Workshop on Virtual Reality Interaction and Physical Simulation VRIPHYS (2012)
J. Bender, A. Kuijper, D. W. Fellner, and É. Guérin (Editors)

High-Resolution Simulation of Granular Material with SPH

Markus Ihmsen Arthur Wahl Matthias Teschner

University of Freiburg

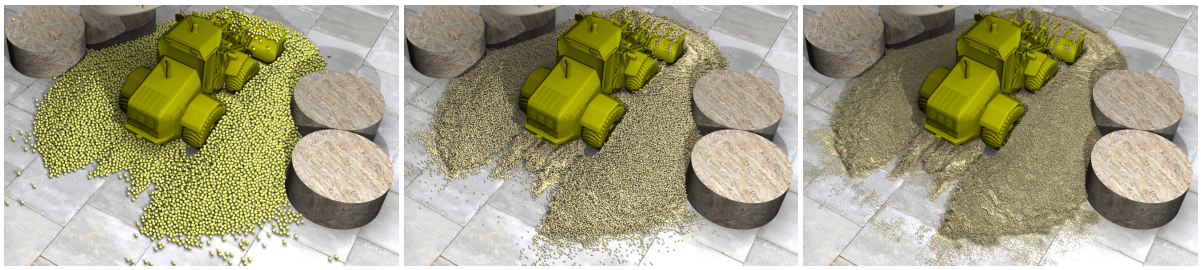

Figure 1: *The proposed framework simulates mechanical behavior at a coarse scale (left). The base simulation is significantly refined with a secondary simulation (middle and right). Simulation particles (38K, 1.4M and 19.4M, left to right) are rendered.*

Abstract

We present an efficient framework for simulating granular material with high visual detail. Our model solves the computationally and numerically critical forces on a coarsely sampled particle simulation. We incorporate a new frictional boundary force into an existing continuum-based method which enables realistic interactions and a more robust simulation. Visual realism is achieved by coupling a set of highly resolved particles with the base simulation at low computational costs. Thereby, visual details can be added which are not resolved by the base simulation.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Animation

1. Introduction

Granular materials, such as sand, rice or coffee beans are conglomerations of discrete solid elements which show unique physical behavior. They settle in stable piles and act like a solid if the average energy is low. When freely flowing, they have similar characteristics as ordinary Newtonian fluids, but unlike fluids, granular material dissipates energy quickly. The complex dynamics arise from the interplay of contact forces between the elements.

In order to capture this behavior, various simulation methods have been developed in the engineering field, e. g. [Bic04, JLL04], and also in computer animation, e. g. [LM93, SOH99, HBK09, NGL10]. In contrast to computational mechanics, the main focus in animation is set on efficient techniques that achieve visually plausible results. In order to allow for efficient implementation, simplifying assumptions and a coarse discretization are employed. While this avoids simulating each physical grain, it poses a new major challenge, namely how to achieve rich visual detail.

The main contribution of this work is a Lagrangian simulation framework which captures granular dynamics realistically and uncovers high visual detail at low computational costs. Appropriate mechanical behavior is modeled by computing frictional and pressure forces on a coarse scale. High visual detail is obtained in a post-process, where a spatially fine-scaled set of particles is coupled to the base simulation. Our model aligns high-resolution particles well with the surface of the base simulation without perceivable clumping. Since external forces are also acting on the fine resolution,

DOI: 10.2312/PE/vriphys/vriphys12/053-060

they can depart freely from the base simulation. The presented upsampling method increases the visual quality dramatically without uncovering the base simulation, even for large upscaling factors and scenarios with dynamic objects, see Fig. 1.

2. Related Work

In computer graphics, granular media is simulated using either discrete [Bic04, LHM95] or continuum methods [ZB05]. In discrete models, the material is discretized into a distinct set of elements and the behavior is captured by directly simulating the interaction between these elements. In [BYM05], grains are modeled by rigid compounds of spheres. Due to the non-spherical shape of the compound structures, stable sand piles can be simulated. By generalizing this approach to rigid bodies, compelling two-way coupling between granular material and solid objects is achieved. Discrete models require a large number of particles to model a fine-grained material. This imposes high computational costs for solving contact dynamics and limits the time step for highly dynamic simulations.

In continuum methods, the grain size is decoupled from the resolution of the simulation. In such models, internal forces are typically computed on a coarse grid. The resulting velocity field is then used to advect a set of fine-scale particles. This was first demonstrated by Zhu and Bridson [ZB05], who simulated sand as an incompressible fluid using the Fluid Implicit Particle (FLIP) method. In order to simulate stable piles, this method classifies the sand domain into regions which are either rigidly moving or flowing. Later, Lenaerts and Dutré [LD09] incorporated this concept to simulate granular material with the Smoothed Particle Hydrodynamics (SPH) method [MCG03]. However, the incompressibility assumption results in undesired cohesive behavior which prevents plausible animations of freely dispersing material. Narain et al. [NGL10] addressed this problem by employing an unilateral incompressibilty constraint, i. e. negative flow divergence is not counteracted by pressure. In the sense of FLIP, they solve the internal forces on an Eulerian grid. The resulting velocities are then used to advect the particles, representing the material. Recently, Alduan and Otaduy [AO11] adapted unilateral incompressibility to the predictive-corrective incompressible SPH (PCISPH) method which was originally designed for incompressible fluid simulation [SP09].

