

A Lagrangian Framework for Simulating Granular Material with High Detail

Markus Ihmsen^{a,*}, Arthur Wahl^b, Matthias Teschner^a

^aUniversity of Freiburg, Germany

^bRWTH Aachen, Germany

Abstract

We present an efficient Lagrangian framework for simulating granular material with high visual detail. Our model solves the computationally and numerically critical forces on a coarsely sampled particle simulation. Pressure and friction forces are expressed as constraint forces which are iteratively computed. We realize stable and realistic interactions with rigid bodies by employing pressure and friction-based boundary forces. Stable formations of sand piles are realized by employing the concept of rigid-body sleeping. Furthermore, material transitions from dry to wet can be modeled. Visual realism is achieved by coupling a set of highly resolved particles with the base simulation at low computational costs. Thereby, detail is added which is not resolved by the base simulation. The practicability of the approach is demonstrated by showing various high-resolution simulations with up to 20 million particles.

Keywords: granular material simulation, SPH, multi resolution

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. [1, 2], and also in computer animation, e.g. [3, 4, 5, 6]. 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.

This paper is an extended version of [7]. The main contribution of [7] 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. Since external forces are also acting on the fine resolution, secondary particles 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. As an extension to [7], we elaborate more clearly on how to initialize and implement the base simulation. In this context, we discuss that the constraint forces are highly interdependent and that this interdependency has a negative impact on the convergence of the solver. As we propose, this issue can be addressed by adapting the concept of rigid-body sleeping. Furthermore, we extend the base simulation to allow animations of dry and wet material. The new model is designed such that transitions from dry to wet material can be simulated. Finally, we analyze the scaling of quality and performance for different parameter sets and resolutions.

2. Related Work

In computer graphics, granular media is simulated using either discrete [1, 8] or continuum methods [9]. 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 [10], 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

*Corresponding author.

Email addresses: ihmsen@informatik.uni-freiburg.de (Markus Ihmsen), wahl@mmi.rwth-aachen.de (Arthur Wahl), teschner@informatik.uni-freiburg.de (Matthias Teschner)

URL: <http://cg.informatik.uni-freiburg.de> (Markus Ihmsen)



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.

59 is then used to advect a set of fine-scale particles. This was first
60 demonstrated by Zhu and Bridson [9], who simulated sand as
61 an incompressible fluid using the Fluid Implicit Particle (FLIP)
62 method. In order to simulate stable piles, this method classifies
63 the sand domain into regions which are either rigidly moving
64 or flowing. Later, Lenaerts and Dutr e [11] incorporated this
65 concept to simulate granular material with the Smoothed Particle
66 Hydrodynamics (SPH) method [12]. However, the incompressibility
67 assumption results in undesired cohesive behavior
68 which prevents plausible animations of freely dispersing material.
69 Narain et al. [6] addressed this problem by employing an
70 unilateral incompressibility constraint, i.e., negative flow divergence
71 is not counteracted by pressure. In the sense of FLIP,
72 they solve the internal forces on an Eulerian grid. The resulting
73 velocities are then used to advect the particles, representing
74 the material. Recently, Alduan and Otaduy [13] adapted unilateral
75 incompressibility to the predictive-corrective incompressible
76 SPH (PCISPH) method which was originally designed for
77 incompressible fluid simulation [14].

78 The purely Lagrangian framework proposed in [13] does
79 not suffer from grid artifacts which is a major benefit compared
80 to the Eulerian framework described in [6]. However, it also
81 introduces new challenges. Following the SPH concept, the
82 pressure is computed based on local density values and not according
83 to the divergence of the velocities. This makes the approach
84 very sensitive to sudden increases in the density. Oscillations
85 in the density field particularly occur at interfaces with
86 dynamic solid objects due to particle deficiency. In the context
87 of fluid simulations, this has been addressed in [15, 16]
88 by treating the boundary as an interface to the fluid simulation.
89 Thereby, spatial and temporal discontinuities of physical properties
90 are avoided, resulting in a smoother and a more robust
91 simulation compared to commonly employed distance-based
92 penalty methods, e.g. [17, 18]. Furthermore, in [13], internal
93 forces and advection are computed using the same discretization
94 scale. Thus, in order to compute fine-grained material,
95 significantly higher computational costs are imposed compared
96 to [6].

97 In this paper, we address these challenges by extending the
98 pure Lagrangian, continuum framework [13] in two ways. First,
99 we show that the physically-based rigid-fluid coupling presented
100 in [16] can be easily adapted to handle smooth interactions with
101 granular material. This eliminates oscillations in the pressure

102 field. As a consequence, less iterations are required when solving
103 for the internal forces even at larger time steps. The second
104 extension enables the simulation of fine-grained material
105 at low computational costs by refining the simulation in a post-
106 process.

107 The general idea of refining the coarse simulation for rendering
108 is not new. Most authors propose to sample the simulation particles
109 with a finer set of pseudo-random particles [11, 6]. High-resolution
110 (HR) particles are fixed to the base particles in order to avoid
111 temporal flickering. As HR particles can not disperse freely,
112 such an approach results in clumping artifacts, perceived as
113 staircase or spherical patterns and a distorted distribution of
114 the material. Narain et al. [6] addressed this problem by
115 inserting additional anisotropic particles in regions where the
116 material diverges. This step is already performed during simulation.
117 In contrast, [19, 20] employ a spatial decomposition
118 for the computation of internal and external forces. In sparse
119 regions, HR particles are not passively advected with the base
120 simulation, but respond to external forces. Thereby, clumping
121 artifacts are significantly reduced.

122 Our refinement method is inspired by the decomposition
123 idea, but departs significantly from previous work. In [20],
124 HR particles are either advected according to external forces
125 or along the base flow. In contrast, our model does not rely
126 on two complementary cases, but blends external forces with
127 the interpolated base velocity. This yields more natural looking
128 results, particularly for large upscaling factors. Furthermore,
129 as we take velocities of boundaries into account, HR particles
130 interact smoothly with complex moving objects.

