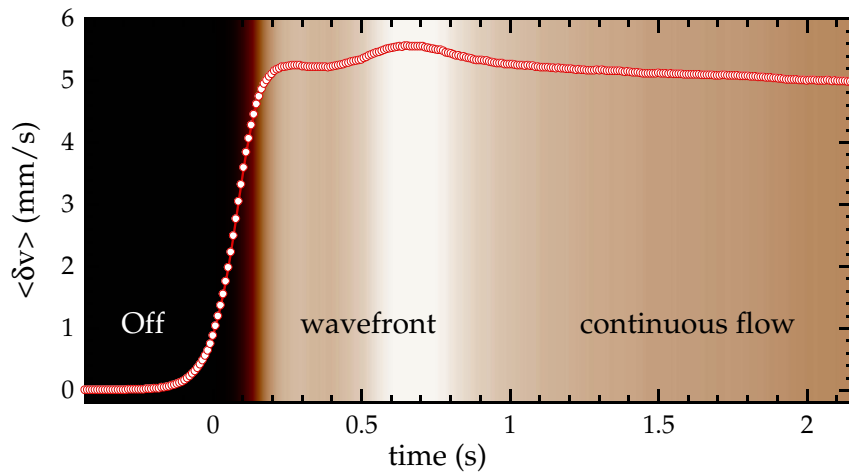


Figure 6.7: [AverageOn] A blow up of turning on region of the average $\delta v(t)$ with its contour in the background. A wavefront is at the head of each avalanche. After the wavefront passes the sand flows smoothly and gradually slows down until finally turning off.

at which $\delta v = 0.01$ mm/s. Then we compute the mean and standard deviation over all events, as plotted.

According to the average avalanche shape fig 6.6, we see that avalanches tend to turn on quickly and shortly thereafter attain their peak fluctuation speed. From that point forward they gradually slow until finally coming to rest. The standard deviation of different events shows that there's a wide variation in how avalanches turn off. As the same supply is exhausted and the avalanche approaches its end the pile stutters and eventually stops. The precise moment that an avalanche ends depends on the precise details of the pile and the particular dynamics of the flow. As a result, the turning off time varies widely among different events, and simply averaging them together results in a distorted average shape. For example, the high tail at long times is an artifact of averaging short and long lived avalanches together. To remove this artifact and uncover the true avalanche shape we normalize each event by its characteristic observables, δv_{max} and t_{on} . Plotted this way the different stages of an avalanche's lifetime become apparent. In the first few moments after the avalanche turns on it achieves its maximum fluctuation speed. From that point forward it slows with nearly constant deceleration as it exhausts its sand supply. At an age of about 0.8 the fluctuation speed takes a sharp turn downwards in the avalanche ends.




Figure 6.8: [AverageOff] The standard deviation normalized by the mean for 1500 separate avalanche events. Inset shows the same quantity calculated for the normalized avalanche shape.

Zooming in on the first stages of an avalanche's life, we see that there is a structure in how an avalanche turns on, fig. 6.7. The average avalanche turns on in about 0.3 seconds, which is related to the flow speed and determined by how long it takes for the avalanche front to traverse the beam spot. Over this time the pile goes from being completely static to attaining its maximum fluctuation speed. One can think of the $\delta v(t)$ time trace as relating to the spatial composition of an avalanche. Early times correspond to the avalanche head and later times to regions farther up the pile. In previous studies [11] it has been discovered that the front of an avalanche has a complex structure that depends on the properties of the flowing grains. For example, one can see convective rolls either foreword or backward depending on how coarse and polydisperse the grains are. This rolling wavefront may be the reason for the sharpness of δv_{max} just after turning on, but without direct imaging of the front in this experiment it is impossible to say for certain. After an avalanche exhausts the majority of its sand supply it reaches a stage of cessation. At this stage the avalanche begins to start-and-stutter before stopping. As a result, there is much variation among different avalanches regarding how and when they turn off. To highlight this point, we plot the normalized variance of different avalanches in fig. 6.8, which diverges at the turning off time. To remove




Figure 6.9: [PSFull] The power spectrum of $\delta v(t)$ for the full 35 hour data set. At very high frequencies the power spectrum saturates to a constant demarcating our measurement accuracy. Frequencies from 0.1 to 100 Hz are timescales of the sand fluctuation speed during avalanche flows and show $1/f^2$ noise. There is a peak at the average avalanche frequency and $1/f$ noise for lower frequencies having to do with spectrum many avalanche cycles.

discrepancies due to differences in avalanche duration and fluctuation speed, we also compute the normalized variance of the normalized shapes as before, shown in the inset. Whichever way we plot it, the normalized variance diverges at the turning off time.

6.5 Avalanche Dynamical Scales

To conclude our study of avalanches we plot the power spectrum of the full 35 hr $\delta v(t)$ time trace, fig. 6.9. The enormous dynamical range afforded us with speckle-visibility spectroscopy is necessary since we see spectral power over nearly all frequencies observed. At the very highest frequencies, we resolve the noise floor. At lower frequencies, between 0.01 and 100 Hz, there is a large contribution to the spectral power from velocity fluctuations of grains in the flowing state. The spectrum peaks between 0.01 and 0.1 Hz, corresponding to the timescale of avalanche occurrence. On very long timescales and very low frequencies the power spectrum approaches $1/f$

implying that there is structure to how often avalanches occur.

In this study of avalanches we've characterized both the statistics of avalanche occurrence, how long they are on and how long there off, and also the dynamics of avalanches in the flowing state, and we've done so using a single light-scattering technique. Speckled-visibility spectroscopy permits us to observe avalanches over ten orders of magnitude in time so that we are able to capture the complete dynamics of 1500 separate events. With this data, we uncover a characteristic avalanche dynamical structure.

Chapter 7

Conclusions

This work has focused on two different classes of experiment. Both have to do with collections of grains driven close to jamming. In the first class of experiment we have studied grains fluidized by air moving in two dimensions at the various densities and driving amplitudes. In the second class of experiment we have studied avalanche flows using light scattering.

The primary result of the first class of experiment was to uncover striking similarities of granular systems driven close to jamming and thermal systems undergoing a glass transition. Structurally, the pair correlation function and circularity factor distribution of Voronoi cells show hallmark features that increase in prominence with proximity to the jamming transition. Each of these functions developed characteristics that were a direct consequence of the spheres being packed tightly together, essentially in hard-core contact. As a result, local packing symmetries which are incommensurate with crystallization nonetheless affected the average shape of Voronoi cells and the likelihood of finding neighbors at noninteger distances. Dynamically, the tight packing resulted in a subdiffusive plateau at intermediate times in the mean-squared displacement. From a rheological standpoint the system is rigid over a certain range of intermediate times (or large forces) and deforms viscously at long times (small forces). The similarity goes beyond ensemble averages over single particle quantities such as the pair correlation function and the mean-squared displacement. Looking at the displacement fields in space, dynamical strings became evident that were reminiscent of similar dynamical structures seen in simulations of supercooled liquids and colloidal glasses. These strings of correlated motion represent a dynamical length scale that in the case of supercooled liquids diverged in a way consistent with Mode-Coupling Theory and the

Vogel-Tammann-Fulcher Equation. Similarly, the dynamical length scale of this granular system also diverged consistent with MCT and the VTF equation, and in fact a simple fit to the latter extrapolated to the jamming transition density independently measured at 83%.

At low-density the system resembled a simple liquid. To test the extent of this analogy and its relation to jamming we also measured three effective temperatures that in thermal equilibrium are all equal. First, we measured the average kinetic energy of the spheres. Second, we measured the average energy of a granular thermometer, which is a weighted sphere oscillator. Third, we measured the Einstein temperature, the ratio of diffusion to mobility. The kinetic energy of the balls and the granular thermometer are instantaneous and local quantities. For the energy of the thermometer to have any thermodynamic meaning its potential and kinetic energy must be set by the bath temperature and they must therefore be equal. We found this to be the case and were thus able to define an effective temperature from the total average energy of the thermometer. However, this effective temperature is only representative of the system if its time average is equal to the ensemble average of the kinetic energy of the grains. Thus, the system must be ergodic. Going one step further, the Einstein temperature relates the short time fluctuations of the spheres to the longtime relaxation of the system, and this ratio is precisely set by the bath temperature, which connects it to the other two effective temperatures. Each of these effective temperatures depends on a different tenet of statistical mechanics in order for them all to be equal. We have therefore measured them on approach to jamming to verify their consistency. However we cut phase space, either by holding driving constant and increasing density or by holding density constant and decreasing driving, the instantaneous localized energies of the spheres and the thermometer were significantly larger than the Einstein temperature, suggesting that spheres can vibrate in their local environment without moving far and rearranging. In this sense, there is a disconnect between the short time random vibration of the particles in the sampling of configuration space.

The second class of experiment we've studied has been granular avalanches flowing down a channel using Speckle-Visibility Spectroscopy and high-speed particle tracking. The primary result of this experiment was to record avalanche dynamics with unparalleled resolution over ten orders of magnitude in time, which allowed us to observe both the short time microscopic fluctuations of the sand particles and the long time behavior of thousands of avalanche events. This dynamical range may seem excessive but avalanches span an enormous range of time and length scales. The

shortest time fluctuations are components of avalanches and influence the long-time dynamics. SVS provides us with the spatial resolution and dynamical range that is therefore necessary to fully characterize this complex granular flow. But with this experiment we've only scratched the surface. We have not yet completely characterized how changes at the microscale couple to macroscale avalanche dynamics. One question is whether reducing the grain size is equivalent to increasing the channel size. There are indications that a universal scaling will connect these two system perturbations, but as of yet no experimental evidence has been obtained.

Appendix A

Crystallized Steel Ball Bearings

In previous studies of air-fluidized grains some systems were seen to resemble thermal systems. In the case of a single ping-pong ball, the system behaved exactly like a Brownian particle in two-dimensions trapped in a harmonic potential, and the motion could even be described by statistical mechanics [1, 64]. We have also studied dense bidisperse packings of steel ball bearings as a function of increasing area fraction to observe how the structure and dynamics change as we approach jamming from below. This is the subject matter of Chapter 2. In this appendix we discuss an auxiliary experiment performed with monodisperse packings of steel ball bearings in a hexagonal geometry. In fig. A.1 we show two pictures of the system for two different area fractions. The hexagonal boundary is commensurate with crystallization and minimizes boundary-induced defects. An insert is used both above and below the sieve to ensure that the air flows uniformly upward near the walls. We observe the system as usual with a 120 Hz camera and track all particles over the entire experiment.

The first measurement we present is the mean-square displacement shown in fig. A.1. At low densities the mean-square displacements are essentially identical to what was seen for bidisperse steel balls and resembles a simple liquid. At short times the mean-squared displacement is ballistic and at long times diffusive. At very long times it saturates because particles diffuse to the boundary and can go no further. At higher densities things are quite different. At short-times motion is still ballistic but a plateau develops at intermediate times. This plateau is reminiscent of the plateau seen previously in disordered bidisperse systems at high density, and was interpreted in terms of the “cage” effect, in which particles are trapped in a cage formed of their nearest neighbors. Unlike




Figure A.1: [CrystalMSD] The images show the system from above for low density and high-density snapshots. At low densities the system is liquid-like and at high densities it is crystalline. In the figure below we plot the mean square displacement for the four densities studied. At low densities the mean-squared displacement is diffusive at long times and high densities when the system crystallizes, the the mean-squared displacement saturates at long times. Because of turbulent fluctuations, all densities are ballistic at short times.




Figure A.2: [CrystalDOS] Density of phonon states computed from the Fourier transform of the velocity autocorrelation function, averaged over all particles. At low densities the system is liquid-like and the density of states approaches a finite value at zero frequency, so that there are translational modes. At high density when the system is crystalline, the density of states goes like $1/f$ as $f \rightarrow 0$, which is consistent with that of an atomic crystal in two dimensions.

bidisperse packings, this monodisperse system can crystallize at high density, as evident in the image, and in such instances the mean-square displacement becomes even more flat at intermediate times with a slope that is nearly zero over three decades. Thus, the particles are localized for long periods and vibrate about equilibrium positions. The crystal is not perfect and at the very longest times there is a slight upward subdiffusive trend representative of a few beads moving out of their unit cells.