The purely Lagrangian framework proposed in [AO11] does not suffer from grid artifacts which is a major benefit compared to the Eulerian framework described in [NGL10]. However, it also introduces new challenges. Following the SPH concept, the pressure is computed based on local density values and not according to the divergence of the velocities. This makes the approach very sensitive to sudden increases in the density. Oscillations in the density field particularly occur at interfaces with dynamic solid objects due to particle deficiency. In the context of fluid simulations, this has been addressed in [IAGT10, AIA*12] by treating the boundary as an interface to the fluid simulation. Thereby, spatial and temporal discontinuities of physical properties are avoided, resulting in a smoother and a more robust simulation compared to commonly employed distance-based penalty methods, e. g. [MST*04, BTT09]. Furthermore, in [AO11], internal forces and advection are computed using the same discretization scale. Thus, in order to compute fine-grained material, significantly higher computational costs are imposed compared to [NGL10].

In this paper, we address these challenges by extending the pure Lagrangian, continuum framework [AO11] in two ways. First, we show that the physically-based rigid-fluid coupling presented in [AIA*12] can be easily adapted to handle smooth interactions with granular material. This eliminates oscillations in the pressure field. As a consequence, less iterations are required when solving for the internal forces even at larger time steps. The second extension enables the simulation of fine-grained material at low computational costs by refining the simulation in a post-process.

The general idea of refining the coarse simulation for rendering is not new. Most authors propose to sample the simulation particles with a finer set of pseudo-random particles [LD09, NGL10]. High-resolution (HR) particles are fixed to the base particles in order to avoid temporal flickering. As HR particles can not disperse freely, such an approach results in clumping artifacts, perceived as staircase or spherical patterns and a distorted distribution of the material. Narain et al. [NGL10] addressed this problem by inserting additional anisotropic particles in regions where the material diverges. This step is already performed during simulation. In contrast, [ABC*07, ATO09] employ a spatial decomposition for the computation of internal and external forces. In sparse regions, HR particles are not passively advected with the base simulation, but respond to external forces. Thereby, clumping artifacts are significantly reduced.

Our refinement method is inspired by the decomposition idea, but departs significantly from previous work. In [ATO09], HR particles are either advected according to external forces or along the base flow. In contrast, our model does not rely on two complementary cases, but blends external forces with the interpolated base velocity. This yields more natural looking results, particularly for large upscaling factors. Furthermore, as we take velocities of boundaries into account, HR particles interact smoothly with complex moving objects.

3. Granular Simulation Framework

Our framework builds on the SPH-based continuum method proposed in [AO11] which is described in Section 3.1. Section 3.2 discusses how the convergence of the algorithm can be improved by incorporating a more versatile treatment

of the granular-solid interface. Subsequently, Section 4 describes how the base simulation is refined, in order to efficiently simulate and render fine-grained material.

3.1. Coarse Scale Simulation

In SPH, the pressure at a position $\mathbf{x}_i$ is typically determined according to the local density fluctuation $\rho_{err}(\mathbf{x}_i) = \rho(\mathbf{x}_i) - \rho_0$, where ρ_0 is the reference density of the material. The density $\rho_i \equiv \rho(\mathbf{x}_i)$ is computed via the SPH interpolation concept:

$$\begin{aligned}\rho_i &= \sum_j V_j \rho_j W(\mathbf{x}_i - \mathbf{x}_j, h) \\ &= \sum_j m_j W(\mathbf{x}_i - \mathbf{x}_j, h),\end{aligned} \tag{1}$$

where m_j and $V_j = \frac{m_j}{\rho_j}$ denote mass and volume represented at $\mathbf{x}_j$. $W_{ij} \equiv W(\mathbf{x}_i - \mathbf{x}_j, h)$ is a kernel function with support h.

