

Particle-Based Simulation of Granular Materials

Nathan Bell¹, Yizhou Yu¹ and Peter J. Mucha²

¹University of Illinois at Urbana-Champaign

²Georgia Institute of Technology

Abstract

Granular materials, such as sand and grains, are ubiquitous. Simulating the 3D dynamic motion of such materials represents a challenging problem in graphics because of their unique physical properties. In this paper we present a simple and effective method for granular material simulation. By incorporating techniques from physical models, our approach describes granular phenomena more faithfully than previous methods. Granular material is represented by a large collection of non-spherical particles which may be in persistent contact. The particles represent discrete elements of the simulated material. One major advantage of using discrete elements is that the topology of particle interaction can evolve freely. As a result, highly dynamic phenomena, such as splashing and avalanches, can be conveniently generated by this meshless approach without sacrificing physical accuracy. We generalize this discrete model to rigid bodies by distributing particles over their surfaces. In this way, two-way coupling between granular materials and rigid bodies is achieved.

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Physically based modeling I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Animation I.6.8 [Simulation and Modeling]: Types of Simulation—Animation

1. Introduction

Granular materials, such as sand, soil, powders and grains, present an indispensable part of the real world. They are simply very large ensembles of macroscopic particles [JNB96]. Given their ubiquity, there is an inevitable need for granular material simulation for graphics-related applications such as building virtual environments and animating natural phenomena. In fact, there have been a few inspirational techniques on this topic in graphics [LM93, CLH96, SOH99, ON03]. However, these techniques usually do not have sufficient spatial or physical accuracy to describe important granular phenomena. More mature techniques for 3D environments still need to be developed. In this paper, we aim to faithfully perform 3D dynamic simulations of granular materials. In such simulations, we not only need to simulate the interactions, such as collisions and friction, among the particles, but also the interactions between granular particles and other larger-scale objects.

In addition to their role in graphics, granular materials have profound industrial significance. Indeed, the grinding of particles and ores alone accounts for an estimated 1.3% of U.S. electrical power consumption [JGD94]. Hence there

is considerable interest in developing models to better understand granular behavior. Different from fluids and other materials, granular materials exhibit unique physical properties which demand novel simulation techniques. For example, a sand pile at rest with a slope lower than the angle of repose behaves like a solid, i.e., the material remains at rest even though gravitational forces create stresses on its surface and in its interior. If the pile is tilted a few degrees above the angle of repose, grains start to flow. However, the main flow exists in a boundary layer at the pile's surface with only limited motion in the interior. As another example, when a granular material is held in a cylindrical container, the pressure at the bottom of the container does not necessarily—as in fluids—increase indefinitely as the height of the material is increased [dG99]. Instead, for a sufficiently tall column, the pressure saturates to a maximum value independent of the height. Because of the contact forces among grains and static friction with the sides of the container, the container walls support the extra weight. It is this feature that allows the sand in an hourglass to flow at a nearly constant rate. Clearly, an accurate model for granular materials must be able to reliably reproduce such phenomena.



Figure 1: Friction forces produce a constant rate of flow through the hourglass.

In this paper, we present a simple and effective simulation method built upon both theoretical and experimental results in physics. We adopt the practice of modeling granular material with discrete elements represented by particles [CS79, HL98], while selecting interparticle interactions based in part on computational cost. This approach permits faithful reproduction of a wide range of both static and dynamic granular phenomena. Furthermore, the discrete nature of our method obviates the need for an explicit surface representation, thereby allowing the simulation of sparse events. Interactions among particles are governed by a molecular-dynamics based contact model. Contact forces are obtained explicitly from the relative velocity and overlap between particles. The formulations of these forces are derived from elasticity theory and experimental results. One major advantage of this contact model is that large numbers of bodies in persistent contact are handled efficiently.

We extend this contact model to rigid bodies by covering their surfaces with particles. A signed distance representation is used to produce a uniform distribution of particles and allow particles to be placed at an offset from the original mesh. The force and torque accumulated from all these surface particles are used to integrate the object through rigid body motion. This produces true two-way coupling between granular particles and rigid bodies. As a by-product, interactions between different rigid bodies can be simulated using the same particle-based approach. Compelling demonstrations for these types of interactions have been produced.

1.1. Related Work in Graphics

Using particles to simulate natural phenomena has a long history in computer graphics. Early work, such as [Ree83],

developed techniques to animate and render irregular objects with particle systems. However, it was the introduction of particle systems with pairwise interactions that made it possible to simulate fluids and deformable objects. [MP89] simulated viscous fluids with particle interactions based on the Lennard-Jones potential force. [TPF89] paired particles to better simulate deformable objects. [Ton91] improved particle motion by adding additional particle interactions based on heat transfer among particles. There have also been successful applications of more accurate particle-based fluid models in graphics. [DC99, MCG03] adapted *smoothed particle hydrodynamics* (SPH) [Mon94] to compute the motion of deformable substances and liquids. [HMT01] modified the formulation of SPH to simulate hair-hair interactions by modeling hair as a fluid-like continuum. Recently, another particle method, Moving Particle Semi-Implicit (MPS), was also developed [YKO96] along with its application to incompressible fluid simulation in graphics [PTB*03].