131 3. Coarse Scale Simulation

132 For computing the coarse simulation, we build on the SPH-based
133 continuum method proposed in [13] which is described in
134 Section 3.1 and Section 3.2. Section 3.3 discusses how the convergence
135 of the algorithm can be improved by incorporating a more versatile
136 treatment at the interface with solid boundaries. The description
137 of the base simulation is completed by explaining how to simulate
138 transitions from dry to wet material in Section 3.4. Finally,
139 Section 4 describes how the base simulation is refined, in order
140 to efficiently simulate and render fine-grained material.
141

142 3.1. Forces

143 In SPH, the pressure at a position \mathbf{x}_i is typically determined
144 according to the local density fluctuation $\rho_{err}(\mathbf{x}_i) = \rho(\mathbf{x}_i) - \rho_0$,
145 where ρ_0 is the reference density of the material. The density
146 $\rho_i \equiv \rho(\mathbf{x}_i)$ is computed via the SPH interpolation concept:

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

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

149 Granular flow is governed by unilateral incompressibility
150 which is described by two inequality constraints $\rho \leq \rho_0$ and
151 $p \geq 0$. Accordingly, pressure values p are not negative and
152 the material can not be compressed beyond ρ_0 . These con-
153 straints can be directly plugged into any incompressible SPH
154 solver by clamping negative pressures to zero. The resultant
155 pressure field is then used to compute the pressure forces as

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

156 For simulating granular material, [13] introduces a friction
157 model which minimizes relative velocities measured by the strain
158 rate $\dot{\epsilon}$. The strain rate is computed as $\dot{\epsilon} = 0.5 (\nabla \mathbf{v} + \nabla \mathbf{v}^T)$, where
159 the gradient of the velocity \mathbf{v}_i is given by

$$\nabla \mathbf{v}_i = \sum_j V_j \nabla W_{ij} \mathbf{v}_j^T. \quad (3)$$

Then a frictional stress tensor $\hat{\mathbf{s}}$ is computed that dissipates the
strain rate. In order to simulate material with different angles of
repose θ , the Drucker-Prager yield criterion

$$\sqrt{\sum \hat{s}_{ij}^2} \leq p \sqrt{2} \sin \theta \quad (4)$$

160 is employed. Thereby, friction is limited by pressure, i.e., in the
161 absence of pressure, the frictional stress is zero. The frictional
162 forces are computed with

$$\mathbf{F}_i^f = -m_i \sum_j m_j \left(\frac{\hat{s}_i}{\rho_i^2} + \frac{\hat{s}_j}{\rho_j^2} \right) \nabla W_{ij}. \quad (5)$$

163 3.2. Implementation

164 In order to compute (2) and (5), stress and pressure val-
165 ues have to be determined. For this, we employ the predictive-
166 corrective method presented in [13], which is an adaption of
167 the PCISPH method [14]. In this method, the constraint forces,
168 i.e., pressure and friction forces, are iteratively computed in a
169 Jacobi-like manner as described in the following.

170 Pressure

171 In each iteration l , pressure is related to the predicted den-
172 sity deviation at time $t + \Delta t$ using an equation of state (EOS)
173 as

$$p_i^l(t) = \sum_l \gamma \left(\max \left(0, \rho_i^l(t + \Delta t) - \rho_0 \right) \right), \quad (6)$$

174 where γ is a scaling factor which can be precomputed for a
175 prototype particle with full neighborhood i as

$$\gamma = \frac{\rho_0^2}{2m_i^2 \Delta t^2 \sum_j \nabla W_{ij}^T \nabla W_{ij}}. \quad (7)$$

176 $\rho_i^l(t + \Delta t)$ denotes the predicted density of particle i in iteration
177 l which is computed as

$$\rho_i^l(t + \Delta t) = \sum_j m_j W(\mathbf{x}_i^l(t + \Delta t) - \mathbf{x}_j^l(t + \Delta t), h). \quad (8)$$

178 The predicted positions $\mathbf{x}_i^l(t + \Delta t)$ are computed with

$$\mathbf{x}_i^l(t + \Delta t) = \mathbf{x}_i(t) + \Delta t \mathbf{v}_i(t) + \Delta t^2 \frac{\mathbf{F}_i^{adv}(t) + \mathbf{F}_i^{p^{l-1}}(t) + \mathbf{F}_i^{f^{l-1}}(t)}{m_i}, \quad (9)$$

179 where $\mathbf{F}_i^{p^{l-1}}(t)$ is the pressure force and $\mathbf{F}_i^{adv}(t)$ denotes non-
180 pressure forces such as gravity and cohesion. Please note that
181 $\mathbf{F}_i^{p^{l-1}}(t)$ is a constraint force which uses implicit information,
182 $\mathbf{F}_i^{adv}(t)$ uses exclusively information available at time t and does
183 not change during the iterations. Employing the momentum
184 preserving discretization (2), we can compute the pressure force
185 in iteration l as

$$\mathbf{F}_i^p(t) = -m_i \sum_j m_j \left(\frac{p_i(t)^l}{\rho_i(t)^2} + \frac{p_j(t)^l}{\rho_j(t)^2} \right) \nabla W_{ij}(t). \quad (10)$$

186 Friction

187 Frictional forces constrain the strain rate to be below a user-
188 defined threshold. The discretization (5) is a momentum pre-
189 serving SPH approximation of $\nabla \cdot \hat{\mathbf{s}}_i$. Like for pressure, the
190 unknown stress tensors are iteratively computed in a predictive-
191 corrective manner. Thereby, stress is locally related to the pre-
192 dicted strain rate $\dot{\epsilon}_i^l(t + \Delta t)$ with

$$\mathbf{s}_i^l = \mathbf{s}_i^{l-1} + \mathbf{D}^{-1} \dot{\epsilon}_i^l(t + \Delta t), \quad (11)$$