Granular flow is governed by unilateral incompressibility which is described by two inequality constraints $\rho \leq \rho_0$ and $p \geq 0$, where either $\rho = \rho_0$ or $p = 0$. Accordingly, pressure values p are not negative and the material can not be compressed beyond ρ_0. These constraints can be directly plugged into any incompressible SPH solver by clamping negative pressures to zero. The resultant pressure field is then used to compute the pressure forces as

$$\mathbf{F}_i^p = -m_i \sum_j m_j \left(\frac{p_i}{\rho_i^2} + \frac{p_j}{\rho_j^2} \right) \nabla W_{ij}. \tag{2}$$

For simulating granular material, [AO11] introduces a friction model which minimizes relative velocities measured by the strain rate $\dot{\varepsilon}$. The strain rate is computed as $\dot{\varepsilon} = 0.5 \left(\nabla \mathbf{u} + \nabla \mathbf{u}^T \right)$, where the gradient of the velocity $\mathbf{u}_i$ is given by

$$\nabla \mathbf{u}_i = \sum_j V_j \nabla W_{ij} \mathbf{u}_j^T. \tag{3}$$

Then a frictional stress tensor $\mathbf{s}$ is computed that dissipates the strain rate. In order to simulate material with different angles of repose θ, the Drucker-Prager yield criterion, $\sqrt{\sum s_{ij}^2} \leq p\sqrt{2}\sin\theta$ is employed. Thereby, friction is limited by pressure, i. e. in the absence of pressure, the frictional stress is zero. The frictional forces are computed with

$$\mathbf{F}_i^f = -m_i \sum_j m_j \left(\frac{\mathbf{s}_i}{\rho_i^2} + \frac{\mathbf{s}_j}{\rho_j^2} \right) \nabla W_{ij}. \tag{4}$$

In order to compute (2) and (4), pressure and stress values need to be determined. Therefore, we employ the iterative predictive-corrective PCISPH method [SP09] as proposed in [AO11]. Alternatively, an equation of state [MCG03, BT07] or a global projection scheme [PTB*03] could be used.

3.2. Solid Body Interaction

The forces given in the previous section are internal forces, acting inside the material. At interfaces with solid objects, additional considerations have to be employed in order to guarantee non-penetration of objects and to simulate two-way coupled interactions with variable friction.

The direct forcing method [BTT09] used in [AO11] corrects predicted penetrations by constraining positions to the rigid surface. Different slip conditions are realized by directly manipulating particle velocities. This model does not conserve momentum and leads to significant oscillations in the pressure field as shown for fluids in [AIA*12]. Additionally, for the granular model described in Sec. 3.1, the noisy pressure field results in discontinuities of the frictional forces (4). In order to avoid perceivable artifacts, a small time step has to be employed and/or the number of iterations for correcting density errors has to be set high, i. e. larger than five.

We improve the robustness and the quality of the simulation by adapting the boundary handling method of [AIA*12] to interactions with granular material. Thereby, the surface of rigid objects is sampled with particles as in [BYM05]. Boundary particles contribute to the density of a granular particle. Accordingly, (1) is rewritten as

$$\rho_i = \sum_j m_j W(\mathbf{x}_i - \mathbf{x}_j, h) + \sum_b \frac{\rho_0}{\delta_b} W(\mathbf{x}_i - \mathbf{x}_b, h), \tag{5}$$

where $\delta_b = \frac{1}{\sum_j W_{bj}}$ is the number density of a boundary particle b, summed up over all boundary particle neighbors j. $\Psi_b(\rho_0) \equiv \frac{\rho_0}{\delta_b}$ scales the contributions of a boundary particle according to the local boundary sampling and the reference density of the granular material. Thereby, (5) estimates the density correctly at boundaries for arbitrary samplings and material densities. Non-penetration is realized via pressure forces, acting from a boundary particle b on a granular particle i with

$$\mathbf{F}_{i \leftarrow b}^p = -m_i \Psi_b(\rho_0) \frac{p_i}{\rho_i^2} \nabla W_{ib}. \tag{6}$$

A detailed derivation can be found in [AIA*12].

We extend this model to interactions with granular materials by adding contributions of the boundary when computing frictional stresses. Therefore, we modify (3) to

$$\nabla \mathbf{u}_i = \sum_j V_j \nabla W_{ij} \mathbf{u}_j^T + \sum_b \frac{1}{\delta_b} \nabla W_{ib} \mathbf{u}_b^T \tag{7}$$

The reformulation of (4) can be derived similar to (6) as

$$\mathbf{F}_{i \leftarrow b}^{f_s} = -m_i \Psi_b(\rho_0) \frac{\mathbf{s}_i}{\rho_i^2} \nabla W_{ib}. \tag{8}$$

However, (8) depends on the properties of the granular material, i. e. angle of repose and local pressure, but not on the material properties of the rigid object. In order to model