Although being effective for fluids and deformable objects, these techniques are not appropriate for granular materials for the following reasons. First, the formulations for particle collision were not specifically designed according to the real material properties of granular particles. Some of them were designed to reflect fluid properties, such as zero divergence (though the divergence in granular materials is typically small). Second, static friction among particles is missing from these methods. Only cohesive forces were used to keep particles together. However, static friction and cohesion give rise to different effects. Fluid viscosity in some of the methods only resembles dynamic friction.

Nevertheless, there have been a few inspirational techniques focusing on granular materials in graphics [LM93, CLH96, SOH99, ON03]. These techniques usually do not have sufficient spatial or physical accuracy. The techniques in [LM93, CLH96] were derived from physics, but they model soil as a height field which precludes the simulation of many interesting 3D effects, such as splashing. While the methods in [SOH99, ON03] can produce interesting granular behavior for footprints or sandboxes, their empirical nature prevents them from faithfully simulating high-speed scenarios where significant energy is exchanged among particles and objects.

While the proposed method is in many ways similar to [LHM95], there are a number of crucial differences. Firstly, our model for particle interaction is more rigorously developed than the damped linear spring force used in their paper. Secondly, the non-spherical particles in their method are assembled using stiff springs whereas we consider each grain as a rigid object. Thirdly, we extend the particle-based method to handle two-way coupling with rigid bodies. While their method can produce compelling animations using a small number of particles, we remark that there is no natural extension of their multi-scale aspect to three dimensions,

since particles cannot be assumed to be in contact with a ground plane.

1.2. Physical Models of Granular Materials

There are a number of computational models used to investigate granular materials. Event-driven (ED), or hard-sphere algorithms, are based on the calculation of changes from distinct collisions between particles (typically spheres or polyhedra). The efficiency and utility of ED methods rely on the assumptions of binary collisions, short contact durations, and relatively long intervals between consecutive events. Hence such methods are well suited to simulate rapid flows, whose behavior resembles that of a gas. However in many cases of interest some or all of these assumptions are violated. In systems with strong dissipation the time between collisions may vanish, resulting in so-called *inelastic collapse* [HL98]. Likewise, ED methods are inappropriate for systems with persistent contacts and frequent collisions such as heaps, hoppers, and shaken materials.

Since modeling individual particles explicitly is often infeasible (e.g. fine powders), there is strong interest in continuum models for granular media. Avalanches, landslides, and other large-scale phenomena can in principle be simulated using continuum techniques [AT01]. Unfortunately these techniques are relatively new and have not yet reproduced the rich small-scale details available in the results of their discrete counterparts.

Molecular Dynamics (MD) or soft-sphere methods are another common way to simulate granular materials. Like ED, MD methods are discrete, with particles also represented by spheres or polyhedra. In ED methods, contacts are resolved through an instantaneous change in the momentum of colliding particles, thereby enforcing non-interpenetration. In contrast, MD methods permit small overlap between particles to simulate the elastic deformation of material elements. This overlap is then used along with other physical quantities to determine an appropriate contact force. This distinction allows MD methods efficiently model systems with persistent contacts and frequent collisions. Moreover, MD methods have demonstrated a wide variety of both static and dynamic granular phenomena such as heap formation, patterns in vibrating layers, force chains, and particle size segregation [Lee94, SvHSvS04, DR99]. Therefore, the simulation technique in this paper is based on MD methods.

2. Contact Forces

MD based models are characterized by their handling of contact forces. Since our objects will be composed of collections of round particles, it suffices to describe the contact force between a single pair of spheres, computing the total force on a given sphere as the sum over its pairwise interactions. We first define the relevant physical quantities used in the contact models and then examine a few common force schemes.




Figure 2: Quantities used in collision modeling.

We divide the contact force $\vec{F} = \vec{F}_n + \vec{F}_t$ into normal and tangential components.