The mean-squared displacement shows that as density goes up and the system crystallizes particle motion becomes localized over a wide range of timescales. Viewed from a different perspective, this crystalline system can be thought of as a material in which the correlated motion of the beads are phonons excited by turbulent fluctuations from the air. These phonons can be quantified by measuring the crystal's density of vibrational states fig. A.2, which is calculated from the velocity time traces [1]. For low densities, the density of states is liquid-like and increases with




Figure A.3: [High] Three representations of the system for $\phi = 0.838$. The first shows a raw image of the system taken from above, showing that the system is clearly crystallized. The second image shows the Voronoi diagram for the configuration with each cell colored according to its number of sides. Six sided cells dominate the diagram, with a few 5-7 dislocations strewn throughout. The third image shows a snapshot of the average velocity vectors for a time interval 10 seconds. The velocity vectors tend to align along the crystal axis and are reminiscent of phonons.

decreasing f so that it approaches the y-axis at a finite value. Thus, there are low-frequency translational modes and the system does not remain in the same configuration. The downward turn at low frequencies for the three lowest area fractions is a finite size effect since particles diffuse to the wall. The frequency at which the density of states turns down is equal to the reciprocal of the time that the mean-square displacement saturates at the wall. For high area fractions the density of states is very different. Beyond a certain area fraction the density of states develops a peak that both broadens and shifts to high frequencies for high area fraction. Evidently, the system's energy is distributed among higher frequency modes when the balls are packed together in a crystalline array. This is true even for high density disordered systems with the difference that the peak is broader in the ordered case. At low frequencies for this monodisperse steel-bead crystal the density of states approaches $1/f$, which is exactly what one would expect for crystals composed of atoms in two-dimensions driven by thermal fluctuations.

The density of states is a way of quantifying how energy is distributed in the system. It is also useful to look at the makeup of these phonons in space. In fig. A.3 we show a few images for $\phi = 0.838$. The first image shows a raw picture of the system from above, and the crystallization

is evident. The next image shows of Voronoi diagram for the configuration, with each cell colored according to its number of sides. Six sided cells dominate the diagram with a few 5-7 dislocations strewn throughout. The final image shows average velocity vectors for an averaging time of 10 seconds, which is within the plateau region of the mean square displacement and corresponds to a frequency when the density of states has already approached $1/f$. Even though motion is driven by random turbulent fluctuations, the beads move in a correlated way so that such phonons tend to align along the crystal axis.

Appendix B

Clustering in the Bed

This work was done in collaboration with Florian Ebert.

For some time we have sought to more thoroughly characterize the interaction potential between dense bidisperse packings of beads in the bed. We had no reason to expect that the potential measured previously for two identical beads would apply in the case of many beads since the interaction is supplied by the air, which is greatly perturbed by dense packings. What's more, in previous experiments (Chapter 1, Ref. [1]) we learned that the interaction for differently sized grains was inherently different to that of identically sized grains. In the case of two identically sized grains we found the interaction was purely repulsive, but for grains of different size it was repulsive at long-distances and attractive at short distances.

A recent study [36] bidisperse superparamagnetic colloids confined in two-dimensions and exposed to a magnetic field, found that the small beads tended to cluster together. The difference in the size of the colloids resulted in a difference in the way they interacted. Specifically, the interaction was negative non-additive, which is to say that the sum of the small-big interactions was less than the sum of the big-big and small-small interactions. It was then shown theoretically and in simulation that clustering of the small beads would occur naturally in equilibrium.

Bidisperse packings of beads in the air-fluidized bed show clustering of a similar kind among the small beads, fig. B.1. In the figure are two representative snapshots of two area fractions, a low and a high, and clustering is evident in both. It is natural to wonder whether such clustering in the bed is related to that seen in the colloid experiment, because if so it may imply that the air-ball interaction is negative non-additive as well, and further strengthen the analogy of this




Figure B.1: [CircleImages] Representative snapshots of two different area fractions, as labelled. The big grains are colored light gray and small grains are colored dark gray and have a size that is proportional to the relative size of the grain. Clustering is evident at both area fractions and may be a consequence of negative non-additivity of the bead-air-bead interaction.

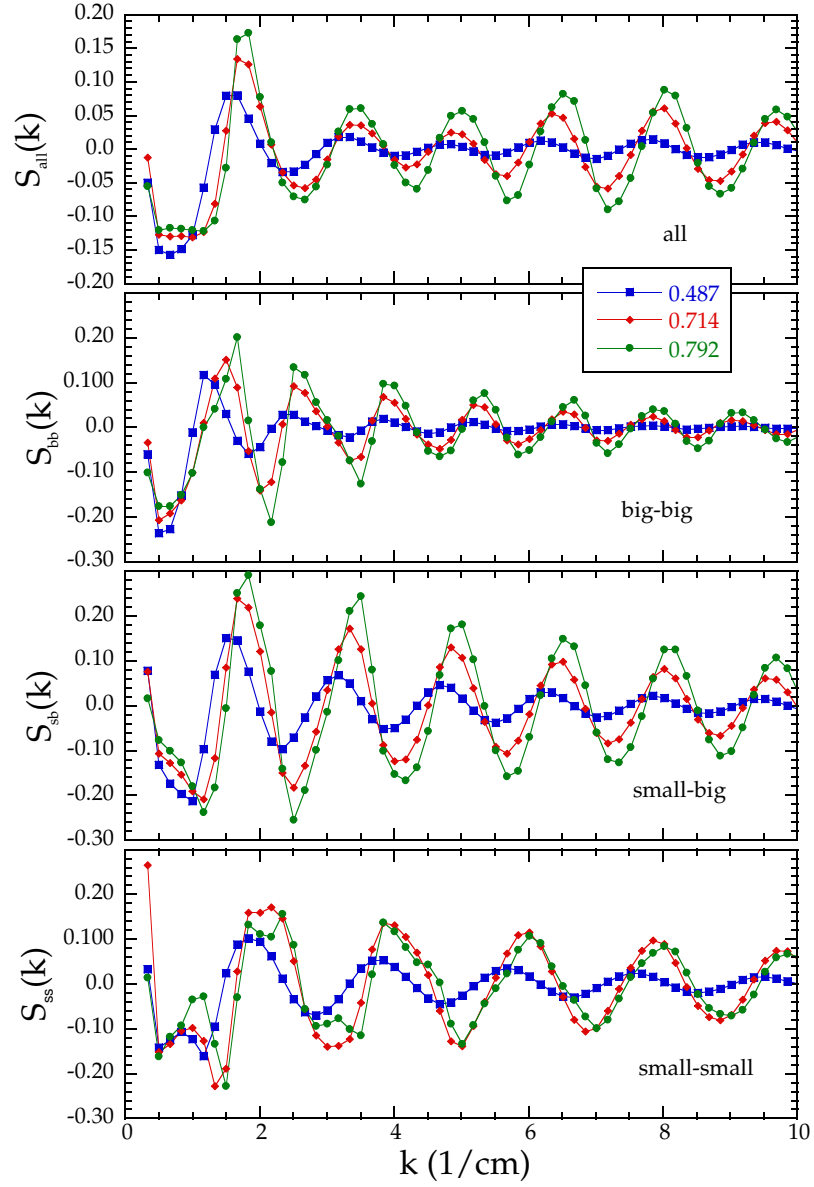


Figure B.2: [StructureFactor] The structure factor averaged over time and space as a function of wave vector for three different area fractions. Each component of the structure factor is measured separately, as labelled. (a) Taking all beads, (b) big-big, (c) small-big, and (d) small-small contributions. The expectation based on previous negative non-additive systems is a small wave vector peak in the small-small component [36].




Figure B.3: [EulerCharacter] The Euler characteristic measured for (a) the big beads and (b) the small beads, as a function of covering, or effective particle size. The curves have been displaced by a constant for clarity.

driven-dissipative steady-state to thermal equilibrium.

The first step then is to measure quantitatively the clustering that is evident in the images. Following the lead of the colloid experiment we measure the structure factor of each configuration for different components separately and average over time, fig. B.2(a-d). The structure factor is the fourier transform of the pair-correlation function and relates information about characteristic distances between particles. We have measured the structure factor for systems of different density with similar results and plot them all in the figure. Each figure presents a different component of the structure factor for different bead combinations. In the colloid experiment the effect of small bead clusters is a low wave-vector peak in the small-small structure factor with no corresponding peaks in the other structure factors. This is exactly what we see, fig. B.2d. The wave-vector of the peak corresponds to the reciprocal lengthscale of the clusters, which we find to be at $k \simeq 1 \text{ cm}^{-1}$, about 3 bead diameters.

The structure function is therefore one way of viewing clusters. Another way is based on a method from integral geometry that can characterize cluster morphology. The first step is to parse the image into a grid and the gradually expand the “covering” of each particle, keeping track of the number of connected domains C , and holes H . The Euler characteristics is then, $E = C - H$. We

measure the Euler characteristic for the big and small beads separately for the same three densities as before, fig. B.3. The curves have been displaced by $E = 2$ for clarity. The salient feature is the second plateau at about $E = 0.5$, which is representative of the same spongy topology of the small-bead clusters that was evident in the colloid experiment.

In this appendix we have shown that clustering of small beads in the gas-fluidized bed is quantitatively similar to clustering seen for superparamagnetic colloids at equilibrium. Simulation and theory has traced the origin of this clustering to negative non-additivity of the interaction. In the case of two spheres of different size in the gas-fluidized bed we found that the potential was attractive at short distances, and speculate that this may be the source of negative non-additivity in this air-driven granular system.

Appendix C

Voronoi Shape Fluctuations

In this appendix we present additional analysis performed on the driven bidisperse system of steel beads discussed in Chapter 2. For each configuration that the system is in at each instant in time, one can generate a unique two-dimensional Voronoi diagram. Such diagrams relate a multitude of information about the system's instantaneous structure. For example, one can immediately tell how ordered the system is, if it is crystalline or amorphous, and how many nearest neighbors each particle has based on the Voronoi criterion. Because the system is driven in a steady state it continuously evolves into new configurations, and associated with each of these configurations is a Voronoi diagram. Recent theoretical and simulation study of two-dimensional monodisperse atoms undergoing crystallization showed that the fluctuations in Voronoi structure depend sensitively on the degree of crystallization [73]. Additional work from the same group also showed that in simulations the transition from a Wigner liquid to a Wigner glass could also be detected in the noise spectrum of these topological fluctuations [74]. Specifically, the power spectrum of the average number of six-sided Voronoi cells has a functional form that is constant (white noise) for a liquid and $1/f$ for a glass.

Our system of driven bidisperse steel beads is similar to the simulation of bidisperse particles. We track all particles at each instant over a wide range of densities, and thus are ideally positioned to test the generality of this theoretical and simulation study. In this appendix, we present the results of the test, and find that the noise spectrum of Voronoi fluctuations provides an alternative formalism for identifying the “pre-jammed” steady-state in our system, which is similar to the




Figure C.1: [Timetraces] The top figure is for a low density system, $\phi = 0.487$ and the bottom figure is for a high density system, $\phi = 0.792$. The light blue curve shows the percentage of cells that have six sides at given instant for each configuration. The purple curve shows the percentage of cells that have circularity within half a standard deviation of the average at each instant for each configurations. These topological averages fluctuate in a way that depends on the grain dynamics and structure.

supercooled state of thermal system. However, we note that this formalism like all others is inherently dynamical since instantaneous Voronoi structure is not sufficient to distinguish the liquid and glass states.

The method begins with the Voronoi diagram which is a mathematical transformation of the instantaneous configuration. Next, we calculate the percentage of six-sided Voronoi cells at each instant, which reduces the instantaneous configuration to a single number, $P_6(t)$. Thus, we obtain a time trace for $P_6(t)$ that fluctuates as the system evolves into new configurations [73,74]. However, we remark that because the number of sides of a Voronoi cell is a discrete quantity that changes abruptly only as a result of large changes in particle locations, the noise spectrum will have a limited range of sensitivity and may be prone to statistical error. We introduce another observable based on the circularity of the Voronoi cells, a continuous quantity that varies smoothly and is sensitive to small changes in structure. Specifically, we compute the average number of cells with circularity within half a standard deviation of the average, which we dub $P_C(t)$. Both of these quantities are designed to capture fluctuations in Voronoi structure, and are plotted in fig. C.1(a,b). Previous experiments (Chapter 2) showed that at low density the system resembles a simple liquid and particles are free to diffuse and at high density it resembles a glass and particles are trapped in a cage formed of their nearest neighbors over a wide range of timescales. These hallmark dynamical differences are manifest in the $P_6(t)$ and $P_C(t)$ time traces, which fluctuate wildly for low density and change more slowly for high density.

A way to quantify the noise of these topological time traces is to compute the power spectrums of $P_6(t)$ and $P_C(t)$, fig. C.2. For all densities and for both measures, the power spectrums approach $1/f^2$ for high frequency. This was not seen in the simulations and we believe is because we are able to observe the short time ballistics of the beads. The power spectrums change slope at the average bead collision frequency and at lower frequencies the beads break through their cages and move into new locations. At these frequencies the form of the noise-power depends on the combination of structure and dynamics. For low density liquid-like systems the noise spectrum is constant (white noise) and for high density glass-like systems, the noise spectrum approaches $1/f$ for both $P_6(t)$ and $P_C(t)$. This is in perfect agreement to the studies of Refs. [73,74] and shows that the noise spectrum of Voronoi fluctuations may be able to distinguish the glass-like behavior of this driven-dissipative system.