193 where \mathbf{D} denotes a diagonal matrix which can be precomputed
194 for a prototype particle with full neighborhood

$$\mathbf{D} = \frac{\rho_0^3}{2m_i^2 \Delta t} \sum_j \nabla W_{ij} \nabla W_{ij}^T. \quad (12)$$

195 In (11), the full stress tensor \mathbf{s} is given. However, as in [6,
196 13], friction is expressed by the traceless deviatoric part $\hat{\mathbf{s}}$, only.
197 Furthermore, in order to simulate different angles of repose, the
198 yield criterion (4) has to be satisfied. Employing a piecewise

199 approximation of (4) on each matrix component $s_i(ab)$ as in [6],
 200 the frictional stress is computed as

$$\hat{s}_i^l = s_i^l - \frac{1}{3} \text{trace}(s_i^l) \mathbf{I}_{3 \times 3} \quad \text{with} \quad \|\hat{s}_i^l(ab)\| \leq p_i^l \sqrt{2/3} \sin \theta, \quad (13)$$

201 where $\mathbf{I}_{3 \times 3}$ is the identity matrix. In each iteration, the predicted
 202 strain rate is computed as

$$\dot{\epsilon}_i^l(t + \Delta t) = 0.5 \left(\mathbf{v}_i^l(t + \Delta t) + (\mathbf{v}_i^l(t + \Delta t))^T \right), \quad (14)$$

203 where the predicted velocities in each iteration are evaluated as

$$\mathbf{v}_i^l(t + \Delta t) = \mathbf{v}_i(t) + \Delta t \frac{\mathbf{F}_i^{adv}(t) + \mathbf{F}_i^{p^{l-1}}(t) + \mathbf{F}_i^{f^{l-1}}(t)}{m_i}. \quad (15)$$

204 Algorithm 1 outlines the steps performed in each simulation
 205 update of the base solver. Discrete particle forces are computed
 206 to limit inter-particle distances as in [13].

207 Discussion

208 In this implementation, pressure and frictional constraint
 209 forces are computed simultaneously using a local approach.
 210 The iterative loop can be terminated when the maximum density
 211 $\max_i \rho_i(t + \Delta t)^l$ is below a threshold η and the maximum
 212 variation of the stress tensor is below the threshold ζ . How-
 213 ever, the constraint forces are interdependent. This interdepen-
 214 dency influences the convergence in low energy regions which
 215 prevents stable formations of sand piles, i.e., sliding does not
 216 stop even for large angles of repose. This effect is even more
 217 significant if a third constraint force is employed such as direct-
 218 forcing [18] to handle interactions with boundaries. In order
 219 to simulate stable formations of sand piles, we suggest to ap-
 220 ply the well-known concept of rigid-body sleeping [21] to SPH
 221 particles. Thereby, we do not integrate SPH particles in time as
 222 long as their velocity is lower than a user-defined *sleep velocity*.
 223 Particles wake up automatically when their net-acceleration ex-
 224 ceeds a critical value. We found that this is already sufficient to
 225 stop unwanted sliding at the free surface, but not at the contact
 226 region with solid boundaries. At interfaces with solid objects,
 227 additional considerations have to be employed. The next sec-
 228 tion addresses this issue.

229 3.3. Solid Body Interaction

230 The direct forcing method [18] used in [13] corrects pre-
 231 dicted penetrations by constraining positions to the rigid sur-
 232 face. Different slip conditions are realized by directly manip-
 233 ulating particle velocities. This model does not conserve mo-
 234 mentum and leads to significant oscillations in the pressure field
 235 as shown for fluids in [16]. Additionally, for the granular model
 236 described in Sec. 3.1, the noisy pressure field results in disconti-
 237 nuities of the frictional forces (5). In order to avoid perceivable
 238 artifacts, a small time step has to be employed and/or the num-
 239 ber of iterations for correcting density errors has to be set high,
 240 i.e., larger than five.