Let \vec{x}_1 and \vec{x}_2 be the centers of particles of radius R_1 and R_2 respectively. The forces described here are applied to the particle centered at x_1 . An equal and opposite force is applied to the particle centered at x_2 . Define the overlap ξ of the particles and the normalized line of centers \vec{N} as follows:

$$\xi = \max(0, R_1 + R_2 - \|\vec{x}_2 - \vec{x}_1\|), \quad (1)$$

$$\vec{N} = \frac{\vec{x}_2 - \vec{x}_1}{\|\vec{x}_2 - \vec{x}_1\|}. \quad (2)$$

Let \vec{v}_1 and \vec{v}_2 denote the velocity of each particle at the contact point and define the relative velocity \vec{V} , relative velocity in the normal direction $\dot{\xi}$, and tangential velocity \vec{V}_t :

$$\vec{V} = \vec{v}_1 - \vec{v}_2, \quad (3)$$

$$\dot{\xi} = \vec{V} \cdot \vec{N}, \quad (4)$$

$$\vec{V}_t = \vec{V} - \dot{\xi} \vec{N}. \quad (5)$$

2.1. Normal Forces

In this section we describe a number of frequently used normal forces and their properties. We restrict our attention to force rules of the form

$$\vec{F}_n = f_n \vec{N}, \quad (6)$$

$$f_n + k_d \xi^\alpha \dot{\xi} + k_r \xi^\beta = 0, \quad (7)$$

with viscous damping coefficient k_d and elastic restoration coefficient k_r . Particle stiffness is controlled by k_r and dissipation during collisions by k_d . In the simplest case we choose $\alpha = 0, \beta = 1$ and obtain a damped linear spring force,

$$f_n + k_d \dot{\xi} + k_r \xi = 0. \quad (8)$$

Here k_r and k_d can be chosen to match experimental data, or manually tuned to produce the desired response. Given a coefficient of normal restitution ϵ_n and contact duration t_c the corresponding coefficients can be determined:

$$k_d = 2 m_{\text{eff}} \frac{-\ln \epsilon}{t_c}, \quad (9)$$

$$k_r = \frac{m_{\text{eff}}}{t_c^2} (\ln^2 \varepsilon + \pi^2), \quad (10)$$

$$\frac{1}{m_{\text{eff}}} = \frac{1}{m_1} + \frac{1}{m_2}. \quad (11)$$

Here m_{eff} represents the reduced mass, given particle masses of m_1 and m_2 respectively. While the linear model is adequate for many purposes, more accurate options exist. Drawing from elasticity theory, we can choose $\beta = 3/2$ which corresponds to the Hertz contact force between spheres [LL86],

$$f_n + k_d \xi^\alpha \dot{\xi} + k_r \xi^{3/2} = 0. \quad (12)$$

Hertz theory relates the stiffness coefficient k_r to the material properties of the spheres in contact. Let E_1 and E_2 be the Young's moduli and μ_1 and μ_2 the Poisson ratios of the two spheres. Now define the effective radius R_{eff} , Young's modulus E_{eff} , and corresponding stiffness coefficient k_r :

$$\frac{1}{R_{\text{eff}}} = \frac{1}{R_1} + \frac{1}{R_2}, \quad (13)$$

$$\frac{1}{E_{\text{eff}}} = \frac{1 - \mu_1^2}{E_1} + \frac{1 - \mu_2^2}{E_2}, \quad (14)$$

$$k_r = \frac{4}{3} E_{\text{eff}} \sqrt{R_{\text{eff}}}. \quad (15)$$

So far, we have left α unspecified in the non-linear model. In the linear model ($\alpha = 0, \beta = 1$), the coefficient of normal restitution ε_n was constant for a given choice of parameters. However both experimental and theoretical results demonstrate a dependence of ε_n on the impact velocity [RPBS99, SP98]. These results suggest a relationship like $(1 - \varepsilon_n) \propto v^{1/5}$ [SDW96]; qualitatively ε_n decreases with increasing impact velocity. In the non-linear model, different values of α give rise to different dependencies. Figure 3 demonstrates this relationship for common parameter pairs. Constant ε_n are obtained in both the linear model and the non-linear model with $\alpha = 1/4$. Using $\alpha = 0, \beta = 3/2$ has the unnatural effect of producing more elastic collisions with increasing velocity. Of the four, the model that best agrees with experimental results is $\alpha = 1/2, \beta = 3/2$. This choice is comprehensively studied in [SP98] and [RPBS99].

2.2. Shear Forces

If desired, one can introduce a force to account for shear friction between particles. The most basic approach simply defines a force opposing the tangential velocity.

$$\vec{F}_t = -k_t \vec{V}_t \quad (16)$$

Here k_t is a viscous damping term not unlike k_d . This rule can be problematic since $\|\vec{F}_t\|$ is not limited by the normal force. A contrasting option accounts for this Coulomb law