Figure C.2: [PSP6PC] The power spectra for the time traces $P_6(t)$ and $P_c(t)$. Previous simulations of atomic or colloidal systems undergoing a glass transition found that the power spectra is constant for a liquids and approached $1/f$ at low frequency for a crystal or glass.

Appendix D

Sensitivity of the Four-Point Susceptibility on the Overlap Threshold Size

The four-point susceptibility χ_4 , is based on the overlap order parameter and has developed a strong following among persons seeking to understand anomalous dynamics in glass-forming liquids and recently in granular materials (Chapters 3 and 4). χ_4 is useful because it is sensitive to correlated motion. But in this appendix we will demonstrate that the overlap order parameter and four-point susceptibility are tricky functions to properly measure, and one must be mindful of systematic errors introduced by arbitrary parameter choices.

The overlap order parameter, defined in Chapters 3 and 4, depends on the specified threshold size a_0 . Fixing a_0 is equivalent to choosing a q vector in the intermediate scattering function and corresponds to the overlap “size” of each particle. This overlap size is not the same as the particle size. It can be either bigger or smaller depending on the length scale of interest. The parameter size is a necessary input that is used to calculate the fraction of overlaps between different configurations. Depending on the size of the overlap function, one can greatly influence the fraction of particles that overlap between two different configurations. To clarify this point we construct fig. D.1a, which shows the self-part of the overlap order parameter for different parameter choices for the same data. $Q(\tau)$ is the fraction of particles that have remained within a radius of a_0 of their

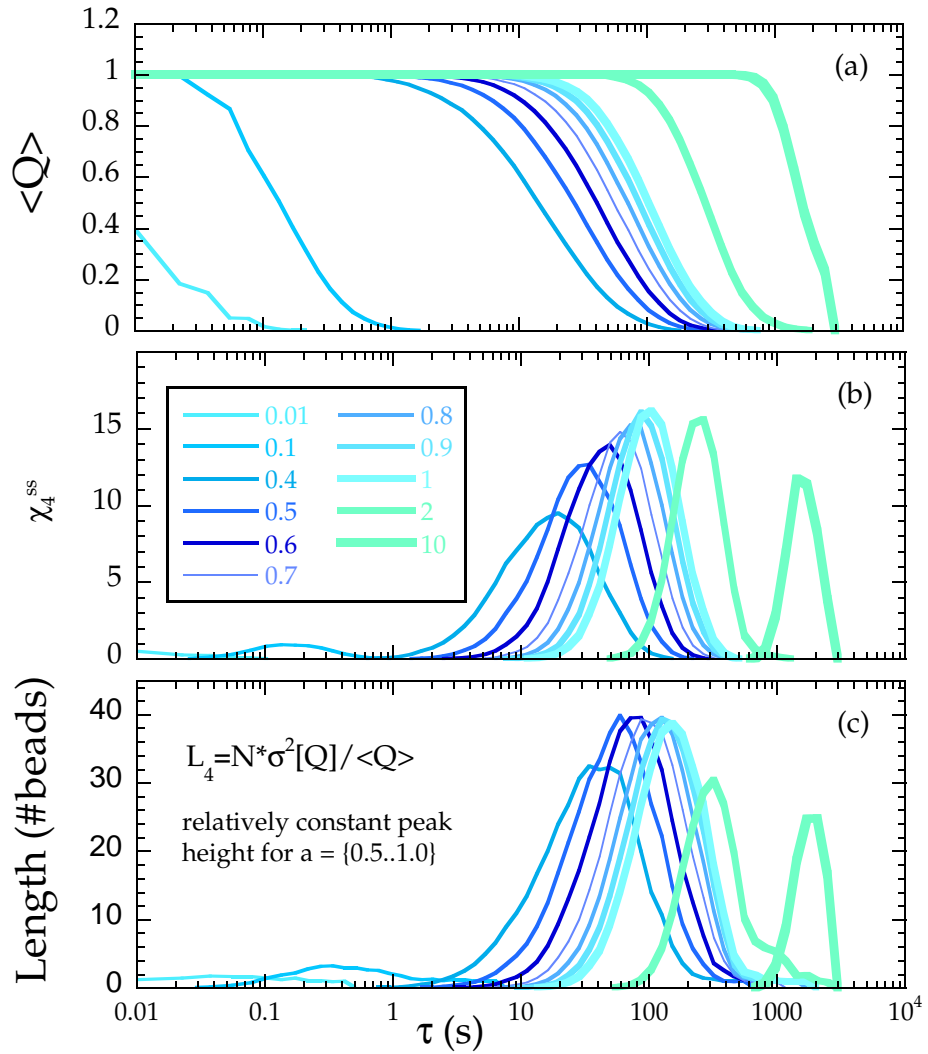


Figure D.1: [OOPPparameterChoice] (a) The overlap order parameter for a single area fraction for different choices of the overlap threshold size. (b) The four-point susceptibility for different choices of overlap threshold size. (c) The size of the correlated motion for different choices of overlap threshold size. All measurements are sensitive to the overlap threshold size.

original location as a function of time, averaged over all initial configurations. The plot shows that even for exactly the same data, the choice of a_0 , shown in the legend of fig. D.1b, greatly influences $Q(\tau)$. If a_0 is too small, then $Q(\tau)$ decays rapidly to zero due to short-time random ballistic motion. If a_0 is too large, then $Q(\tau)$ is insensitive to all but the largest displacements. These two extremes may be easy to avoid and thus overcome, but there is a vast range of choices in between and at some point the proper choice is not so obvious.

Since a_0 influences $Q(\tau)$, it also influences χ_4 , which is defined in terms of the overlap order parameter (Chapters 3 and 4). Fig. D.1b bears this out, showing χ_4 for the same data for the different a_0 choices. Such sensitivity to the choice of a_0 is undesirable since one effectively specifies the length scale of interest ahead of time, and may unwittingly limit the range of length scales that will be observed. The interpretation and universality of the χ_4 measured is also an important question that must be carefully answered. In our test we show that by varying a_0 over three orders of magnitude, we shift the time of the peak of χ_4 , which is a characteristic time of the dynamical heterogeneities, by six orders of magnitude.

The natural question to ask then, is how one should go about choosing a_0 . A somewhat arbitrary but usually satisfactory way of handling the problem is to choose the overlap threshold size at which the height of the peak of χ_4 is a maximum as a function of a_0 . For our system this turns out to be about 0.1 to 1 ball diameters. As seen in the figures, over this range of overlap sizes, the peak time and height of χ_4 doesn't change very much. The characteristic length scale associated with the dynamical heterogeneities, fig. D.1c, is also relatively constant over this range of a_0 , so that there is some flexibility in the choice.

To sum up, one must be careful when measuring the overlap order parameter and related four-point susceptibility. Both are sensitive to the choice of the overlap threshold size, a_0 which must be specified ahead of time. One way of dealing this is to check the dependence of these functions on a_0 and see whether there is a range over which they do not depend too strongly. There is no guarantee that such a range will exist. Even then, one must choose a single value. Another possibility suggested by the trends of the curves in fig. D.1(a-c), is it to average over many choices of a_0 . We have developed two alternative methods, the persistent area and persistent bond, which are more rigorous and straightforward to measure and also less prone to error, which is the subject of Chapter 4.

Appendix E

Energy Cascade in Driven Beads

Much work of late has uncovered so-called spatially heterogeneous dynamics (SHD) in a wide range of systems. They've been seen in simulations of supercooled liquids, measured indirectly in atomic and polymeric glass formers, and even in dense granular packings. The discovery of SHD in granular systems is of particular interest since it implies that they may be a generic feature of systems close to jamming. There is therefore a pressing need to discover an appropriate way of quantifying SHD in all these systems, and there are already a few candidates. One that has been around for some time in the glass community is the so-called four-point susceptibility, which is related to the self-part of the intermediate scattering function and can be indirectly measured using scattering techniques atomic systems [9]. Other candidates, which are the subject matter of chapter 4, are the persistent area and persistent bond. The advantage of these functions is that they can capture SHD over a wide range of time and length scales without the need to specify input parameters or choose a wave-vector of interest. Recently, a very simple function was measured in colloidal glasses [94], which was the the transverse and longitudinal velocity crosscorrelation function as a function of distance.

In this appendix we introduce measurements of a different variety that are most closely related to the colloid crosscorrelation function. Rather than computing the dot-product, we compute the longitudinal velocity difference between particles. This function has a long track record the study of turbulent fluids, and an enormous advantage is that it comes with a depth of theoretical understanding from the field of turbulence. A famous theoretical result by A. N. Kalmagorov showed that the third moment of the structure function in a turbulent fluid follows the famous $-4/5$ law.




Figure E.1: [MSDs] The mean-squared displacement for the two area fractions of interest. The blue curve is for low area fraction and the red curve is high area fraction. The low-density system saturates at long time due to finite system size, and is super diffusive between 10 and 100 seconds due to a convection current.

That this moment is negative implies that energy cascades from large scales, in which it is injected, down to small scales, where it has dissipated. The purpose of this section is to present results in which the third moment of the longitudinal structure function, $S_3(r)$, was measured for a dense disordered collection of steel ball bearings in two-dimensions driven by air. In previous studies the system was seen to have spatially heterogeneous dynamics.

We have measured the structure function for two different steel ball bearing systems. Both are 50-50 bidisperse mixtures of two different sized grains. One consists of approximately 400 larger grains, and the other of about 2000 smaller grains, although both are bidisperse. The results are essentially identical except that the far larger number of grains in the small bearing system means that we have better spatial resolution and better statistics. Therefore, we present those results only.

For orientation we present the mean-squared displacement for the two area fractions under investigation, fig. E.1. As found previously (chapters 2-4), the red curve is for a low-density liquid-like system and the blue curve is for a high density glass-like system. Both area fractions

are ballistic at short times and diffuse a long times, with the high density system having a long-lived plateau at intermediate times. This plateau is symptomatic of the so-called “cage effect” in which particles are trapped in a cage of their nearest neighbors. Whatever the origin of the plateau, this mean-squared displacement for this driven system of steel beads is essentially identical to that of colloidal glasses and other glassy systems with spatially heterogeneous dynamics.

In the measurement of the structure function there are two relevant variables that must be considered. There is the lag-time time, over which the average velocity is computed, and the distance of separation between beads. In turbulence the structure function is presented as a function of distance for a particular delay time interval. Following that procedure, we measure $S_3(r)$ for a few delay-time intervals as a function of separation distance, fig. E.2. The intervals chosen are labelled in the figures in terms of number of a frames, which are captured at a rate of 120 Hz. In previous systems spatially heterogeneous dynamics are evident over a certain range of timescales. At short times when the mean-square displacement is ballistic the velocity vector field is random in space and there are no velocity correlations. At long times when the behavior is diffusive, the velocity vectors are also random in space and again there are no correlations. However, at intermediate times there is much heterogeneity in the velocity vector field and strong correlations exist over multiple shells of the nearest neighbors. In the experiments presented here, according to the mean-squared displacement, spatially heterogeneous dynamics should be evident from 1 to 100 s in the low-density system and from 0.1 to 1000 s in the high-density system.