Algorithm 1: Simulation update of the base solver.

```
foreach particle i do
    find neighbors
foreach particle i do
    apply gravity and material viscosity
    reset pressure and stress
    reset pressure and friction force
k = 0
while k < 3 do
    foreach particle i do
        predict velocity and position
    foreach particle i do
        predict density ρ_i^*(t + Δt)
        if ρ_i^*(t + Δt) > ρ_max then
            predict density error and strain rate
            increment stress and pressure
            apply yield condition
        else
            add discrete particle forces
    foreach particle i do
        compute pressure force (2) and (6)
        compute friction force (4) and (10)
    k+ = 1
foreach particle i do
    update velocity and position
```

external friction forces, we additionally employ the artificial viscosity model given in [AIA*12] which is written as

$$\mathbf{F}^{v}_{i\leftarrow b} = -m_i \Psi_b (\rho_0) \Pi_{ib} \nabla W_{ib}, \tag{9}$$

where $\Pi_{ib} = -\frac{\sigma_{ib} h c_s}{2\rho_i} \left(\frac{\min(\mathbf{v}_{ib} \cdot \mathbf{x}_{ib}, 0)}{|\mathbf{x}_{ib}|^2 + \varepsilon h^2} \right)$. Here, c_s denotes the speed of numerical propagation and σ_{ib} controls the friction between the rigid object and the granular material. Applying the sum of (8) and (9) might overshoot the desired dissipative effect. Therefore, we combine both forces by picking the one which maximizes dissipation. Accordingly, the dissipative contribution of a boundary particle b is computed as

$$\mathbf{F}^{f}_{i\leftarrow b} = \max(\left\|\mathbf{F}^{v}_{i\leftarrow b}\right\|, \left\|\mathbf{F}^{f_s}_{i\leftarrow b}\right\|), \tag{10}$$

Algorithm 1 outlines the steps performed in each simulation update of the base solver. Two-way coupling is easily realized by applying pressure (6) and friction (10) forces symmetrically to the boundary, see Fig. 2.

4. Fine Scale Simulation

The simulation model described in the previous section employs a continuum approach which describes the granular flow at a macroscopic scale. Thereby, a simulation particle should be interpreted as a clump of matter and not as a single grain. Indeed, setting up the particle size to the real size of a sand grain, which is only a fraction of a millimeter, is prohibitive. This would not only explode the memory and computational costs, but also restrict the time step due to the CFL condition. Instead, we propose to simulate the material on a coarse scale and apply a secondary simulation with a set of highly resolved particles which can be directly used for rendering.

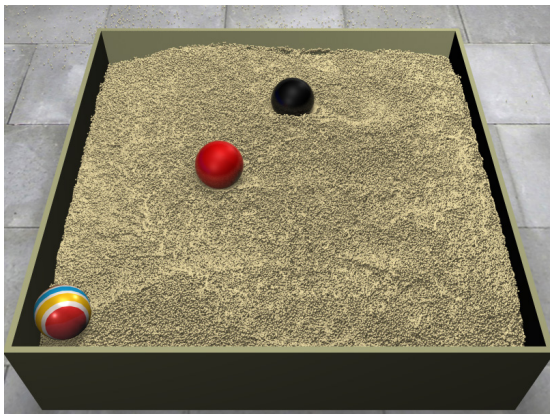

Figure 2: *Demonstration of the proposed two-way coupling. Three spheres with different masses are tossed into a sand pool. Due to the friction-based coupling, the sand supports also the heaviest sphere (black) which is three times more dense than the sand.*

4.1. Sampling

Each time a low-resolution (LR) particle with radius r_{LR} is added, we sample its volume with HR particles. The initial sampling is crucial as it could easily introduce aliasing or distortions. Sampling the spherical volume of a particle leads to gaps while sampling the bounding box might cause staircase patterns. In order to avoid aliasing, we do not only generate HR samples inside the bounding box of the LR particle but also slightly outside, employing a distribution that prefers samples that are inside the LR volume.

Therefore, we divide the bounding box of each LR particle into seven support points, one at the particle center and one at each intersection point of the bounding box with the spherical particle volume. HR particles are randomly sampled around each support point in a cubical volume with length $2r_{LR}$. As the seven sample volumes overlap inside the LR volume, this strategy generates three times more HR samples inside the bounding box of the base particle than outside.

4.2. Advection

We derive the advection of HR particles from the following principles: HR particles should follow the mechanical flow

that is given by the base simulation, but also should be allowed to disperse freely. Further, they should smoothly align with the surface of the base simulation without forming perceivable clumps. Finally, in order to guarantee efficient updates at large time steps, the advection method should not compute internal forces or perform collision tests between HR particles.

Alduan et al. [ATO09] set similar requirements to their model. They map the mechanical behavior by interpolating the velocity of LR particles to HR particles for advection. In order to avoid clumping, particles having one or no LR particle neighbors within the influence radius h_{HR} are only influenced by external forces. However, this simple distinction avoids clumping only if $h_{HR} \approx r_{LR}$. On the other hand, larger values of h_{HR} are required for smooth interpolations of the velocities. In all our experiments, we set $h_{HR} = 3r_{LR}$.