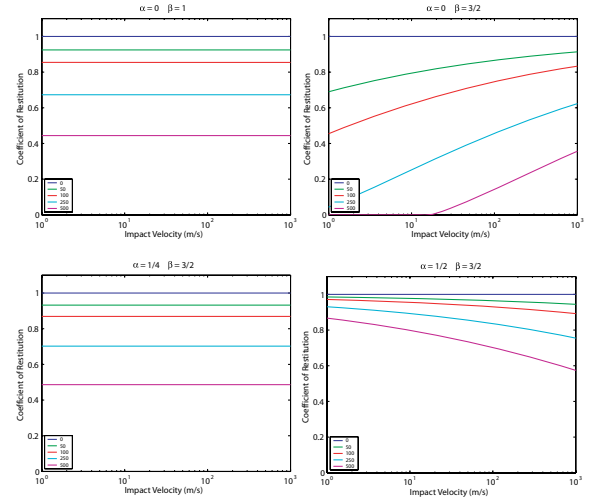


Figure 3: The coefficient of normal restitution is independent of the velocity in both the linear model $\alpha = 0, \beta = 1$ and non-linear model with $\alpha = 1/4$. The non-linear model with $\alpha = 1/2$ becomes more dissipative with increasing velocity, in agreement with experimental results. However choosing $\alpha = 0, \beta = 3/2$ produces the opposite behavior. In all tests $k_r = 1.0E6$ and $k_d = \{0, 50, 100, 250, 500\}$.

for friction coefficient μ and normal force f_n ,

$$\vec{F}_t = -\mu f_n \frac{\vec{V}_t}{\|\vec{V}_t\|}. \quad (17)$$

The numerical difficulties encountered by the discontinuity at $\|\vec{V}_t\| = 0$ can be alleviated by combining (16) with (17):

$$\vec{F}_t = -\min(\mu f_n, k_t \|\vec{V}_t\|) \frac{\vec{V}_t}{\|\vec{V}_t\|}. \quad (18)$$

Here the magnitude of k_t affects oblique impacts and should be large enough that the behavior of force (18) resembles that of (17) [SDW96].

The shear friction forces discussed above can only slow movement in the tangent direction, not stop or reverse it. These rules are inadequate for applications that require truly static friction, such as heap formation, angles of repose, and the quasi-static stages of avalanches. In such situations, there is a threshold force below which the grains do not move at all, opposed by static friction originating in the even smaller scale surface interactions of particle contact. However, implementation of a simple threshold rule is history dependent and computationally intensive. A common solution is to introduce a virtual spring when particles first touch [LH93]. While the particles remain in contact, the *total* displacement \vec{D} from the initial contact point is used to compute the shear

force:

$$\vec{F}_t = -\min(\mu f_n, k_t \|\vec{D}\|) \frac{\vec{D}}{\|\vec{D}\|}. \quad (19)$$

Another approach purely relies on particle geometry, using non-spherical particles to create the necessary roughness for static friction. Non-spherical particles are often constructed of polygons [PB95] or spheres connected through springs [PB93], though the latter can become costly for sufficiently stiff springs. The spatial configuration of these spheres or polygons represents the microstructure of non-spherical particles. No explicit static friction force need be used for these compound non-spherical objects. When two non-spherical particles are in contact with each other, the normal forces between their rough structures produce a resistance to relative tangential motion, thus creating actual effects of static friction. In order to draw upon the existing theory for spherical particles, we model non-spherical particles with collections of spheres, avoiding overly stiff interactions by replacing internal springs with rigid motion constraints. With static friction obtained via normal forces between compounds of spheres, we can then limit our attention to dynamic shear forces as expressed in (18).

3. Granular Material Simulation

We directly model a granular material as a large collection of granular particles which may be in persistent contact when the material is static. However, new external forces may cause some or all of the particles to have relative motion, leading to momentum and energy exchange through collisions. The aforementioned contact forces are used for computing interactions between both static particle configurations and dynamic assemblies.

3.1. Granular Particles

Our model of granular particles is inspired by the non-spherical particles in [PB93]. However, instead of using springs connecting round particles, we more simply model a grain as multiple round particles that are constrained to move together as a single rigid body. Such rigid collections




Figure 4: Non-spherical particles are used to accurately model static granular phenomena. Here tetrahedron and cube shapes are constructed from spheres.

of spheres have been considered in discrete element simulations before [WB93], but were unable with then-present computers to simulate the broad scope of granular phenomenon which can now be simulated.

The positions of the round particles are given as offsets from the body's center of mass. Let \vec{o}_i denote the offset of a given particle, R the current orientation matrix, and \vec{x} the center of mass of the grain. The current position of the particle is simply $\vec{x} + R \vec{o}_i$. Dissipation in the contact force model depends on the relative velocity at the point of contact \vec{p} . We can determine the velocity at the contact point from the velocity \vec{v} and angular velocity $\vec{\omega}$ of the grain using $\vec{v} + \vec{\omega} \times (\vec{p} - \vec{x})$. A contact force \vec{F} applied at a contact point \vec{p} contributes a torque $(\vec{p} - \vec{x}) \times \vec{F}$ to the body. The net force on a grain is simply the sum of the individual contact forces applied to its particles.