In fig. E.2 we show the structure function results and corresponding average velocity vector fields for the low-density system for different averaging times as labelled. It is clear that when particles are trapped and the mean-squared displacement plateaus, there are spatially heterogeneous dynamics. They first become apparent at one second (100 frames) and last till about 100 seconds (10,000 frames). Longer than 10 seconds the mean-squared displacement is super diffusive and we start to see curl in the average velocity vector field, which is due circulation in the driving airflow. This curl is constant in time and therefore does not influence functions like the four-point susceptibility, persistent area, and persistent bond, but it clearly affects the structure function. It is also evident that it sure times when the particles have not move very far, the structure function oscillates due to the discrete nature of the beads. In turbulence such discreteness occurs at the atomic level

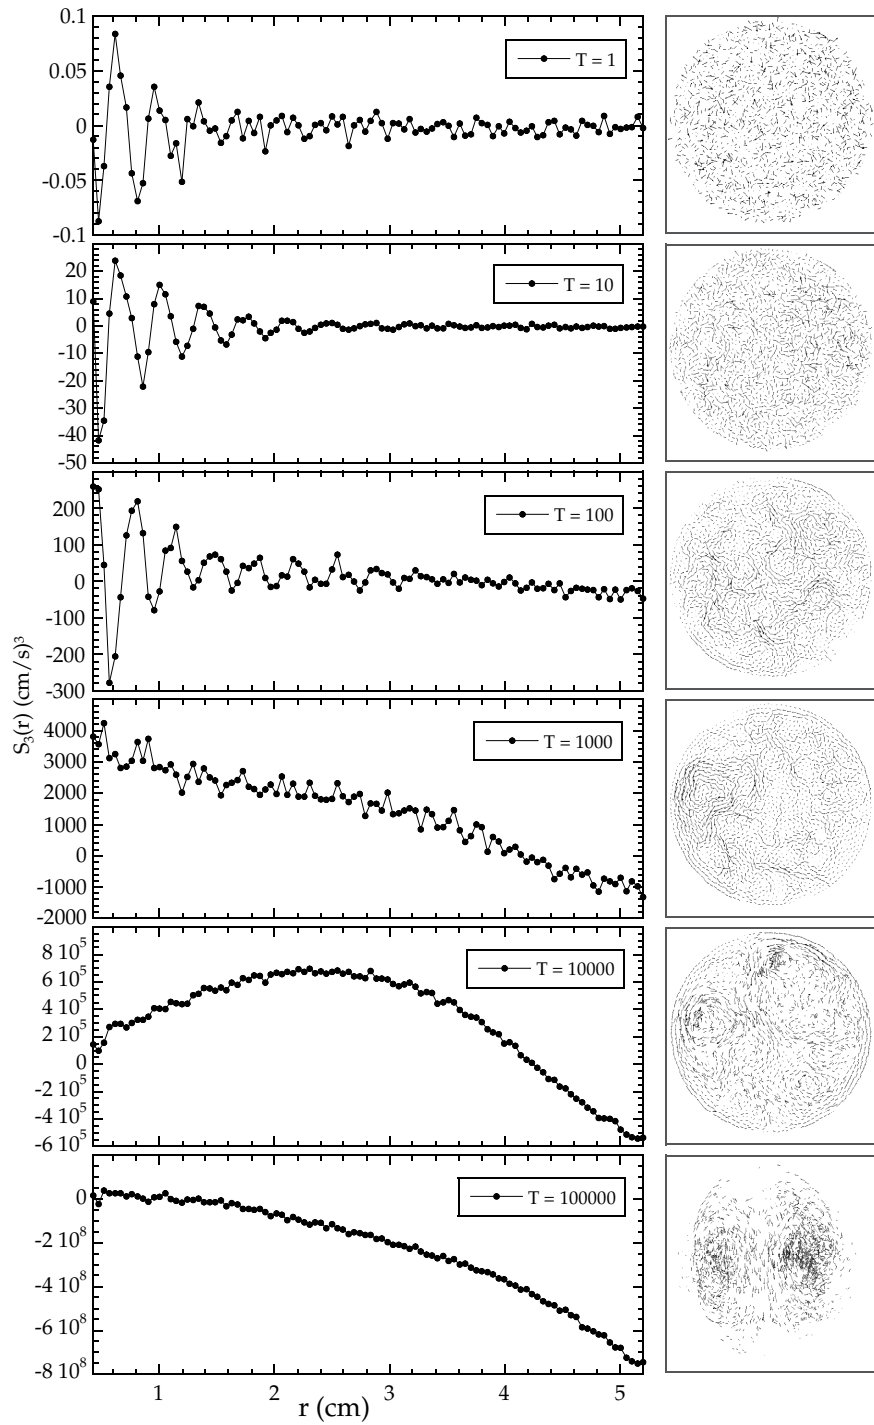


Figure E.2: [S3rAF737] $\phi = 0.737$. The left-hand side shows the third moment of the structure function and the right hand side shows the average velocity vector fields for different time intervals as labelled. The time labels are in number of frames which are intervals of 1/120 of a second.

and is therefore invisible, and this may be an important difference of granular systems. As the average time gets larger the oscillations increase in magnitude and become more negative. Eventually, most of the function is negative, and for the images $T > 1000$ fr there is a change in slope in the structure function at 2 cm, which is close to the length of the correlated velocity. In turbulence, the fact that third moment is negative is an indication of energy cascade. In turbulent fluids it was worked out theoretically that this moment follows the “-4/5 law” so that energy cascades from large to small scales. Unfortunately, in this granular system the law cannot be properly verified because the functional form is not the same. Still, the fact that this moment is negative may still imply energy cascade to smaller scales.

For high-density things are quite different. According to the mean-squared displacement this system is much more glass-like with a plateau that lasts for four decades. Dynamical differences are also evident in the spatial makeup of the velocity vectors, shown in the images for different delay intervals, fig. E.3. According to the MSD, caging lasts from 0.1 to 100 seconds, and comparison with the images shows that spatially heterogeneous dynamics first become evident for $T = 10$ fr and last to $T = 10000$ fr. The heterogeneities are also larger than for the low-density system and increase in both length and coherence with increased averaging time. We also note that because of the high density of this particular system we do not see convection at our longest times. This could be because the strong hard-core interaction of the beads far exceeds the small DC circulation in the air, or it could be that with three more decades at long time such circulation would be evident. Regardless, it is clear that there is strong spatial heterogeneity in the dynamics over the timescales we do observe the system and they suffice to test the structure function. As with short times in the low-density system, there are positive and negative oscillations in the structure function, which again is a consequence of the discrete particle size. Interestingly, these oscillations persist for large separation distances, well beyond oscillations in the pair correlation function. Unlike low-density, the high density structure function oscillates symmetrically around zero with only a slight preference for negativity, and we do not see obvious indication of energy cascade. The only obvious difference between the velocity vector fields is the long time DC convection current. It is also possible that we do not have sufficient long time resolution to see the cascade of energy since the relaxation time of this glass-like system is orders of magnitude larger.

To conclude, the longitudinal structure function is certainly an interesting and novel way of

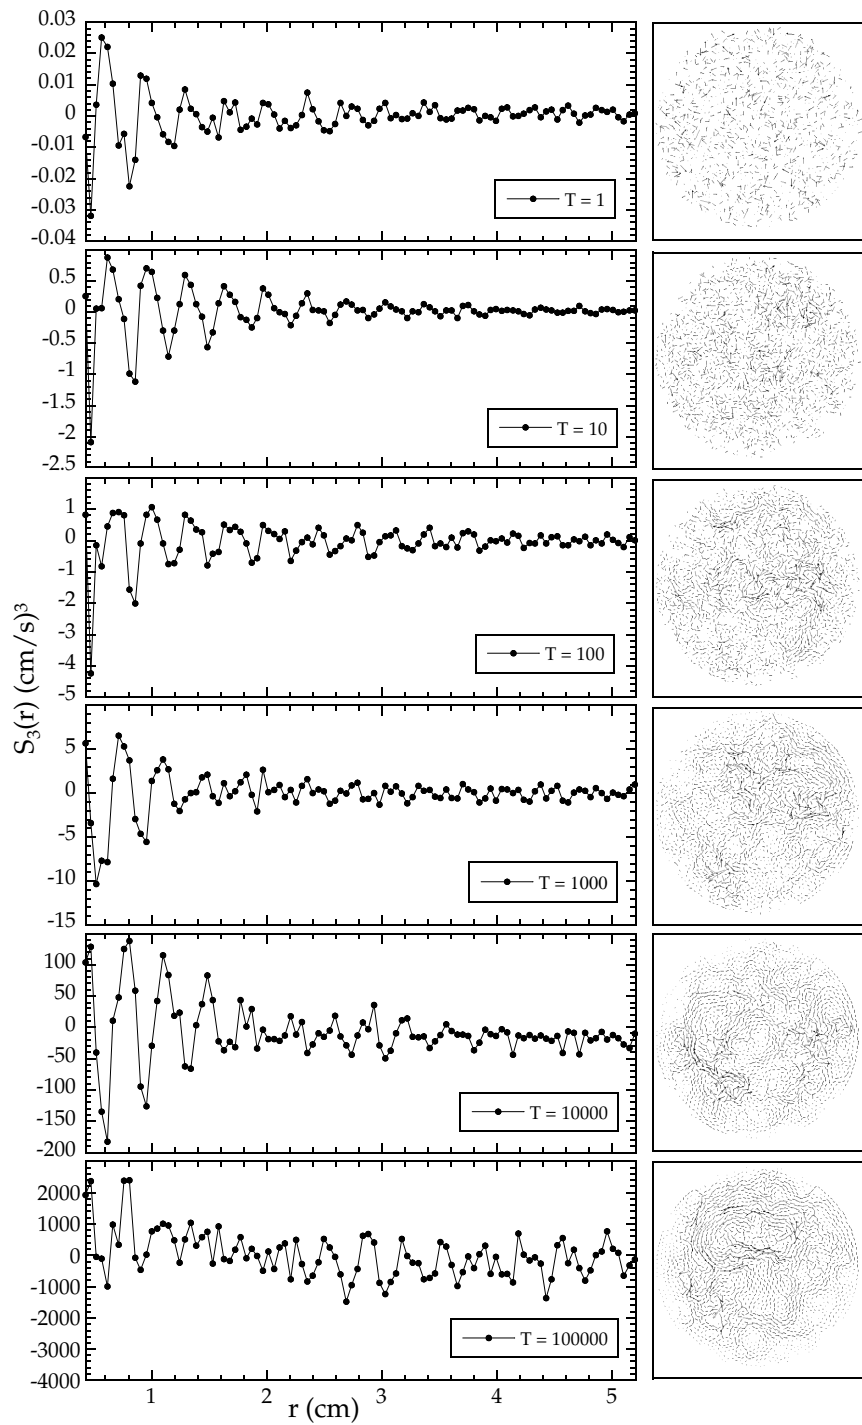


Figure E.3: [S3rAF802] $\phi = 0.802$. The left-hand side shows the structure functions and right-hand side shows the average velocity vector fields exactly as in fig. E.2.

analyzing granular flows. However, it does a poor job of quantifying spatially heterogeneous dynamics that are obvious in the images and it is clear that existing functions such as the four-point susceptibility, persistent area, and persistent bond do a better job. Nonetheless, the deep theoretical understanding that can be borrowed from the field of turbulence may inspire granular physicists to study it more closely, since it offers an entirely unique perspective. For example, the fact that there is negativity $S_3(r)$ could signal an energy cascade from large to small scales in granular systems such as this one. Such a cascade is something that granular physicists accept as an ideological truth but have never rigorously verified.

Appendix F

Small Sand Avalanches

The experiment that we present in this appendix is similar to that of Chapter 6 with the exception that here are used smaller sand particles. The diameter of the sand used in this experiment is about $100\ \mu\text{m}$ with polydispersity of about 10%, so that these “small” sand grains are about a third the diameter of the grains used in the experiment of Chapter 6. Our purpose for repeating the experiment with smaller grains is twofold. First, we would like to know how changes at microscale effect macroscopic avalanche behavior. Second, in the previous experiment we saw that the velocity profile across the channel depended on grain size. In particular, it was found to be parabolic for the big grains and more plug-like for the small grains. We would like to determine whether these difference in the profile are symptomatic of other differences in the avalanche cycle. For example, it is possible that a parabolic profile is a consequence of force chains spanning the channel which could interfere with the natural behavior of avalanches.

We have observed the small sand system for 10 hrs at a flow rate of 0.07 g/s. For this observation duration and flow rate we witnessed 450 small sand events, which is about a third as many events as were witnessed by observing the big sand for 35 hrs. In fig. F.1a we show a snippet of the full $\delta v(t)$ time trace. Each spike in the time trace represents an avalanche and it is evident that the avalanches are approximately periodic and that the on-state comprises a small fraction of the total cycle. One advantage of having such a short on-state is that we can be assured that we are in the quasi-static flow regime, which is to say that increasing the flow rate only decrease the time between events and does not affect the time of events. In fig. F.1b we show a zoom-in of the grayed region in fig. F.1a. Even in this one event is evident that the small sand avalanches have a much




Figure F.1: [SmallAvalanches] (a) Example time trace of $\delta v(t)$ for intermittent avalanche flow. Each spike corresponds to a single avalanche event. (b) Zoom of the grayed region in figure (a) showing the time trace for a single small sand avalanche.




Figure F.2: [SmallStats] (a) Scatter plot of the time-on and time-off durations. (b) Scatter plot of the maximum fluctuation speed and the time-off duration (c) Scatter plot of the time-on duration and the maximum fluctuation speed. (d) Histogram of the duration of avalanche events. (e) Histogram of the duration between avalanche events. (f) Histogram of the maximum fluctuation speed obtained by the avalanche.

peakier δv_{max} and that the abruptness of the cessation of flow is much sharper, with little of the stick-slip motion that was apparent with the big sand events. Direct comparison of the δv show that the small sand particles tend to fluctuate with smaller δv , and this is probably a consequence of having reduced an important length scale, the diameter of the grains. Nevertheless, while the fluctuation speed is smaller the flow speed down the channel is actually larger, as was measured using superficial particle tracking.