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
 $l = 0$ 
while  $\max_i \rho_i(t + \Delta t)^l > \eta \vee \max_i \|\mathbf{s}_i^l - \mathbf{s}_i^{l-1}\| > \zeta$  do
  | foreach particle  $i$  do
  | | predict velocity and position (15) and (9)
  | | foreach particle  $i$  do
  | | | predict density  $\rho_i^l(t + \Delta t)$ 
  | | | if  $\rho_i^l(t + \Delta t) > \rho_{max}$  then
  | | | | predict density and strain rate (8) and (14)
  | | | | update pressure and stress (6) and (13)
  | | | else
  | | | | add discrete particle forces
  | | | foreach particle  $i$  do
  | | | | compute pressure force (10) and (17)
  | | | | compute friction force (5) and (21)
  | |  $l+ = 1$ 
  | foreach particle  $i$  do
  | | update velocity and position

```

241 We improve the robustness and the quality of the simula-
 242 tion by adapting the boundary handling method of [16] to in-
 243 teractions with granular material. Thereby, the surface of rigid
 244 objects is sampled with particles as in [10]. Boundary particles
 245 contribute to the density of a granular particle. Accordingly, (1)
 246 is rewritten as

$$\rho_i = \sum_j m_j W_{ij} + \sum_b \frac{\rho_0}{\delta_b} W_{ib}, \quad (16)$$

247 where $\delta_b = \sum_j W_{bj}$ is the number density of a boundary particle
 248 b , summed up over all boundary particle neighbors j . $\Psi_b(\rho_0) \equiv$
 249 $\frac{\rho_0}{\delta_b}$ scales the contributions of a boundary particle according to
 250 the local boundary sampling and the reference density of the
 251 granular material. Thereby, (16) estimates the density cor-
 252 rectly at boundaries for arbitrary samplings and material den-
 253 sities. Non-penetration is realized via pressure forces, acting
 254 from a boundary particle b on a granular particle i with

$$\mathbf{F}_{i \leftarrow b}^p = -m_i \Psi_b(\rho_0) \frac{p_i}{\rho_i^2} \nabla W_{ib}. \quad (17)$$

255 A detailed derivation can be found in [16].

256 We extend this model to interactions with granular materi-
 257 als by adding contributions of the boundary when computing
 258 frictional stresses. Therefore, we modify (3) to

$$\nabla \mathbf{v}_i = \sum_j V_j \nabla W_{ij} \mathbf{v}_j^T + \sum_b \frac{1}{\delta_b} \nabla W_{ib} \mathbf{v}_b^T \quad (18)$$

259 The reformulation of (5) can be derived similar to (17) as

$$\mathbf{F}_{i \leftarrow b}^{fs} = -m_i \Psi_b(\rho_0) \frac{S_i}{\rho_i^2} \nabla W_{ib}. \quad (19)$$

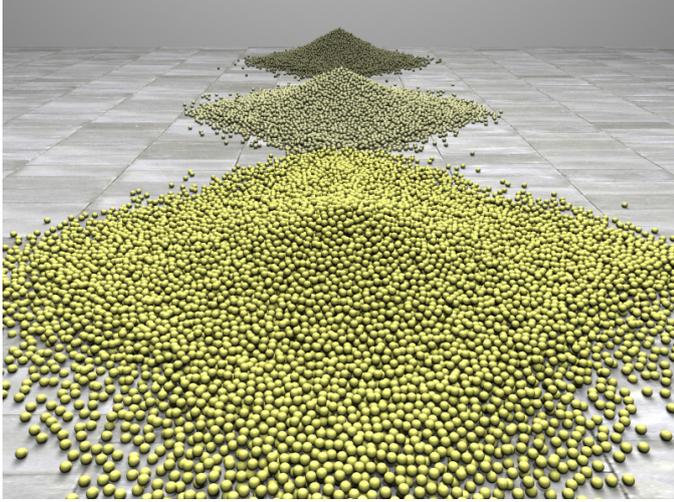


Figure 2: Sand piles with different angles of repose 30° , 45° and 60° (front to back).

260 However, (19) 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 external friction forces, we additionally employ the artificial viscosity model given in [16] which is written as

$$\mathbf{F}_{i \leftarrow b}^v = -m_i \Psi_b(\rho_0) \Pi_{ib} \nabla W_{ib}, \quad (20)$$

265 where $\Pi_{ib} = -\frac{\sigma_{ib} h c_s}{2\rho_i} \left(\frac{\min(\mathbf{v}_{ib} \cdot \mathbf{x}_{ib}, 0)}{|\mathbf{x}_{ib}|^2 + \epsilon h^2} \right)$. 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 (19) and (20) 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

$$\mathbf{F}_{i \leftarrow b}^f = \max(\|\mathbf{F}_{i \leftarrow b}^v\|, \|\mathbf{F}_{i \leftarrow b}^s\|), \quad (21)$$

272 Two-way coupling is easily realized by applying pressure (17) and friction (21) forces symmetrically to the boundary.

274 3.4. Material Properties

275 The simulation model described so far is applicable to simulations of granular material with different angles of repose, see Fig. 2. We further extend the granular simulation model to allow for simulations of dry and wet sand with controllable scaling. In dry sand volumes, air voids can form between grains, which reduces friction and allows the material to disperse freely. In contrast, moist sand volumes exhibit a higher friction between grains which results in a sticky behavior. As a consequence, compared to dry material, a larger force needs to be exerted to make the wet material flow. This phenomenon can not be simulated by the frictional force (5) alone as this force acts only in tangential direction. In [13], the model of Gascon et al. [22] is adapted in order to simulate cohesive effects as a function of the norm of the strain rate. However, this model imposes a severe restriction on the time step as mentioned in [13].

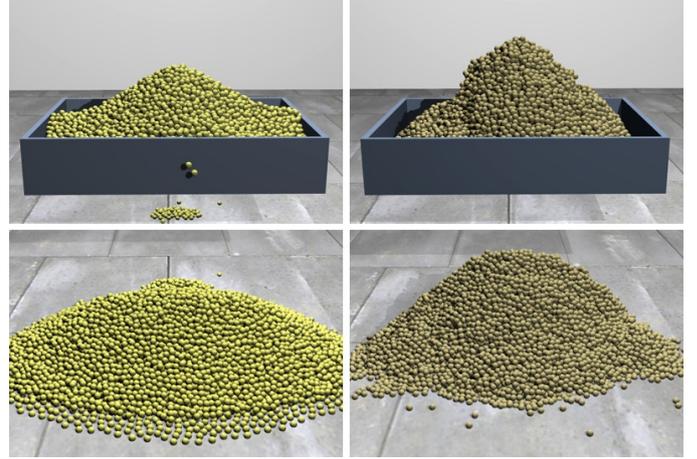


Figure 3: Simulation of dry sand (left) and wet sand (right). For both volumes the angle of repose is 50° . The materials differ in the cohesion intensity and the material viscosity. The wet volume has a cohesion intensity of 0.7 and the viscosity is set to 0.3. The cohesion intensity of the dry volume was set to zero and the viscosity was set to a minimal value of 0.01.