The motivations for such a non-spherical grain model are as follows. First, non-spherical particles produce a considerable angle of repose and exhibit stick-slip behavior, in accordance with experimental results [PB93]. Spherical particles can produce neither the same angle of repose nor stick-slip behavior without some explicit static friction force rule. While one can partially remedy this problem by applying equation (19) to round particles [LH93], non-spherical granular particles offer an arguably more natural solution. Second, we find that the bookkeeping associated with the virtual-springs in (19) is computationally expensive. Third, the idea of approximating objects with round particles extends naturally to general rigid bodies.

3.2. Contact Detection

In this section we briefly discuss an efficient method for identifying particle contacts. Since all our objects are compounds of spheres, contact detection reduces to finding overlapping spheres.

Clearly the naive $O(n^2)$ method of considering all pairs of particles is impractical for even small simulations. Assuming all particles are of similar size, we can apply spatial hashing to more efficiently resolve contacts. A voxel size two times the maximum particle radius guarantees that only particles belonging to adjacent voxels can be in contact. Hence the problem of finding all particle contacts reduces to exploring the 27 voxels surrounding a given particle. To exploit temporal coherence, the hash structure is kept between time steps and updated whenever a particle moves between voxels.

While this scheme provides $O(n)$ lookup of all potential contacts, performing 27 individual lookups per particle can dominate the overall computational cost. Further exploiting temporal coherence, we can reduce the cost to one lookup by inserting the particle into all 27 neighboring voxels. This approach biases the cost of maintaining the hash structure to updates, rather than queries. Justifying this preference, we note that particles typically move short distances during a




Figure 5: A steel ball collides with a sand pile producing a large splash.

single adaptively-integrated time step. Moreover, it is clear that particles in persistent contact in a slowly-moving collection are restricted by numerical stability to moving a fraction of a radius in a single time step. Hence it is very likely for a particle to remain within a voxel through many time steps. When using this method we observe a significant reduction in compute time. In simulations with large stagnant regions, the cost of contact detection is negligible.

3.3. Interaction with Rigid Bodies

In many situations, granular materials need to interact with rigid bodies. For example, they are excavated by rigid tools and transported in rigid containers. In a landslide, they destroy large structures while at other times, we may simply throw small objects at them to see a splash. Realistic interaction between granular materials and rigid bodies is thus an integral part of granular material simulation. In our method, we generate particles to completely cover the surface of a rigid object and simulate the interaction between granular particles and these surface particles using the contact forces described in Section 2. With this method we achieve efficient two-way coupling between granular particles and rigid bodies.

We create our granular particles from spheres, prescribing a few positions to create the desired shape. In the case of an arbitrary shape requiring hundreds or thousands of particles, such manual placement is not an option. We require a general method for representing shapes with particles. Given their ubiquity in graphics, we introduce a way to utilize bodies defined by triangle meshes.

3.3.1. Surface Particle Generation for Triangle Meshes

While one could place particles using the triangles directly [Tur92], we prefer to first construct a signed distance representation and then apply ideas from Witkin and Heckbert's [WH94] method for sampling implicit surfaces with particles. This approach has the advantage that one can easily place particles completely inside, outside, or at an arbitrary offset from the original mesh. The signed distance field is computed on a voxel grid of resolution comparable to the radius

of the particles to be placed. Construction of the signed distance field follows the algorithm described in [BA02]. Since only a narrow band of signed distances near the surface is required, this procedure is fast. The signed distance values are then stored in a hash table.

An optional offset is then taken and an iso-surface extracted using Marching Tetrahedra [TPG99]. This offset triangle mesh is only used for initializing particle positions. The initial positions are chosen randomly within each triangle and the number of particles is controlled by the triangle's area A , the particle radius R , and a user defined density value D . We first place $\lfloor DA/(\pi r^2) \rfloor$ particles inside the triangle and then one additional particle with probability $DA/(\pi r^2) - \lfloor DA/(\pi r^2) \rfloor$.

Once initialized, the particles are then permitted to float




Figure 6: Rigid bodies defined by triangle meshes are covered with particles. The signed distance representation permits placement of particles at an arbitrary offset from the mesh surface. Here a camel model is represented by 4147 spheres.



Figure 7: Structures swept away by an avalanche.



Figure 8: A bulldozer clears the remains of the avalanche.

along the signed distance field according to a simplified version of the interaction in [WH94]. Each of these "floaters" has a position \vec{P} and a velocity $\vec{V} = \vec{V}_F + \vec{V}_R$. Floaters are constrained to the zero-isosurface using a feedback term, $\vec{V}_F = -\phi(\vec{P})\vec{N}$, where $\phi(\vec{P})$ and \vec{N} are the signed distance and normal at \vec{P} respectively. These values are obtained from the signed distance field through trilinear interpolation. Since a uniform distribution of particles is desired, a simple repulsion model is also used.

$$\vec{V}_R = \sum_{P_i \in S \setminus P} K(\|\vec{P} - \vec{P}_i\|) \frac{\vec{P} - \vec{P}_i}{\|\vec{P} - \vec{P}_i\|}. \quad (20)$$

Here K is a kernel function with compact support and S is the set of floaters where K is non-zero. The floaters are integrated through time until the convergence criteria are satisfied. The final positions of the floaters are then fixed on the rigid body during subsequent simulations with granular particles. We find that this approach produces natural coverings for a wide range of models. Figure 6 shows the result of this procedure applied to a camel model.