In contrast to [ATO09], we do not employ an explicit distinction of two cases, but propose a weighting that automatically and smoothly blends the contributions of the base simulation and external forces. Therefore, for each HR particle at position $\mathbf{x}_i$, we first compute distance-based weights w as

$$w(d_{ij}) = \max\left[0, \left(1 - \frac{d_{ij}^2}{h_{HR}^2}\right)^3\right] \tag{11}$$

where $d_{ij} = |\mathbf{x}_i - \mathbf{x}_j|$ is the distance between HR particles and LR simulation or boundary particles j in the support radius h_{HR}. (11) is a well-shaped kernel function which smoothly drops to zero. It is typically applied for reconstructing smooth surfaces of particle data [ZB05, AIAT12]. We employ it for computing the average velocity as

$$\mathbf{v}_i^*(t+\Delta t) = \frac{1}{\sum_j w(d_{ij})} \sum_j w(d_{ij})\mathbf{v}_j \tag{12}$$

As long as the number of LR samples is sufficient, (12) interpolates the velocity well. However, in sparsely sampled regions it results in visual clumping. In order to achieve smooth alignments of HR particles without perceivable clumps, we compute the velocities of HR particles as the sum of weighted external forces $\mathbf{F}^g$ and the interpolated velocity with

$$\mathbf{v}_i(t+\Delta t) = (1-\alpha_i)\mathbf{v}_i^*(t+\Delta t) + \alpha_i\left(\mathbf{v}_i(t) + \Delta t\frac{\mathbf{F}^g}{m}\right), \tag{13}$$

where α is non-zero in sparse regions only and increases with higher distances of $\mathbf{x}$ to the center of LR particles. It is defined as

$$\alpha_i = \begin{cases} 1 - \arg\max_j w(d_{ij}) & \frac{\arg\max_j w(d_{ij})}{\sum_j w(d_{ij})} \geq 0.6, \\ 1 - \arg\max_j w(d_{ij}) & \arg\max_j w(d_{ij}) \leq w(r_{LR}), \\ 0 & \text{otherwise.} \end{cases}$$

Here, the constant 0.6 is an empirically tested value which gives the best results when $h_{HR} = 3r_{LR}$. Finally, the position

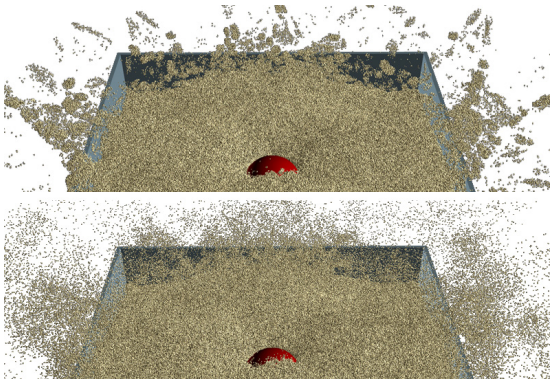

Figure 3: *Upsampling comparison. In [ATO09], HR particles can not disperse freely if two or more base particles are in close proximity (top). Our method avoids clumping by weighting external and internal forces for all HR particles (bottom).*

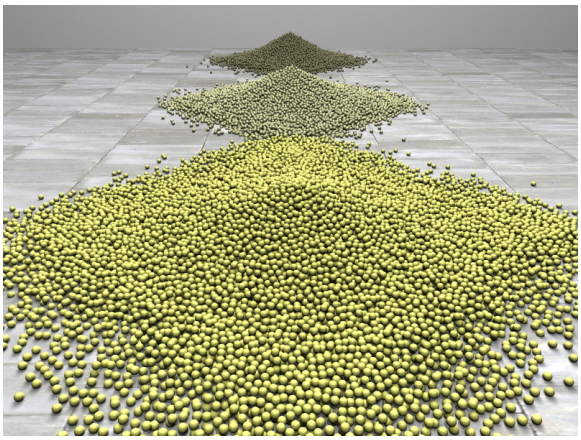

Figure 4: *Sand piles with different angles of repose* 30°, 45° *and* 60° *(front to back).*

is integrated as

$$\mathbf{x}_i(t+\Delta t) = \mathbf{x}_i(t) + \Delta t\mathbf{v}_i(t+\Delta t). \tag{14}$$

Accordingly, external forces are automatically applied in regions where clumping potentially occurs. Contributions of external forces are smoothly faded in and out, which on one hand allows HR particles to disperse freely and on the other hand results in smooth alignment with the materials surface, see Fig. 3. Our model does faithfully upscale scenes with dynamic objects, as we take the positions and velocities of moving objects into account when interpolating the velocities.