290 In contrast, we found that the surface tension model for SPH fluids presented in [23] is well suited to model cohesive effects for granular material without imposing restrictions on the time step. Thereby, surface tension is modeled as a sum of pairwise forces defined as

$$\mathbf{F}_i^c = -\kappa \sum_j m_j W_{ij} \mathbf{x}_{ij} \quad (22)$$

295 where κ controls the cohesion intensity. In order to model higher friction for wet volumes, we propose to increase the artificial viscosity constant of the granular material. We found that this gives better control and is numerically more robust than increasing the angle of repose.

300 The effectiveness of this model to simulate wet material is demonstrated on a simple scenario, where we simulated a dry and a wet volume, see Fig. 3. The volumes differ only in the cohesion intensity and the viscosity constant. The dry material is simulated with no cohesion and an artificial viscosity constant of 0.01. For the wet volume, the cohesion intensity is set to 0.7 and the viscosity to 0.3. In order to simulate wetting processes when coupling the granular material with an SPH fluid, the transition from dry to wet sand can be realized by just increasing the cohesion intensity and the viscosity constant of the granular material.

311 3.5. Parameters

312 For the SPH interpolations, we use the cubic spline kernel [24]. The support radius h is chosen as two times the average particle distance d_0 for the material at rest. Accordingly, we precompute the mass of each SPH particle with $m = V\rho_0$, where $V = \frac{1}{6}\pi d_0^3$ is the particle volume of a sand grain. The yield density ρ_0 is user given and can vary for different materials, e.g., for dry sand, we set ρ_0 to 1602 kg/m^3 .

319 4. Fine Scale Simulation

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

331 4.1. Sampling

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

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

349 4.2. Advection

350 We derive the advection of HR particles from the follow-
 351 ing principles: HR particles should follow the mechanical flow
 352 that is given by the base simulation, but also should be allowed
 353 to disperse freely. Further, they should smoothly align with
 354 the surface of the base simulation without forming perceivable
 355 clumps. Finally, in order to guarantee efficient updates at large
 356 time steps, the advection method should not compute internal
 357 forces or perform collision tests between HR particles.

358 Alduan et al. [20] set similar requirements to their model.
 359 They map the mechanical behavior by interpolating the veloc-
 360 ity of LR particles to HR particles for advection. In order to
 361 avoid clumping, particles having one or no LR particle neigh-
 362 bors within the influence radius h_{HR} are only influenced by ex-
 363 ternal forces. However, this simple distinction avoids clumping
 364 only if $h_{HR} \approx r_{LR}$. On the other hand, larger values of h_{HR}
 365 are required for smooth interpolations of the velocities. In all our
 366 experiments, we set $h_{HR} = 3r_{LR}$.

367 In contrast to [20], we do not employ an explicit distinction
 368 of two cases, but propose a weighting that automatically and
 369 smoothly blends the contributions of the base simulation and

370 external forces. Therefore, for each HR particle at position \mathbf{x}_i ,
 371 we first compute distance-based weights w as

$$w(d_{ij}) = \max \left[0, \left(1 - \frac{d_{ij}^2}{h_{HR}^2} \right)^3 \right] \quad (23)$$

372 where $d_{ij} = |\mathbf{x}_i - \mathbf{x}_j|$ is the distance between HR particles and
 373 LR simulation or boundary particles j in the support radius h_{HR} .
 374 (23) is a well-shaped kernel function which smoothly drops to
 375 zero. It is typically applied for reconstructing smooth surfaces
 376 of particle data [9, 25]. We employ it for computing the average
 377 velocity as

$$\mathbf{v}_i^*(t + \Delta t) = \frac{1}{\sum_j w(d_{ij})} \sum_j w(d_{ij}) \mathbf{v}_j \quad (24)$$

378 As long as the number of LR samples is sufficient, (24) in-
 379 terpolates the velocity well. However, in sparsely sampled re-
 380 gions it results in visual clumping. In order to achieve smooth
 381 alignments of HR particles without perceivable clumps, we com-
 382 pute the velocities of HR particles as the sum of weighted ex-
 383 ternal forces \mathbf{F}^g and the interpolated velocity with

$$\mathbf{v}_i(t + \Delta t) = (1 - \alpha_i) \mathbf{v}_i^*(t + \Delta t) + \alpha_i \left(\mathbf{v}_i(t) + \Delta t \frac{\mathbf{F}^g}{m} \right), \quad (25)$$

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

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

387 Here, the constant 0.6 is an empirically tested value which gives
 388 the best results when $h_{HR} = 3r_{LR}$. Finally, the position is inte-
 389 grated as

$$\mathbf{x}_i(t + \Delta t) = \mathbf{x}_i(t) + \Delta t \mathbf{v}_i(t + \Delta t). \quad (26)$$