3.4. Time Integration

Integration through time is accomplished through rigid body evolution [Bar97]. Mass properties for granular particles are obtained from their constituent spheres. In the case of rigid bodies, the defining triangle mesh is used to derive the relevant quantities [Mir96]. Since the time step required for numerical stability can vary greatly during simulation, an adaptive ODE integrator is a necessity. We choose to solve the equations of motion using an adaptive Runge-Kutta-Fehlberg method (RKF45) [Lam91]. This choice would perhaps seem unnatural, given the discontinuities in the lower-order derivatives of motion when sphere pairs transition into and out of contacts. Nevertheless, test simulations on different flows found that RKF45 was more efficient than lower-order methods at typically-specified error tolerances, perhaps because these discontinuities are at isolated points in both space and time. Before computing contact forces, the hash structure described in 3.2 is updated to reflect the cur-

rent particle positions. After forces have been computed the relevant derivatives are returned to the integrator.

4. Results and Discussion

We have applied our method to several large-scale simulations. Figure 5 shows a steel ball colliding with a sand pile, represented by 45,494 tetrahedron-shaped particles. Blocked by a glass wall, the wave of granular material is halted, and produces a heap near the base of the wall. Here the importance of accurate modeling of both static and dynamic phenomena is made clear as the granular material transitions rapidly from static to dynamic and then back to static regimes. The two-way coupling between granular materials and rigid bodies is highlighted in Figures 7 and 8. In the first sequence an avalanche destroys structures composed of rigid bodies. At first the bodies resist the oncoming mass, producing small splashes against the forward structures and diverting the flow. However the structures are soon overwhelmed, and become part of the flow themselves. After the avalanche has laid waste to this virtual environment, a bulldozer passes through to push aside the remnants. Rigid bodies and granular material are swept away by the bulldozer's blade. In all sequences the ground surface is covered with particles using the same technique as the rigid bodies. In the splash and hourglass, an additional layer of smaller particles are added to increase the roughness of the surface.

Since our treatment of rigid bodies also handles their mutual interactions, we produced a sequence with 1000 rings to mimic the simulation in [GBF03]. Two sequences were created to demonstrate varying quantities of friction and dissipation. Using low friction and dissipation causes the rings to scatter, collecting near the edges of the containing box. When greater friction and dissipation are used, the rings clump closer to the rods at the center of the box. To this point we have distinguished between rigid bodies and granular particles, however we note that in such large-scale simulations we may regard individual rigid bodies as complex grains of granular material. Indeed, in many such cases the distinction is only a matter of scale and quantity. For ex-



Figure 9: 1000 rings after being dropped onto 25 poles. Each ring is approximated by 110 particles.

ample, the individual characteristics of a large boulder are important when considering the dynamics of a few objects. However, when considering an avalanche consisting of many thousands of boulders, only the large-scale behavior of the flow is significant. Therefore, we may model large-scale phenomena by replacing detailed representations with suitable approximations which capture the same essential characteristics. While each of our rings are approximated by only 110 round particles, the resulting simulation is quite realistic. Moreover, the runtimes of our simulation compare favorably to those in [GBF03].

Figure 1 shows an hourglass simulation with 27,427 non-spherical particles, an order of magnitude greater than the optimization-based simulation in [MS01] which considered only 1,000 spheres. The optimization-based approach emphasizes the avoidance of small integration steps caused by stiffness. However, this emphasis may be misplaced in large-scale simulation, since the cost of many small time steps of $O(n)$ complexity is asymptotically preferable to significantly larger time steps of superlinear complexity. While this optimization-based algorithm can rapidly simulate small systems, its $O(n^2)$ scaling [MS01] limits the method's applications.

The computational cost of our method is determined by the choice of physical parameters and the dynamics of a simulation. While the bulldozer and avalanche simulations are of comparable size, the highly dynamic avalanche animation required a smaller average time step. A Young's modulus of 10 MPa was used for all simulations, except the ring sequences. Here a 100 MPa was used to reduce overlap during collision.

5. Conclusions and Future Work

We have demonstrated an efficient discrete element method to accurately capture granular phenomena and two-way coupling with rigid bodies, in terms of simple interactions between rigid-motion-constrained collections of spheres. Large-scale simulations of both granular materials and rigid-bodies have been produced using the same framework. We hope that this paper will inspire much more future development in this direction. In the following, we present a few topics for future research. Our treatment of granular interaction neglects the effects of cohesive forces between particles. In the case of moist sand, such forces produce a considerable increase in the sand's angle of repose. A related challenge is to model the two-way coupling between granular materials and fluids. The third challenge is to generate textures for granular material simulations. Such textures could be used to hide the underlying granularity of the physical simulation or to describe a spatially varying property such as moisture. An evolving granular material represents a dynamic volume texture, however it is unclear how to generate such a texture from a collection of discrete elements.