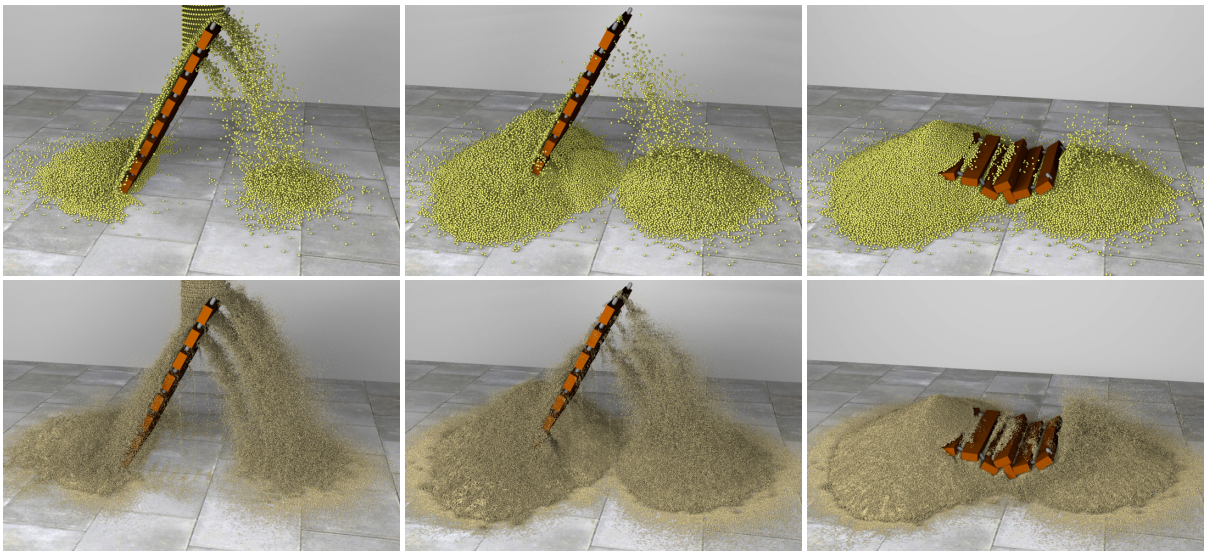

Figure 5: *Sand is poured over a fixed rope-ladder. LR particles (top) do not slip through all gaps, in contrast to the finer-scaled HR particles (bottom). The upsampling captures a nice pattern at the ground which is not resolved by the base simulation (left). At the end, the constraints of the ladder are released.*

5. Results

We present several scenarios which show different aspects of our approach including one-way and two-way solid body interaction. Comparison to previous work and performance evaluation are also provided. Timings are given for a 12-core 3.46 GHz Intel i7 with 24 GB of RAM, see Table 1.

Implementation. For the SPH interpolations, we use the cubic spline kernel [Mon05]. We optimize the base simulation by applying the well-known concept of rigid-body sleeping [Sch02] to SPH particles. Thereby, we do not integrate SPH particles in time as long as their velocity is lower than a user-defined *sleep velocity*. Particles wake up automatically when their net-acceleration exceeds a critical value. The efficiency is further improved by parallelizing the algorithm with the data structures described in [IABT11] using OpenMP. Images were rendered with mental ray v3.9.4 [NVI11].

5.1. Solid Body Interaction

We demonstrate the effectiveness of the proposed friction and pressure forces applied from and to the boundaries on a simple test scenario, see Fig. 4. In this scene, we simulated three stable piles with different angles of repose. We set the time step to 1 millisecond and used a fixed number of three iterations to compute pressure and stress values. For this setting, the proposed boundary handling enforces smooth pressure gradients at the boundary. In contrast, if we employ direct forcing as in [AO11], the density field oscillates which results in unnatural accelerations perceived as 'popping of particles' as is shown in the accompanying video.

Two-way coupling between rigid objects and granular material is demonstrated in Fig. 2. Rigid objects interact differently with the granular material according to their density since the dissipation of energy for lighter objects is much less than for heavy objects. In contrast to fluids, granular material supports even objects with much higher material density than its own. This behavior is faithfully captured as the employed coupling is based on frictional stresses.

5.2. Upsampling

In order to show the benefit of the proposed refinement model, we compare it to [ATO09] on a scene where a heavy sphere hits the sand surface with high velocity, creating a splash. The base simulation sampled the material with 100K particles while the secondary simulation used 20.8M particles. Since in [ATO09], HR particles with more than one LR neighbor are not influenced by external forces, cluster of HR particles move uniformly in splash regions and at the sand surface, see Fig. 3-top. In contrast, our method does not rely on complementary cases to compute the velocity, but weights external forces and base velocities using a well shaped kernel function. Accordingly, HR particles are allowed to disperse freely, see Fig. 3-bottom.

Opposed to previous work, we show that HR particles interact smoothly with moving objects, e. g. a bulldozer, see Fig. 1. This is realized by taking positions and velocities

of boundary particles into account when interpolating the velocity field. We further demonstrate that our method can cope with very large upscaling factors. Fig. 1-middle shows the refinement with an upscaling factor of 38 which corresponds to the largest value used in [ATO09]. The realism is significantly improved by setting the refinement factor to 500, Fig. 1-right.

We allow particles to depart freely from the base simulation by computing the velocities of HR particles as the weighted sum of external forces and the interpolated velocity field. Thereby, HR particles uncover details that are not captured by the base simulation, see Fig. 5.