F.1 Small Sand Avalanche Statistics

As in Chapter 6, we begin the study of the avalanche cycle by compiling statistics of time-on, time-off, and δv_{max} . In the bottom row of fig. F.2 we plot histograms of the three observables and in the top row we plot correlations between the observables. Fig. F.2d shows a histogram of

the time-on durations. The small sand avalanches have an average duration of 2.4 seconds with a standard deviation of only a half second. The small sand avalanches last only about a tenth as long as the big sand avalanches. In fig. F.2b we plot the time-off distribution, showing that on average the small sand avalanches are off for about 80 s with the standard deviation of 20 s. This time off distribution has only a single peak, suggesting that the double peak in the big sand experiment may be an artifact of variation in the flow rate. Similar to the big sand histograms we see that there are high tails that correspond to uncommonly long lasting events. These wide variations in how long events last and how long they are off is related to the wide range of stable packing densities and angles of the granular pile. Whether an avalanche will flow or come to rest does not only depend on the volume of sand that has collected on the pile. It also depends on the particular arrangement of the grains, which is a distinguishing feature of granular systems in general. In fig. F.2f we plot the histogram of δv_{max} . It is symmetrically distributed about an average of 4.5 mm/s with a small standard deviation of 0.2 mm/s, suggesting that while small sand avalanches may vary widely in size, they all tend to fluctuate at the same speed.

To take the analysis further we also consider correlation plots between the observables, fig. F.2(a-c). Fig. F.2a shows the correlation between the time-off and time-on. Long time-offs seem to correspond to long time-ons, suggesting that, as with the big sand avalanches and as suggested by variations in the on and off time distributions, small sand avalanches come in different sizes. After all, the time-off depends on how much sand must be replenished before the maximum angle stability is surpassed. The time-on depends on how much excess sand is available above the angle of repose. Both durations should increase for larger avalanches at fixed flow rate. In fig. F.2(b, c) we plot correlations of δv_{max} with time-on and time-off and, similar to the big sand avalanches, find little correlation.

F.2 Anatomy of a Small Sand Avalanche

In this section we discuss the anatomy of the average avalanche dynamical shape. Avalanches tend to occur roughly periodically with an average time off of about 80 s and an average time on of about 2 s. To get a better sense of the dynamical shape of the on-state we average together 450 separate events witnessed over the 10 hr observation time, fig. F.3. To compute the average, we




Figure F.3: [SmallShape] Average dynamical avalanche shape of 450 separate small sand events. Similar to the big avalanches, there is a peak just after the turning on time followed by a gradual deceleration as the sand supply is exhausted. There is a high tail at long times $t > 3$ s, which is an artifact of averaging together events of different duration. Inset shows the average avalanche normalized shape, in which each event is normalized by its maximum fluctuation speed and its total time-on. Normalization removes the long time tail and uncovers the universal avalanche shape and the abruptness of jamming at turning-off.

line up the avalanches according to the time that they turn on, which we measure by thresholding the $\delta v(t)$ data at 0.01 mm/s. Because the transition from off to on is very quick the threshold can be changed by a factor of ten without changing the results of the average. We compute the average avalanche dynamical shape by computing the mean and standard deviation over the different events at each delay-time instant. As for the big sand, just after turning on the small sand events attain their maximum fluctuation speed. Then the avalanche ages, constantly decelerating until reaching the stage of cessation when the flow abruptly stops. There are a few notable differences between the big and small sand avalanches. The peak of $\delta v(t)$ is peakier and the constant deceleration region that follows is very flat. Similar to the big sand avalanches, by computing a simple average over all avalanches we end up with high tail at long delay. This is also an artifact of averaging short and long-lived avalanches together. Therefore, we compute the average normalized shape as before, by normalizing each event by its maximum fluctuation speed and total duration. This way each avalanche begins at normalized time of 0 and ends at normalized time of 1, inset fig. F.3. The average normalized shape does not have the long time tail and is a more representative picture of the universal avalanche dynamical shape. The relatively small standard deviation of different events that differ greatly in total flow duration suggests that there is a universal shape that is independent of the total flow time and maximum fluctuation speed. Still, it is evident that the variation among different events blooms just before turning off so that the fluctuation dynamics are poor predictors of the precise jamming time. If avalanches tended to end the same way one would expect the opposite, that the variation would decrease as the jamming time is approached.

The average avalanche shape shows that there is dynamical structure to the spatial composition of avalanches. In fig. F.4 we show a zoom in of the turn-on region. In the background is a contour plot of $\delta v(t)$ for the zoom in, which clearly partitions the turn-on of the average avalanche into three distinct regions. Initially the pile is completely static so that $\delta v(t) = 0$. Then $\delta v(t)$ reaches a maximum as the avalanche front traverses the beam spot. The flow then drops to a value that corresponds to fluctuation dynamics farther back in the pile, as the flow constantly decelerates. There is evidence in the literature that the front of an avalanche has complex structure that depends on the composition of the grains. In the frame of the avalanche, for example, one can see convective rolls either foreword or backward depending on grain roughness, shape, and polydispersity. At the very least, we see that the fluctuation speed spikes just after turning on, and that the width of the




Figure F.4: [SmallOn] Zoom in of the average turning on of small sand avalanches. In the background is a contour plot of $\delta v(t)$ showing that there are three stages in how avalanches turn on.

spike is proportional to the time it takes for the front of the avalanche to traverse the beam spot.

To get a better sense of how avalanches turn off we compute the power spectrum of the full data set, fig. F.5. First, we coarse-grain $\delta v(t)$ which allows us to more easily handle the data in memory but also removes 2-1/2 decades of high frequency information, which is mostly noise-floor saturation as in the power spectrum of Chapter 6. Nevertheless, it is clear that during avalanche flows the power spectrum goes like $1/f^2$. Similar to the big sand spectrum, there is a gentle hump close to the reciprocal of the avalanche average period. There are also spurious oscillations at lower frequencies which we believe are transform artifacts since the on-time comprises only 3% of the avalanche cycle and $\delta v = 0$ a majority of the time, fig. F.1a.

In this appendix we present a similar experiment to that of Chapter 6 except that the grains here are about a third the size. The important question that we seek to answer is whether changes at the microscale affect the macroscale flow dynamics, and in particular if there are artifacts due to finite size of the grains. Indeed, the data suggest that avalanche dynamics are qualitatively independent of the grain size for our experimental geometry, although there are some quantitative




Figure F.5: [PSSsmall] Power spectrum of the full $\delta v(t)$ time trace for the small sand avalanches. As with the big sand avalanches of Chapter 6, the small sand power spectrum is $1/f^2$ for timescales corresponding to the avalanche on-state, and has a gentle peak at the inverse average avalanche period. At low frequencies there are spurious oscillations which may be due to the on-state comprising a relatively small portion of the avalanche cycle so that $\delta v = 0$ most of the time.

differences and possible indications of finite size effects in the big sand experiment. A few the universal features are nonetheless clear. First, avalanches tend to have an intrinsic variation in their size, which is a consequence of the intrinsic variation of stable packing densities and angles of stability of granular piles. Second, even though avalanches may differ greatly in size and duration, there is a universal avalanche shape. This is true both for the big and small sand avalanches and is best seen in the average normalized avalanche shape. Third, avalanches tend to turn on in a similar way, going from completely static and to quickly reaching their maximum fluctuation speed as the wavefront passes. Fourth, there is a wide variation in avalanche dynamics prior to the cessation of flow, so that one cannot predict the precise jamming time. This is reminiscent of the unpredictable and often catastrophic behavior of many granular flows, whether in a hopper, grain elevator, or down the side of the mountain.

Appendix G

Technical Appendix

G.1 Digital Video Imaging of Particles in the Gas-Fluidized bed.

G.1.1 Gas-fluidized Bed and Its Components

The apparatus was largely built by my predecessor Rajesh, although it has been altered over my reign. For more information about the apparatus I refer you to his thesis. There is also much information about gas-fluidization technology generally, and a beautiful compendium of our early results.

The Gas-Fluidized Bed Apparatus

The apparatus that we call the gas fluidized bed is a $1.5 \times 1.5 \times 4 \text{ ft}^3$ box with two holes cut out of it, one at the bottom and one at the top. To the hole at the bottom is connected an air blower and into the hole of the top we place sieves. The innards of the box contain air filters sandwiched between perforated metal sheets which help even-out the airflow and reduce circulation and nonuniformity. There are many things that can cause circulation and nonuniformity in the air and is is extremely difficult to get it completely even. This can have undesirable effects on the behavior of particles fluidized in the sieve above and it is therefore important to ensure that the windbox is clear and that the sieve is defect free.

The blower is a standard two-speed carpet dryer connected to the outlet through a Variac auto-transformer. Using the autotransformer we can finely control the voltage applied to the blower so

that even at fixed blower setting we can finely control the airspeed over a range from 0 to 1000 cm/s with a resolution of about 5 cm/s. We verify the airspeed directly using a hotwire anemometer. The anemometer is a delicate device which uses a thin wire and resistor to measure the airflow passing over it. If the probe is accidentally banged the wire can be easily snapped and costs about \$300 to replace. If the anemometer's readings are erratic, likely it has run out of batteries which can be replaced by screwing off the back panel. The resolution of the anemometer is 10 cm/s, although it is possible to get finer resolution by measuring airspeed as a function of applied voltage. Then the airspeed can be interpolated for a given applied voltage without having to measure it directly.

Sieves can be purchased from McMasterCarr's web site. It is important to choose a sieve with a small enough mesh size, typically 100 μm , so that balls rolling on it are not strongly influenced by the grain of the mesh. It is also important to ensure that there are no dents in the mesh for the same reason. If the mesh is dented or the walls are not circular, it can affect the behavior of particles. The sieve hole is 12 in in diameter and so we use 12 in sieves from McMaster. Often this is too large a diameter for the experiment we want to perform and so we have machined inserts to place above and below the mesh. The inserts must be machined with an outer diameter that allows them to fit snugly into the walls of the sieve. Another advantage to using inserts is that you can choose the particular shape of the boundary. In past we've used circular, square, diagonal, and even stadium shaped boundaries specific to the experiment. The inserts must be inserted both above and below sieve. If an insert is inserted only above, the air will have a horizontal component in the plane of the sieve near the boundary, and this can greatly affect the behavior of particles nearby. Sometimes this effect is desirable, but most of the time it is not.

The last component of the apparatus is the base, which houses a swivelling platform and also a motor that can tilt the platform. The base sits on a tripod of three adjustable feet. The feet are able to adjust the level of the base perpendicular to the swivel axis. The swivelling platform can be tilted to a maximum angle of about 2 deg. To tilt we use an electric motor connected through an adjustable shaft to one side of the platform. The motor is controlled by an Electrocraft 2000 controller which is connected through a serial port to a computer. The software that controls the motor is called "IQ Master" and can run on Windows XP, although is designed for older versions of Windows. It is extremely buggy and cumbersome but when you get it working it is reliable. If you need a new copy of the software or instruction manual you can call the company and ask

for “Jim” who is the only person alive who knows how this thing works. Unfortunately, tracking Jim down can be a difficult matter, as the manufacturer or who built the controller and developed the software has been bought out, and you will have the call around to find the appropriate phone number.

To use the IQ software you must first make sure that the Electrocraft 2000 is enabled. To enable it there is a switch that sends 10 volts into the appropriate pin. If there is a fault the LED will shine red rather than green, and it can be cleared by flipping the switch off and then back on after a few seconds. You should see the LED flip back at green to indicate that the fault has been cleared. Once the Electrocraft is enabled start the IQ software on the computer and go to the “Run Control” menu. Select the velocity option and specify the rotation speed in revolutions per second. It is recommended that you not exceed 20 rev/s since the apparatus is not designed to swivel very fast. To start motion click the “start move” button. Do not click the “run program” button as it seems to be a decoy that will cause the motor to behave erratically and spin out of control. Often the IQ Master software will hang when you try to open the run control menu. If this happens there is no logical way to gain control over the motor. I have found that the only thing to do is force quit and alternate restarting the computer and motor until you can establish control without hanging. Sometimes even this will not work and in these cases connect the Electrocraft controller to a different computer and attempt to establish control with it. The IQ software was not written to run under Windows XP and it often conflicts with other network programs running in the background. Therefore, attempting to establish control with a different computer often works because it usually has different tasks running in the background. If you gain control with the other computer switch back to the original computer and run the IQ master software and run control again. In most cases it will work. Once you have control the connection is typically very liable provided that you do not close IQ Master or the run control window. If you want to stop motion simply set velocity to 0 rev/s. The IQ Master is content in this state indefinitely and will allow you to control the motor reliably until you close it.