390 Accordingly, external forces are automatically applied in re-
 391 gions where clumping potentially occurs. Contributions of ex-
 392 ternal forces are smoothly faded in and out, which on one hand
 393 allows HR particles to disperse freely and on the other hand re-
 394 sults in smooth alignment with the materials surface, see Fig. 4.
 395 Our model does faithfully upscale scenes with dynamic objects,
 396 as we take the positions and velocities of moving objects into
 397 account when interpolating the velocities.

398 5. Results

399 We present several scenarios which show different aspects
 400 of our approach including one-way and two-way solid body in-
 401 teraction. Comparison to previous work and performance eval-
 402 uation are also provided. Timings are given for a 12-core 3.46
 403 GHz Intel i7 with 24 GB of RAM, see Table 1. We make use of
 404 all threads by incorporating the parallel algorithms and cache-
 405 efficient data structures described in [26] using OpenMP. Im-
 406 ages were rendered with mental ray v3.9.4 [27]. The video se-
 407 quences are encoded with 50 frames per second. Thus, one
 408 frame corresponds to 0.02s in the following discussions.

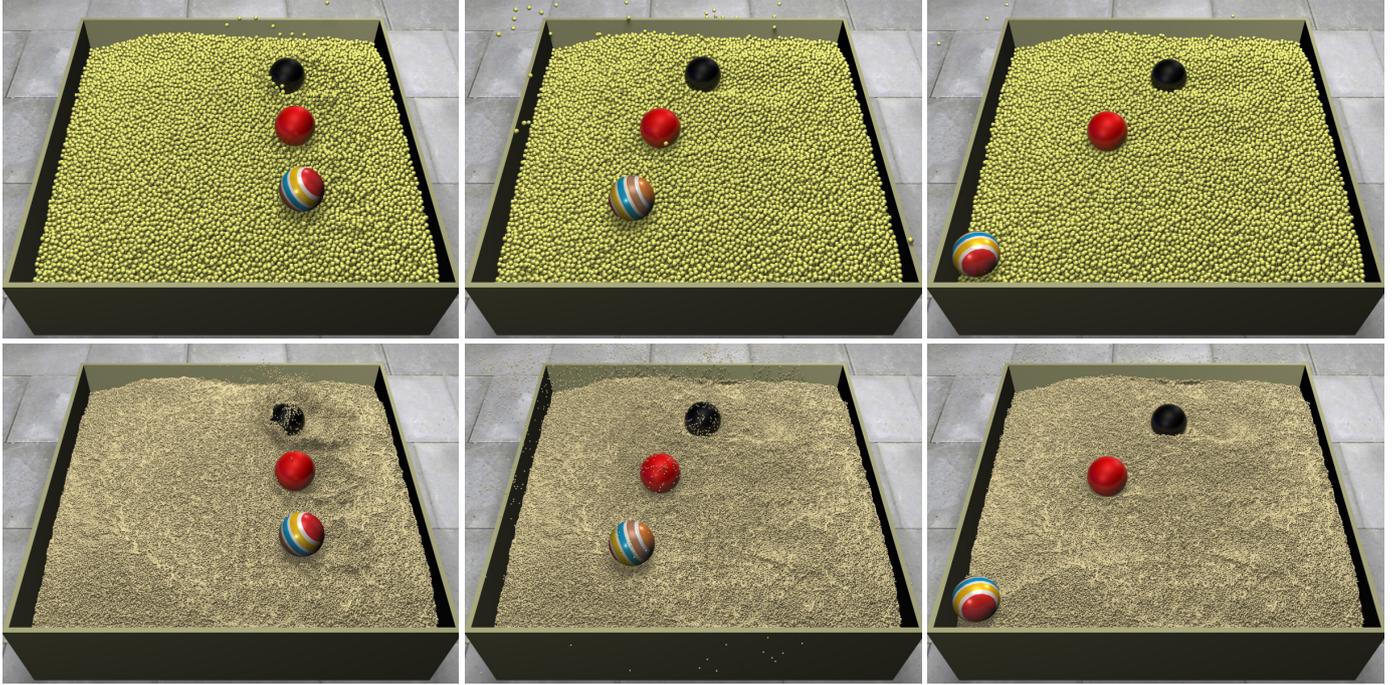


Figure 5: Demonstration of the proposed two-way coupling. Three spheres with different masses are tossed into a sand pool. Due to the friction-based coupling, the sand supports also the heaviest sphere (black) which is three times more dense than the sand. The top row shows the base simulation with 96K particles, the bottom row shows the same frames for the secondary simulation with 11.3M particles.

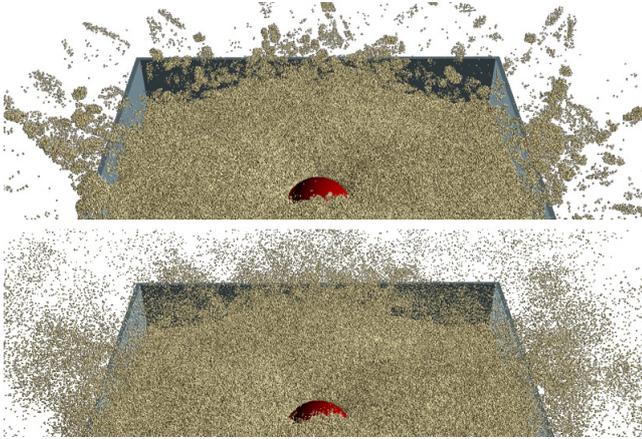


Figure 4: Upsampling comparison. In [20], HR particles can not disperse freely if two or more base particles are in close proximity (top). Our method avoids clumping by weighting external and internal forces for all HR particles (bottom).

420 the accompanying video.

421 Two-way coupling between rigid objects and granular material is demonstrated in Fig. 5. Rigid objects interact differently
 422 with the granular material according to their density since the
 423 dissipation of energy for lighter objects is much less than for
 424 heavy objects. In contrast to fluids, granular material supports
 425 even objects with much higher material density than its own.
 426 This behavior is faithfully captured as the employed coupling
 427 is based on frictional stresses.
 428

429 5.2. Upsampling

430 In order to show the benefit of the proposed refinement
 431 model, we compare it to [20] on a scene where a heavy sphere
 432 hits the sand surface with high velocity, creating a splash. The
 433 base simulation sampled the material with 100K particles while
 434 the secondary simulation used 20.8M particles. Since in [20],
 435 HR particles with more than one LR neighbor are not influenced
 436 by external forces, cluster of HR particles move uniformly in
 437 splash regions and at the sand surface, see Fig. 4-top. In con-
 438 trast, our method does not rely on complementary cases to com-
 439 pute the velocity, but weights external forces and base velocities
 440 using a well shaped kernel function. Accordingly, HR particles
 441 are allowed to disperse freely, see Fig. 4-bottom.

442 Opposed to previous work, we show that HR particles inter-
 443 act smoothly with moving objects, e.g., a bulldozer, see Fig. 1.
 444 This is realized by taking positions and velocities of boundary
 445 particles into account when interpolating the velocity field. We
 446 further demonstrate that our method can cope with very large
 447 upscaling factors. Fig. 1-middle shows the refinement with an
 448 upscaling factor of 38 which corresponds to the largest value

409 5.1. Solid Body Interaction

410 We demonstrate the effectiveness of the proposed friction
 411 and pressure forces applied from and to the boundaries on a
 412 simple test scenario, see Fig. 2. In this scene, we simulated
 413 three stable piles with different angles of repose. We set the
 414 time step to 1 millisecond and used a fixed number of three it-
 415 erations to compute pressure and stress values. For this setting,
 416 the proposed boundary handling enforces smooth pressure gra-
 417 dients at the boundary. In contrast, if we employ direct forcing
 418 as in [13], the density field oscillates which results in unnat-
 419 ural accelerations perceived as 'popping of particles' as is shown in

	# 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. 5)	96K	11.3M	3.1 s	5.8 s
Splash (Fig. 4)	83K	20.8M	2.9 s	10.3 s
Sand-Piles (Fig. 2)	45K	-	1.6 s	-
Rope-Ladder (Fig. 6)	138K	17.3M	4.6 s	8.3 s

Table 1: Performance measurements for the given scenarios.

used in [20]. 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. 6.

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 [13]. 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 [13].