5.3. Performance

The presented framework advances the efficiency compared to previous work in two ways. First, the employed boundary handling results in smooth pressure gradients which improves the robustness compared to [AO11]. Thus, larger time steps can be handled at a smaller number of iterations for computing pressure and frictional forces. In all presented scenarios, the primary simulation was performed with a time step of 1 ms and a fixed number of three iterations. Thereby, we measured a speed up of up to 6 compared to [AO11].

Second, we employ the refinement as a post-processing step at a different temporal resolution. In all presented scenarios, the time step for the secondary simulation was set to 10 ms. As no interactions between HR particles are computed, each particle can be updated independently which permits a straightforward parallelization. It should be noticed that our implementation took on average 11 seconds (1.4 s LR + 9.6 s HR) per frame for a scene with 19.4 million particles, see Table 1. In contrast, Alduan et al. [ATO09] reported an update rate of 5.5 minutes for 1.6 million HR particles and the same number of LR particles. Although this comparison does not take the respective hardware configurations into account, it indicates the efficiency of the proposed model.

6. Conclusion

We have presented an efficient framework for computing high-resolution simulations of granular material using a pure Lagrangian method. Performance-critical forces are computed in a primary simulation on a relatively coarse scale. Here, the proposed friction-based coupling models interactions with rigid objects realistically and handles comparable large time steps. Visual detail is added in a secondary simulation where high-resolution particles are coupled to the base flow. For advecting secondary particles, we propose a smooth weighting of external forces with the velocity field of the base simulation. This technique adds detail that is not captured otherwise while clumping of high-resolution particles is avoided, even for very large upscaling factors.

	# Particles		Time / Frame	
	LR	HR	LR	HR
Bulldozer (Fig. 1)	38K	1.4M 19.4M	1.4 s	1.1 s 9.6 s
Spheres (Fig. 2)	96K	11.3M	3.1 s	5.8 s
Splash (Fig. 3)	83K	20.8M	2.9 s	10.3 s
Sand-Piles (Fig. 4)	45K	-	1.6 s	-
Rope-Ladder (Fig. 5)	138K	17.3M	4.6 s	8.3 s

Table 1: *Performance measurements for the given scenarios.*

6.1. Limitations and Future Work

The proposed coupling does not model static friction correctly which prevents us from animating some interesting scenes, e. g. an accelerating dumper truck filled with sand. In such a scenario, our model fails to keep the sand pile at rest when the truck starts moving.

Furthermore, in our refinement model HR particles never interact with each other which makes it very efficient to compute on one hand. However, this might cause compression artifacts in sparsely sampled LR regions, e. g. at the edges of a sand pile, as HR particles are attracted to the ground by gravity. As a simple solution, interactions between HR particles could be computed in relevant regions employing either a discrete or a continuum model. Alternatively, we believe that this issue can be addressed more efficiently by finding an adequate extension to the proposed interpolation heuristic.

Massive conglomerations of granular material, e. g. a beach, are more efficiently represented by height-fields [HBK09, KBKS09]. However, in these methods, the level of detail is limited as only two-dimensional information is mapped onto the three-dimensional space. Dispersion effects or the sliding of single grains can not be captured with such a representation.

In future work, the granular material could be coupled with an SPH fluid simulation in order to animate erosion effects and transitions from dry to moist sand and mud. In this context, we plan to investigate how SPH fluid simulations can benefit from the proposed refinement model.

Acknowledgements

We thank the reviewers for their helpful comments. This project is supported by the German Research Foundation (DFG) under contract numbers SFB/TR-8 and TE 632/1-2. We also thank NVIDIA ARC GmbH for supporting this work.

References

[ABC*07] AMMANN C., BLOOM D., COHEN J. M., COURTE J., FLORES L., HASEGAWA S., KALAITZIDIS N., TORNBERG T., TREWEEK L., WINTER B., YANG C.: The birth of sandman. In *ACM SIGGRAPH 2007 sketches* (2007). 2

[AIA*12] AKINCI N., IHMSEN M., AKINCI G., SOLENTHALER B., TESCHNER M.: Versatile Rigid-Fluid Coupling for Incompressible SPH. *ACM Trans. on Graphics (Proc. SIGGRAPH)* (2012). 2, 3, 4

[AIAT12] AKINCI G., IHMSEN M., AKINCI N., TESCHNER M.: Parallel Surface Reconstruction for Particle-Based Fluids. *Computer Graphics Forum* (2012). doi:10.1111/j.1467-8659.2012.02096.x. 5

[AO11] ALDUÁN I., OTADUY M. A.: Sph granular flow with friction and cohesion. In *Proc. of the 2011 ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (2011), pp. 25–32. doi:10.1145/2019406.2019410. 2, 3, 6, 7

[ATO09] ALDUÁN I., TENA A., OTADUY M. A.: Simulation of High-Resolution Granular Media. In *Proc. of Congreso Español de Informática Gráfica* (2009), pp. 1–8. 2, 5, 6, 7