We are also interested in developing a more rigorous model for rigid body dynamics based on surface particles. While our current method produces plausible animation, the different numbers of spherical contacts across various collision configurations affects properties such as the coefficient of restitution. Also, closer examination of the relationship between the geometric roughness and friction is required. Ideally, one would specify an angle of repose or friction coefficient and produce grains that exhibit those characteristics.

Lastly, we are interested in different force schemes and integration methods. Since our implementation of static friction relies on particle geometry there is little static friction between rigid bodies. Unlike granular particles, the surfaces of rigid bodies are relatively smooth and prohibit significant interlocking. One potential solution is to apply friction impulses as described in [BFA02] and [ITF04]. Also, a semi-implicit integrator [BW98] would handle stiff particle interactions more efficiently than our explicit method.

Simulation	Round Particles	Frames	Min. / Frame
Hourglass	109,708	1600	3.18
1000 Rings	110,000	460	3.73 3.09
Splash	186,892	480	3.41
Avalanche	294,820	720	26.40
Bulldozer	310,149	300	17.40

Table 1: Timing data collected on a set of 3Ghz class PCs.

Acknowledgments

We wish to thank the anonymous reviewers for their valuable comments. This work was partially supported by a grant from UIUC Research Board.

References

- [AT01] ARANSON I. S., TSIMRING L. S.: Continuum description of avalanches in granular media. *Physical Review E* 64 (2001), 020301.
- [BA02] BAERENTZEN J. A., AANAES H.: Generating signed distance fields from triangle meshes. *IMM Technical Report*, 21 (2002).
- [Bar97] BARAFF D.: An introduction to physically based modeling, course notes. ACM SIGGRAPH, 1997.
- [BFA02] BRIDSON R., FEDKIW R., ANDERSON J.: Robust treatment of collisions, contact and friction for cloth animation. In *SIGGRAPH 2002 Proceedings* (2002), pp. 594–603.
- [BW98] BARAFF D., WITKIN A.: Large steps in cloth simulation. In *SIGGRAPH 1998 Proceedings* (1998), pp. 43–54.
- [CLH96] CHANLOU B., LUCIANI A., HABIBI A.: Physical models of loose soils dynamically marked by a moving object. In *Computer Animation* (1996), pp. 27–35.
- [CS79] CUNDALL P. A., STRACK O. D. L.: A discrete numerical model for granular assemblies. *Geotechnique* 29, 1 (1979), 47–65.
- [DC99] DESBRUN M., CANI M.-P.: *Space-time adaptive simulation of highly deformable substances*. Tech. rep., INRIA, 1999.
- [dG99] DE GENNES P. G.: Granular matter: a tentative view. *Reviews of Modern Physics* 71 (1999), S374–S382.
- [DR99] DURY C. M., RISTOW G. H.: Competition of mixing and segregation in rotating cylinders. *Physics of Fluids* 11, 6 (1999), 1387–1394.
- [GBF03] GUENDELMAN E., BRIDSON R., FEDKIW R.: Nonconvex rigid bodies with stacking. *ACM Trans. Graph.* 22, 3 (2003), 871–878.
- [HL98] HERRMAN H. J., LUDING S.: Modeling granular media on the computer. *Continuum Mech. Thermodyn.* 10 (1998), 189–231.
- [HMT01] HADAP S., MAGNENAT-THALMANN N.: Modeling dynamic hair as continuum. In *Eurographics Proceedings. Computer Graphics Forum, Vol.20, No.3* (2001).
- [ITF04] IRVING G., TERAN J., FEDKIW R.: Invertible finite elements for robust simulation of large deformation. In *SCA '04: Proceedings of the 2004 ACM SIGGRAPH/Eurographics symposium on Computer animation* (2004), pp. 131–140.
- [JGD94] JENNIS B., GREEN J., DAVIES R.: The legacy of neglect in the u.s. *Chemical Engineering Progress* 90 (1994), 32–43.
- [JNB96] JAEGER H. M., NAGEL S. R., BEHRINGER R. P.: Granular solids, liquids, and gases. *Reviews of Modern Physics* 68, 4 (1996), 1259–1273.
- [Lam91] LAMBERT J. D.: *Numerical Methods for Ordinary Differential Systems: The Initial Value Problem*. Wiley, 1991.
- [Lee94] LEE J.: Heap formation in two-dimensional granular media. *Journal of Physics A* 27 (1994).