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. [20] 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.

5.4. Resolution Scaling

The effect of the frictional force employed in this framework is not invariant to the spatial and temporal discretization. This is demonstrated on a simple scenario where a dry material with an angle of repose of 55° is simulated at three different resolutions, see Fig. 7. The coarse simulation with 7K particles and a radius of $0.05m$ ran at a time step of $1ms$. The same time step could be set for the medium resolved simulation with 60K particles. The high-resolution simulation with 500K particles required a much smaller time step of $0.1ms$. The rest angles of the sand piles obtained by the simulation were 31° (7K), 35°

(60K) and 45° (500K). Thus, non of the simulations obtained the predefined angle of repose of 55° . Although this series indicates that the simulated angle of repose might converge to the desired value, this could not be verified as we could not find a practical time step for higher resolutions.

Due to the difference in the tolerated time step size, we measured a huge difference in the overall performance. One video frame for the different simulations could be computed in $0.3s$ (7K), $2.6s$ (60K) and $98s$ (500K). In order to evaluate the quality of the proposed refinement scheme, we also compare the simulation result of the highly resolved base simulation with the upsampled coarse simulation with 500K secondary particles, see Fig. 7, bottom. The secondary simulation took only $0.8s$ per frame to refine the 7K simulation. Thus, in total it took 1.1 seconds ($0.3s$ base simulation + $1.1s$ upsampling) to compute a frame with the proposed pipeline. This is almost 90 times faster than the time required for the highly resolved base simulation. The perceivable difference between the two simulations can be attributed to the differences in the base simulation (7K vs 500K) where we measured a deviation of 14° for the obtained angle of repose. On the other hand, this simple scenario is challenging for the secondary simulation as the movement of the granular material is slow and nearly uniform. Furthermore, the base simulation is very coarsely sampled with 7K particles. For this configuration, clustering in the refined simulation is perceivable.

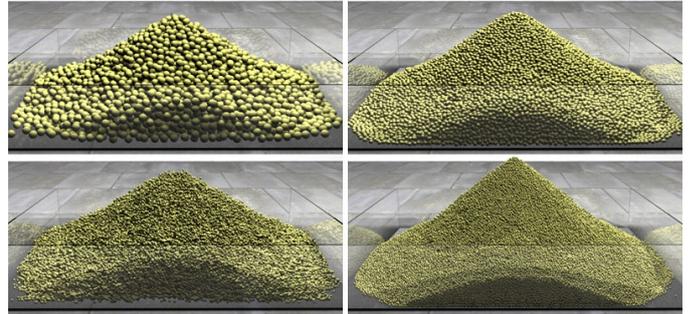


Figure 7: Base simulation with different resolutions. A coarse simulation of 7K particles with radius $r = 0.05m$ (top-left), 60K particles with $r = 0.25m$ (top-right) and a highly resolved simulation using 500K particles with $r = 0.125m$ (bottom right). The bottom left image shows the refined coarse simulation with 500k secondary particles.

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. We showed how to realize stable and realistic interactions with rigid bodies by employing pressure and friction-based boundary forces. The stability of the base simulation can be significantly improved by incorporating the concept of rigid-body sleeping. This also stops unwanted sliding of sand piles. We have further extended the base simulation to allow for simulation of dry and wet sand with controllable scaling.



Figure 6: 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.

528 Visual detail is added in a secondary simulation where high-
 529 resolution particles are coupled to the base flow. For advecting
 530 secondary particles, we propose a smooth weighting of external
 531 forces with the velocity field of the base simulation. This tech-
 532 nique adds detail that is not captured otherwise while clump-
 533 ing of high-resolution particles is avoided, even for very large
 534 upscaling factors. We provided a thorough discussion of the
 535 implementation and of practical parameters.

536 6.1. Limitations and Future Work

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

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

552 Massive conglomerations of granular material, e.g., a beach,
 553 are more efficiently represented by heightfields [5, 28]. How-
 554 ever, in these methods, the level of detail is limited as only two-
 555 dimensional information is mapped onto the three-dimensional
 556 space. Dispersion effects or the sliding of single grains can not
 557 be captured with such a representation.

558 In future work, the granular material could be coupled with
 559 an SPH fluid simulation in order to animate erosion effects and

560 transitions from dry to moist sand and mud. In this context, we
 561 plan to investigate how SPH fluid simulations can benefit from
 562 the proposed refinement model.

563 Acknowledgements

564 This project is supported by the German Research Founda-
 565 tion (DFG) under contract numbers SFB/TR-8 and TE 632/1-2.
 566 We also thank NVIDIA ARC GmbH for supporting this work.

567 References