[Bic04] BICANIĆ N.: Discrete element methods. *Encyclopedia of Computational Mechanics 1* (2004). doi:10.1002/0470091355.ecm006. 1, 2

[BT07] BECKER M., TESCHNER M.: Weakly compressible SPH for free surface flows. In *Proc. of the 2007 ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (2007), pp. 209–217. 3

[BTT09] BECKER M., TESSENDORF H., TESCHNER M.: Direct forcing for lagrangian rigid-fluid coupling. *IEEE Transactions on Visualization and Computer Graphics 15*, 3 (2009), 493–503. 2, 3

[BYM05] BELL N., YU Y., MUCHA P. J.: Particle-based simulation of granular materials. In *Proc. of the 2005 ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (2005), SCA '05, pp. 77–86. 2, 3

[HBK09] HOLZ D., BEER T., KUHLEN T.: Soil Deformation Models for Real-Time Simulation: A Hybrid Approach. In *VRIPHYS 09: Sixth Workshop in Virtual Reality Interactions and Physical Simulations* (2009), pp. 21–30. doi:10.2312/PE/vriphys/vriphys09/021-030. 1, 7

[IABT11] IHMSEN M., AKINCI N., BECKER M., TESCHNER M.: A Parallel SPH Implementation on Multi-core CPUs. *Computer Graphics Forum 30*, 1 (2011), 99–112. doi:10.1111/j.1467-8659.2010.01832. 6

[IAGT10] IHMSEN M., AKINCI N., GISSLER M., TESCHNER M.: Boundary Handling and Adaptive Time-stepping for PCISPH. In *Proc. VRIPHYS* (2010), pp. 79–88. 2

[JLL04] JOSSERAND C., LAGRÉE P.-Y., LHUILLIER D.: Stationary shear flows of dense granular materials: a tentative continuum modelling. *The European Physical Journal E: Soft Matter and Biological Physics 14*, 6 (2004), 127–135. 1

[KBKS09] KRISTOF P., BENEVS B., KRIVANEK J., STAVA O.: Hydraulic Erosion Using Smoothed Particle Hydrodynamics. *Computer Graphics Forum (Proceedings of Eurographics 2009) 28*, 2 (2009). 7

[LD09] LENAERTS T., DUTRE P.: Mixing fluids and granular materials. *Computer Graphics Forum 28*, 2 (2009), 213–218. doi:10.1111/j.1467-8659.2009.01360.x. 2

[LHM95] LUCIANI A., HABIBI A., MANZOTTI E.: A multi scale physical model of granular materials. In *Proc. Graphics Interface* (1995), pp. 136–146. 2

[LM93] LI X., MOSHELL J. M.: Modeling soil: Realtime dynamic models for soil slippage and manipulation. In *In SIGGRAPH 1993 Proceedings* (1993), pp. 361–368. 1

[MCG03] MÜLLER M., CHARYPAR D., GROSS M.: Particle-based fluid simulation for interactive applications. In *Proc. of ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (2003), pp. 154–159. 2, 3

[Mon05] MONAGHAN J.: Smoothed particle hydrodynamics. *Reports on Progress in Physics 68*, 8 (2005), 1703–1759. 6

[MST*04] MÜLLER M., SCHIRM S., TESCHNER M., HEIDELBERGER B., GROSS M.: Interaction of fluids with deformable solids. *Computer Animation and Virtual Worlds 15*, 34 (2004), 159–171. 2

[NGL10] NARAIN R., GOLAS A., LIN M. C.: Free-flowing granular materials with two-way solid coupling. *ACM Trans. Graph. (SIGGRAPH Proc.) 29*, 6 (2010), 173:1–173:10. doi:10.1145/1882261.1866195. 1, 2

[NVI11] NVIDIA: mental ray (version 3.9). *http://www.mentalimages.com* (2011). URL: http://www.mentalimages.com. 6

[PTB*03] PREMOZE S., TASDIZEN T., BIGLER J., LEFOHN A., WHITAKER R.: Particle-Based Simulation of Fluids. *Computer Graphics Forum (Proc. of Eurographics) 22* (2003), 401–410. 3

[Sch02] SCHMIDL H.: *Optimization-Based-Animation*. PhD thesis, University of Miami, 2002. 6

[SOH99] SUMNER R. W., O'BRIEN J. F., HODGINS J. K.: Animating Sand, Mud, and Snow. *Computer Graphics Forum 28* (1999), 1–11. 1

[SP09] SOLENTHALER B., PAJAROLA R.: Predictive-corrective incompressible SPH. *ACM Trans. on Graphics (Proc. SIGGRAPH) 28* (2009), 40:1–40:6. 2, 3

[ZB05] ZHU Y., BRIDSON R.: Animating sand as a fluid. *ACM Trans. Graph. (SIGGRAPH Proc.) 24*, 3 (2005), 965–972. doi:10.1145/1073204.1073298. 2, 5