Lights, Camera

Undoubtedly, the most important aspect of any imaging experiment is lighting and optics. In our case we primarily study spheres which have wonderful geometrical properties for the purposes of

imaging. The rule governing the scattering of photons off a sphere is very simple, and as a result it is simple to concoct lighting that makes imaging easy. For example, we use six incandescent lights three feet above the spheres in a ring that surrounds the camera. Because the lights are clustered near one another, close to the camera, and also far from the spherical particles, the only photons emitted from the lights that reflect into the camera are those that scatter off the tops of the spheres. Thus, one sees a dot at the center of each sphere. These dots are smaller and brighter depending on how shiny the spherical particles are. If the particles are opaque as they often are in our case, then rather than seeing a dot, the entire cross-sectional area of the sphere will be lit up. As a result, when two opaque spheres near one another they will be imaged as touching, which complicates image analysis. It is therefore preferable to use shiny particles whenever possible, even if the particles aren't spheres. If opaque particles must be used, then an image analysis technique that works well to separate touching circles is an erosion. An erosion algorithm is available in LabVIEW's "Gray Morphology VI" in the "Vision" toolkit (for an example of an erosion see fig. 1.1 in the first chapter). It works best when fed a circular erosion tool, which is an array with a circular distribution of 1s.

With suitable lighting the next step is preparing the camera to capture images. We use a rather old Pulnix 6710 120 Hz 8Bit CCD camera. There are two classes of camera settings, hardware settings and software settings. Hardware settings have to do with the lens and must be done manually by tuning the zoom, focus, and aperture. I rarely change zoom lens settings. Software settings must be changed using the "Measurement and Automation Explorer Window" that is installed with LabVIEW. The camera is listed under the IMAQ menu and allows you to select a desired CCD subsection and intensity lookup table, such as logarithmic, linear or binary threshold, in which case you must also specify the desired threshold values. There are a few dead pixels on the camera that usually do not interfere with imaging. If they do, a simple fix is to include a small subroutine in the image capture program that interpolates the dead pixel based on the average of the pixels surrounding it.

To capture images we use a buffered image stream written in LabVIEW that saves them as an AVI movie in "Microsoft RLE" codec. When saving a binary movie it is usually best to choose a "lossless" codec like Microsoft RLE. This codec is rather old and does not have the best compression, but it is a simple algorithm that allows the computer to stream at the camera's full frame

rate, and it is lossless. We use a buffered stream which allocates memory space to temporarily store images. Otherwise, the camera and computer are likely to fall out of sync and frames can be dropped. Occasionally, the buffer will fill if the computer is running another process, and movies will be encoded improperly, which will be obvious when viewed. In these cases it is wise to close all background processes or restart the computer to clear memory. If you desire to write your own buffered stream data capture program, the best thing is to start with the buffer ring example that comes with LabVIEW and alter it as necessary. The same program can be used to store images if you should desire, but because of the large number of images captured in a data acquisition, typically on the order of 100,000, storing images becomes cumbersome since Windows does not deal well with directories with 100,000 files in them.

G.1.2 Image Analysis and Particle Tracking

Locating and Tracking Particles

There are two stages to particle tracking. The first is to locate all particles in an image. This is when image analysis must be used and is usually the part when fatal imaging problems become apparent. The second is particle “tracking,” which means to link the locations of a single particle together as a time trace. It should always be the goal in imaging experiments to solve imaging problems with optics and lighting when ever possible, leaving analysis to do only the simplest tasks.

All particle finding methods typically come down to variations on a similar theme. First an image with bright spots is captured with a camera. Then, the image is thresholded into binary and a binary particle finding algorithm is run. If particles must be separated, this is when one runs an erosion or, if the particles are circular, a circle-finding algorithm. Such algorithms are all available in LabVIEW’s vision toolkit. At this stage if one desires higher center-of-mass accuracy, than one can go back to the original image and perform a weighted average over the pixel intensities in a neighborhood around the binary centers-of-mass, refining the locations. Usually, this unnecessarily complicates the procedure with little additional precision. Moreover, one must capture images in grayscale, not simply in binary, greatly increasing the amount of disk space required and the likelihood of buffer overrun.

Particle locating routines are also available in IDL. They are much more sophisticated, a bit more accurate, and widely used in the scientific community. They can be freely downloaded from

Eric Weeks' web site. The drawback to using such routines is also that they are sophisticated, and one usually does not have a working knowledge of them. This is fine if you only want to track bright spots in a series of images. But often in the course of an experiment one encounters unique problems that have specific needs and standard particle tracking routines may not be up to the task. In these cases, the only real option is to write a program yourself and this is immensely easier if you have a working knowledge of the components. If you are able to tailor a program to your specific needs, you can design the simplest program possible that gets the job done. In image analysis and particle tracking simplicity is tantamount to reliability.

Whichever way particles are located the end result is an array containing the x-position and y-position of each particle that was found in each image. The next step is to track the particles, which is to say, to find each particle's location in all the images and link them together as a single time trace. As with particle finding methods, most tracking methods do the same thing. Since particles move a small distance from one image to the next, particles that are nearest to one another in two consecutive images are identified as the same particle. This procedure is simple, reliable, and fast, and should be used whenever possible. If particles enter or leave the field of view then one must add additional code to account for these possibilities, and this can greatly complicate the tracking program. In such instances one may consider using the freely available code for IDL since it is capable of handling very complicated scenarios out of the box.

Saving and Refining Particle Tracks

There are many datatypes that one can use to save particle tracks. My program written in LabVIEW uses a "datalog" file, which is a unique LabVIEW data type that cannot be easily understood by other programs or languages. The datalog is a general datatype that can store any kind of data structure in LabVIEW. My tracking routine saves "DATFiles" which are datalog files that contain a two-dimensional array of three-element clusters. Each cluster contains the x-position, y-position, and area of a particle. Each column in the DATFile is for a single particle and each row is for single time. Therefore, the top row contains the locations of all particles found in the first frame, and the left-most column contains the locations of the first particle in all the frames. For very long data runs, my tracking program saves multiple DATFiles with a maximum predefined size that keeps the arrays small enough to be conveniently handled in memory.

A different way to save particle tracks is the method used by the standard IDL tracking routines. They save spreadsheet files containing a two-dimensional array of data. Each column in the array corresponds to a particle property. For example, the columns could be, from left to right, x-position, y-position, particle ID number, and frame number. In this architecture, every particle in the first image will have a frame number value of 0, every particle in the second image will have a frame value of 1, and so forth. Each particle identified has a unique ID number in its particle ID column. This architecture is amenable to particles entering and leaving the field-of-view, since it does not need to be altered to handle this scenario – each particle is given a unique ID when it enters that is permanently retired when it leaves. It also puts all the data in a single file which can be a good or a bad thing. On the one hand, it makes the data relatively easy to access. If one desires all of the locations from the first image, then one simply searches and extracts every row that has a frame number equal to 0. Similarly, if one wants the data for the first particle found in the first image than one searches and extracts all the data with a particle ID number column equal to 0. Such a search procedure is reasonably quick for moderate data sets. On the other hand, if one tracks a large number of particles for a long time, this file can become inconveniently large, making it impossible to handle in memory. One can solve this problem by chopping the file into pieces or writing a sophisticated disk read program, but these fixes can largely nullify the inherent advantages of the data structure.

Often, after particle tracks are saved one also need their derivatives. The straightforward thing to do is compute “finite differences,” which is to say, given a time trace for some variable, compute the difference for all consecutive frames and divide each by the time interval. Finite differences are notoriously inaccurate if the signal is noisy or has digitization artifacts. In these cases it may be preferable to fit the input trace with a polynomial and differentiate the polynomial to obtain a derivative – a fast-fourier transform filter can be used with similar results. One can choose both the window size and the order of the polynomial to fit, and also the weight that each data point gets. We use a polynomial routine that typically fits a third-order polynomial over 10 points with a Gaussian weighting that vanishes at the edges of the window. The Gaussian weighting helps to keep the derivatives continuous from one frame to the next. When fitting a polynomial to data, a rule of thumb is to never fit greater than a third order polynomial, so the window size must be set accordingly. In general, I use finite differences first and the polynomial fit routine when I need

accurate derivatives.

G.2 Speckle-Visibility Spectroscopy of Avalanches

G.2.1 Avalanche Apparatus and Its Components

The avalanche apparatus has five major components. The funnel above stores the sand which falls through a sand valve and into the avalanche channel. The avalanche channel is below the sand valve and screwed into an optical table. It is where avalanches build and are generated, and it expunges outflowing sand into a receptacle. The other components of the apparatus have to do with how avalanches are observed. We observed the system using both light scattering and direct imaging of superficial layers. Light scattering is the primary method which we accomplish with a laser and linescan camera. For superficial particle tracking we use a high-speed square CCD camera and halogen lights.

The funnel sits atop the apparatus and holds sand to be flowed into the channel. At the base of the funnel is a sand valve which we have custom built. The sand valve has two stages. The first stage is a simple valve that can be opened or closed to coarsely control the flow. This valve cannot be closed too small or the sand will jam and stopped flowing. Thus, we require a second stage which uses a knife edge to divert only a small portion of the sand into the channel, and the rest goes into a waste tube. This sand valve arrangement is reliable down to 0.05 g/s depending on sand grain size. Below that the flow can fluctuate wildly due to arching above the knife edge, and this can greatly influence avalanche dynamics. If smaller more continuous flow rates are desired it will be necessary to design a better stand valve. One possibility to prevent jamming is to vibrate the sand valve in the hope that it will break arches as they form, but it has been our experience that to break arches the valve must be struck hard, and we are hesitant to add vibration which will in all likelihood affect the light scattering measurements. Another possibility is to pass the sand through a square array of rods. The sand is poured into the top of the array and naturally spreads out as it falls down. Then, one can choose from a wide range of flow rates by selecting a location at the bottom of the array.

The avalanche channel consists of two aluminum bars sandwiched between the glass walls with a non-static coating. The bar at the bottom is the channel floor and the bar at the top obstructs

sand falling from the valve so that it impacts the pile with a minimum amount of momentum. Aluminum and non-static plexiglass are necessary because the sand particles are glass spheres that will accumulate charge if rubbed against an insulating surface. This will cause them to stick to the walls, which is generally undesirable. The concept behind the channel is simple. As sand is added a pile builds between the plates to the maximum angle of stability. When it surpasses the maximum angle of stability an avalanche will flow until it exhausts the excess sand and stops at the angle of repose, allowing the process to repeat indefinitely. The clear glass walls can be easily seen through allowing us to take images with a camera from different angles.

The specifications for the laser are laid out in the user manual. Important safety and maintenance instructions are laid out there as well, and this should always be kept on hand and easily accessible. The laser is powered by a liquid cooled power coupling. This coupling can be either controlled by the computer through a serial connection, or can be controlled manually with the buttons and knobs on front. It is relatively simple and straightforward to control and details are in the manual. The power coupling is cooled by water passed through a temperature controller which must be set to 20 C. To enable the laser it must be powered on with the key turned to the standby mode until the display reads ready. Then the key can be turned which will start lasing in the laser head. Powering a laser on and off repeatedly will greatly increase wear and tear requiring costly maintenance. As a rule, the laser should only be turned off when it will not be used in the next week. Otherwise, it should be set to the standby position with the shutter closed. The power coupling sits below the optical table on the floor and is plugged into a standard wall outlet. A general rule when setting up laser optics is to keep the beam path in a single plane whenever possible. Thus, in our setup, the laser is directly in front of the channel and lined up with the camera. To shoot the laser into the channel it is deflected off of two mirrors aligned vertically on a post, the first to pick the laser up off the table and the second to bounce it down into the channel perpendicular to the plane of the pile. When desired, we expand the beam by passing it through an appropriate lens between the second mirror and the pile.

The beam spot is imaged with a 58 kHz Basler linescan camera. The camera has a 2×1024 array of 8-bit pixels. It is controlled by the computer and plugs into a national instruments data acquisition card. Like any other camera controlled by LabVIEW, it can be tested using the Measurement and Automation Explorer. The great advantage of this camera is that it is fast and allows

us to stream compact data to disk on-the-fly, which allows us to capture data continuously at 58 kHz for many hours. This enormous dynamical range is one of the greatest attributes of the current setup. While we could likely reduce our pixel noise with a more expensive camera, we have not found a camera that will allow us to run continuously streaming 58 kHz data for such long observation times (typically multiple hours).

G.2.2 Avalanche Data Acquisition with LabVIEW

Because the laser light is coherent, one sees an interference pattern when the light scatters off the sand particles. This interference pattern is often called a “speckle pattern.” The primary concept of speckle-visibility spectroscopy is that the sharpness of the speckle pattern for a given camera exposure duration depends on the motion of the light scattering particles. The technique is sensitive to decorrelation of the interference pattern, and therefore to motion on the order of fractions of the wavelength of the scattered light. Sand particles cover this distance quickly and the speckle pattern changes rapidly. Therefore, an extremely fast camera with short exposure duration, and a high-powered laser, are both necessary to see a sharp bright speckled pattern. There are many factors to consider when attempting to optimize the scattering optics, and these are discussed at great length in previous publications [71].