- [LH93] LEE J., HERRMANN H. J.: Angle of repose and angle of marginal stability: molecular dynamics of granular particles. *Journal of Physics A* 26 (1993).
- [LHM95] LUCIANI A., HABIBI A., MANZOTTI E.: A multi-scale physical model of granular materials. In *In Proceedings of Graphics Interface '95* (1995), pp. 136–146.
- [LL86] LANDAU L. D., LIFSHITZ E. M.: *Theory of Elasticity*. Pergamon Books Ltd., 1986.
- [LM93] LI X., MOSHELL J.: Modeling soil: Realtime dynamic models for soil slippage and manipulation. In *SIGGRAPH 1993 Proceedings* (1993), pp. 361–368.
- [MCG03] MÜLLER M., CHARYPAR D., GROSS M.: Particle-based fluid simulation for interactive applications. In *ACM Symposium on Computer Animation* (2003).
- [Mir96] MIRTICH B.: Fast and accurate computation of polyhedral mass properties. *Journal of Graphics Tools* 1, 2 (1996), 31–50.
- [Mon94] MONAGHAN J.: Simulating free surface flows with sph. *Journal of Computational Physics* 110 (1994), 399–406.
- [MP89] MILLER G., PEARCE A.: Globular dynamics: A connected particle system for animating viscous fluids. *Computer and Graphics* 13, 3 (1989), 305–309.
- [MS01] MILENKOVIC V. J., SCHMIDL H.: Optimization-based animation. In *SIGGRAPH 2001 Proceedings* (2001), pp. 37–46.
- [ON03] ONOUE K., NISHITA T.: Virtual sandbox. In *Proceedings of Pacific Graphics 2003* (2003), pp. 252–262.
- [PB93] PÖSCHEL T., BUCHHOLTZ V.: Static friction phenomena in granular materials: Coulomb law versus particle geometry. *Physical Review Letters* 71, 24 (1993), 3963–3966.
- [PB95] PÖSCHEL T., BUCHHOLTZ V.: Molecular dynamics of arbitrarily shaped granular particles. *Journal of Physics I France* 5 (1995), 1431–1455.
- [PTB*03] PREMOZE S., TASDIZEN T., BIGLER J., LEFOHN A., WHITAKER R.: Particle-based simulation of fluids. *Computer Graphics Forum* 22, 3 (2003).
- [Ree83] REEVES W. T.: Particle systems - a technique for modeling a class of fuzzy objects. In *Computer Graphics (Proc. of SIGGRAPH '83)* (1983), vol. 17, pp. 359–376.
- [RPBS99] RAMÍREZ R., PÖSCHEL T., BRILLIANTOV N. V., SCHWAGER T.: Coefficient of restitution of colliding viscoelastic spheres. *Physical Review E* 60 (1999), 4465–4472.
- [SDW96] SCHÄFER J., DIPPEL S., WOLF D. E.: Force schemes in simulations of granular materials. *Journal of Physics I* 6 (1996).
- [SOH99] SUMNER R. W., O'BRIEN J. F., HODGINS J. K.: Animating sand, mud, and snow. *Computer Graphics Forum* 18, 1 (1999), 17–26.
- [SP98] SCHWAGER T., PÖSCHEL T.: Coefficient of normal restitution of viscous particles and cooling rate of granular gases. *Physical Review E* 57, 1 (1998), 650–654.
- [SvHSvS04] SNOEIJER J. H., VAN HECKE M., SOMFAI E., VAN SAARLOOS W.: Packing geometry and statistics of force networks in granular media. *Physical Review E* 70 (2004).
- [Ton91] TONNESEN D.: Modeling liquids and solids using thermal particles. In *Proc. Graphics Interface* (1991), pp. 255–262.
- [TPF89] TERZOPOULOS D., PLATT J., FLEISHER K.: Heating and melting deformable models (from goop to glop). In *Proc. Graphics Interface* (1989), pp. 219–226.
- [TPG99] TREECE G. M., PRAGER R. W., GEE A. H.: Regularised marching tetrahedra: improved iso-surface extraction. *Computers and Graphics* 23, 4 (1999), 583–598.
- [Tur92] TURK G.: Re-tiling polygonal surfaces. In *SIGGRAPH 1992 Proceedings* (1992), pp. 55–64.
- [WB93] WALTON O. R., BRAUN R. L.: Simulation of rotary-drum and repose tests for frictional spheres and rigid sphere clusters. In *DOE/NSF Workshop on Flow of Particulates, Ithaca, NY, 29 Sep. - 1 Oct. 1993* (1993), vol. 29.
- [WH94] WITKIN A. P., HECKBERT P. S.: Using particles to sample and control implicit surfaces. In *SIGGRAPH 1994 Proceedings* (1994), pp. 269–274.
- [YK096] YOON H., KOSHIZUKA S., OKA Y.: A particle gridless hybrid method for incompressible flows. *Int. J. Numer. Meth. Fluids* 29, 4 (1996).