- 568 [1] Bicanic N. Discrete element methods. Encyclopedia of Computational
 569 Mechanics 2004;1. doi:10.1002/0470091355.ecm006.
- 570 [2] Josserand C, Lagrée PY, Lhuillier D. Stationary shear flows of dense
 571 granular materials: a tentative continuum modelling. The European Phys-
 572 ical Journal E: Soft Matter and Biological Physics 2004;14(6):127–35.
- 573 [3] Li X, Moshell JM. Modeling soil: Realtime dynamic models for soil
 574 slippage and manipulation. In: SIGGRAPH 1993 Proceedings. 1993,
 575 p. 361–8.
- 576 [4] Sumner RW, O’Brien JF, Hodgins JK. Animating Sand, Mud, and Snow.
 577 Computer Graphics Forum 1999;28:1–11.
- 578 [5] Holz D, Beer T, Kuhlen T. Soil Deformation Models for Real-Time Sim-
 579 ulation: A Hybrid Approach. In: VRIPHYS 09: Sixth Workshop in Vir-
 580 tual Reality Interactions and Physical Simulations. 2009, p. 21–30. doi:
 581 10.2312/PE/vrphys/vrphys09/021-030.
- 582 [6] Narain R, Golas A, Lin MC. Free-flowing granular materials with
 583 two-way solid coupling. ACM Trans Graph (SIGGRAPH Proc)
 584 2010;29(6):173:1–173:10. doi:10.1145/1882261.1866195.
- 585 [7] Ihmsen M, Wahl A, Teschner M. High-Resolution Simulation of Granular
 586 Material with SPH. In: Proc. VRIPHYS. 2012, p. 53–60.
- 587 [8] Luciani A, Habibi A, Manzotti E. A multi scale physical model of gran-
 588 ular materials. In: Proc. Graphics Interface. 1995, p. 136–46.
- 589 [9] Zhu Y, Bridson R. Animating sand as a fluid. ACM Trans Graph (SIG-
 590 GRAPH Proc) 2005;24(3):965–72. doi:10.1145/1073204.1073298.
- 591 [10] Bell N, Yu Y, Mucha PJ. Particle-based simulation of granular materials.
 592 In: Proc. of the 2005 ACM SIGGRAPH/Eurographics Symposium on
 593 Computer Animation. SCA ’05; New York, NY, USA: ACM. ISBN 1-
 594 59593-198-8; 2005, p. 77–86. doi:10.1145/1073368.1073379.

- 595 [11] Lenaerts T, Dutre P. Mixing fluids and granular materials.
596 Computer Graphics Forum 2009;28(2):213–8. doi:10.1111/j.1467-
597 8659.2009.01360.x.
- 598 [12] Müller M, Charypar D, Gross M. Particle-based fluid simulation for in-
599 teractive applications. In: Proc. of ACM SIGGRAPH/Eurographics Sym-
600 posium on Computer Animation. 2003, p. 154–9.
- 601 [13] Alduán I, Otaduy MA. Sph granular flow with friction and cohesion.
602 In: Proc. of the 2011 ACM SIGGRAPH/Eurographics Symposium on
603 Computer Animation. ISBN 978-1-4503-0923-3; 2011, p. 25–32. doi:
604 10.1145/2019406.2019410.
- 605 [14] Solenthaler B, Pajarola R. Predictive-corrective incompressible SPH.
606 ACM Trans on Graphics (Proc SIGGRAPH) 2009;28(3):1–6.
- 607 [15] Ihmsen M, Akinci N, Gissler M, Teschner M. Boundary Handling and
608 Adaptive Time-stepping for PCISPH. In: Proc. VRIPHYS. 2010, p. 79–
609 88.
- 610 [16] Akinci N, Ihmsen M, Akinci G, Solenthaler B, Teschner M. Versatile
611 Rigid-Fluid Coupling for Incompressible SPH. ACM Trans on Graphics
612 (Proc SIGGRAPH) 2012;.
- 613 [17] Müller M, Schirm S, Teschner M, Heidelberger B, Gross M. Interaction
614 of fluids with deformable solids. Computer Animation and Virtual Worlds
615 2004;15(34):159–71.
- 616 [18] Becker M, Tessendorf H, Teschner M. Direct forcing for lagrangian rigid-
617 fluid coupling. IEEE Transactions on Visualization and Computer Graph-
618 ics 2009;15(3):493–503.
- 619 [19] Ammann C, Bloom D, Cohen JM, Courte J, Flores L, Hasegawa S, et al.
620 The birth of sandman. In: ACM SIGGRAPH 2007 sketches. New York,
621 USA: ACM; 2007,doi:10.1145/1278780.1278812.
- 622 [20] Alduán I, Tena A, Otaduy MA. Simulation of High-Resolution Granular
623 Media. In: Proc. of Congreso Español de Informática Gráfica. 2009, p.
624 1–8.
- 625 [21] Schmidl H. Optimization-Based-Animation. Ph.D. thesis; University of
626 Miami; 2002.
- 627 [22] Gascón J, Zurdo JS, Otaduy MA. Constraint-Based Simulation of Adhe-
628 sive Contact. In: Proc. of the ACM SIGGRAPH / Eurographics Symposi-
629 um on Computer Animation. 2010;.
- 630 [23] Becker M, Teschner M. Weakly compressible SPH for free surface flows.
631 In: Proc. of the 2007 ACM SIGGRAPH/Eurographics Symposium on
632 Computer Animation. 2007, p. 209–17.
- 633 [24] Monaghan J. Smoothed particle hydrodynamics. Reports on Progress in
634 Physics 2005;68(8):1703–59.
- 635 [25] Akinci G, Ihmsen M, Akinci N, Teschner M. Parallel Surface Recon-
636 struction for Particle-Based Fluids. Computer Graphics Forum 2012;doi:
637 10.1111/j.1467-8659.2012.02096.x.
- 638 [26] Ihmsen M, Akinci N, Becker M, Teschner M. A Parallel SPH Implemen-
639 tation on Multi-core CPUs. Computer Graphics Forum 2011;30(1):99–
640 112. doi:10.1111/j.1467-8659.2010.01832.
- 641 [27] NVIDIA . mental ray (version 3.9). <http://www.mentalimages.com>
642 2011;URL <http://www.mentalimages.com>.
- 643 [28] Kristof P, Benevs B, Krivanek J, Stava O. Hydraulic Erosion Using
644 Smoothed Particle Hydrodynamics. Computer Graphics Forum (Proceed-
645 ings of Eurographics 2009) 2009;28(2).