It can be shown that for a scattered Gaussian spectrum the autocorrelation of the electric field, i.e. the interference pattern, can be approximated from the first and second moments of the pixel intensities across the CCD array [71]. Therefore, we do not need to save every pixel value, only the average and mean square average for each exposure. This greatly reduces the total amount of data that must be stored and is the essential feature that allows us to observe the system for such long durations with such high time resolution. Another important factor is to compute the intensity moments across the CCD on-the-fly. This operation must be formed in $17.2 \mu\text{s}$, the capture rate. A program that computes these moments on-the-fly must be extremely efficient. We have developed such a program that is both capable of communicating with the camera and computing the moments on-the-fly at the maximum 58 kHz frame rate. The ability to compute the intensity moments on-the-fly is the greatest hurdle to upgrading the camera. After looking into it for some time, we were unable to find a camera that could capture data at such a high speed while simultaneously computing and storing intensity moments on-the-fly. The vast majority of cameras, while having

superior CCDs, are not designed to stream data this way, but rather store images in a memory buffer that must be downloaded afterwards. This limits the total observation time to a few seconds at best.

The intensity moments are necessary to compute the variance, which is a number that quantifies the sharpness of the speckle pattern as a function of exposure duration and is mathematically related to the dynamics of the scattering sand particles. One technical problem is that the average intensity of the CCD drifts significantly over time and the signal of the sand particles is superimposed on this electronic drift. Because the drift has to do with the laser and CCD, and has nothing to do with sand dynamics, it can be removed by computing the variance ratio. The variance ratio is the ratio of two different variances measured for two different exposure durations. Because the drift is hardware electronic in nature, it is the same however we compute synthetic exposures, whether we compute a synthetic exposure duration of T or $2T$, but the sharpness of the speckle pattern is not. Therefore, by computing the ratio of the variance for two different exposure durations we cancel out electronic drift leaving signal from sand motion intact. This can be easily tested by looking at raw variance time traces for static speckle patterns (a static sand pile). While the variance for a given duration will drift, the variance ratio for two durations computed synthetically is constant.

With the variance ratio in hand we can transform using theory for back scattered light into the fluctuation speed δv [71]. This transformation is a mathematical procedure that we do numerically. It amounts to inverting a rather complicated ratio of exponentials. One method of handling this complicated function is to approximate it by a truncated power series that can be more easily be inverted. We have developed a different numerical algorithm to invert the original function that is much faster than applying the Newton solver. The idea is simple. One specifies the approximate desired accuracy. Then the program plugs a corresponding number of inputs into the original function and obtains a corresponding number of outputs. It then searches the output array for the given variance ratio and thereby finds the δv to which it corresponds. This simple method works because the original function, while mathematically complicated, is monotonic and smooth.

For fast sand fluctuations the variance ratio is typically small and corresponds to the range when the function saturates to a constant, which can make it impossible to properly invert. The error cannot be reduced by using a different inversion method, such as the Newton solver, since it has to do with the form of the function for small variance ratio values. The only option is to

change the form of the function by choosing a larger synthetic exposure ratio. At the moment we use a ratio of $V(2T)/V(T)$, and the effect could be reduced for larger synthetic exposures, such as $V(4T)/V(T)$. Unfortunately, to compute longer synthetic exposures would require the computation of higher order moments on-the-fly from the CCD, which is not easily accomplished for our desired frame rate and with current limitations in computer and camera technology and their communication speed.

G.3 LabVIEW Miscellaneous

LabVIEW is a full-fledged programming language and as a result many habits one may have learned by programming in languages like C are directly applicable to LabVIEW programming. However, because it uses a graphical interface and block diagram programming environment, it is relatively simple to learn, and many novice programmers become proficient before they have an adequate understanding of how the compiler interprets their code. In no place is this more evident than in the handling of memory. Like Java, LabVIEW handles memory automatically and the compiler does its best to minimize the number of operations and data copies that must be made in order to complete a task. Nevertheless, the programmer has the ultimate say over how data flows and the compiler can only do so much to overcome inherent design flaws. Therefore, it is strongly recommend that novice LabVIEW programmers with some working know-how, read the programming to manual available on the National Instrument web site and pay special attention to the section on memory handling. Like any other language, simple changes in programming style can greatly influence the speed, efficiency, and readability of LabVIEW code.

The nitty-gritty details of programming in LabVIEW can only be learned by experience, but it is wise to follow standard production techniques developed for languages like C. Chief among these is the concept of modular design. Modular design has been a longtime priority to the developers of the LabVIEW language. The basic concept is that large complicated tasks should be broken down into an organized set of smaller tasks that can each be accomplished by a relatively simple program. In LabVIEW this amounts to creating “subVIs” whenever applicable. A rule of thumb that was told to me by Paul Dixon when he first introduced me to LabVIEW is, “if you can’t see the whole program on your screen at once then you need to break it down more.” Programs

written with modular design in mind, in any language, are far more likely to be understood by yourself and especially by others attempting to read your code. Thus, such programs are much more likely to be reused and incorporated into larger projects in the future, and this greatly reduces the amount of work that must be done to develop new code to handle new tasks. But undoubtedly the greatest advantage of modular design is that simplifies even fairly sophisticated programs and will make them less buggy and will give you far greater confidence in the results. Such confidence is absolutely essential to independent scientific research because likely you will be the only person who has a working knowledge of your programs and they will be the primary tool that you use to capture and analyze your data.

The only programming detail that I will go into here is how to apply modular design to the directory structure of the programs you write, because I have found it to be an extremely useful technique to writing clean reusable code. The rule is very simple. The directory that contains the program should model in structure the structure of the program itself. Every program or subroutine that you create, has its own folder. I follow the additional convention that this folder be named the same name as the program that owns it. If a program contains many smaller subroutines, which is a hallmark of modular design, then each subroutine gets its own folder in the main program directory, and the tree continues down as far as it must, like a Russian nesting doll. This directory structure greatly simplifies incorporating old programs as a small part of new larger programs, since all one must do is copy the main program folder to the new program directory. In this way, your code will become stackable and you will see that over time your personal code library will grow in sophistication while always remaining flexible and easily understood.

Bibliography

- [1] A. R. Abate and D. J. Durian. The partition of energy for gas-fluidized grains. *Phys. Rev. E* **72**, 031305 (2005).
- [2] A. R. Abate and D. J. Durian. Approach to jamming in an air-fluidized granular bed. *Phys. Rev. E* **74**(3), 031308 Sep (2006).
- [3] E. Achenbach. Vortex shedding from spheres. *J. Fluid Mech.* **62**(2), 209–221 (1974).
- [4] M. Aichele, Y. Gebremichael, F. W. Starr, J. Baschnagel, and S. C. Glotzer. Polymer-specific effects of bulk relaxation and stringlike correlated motion in the dynamics of a supercooled polymer melt. *J. Chem. Phys.* **119**, 5290–5304 (2003).
- [5] R. A. Bagnold. Experiments on a gravity-free dispersion of large solid spheres in a newtonian fluid under shear. *Proc. R. Soc. London Ser. A* **225**, 49–63 (1954).
- [6] G. Baxter and J. Olafsen. Gaussian statistics in granular gases. *Nature* **425**, 680 (2003).
- [7] L. Bellon, S. Ciliberto, and C. Laroche. Violation of the fluctuation-dissipation relation during the formation of a colloidal glass. *Europhys. Lett.* **53**(4), 511–17 (2001).
- [8] C. Bennemann, C. Donati, J. Baschnagel, and S. C. Glotzer. Growing range of correlated motion in a polymer melt on cooling towards the glass transition. *Nature* **399**, 246–249 (1999).
- [9] L. Berthier and J. L. Barrat. Shearing a glassy material: numerical tests of nonequilibrium mode-coupling approaches and experimental proposals. *Phys. Rev. Lett.* **89**(9), 095702 (2002).

- [10] L. Berthier, G. Biroli, J. P. Bouchaud, L. Cipelletti, D. E. Masri, D. L'Hôte, F. Ladieu, and M. Perino. Direct experimental evidence of a growing length scale accompanying the glass transition. *Science* **310**, 1797–1800 (2005).
- [11] T. Börzsönyi, T. C. Halsey, and R. E. Ecke. Two scenarios for avalanche dynamics in inclined granular layers. *Phys. Rev. Lett.* **94**(20), 208001 May (2005).
- [12] G. S. Cargill. Dense random packing of hard spheres as a structural model for noncrystalline metallic solids. *J. Appl. Phys.* **41**(5), 2248–50 (1970).
- [13] M. E. Cates, J. P. Wittmer, J. P. Bouchaud, and P. Claudin. Jamming, force chains, and fragile matter. *Phys. Rev. Lett.* **81**(9), 1841–4 (1998).
- [14] J. C. Crocker and D. G. Grier. Methods of digital video microscopy for colloidal studies. *J. Col. I. Sci.* **179**, 298–310 (1996).
- [15] L. F. Cugliandolo and J. Kurchan. Analytical solution of the off-equilibrium dynamics of a long-range spin-glass model. *Phys. Rev. Lett.* **71**(1), 173–6 (1993).
- [16] G. D'Anna, P. Mayor, A. Barrat, V. Loreto, and F. Nori. Observing brownian motion in vibration-fluidized granular matter. *Nature* **424**, 909–912 (2003).
- [17] O. Dauchot, G. Marty, and G. Biroli. Dynamical heterogeneity close to the jamming transition in a sheared granular material. *Phys. Rev. Lett.* **95**(265701) (2005).
- [18] C. Donati, J. F. Douglas, W. Kob, S. J. Plimpton, P. H. Poole, and S. C. Glotzer. Stringlike cooperative motion in a supercooled liquid. *Phys. Rev. Lett.* **80**, 2338–2341 (1998).
- [19] C. Donati, S. Franz, S. C. Glotzer, and G. Parisi. Theory of non-linear susceptibility and correlation length in glasses and liquids. *J. Non-Cryst. Solids* **307**, 215–224 (2002).
- [20] C. Donati, S. C. Glotzer, P. H. Poole, W. Kob, and S. J. Plimpton. Spatial correlations of mobility and immobility in a glass-forming lennard-jones liquid. *Phys. Rev. E* **60**, 3107–3119 (1999).
- [21] M. T. Dove. *Introduction to Lattice Dynamics*. Cambridge University Press New York (1993).

- [22] J. Duran. *Sands, powders, and grains: An introduction to the physics of granular materials*. Partially ordered systems. Springer New York (2000).
- [23] S. R. N. E. I. Corwin, H. M. Jaeger. Structural signature of jamming in granular media. *Nature* **435**, 1075 – 1078 (2005).
- [24] M. D. Ediger. Spatially heterogeneous dynamics in supercooled liquids. *Annu. Rev. Phys. Chem.* **51**, 99–128 (2000).
- [25] M. D. Ediger, C. Angell, and S. Nagel. Supercooled liquids and glasses. *Journal of Physical Chemistry* **100**(31), 13200–13212 (1996).
- [26] D. A. Egolf. Equilibrium regained: From nonequilibrium chaos to statistical mechanics. *Science* **287**, 101–104 (2000).
- [27] K. Feitosa and N. Menon. Breakdown of energy equipartition in a 2d binary vibrated granular gas. *Phys. Rev. Lett.* **88**(19), 198301 (2002).
- [28] J. P. Garrahan and D. Chandler. Geometrical explanation and scaling of dynamical heterogeneities in glass forming systems. *Phys. Rev. Lett.* **89**, 035704 (2002).
- [29] Y. Gebremichael, F. Starr, T. Schroeder, and S. Glotzer. Spatially correlated dynamics in a simulated glass-forming polymer melt: Analysis of clustering phenomena. *Phys. Rev. E* **64**, 051503 (2001).
- [30] Y. Gebremichael, M. Vogel, and S. C. Glotzer. Particle dynamics and the development of string-like motion in a simulated monoatomic supercooled liquid. *J. Chem. Phys.* **120**, 4415–4427 (2004).
- [31] Y. Gebremichael, M. Vogel, and S. Glotzer. String-like correlated particle motion in a one-component, supercooled liquid. *J. Chem. Phys.* **120**(9) (2004).
- [32] S. C. Glotzer, V. N. Novikov, and T. B. Schroder. Time-dependent, four-point density correlation function description of dynamical heterogeneity and decoupling in supercooled liquids. *J. Chem. Phys.* **112**, 509–512 (2000).

- [33] W. Gotze. Properties of the glass instability treated within a mode-coupling theory. *Z. Phys. B: Condens. Matter* **60**, 195–203 (1985).
- [34] T. S. Grigera and N. E. Israeloff. Observation of fluctuation-dissipation-theorem violations in a structural glass. *Phys. Rev. Lett.* **83**(24), 5038–5041 (1999).
- [35] D. Herisson and M. Ocio. Fluctuation-dissipation ratio of a spin glass in the aging regime. *Phys. Rev. Lett.* **88**(25), 257202 (2002).
- [36] N. Hoffmann, F. Ebert, C. N. Likos, H. Löwen, and G. Maret. Partial clustering in binary two-dimensional colloidal suspensions. *Phys. Rev. Lett.* **97**(7), 078301 Aug (2006).
- [37] P. C. Hohenberg and B. I. Shraiman. Chaotic behavior of an extended system. *Physica D* **37**(1-3), 109–115 (1989).
- [38] M. M. Hurley and P. Harrowell. Non-gaussian behavior and the dynamical complexity of particle motion in a dense two-dimensional liquid. *J. Chem. Phys.* **105**(23), 10521–10526 (1996).
- [39] I. Ippolito, C. Annic, J. Lemaitre, L. Oger, and D. Bideau. Granular temperature: experimental analysis. *Phys. Rev. E* **52**(2), 2072–5 (1995).
- [40] H. M. Jaeger, S. R. Nagel, and R. P. Behringer. Granular solids, liquids, and gases. *Rev. Mod. Phys.* **68**(4), 1259–73 (1996).
- [41] A. Kasper, E. Bartsch, and H. Sillescu. Self-diffusion in concentrated colloid suspensions studied by digital video microscopy of core-shell tracer particles. *Langmuir* **14**(18), 5004–5010 (1998).
- [42] W. K. Kegel and A. van Blaaderen. Direct observation of dynamical heterogeneities in colloidal hard-sphere suspensions. *Science* **287**(5451), 290–293 (2000).
- [43] A. S. Keys, A. R. Abate, S. C. Glotzer, and D. J. Durian. Spatially heterogeneous dynamics in a granular system of steel ball bearings.
- [44] W. Kob, C. Donati, S. J. Plimpton, P. H. Poole, and S. C. Glotzer. Dynamical heterogeneities in a supercooled lennard-jones liquid. *Phys. Rev. Lett.* **79**(15), 2827–30 (1997).

- [45] N. Lacey, F. W. Starr, T. B. Schroder, and S. C. Glotzer. Spatially heterogeneous dynamics investigated via a time-dependent four-point density correlation function. *J. Chem. Phys.* **119**(14) (2003).
- [46] N. Lacey, F. Starr, T. Schroder, and S. Glotzer. Spatially heterogeneous dynamics investigated via a time-dependent four-point density correlation function. *J. Chem. Phys.* **119**(14), 7372–7387 (2003).
- [47] J. S. Langer. Dynamics and thermodynamics of the glass transition. *Phys. Rev. E* **73**(4), 041504 Apr (2006).
- [48] P. A. Lemieux and D. J. Durian. From avalanches to fluid flow: a continuous picture of grain dynamics down a heap. *Phys. Rev. Lett.* **85**(20), 4273–6 (2000).
- [49] A. J. Liu and S. R. Nagel. Jamming is not just cool any more. *Nature* **396**(6706), 21–22 (1998).
- [50] A. J. Liu and S. R. Nagel, editors. *Jamming and Rheology: Constrained Dynamics on Microscopic and Macroscopic Scales*. Taylor and Francis NY (2001).
- [51] W. Losert, L. Bocquet, T. C. Lubensky, and J. P. Gollub. Particle dynamics in sheared granular matter. *Phys. Rev. Lett.* **85**(7), 1428–31 (2000).
- [52] R. B. M. D. P. Y. T. P. S. M. B. Stone, D. P. Bernstein. Stress propagation: Getting to the bottom of a granular medium. *Nature* **427**, 503–504 (2004).
- [53] W. N. Man, A. Donev, F. H. Stillinger, M. T. Sullivan, W. B. Russel, D. Heeger, S. Inati, S. Torquato, and P. M. Chaikin. Experiments on random packings of ellipsoids. *Phys. Rev. Lett.* **94**(198001) (2005).
- [54] A. H. Marcus, J. Schofield, and S. A. Rice. Experimental observations of non-gaussian behavior and stringlike cooperative dynamics in concentrated quasi-two-dimensional colloidal liquids. *Phys. Rev. E* **60**(5), 5725–5736 (1999).
- [55] G. Marty and O. Dauchot. Subdiffusion and cage effect in a sheared granular material. *Phys. Rev. Lett.* **94**, 15701 (2005).

- [56] N. Menon and D. J. Durian. Diffusing-wave spectroscopy of dynamics in a three-dimensional granular flow. *Science* **275**(5308), 1920–1922 (1997).
- [57] F. Moucka and I. Nezbeda. Detection and characterization of structural changes in the hard-disk fluid under freezing and melting conditions. *Phys. Rev. Lett.* **94**(4), 040601 (2005).
- [58] D. M. Mueth, G. F. Debregeas, G. S. Karczmart, P. J. Eng, S. R. Nagel, and H. M. Jaeger. Signatures of granular microstructure in dense shear flows. *Nature* **406**(6794), 385–9 (2000).
- [59] R. M. Nedderman. *Statics and kinematics of granular materials*. Cambridge University Press NY (1992).
- [60] C. S. O’Hern, S. A. Langer, A. J. Liu, and S. R. Nagel. Random packings of frictionless particles. *Phys. Rev. Lett.* **88**(7), 075507 (2002).
- [61] C. S. O’Hern, L. E. Silbert, A. J. Liu, and S. R. Nagel. Jamming at zero temperature and zero applied stress: The epitome of disorder. *Phys. Rev. E* **68**(1), 011306 (2003).
- [62] C. O’Hern, S. Langer, A. Liu, and S. Nagel. Force distributions near jamming and glass transitions. *Phys. Rev. Lett.* **86**, 111–114 (2001).
- [63] R. P. Ojha, A. R. Abate, and D. J. Durian. Statistical characterization of the forces on spheres in an upflow of air. *Phys. Rev. E* **71**(1), 016313 (2005).
- [64] R. P. Ojha, P. A. Lemieux, P. K. Dixon, A. J. Liu, and D. J. Durian. Statistical mechanics of a gas-fluidized particle. *Nature* **427**(6974), 521–523 (2004).
- [65] J. S. Olafsen and J. S. Urbach. Clustering, order, and collapse in a driven granular monolayer. *Phys. Rev. Lett.* **81**(20), 4369–72 (1998).
- [66] J. S. Olafsen and J. S. Urbach. Two-dimensional melting far from equilibrium in a granular monolayer. *Phys. Rev. Lett.* **95**(9), 098002 (2005).
- [67] I. K. Ono, C. S. O’Hern, D. J. Durian, S. A. Langer, A. J. Liu, and S. R. Nagel. Effective temperatures of a driven system near jamming. *Phys. Rev. Lett.* **89**(9), 095703 (2002).
- [68] F. H. S. P. G. Debenedetti. Supercooled liquids and the glass transition. *Nature* **410**, 259–267 (2001).

- [69] B. Pouligny, R. Malzbender, P. Ryan, and N. A. Clark. Analog simulation of melting in two dimensions. *Phys. Rev. B* **42**(1, pt.B), 988–91 (1990).
- [70] J. J. Prentis. Experiments in statistical mechanics. *Am. J. Phys.* **68**(12), 1073–1083 (2000).
- [71] S. S. S. P. K. D. D. J. D. R. Bandyopadhyay, A. S. Gittings. Speckle-visibility spectroscopy: A tool to study time-varying dynamics. *Rev. Sci. Instrum.* **79**, 093110 (2005).
- [72] A. Rahman. Correlations in the motion of atoms in liquid argon. *Phys. Rev.* **136**, A405–11 (1964).
- [73] C. Reichhardt and C. J. Olson Reichhardt. Fluctuating topological defects in 2d liquids: Heterogeneous motion and noise. *Phys. Rev. Lett.* **90**(9), 095504 Mar (2003).
- [74] C. Reichhardt and C. J. Olson Reichhardt. Noise at the crossover from wigner liquid to wigner glass. *Phys. Rev. Lett.* **93**(17), 176405 Oct (2004).
- [75] P. Reis, R. Ingale, and M. Shattuck. Crystallization of a quasi-two-dimensional granular fluid. *Phys. Rev. Lett.* **96**, 258001 (2006).
- [76] G. Reydellet, F. Rioual, and E. Clement. Granular hydrodynamics and density wave regimes in a vertical chute experiment. *Europhys. Lett.* **51**(1), 27–33 (2000).
- [77] P. N. Segre, F. Liu, P. B. Umbanhowar, and D. A. Weitz. An effective gravitational temperature for sedimentation. *Nature* **409**, 594–597 (2001).
- [78] L. E. Silbert, A. J. Liu, and S. R. Nagel. Vibrations and diverging length scales near the unjamming transition. *Phys. Rev. Lett.* **95**(9), 098301 (2005).
- [79] L. E. Silbert, A. J. Liu, and S. R. Nagel. Structural signatures of the unjamming transition at zero temperature. *Phys. Rev. E* **73**(4), 041304 (2006).
- [80] H. Sillescu. Heterogeneity at the glass transition: a review. *J. Non-Cryst. Solids* **243**, 81–108 (1999).
- [81] A. P. Sokolov, W. Steffen, and E. Rossler. Dynamics of glass-forming liquids at high temperatures: A general scenario. *J. Phys.: Condens. Matter* **8**, 9587–9591 (1996).

- [82] C. M. Song, P. Wang, and H. A. Makse. Experimental measurement of an effective temperature for jammed granular materials. *Proc. Nat. Acad. Sci.* **102**(7), 2299–2304 (2005).
- [83] J. D. Stevenson, J. Schmalian, and P. G. Wolynes. The shapes of cooperatively rearranging regions in glass-forming liquids. *Nature Physics* **2**, 268–274 (2006).
- [84] G. K. Suryanarayana and A. Prabhu. Effect of natural ventilation on the boundary layer separation and near-wake vortex shedding characteristics of a sphere. *Exp. Fluids* **29**, 582–591 (2000).
- [85] D. Thirumalai and R. D. Mountain. Activated dynamics, loss of ergodicity, and transport in supercooled liquids. *Phys. Rev. E* **47**(1), 479–489 (1993).
- [86] C. Toninelli, M. Wyart, L. Berthier, G. Biroli, and J. P. Bouchaud. Dynamical susceptibility of glass formers: Contrasting the predictions of theoretical scenarios. *Phys. Rev. E* **71**, 041505 (2005).
- [87] S. Torquato. *Random Heterogeneous Materials: Microstructure and Macroscopic Properties*. Springer NY (2001).
- [88] W. Van Meegen and S. M. Underwood. Tracer diffusion in concentrated colloidal dispersions: 2. non-gaussian effects. *J. Chem. Phys.* **88**(12), 7841–7846 (1988).
- [89] M. Vogel and S. C. Glotzer. Spatially heterogeneous dynamics and dynamic facilitation in a model of viscous silica. *Phys. Rev. Lett.* **92**, 255901 (2004).
- [90] M. Vogel and S. C. Glotzer. Temperature dependence of spatially heterogeneous dynamics in a model of viscous silica. *Phys. Rev. E* **70**, 061504 (2004).
- [91] E. R. Weeks, J. C. Crocker, A. C. Levitt, A. Schofield, and D. A. Weitz. Three-dimensional direct imaging of structural relaxation near the colloidal glass transition. *Science* **287**(5453), 627–31 (2000).
- [92] E. R. Weeks and D. A. Weitz. Subdiffusion and the cage effect studied near the colloidal glass transition. *Chem. Phys.* **284**, 361–376 (2002).

- [93] E. Weeks and D. Weitz. Properties of cage rearrangements observed near the colloidal glass transition. *Phys. Rev. Lett.* **89**(095704) (2002).
- [94] E. R. Weeks, J. C. Crocker, and D. A. Weitz. Short and long range correlated motion observed in colloidal glasses and liquids. (2006).
- [95] H. R. Wendt and F. F. Abraham. Empirical criterion for glass-transition region based on monte-carlo simulations. *Phys. Rev. Lett.* **41**(18), 1244–1246 (1978).
- [96] M. Wyart, S. R. Nagel, and T. A. Witten. Geometric origin of excess low-frequency vibrational modes in weakly connected amorphous solids. *Europhys. Lett.* **72**(3), 486–492 (2005).
- [97] M. Wyart, L. E. Silbert, S. R. Nagel, and T. A. Witten. Effects of compression on the vibrational modes of marginally jammed solids. *Phys. Rev. E* **72**(5), 051306 (2005).
- [98] R. Yamamoto and A. Onuki. Dynamics of highly supercooled liquids: Heterogeneity, rheology, and diffusion. *Phys. Rev. E* **58**, 3515–3529 (1998).