

A Convex Complementarity Approach for Simulating Large Granular Flows

Alessandro Tasora*

Università degli Studi di Parma,
Dipartimento di Ingegneria Industriale, 43100 Parma, Italy

Mihai Anitescu†

Mathematics and Computer Science Division
Argonne National Laboratory,
9700 South Cass Avenue, Argonne, IL 60439, US

Abstract

Aiming at the simulation of dense granular flows, we propose and test a numerical method based on successive convex complementarity problems. This approach originates from a multibody description of the granular flow: all the particles are simulated as rigid bodies with arbitrary shapes and frictional contacts. Unlike the discrete element method (DEM), the proposed approach does not require small integration time steps typical of stiff particle interaction; this fact, together with the development of optimized algorithms that can run also on parallel computing architectures, allows an efficient application of the proposed methodology to granular flows with a large number of particles. We present an application to the analysis of the refueling flow in pebble-bed nuclear reactors. Extensive validation of our method against both DEM and physical experiments results indicates that essential collective characteristics of dense granular flow are accurately predicted.

1 Introduction

During the last decade, research on the dynamical behavior of granular matter has attracted significant attention from the scientific community. Simulations of granular dynamics can be particularly useful in many fields of engineering, such as mining, mineral processing, pharmaceuticals, oil and gas extraction, rock soil dynamics, agriculture, and food handling. Many devices that deal with crushing, transporting, or mixing particles are still difficult or impossible to simulate because of the large number of moving parts.

*Corresponding author: Address: Università degli Studi di Parma Dipartimento di Ingegneria Industriale, V.G.Usberti 181/A, 43100 Parma, Italy; phone: 0039-0521-905895; e-mail:tasora@ied.unipr.it

†e-mail:anitescu@mcs.anl.gov

On the scale of individual grains, granular materials follow the laws of classical mechanics: particles are subject to many complex collisions and contacts. On mesoscopic and macroscopic scales, however, their behaviour is different from classical phases such as fluids, solids, or gases [1]. The discontinuities that can affect granular materials on a microscopic level introduce a level of complexity that is not met in the simulation of other porous media such as foams, catalysts, or sponges, which do not feature grain-level diffusion and mixing [2].

For instance, size segregation devices, shakers, and rotating drums for the pharmaceutical industry exhibit pattern formation phenomena such as segregation and stratification of species that can be observed at a mesoscopic scale. On macroscopic scales other complex nonlinear phenomena exist whose numerical simulation is still challenging, such as in flows of silos and hoppers, in landslides and avalanches, and in ground compaction under the pressure of vehicle tires.

To simplify the problem, many researchers attempted to derive continuum models for dense granular flows. These approaches can be helpful in describing the average behavior of granular materials, but a general theory is still lacking [3] [4].

On the other side are brute-force approaches that are biased toward the microscopic description of the material, favoring a realism over the cost of an increased computational burden. Among these approaches the most used is the discrete element method (DEM) of Cundall and Strack [5], where the granular medium is considered as collection of interacting discrete objects; each object moves in response to the total force that results from body forces, such as gravity and inertia, and repulsive forces caused by contacts with other particles [6]. In most DEM embodiments, the critical part is the description of the contact forces: the Hertz contact theory for contacting spheres can be used to introduce spring-dashpot forces between the objects in contact, and additional force fields can be used to represent the friction forces. In many scenarios, however, the particles feature high stiffness, thus leading to differential problems that require extremely small time steps in order to avoid divergence during time integration [7]. Faster simulations could be achieved at a cost of using less stiff values, but this heuristic shortcut leaves the difficulty of tuning proper values for the contact coefficients [8].

In this paper we propose a multibody method that can simulate a very large number of objects; but, unlike the DEM approach, our method enforces contacts with complementarity constraints, leading to a differential variational inequality (DVI) problem [9]. Although the mathematical formulation of DVIs is more complex than that of the ordinary differential equations (ODEs) embedded in DEM methods, the DVI has the remarkable advantage of allowing the use of larger time steps because contacts are enforced geometrically. Moreover, knowledge of real or heuristic stiffness of the particles is not required, since the bodies can be considered rigid [10] [11] [12].

We solve the original DVI by using a time-stepping scheme. The difficult part is, at each time step, the solution of a nonlinear complementarity problem (NCP) resulting from the optimality conditions of an optimization problem that, given the Coulomb friction model, poses complex nonlinear constraints on contact forces. Previous researches in this area tended to adapt solvers for linear complementarity problems (LCPs), at the cost of approximating the Coulomb friction cones with faceted pyramids. Moreover, most available LCP solvers are of direct type and cannot solve problems with many unknowns [13] [14].

Instead, we recently developed an iterative matrixless method that can solve the NCP, in

form of a cone complementarity problem (CCP), with millions of unknowns and with moderate computational efforts [15, 16]. Like other iterative complementarity solvers, [17], our method requires the repetition of a contraction mapping that is stopped once the convergence is reached within a reasonable accuracy.

We successfully used our approach to simulate systems of bodies with half a million of contacts, thus suggesting its adoption in the simulation of granular flows where thousands or millions of rigid particles with arbitrary shapes move with frictional contacts.

Since the stiffness from the spring and dashpot is absent, our method can take much larger time steps. On the other hand, the CCP problem is clearly more complex than the very few force evaluation that are needed by one time step of the DEM method. Therefore the computational benefit from our method is a function of how fast we can solve the CCP.

For small problems the faster time-to-solution of CCP methods was clearly demonstrated [8]. Our estimates of the CCP timing for the pebble bed simulation result in a factor of 300 speedup compared to results reported with DEM [7]. In addition the reference [7] mentions the fact that computational considerations have required taking the spring constant a factor of 10^5 smaller than supported by Hertzian contact theory. Nevertheless, such comparisons are not yet strong indicators, since the results are run on different architectures, different processors, and quite clearly with different tolerances.

Assessing the relative behavior of DEM and CCP in terms of time-to-solution for given tolerance of the solution methods on the same architecture is an important endeavor with a high development cost. To decide if that path is worth pursuing we first settle on a question that we believe it is of a more immediate importance: *are CCP methods able to predict the behavior of dense granular flow at all?* While theoretical analysis does demonstrate weak convergence to a solution [18], there is the nontrivial risk that the convergence concept is too weak (since non-uniqueness of solutions may occur, as rigid body frictional configurations often exhibit [19]) or that the time steps we take are too large for the regimes we are trying to simulate. In particular, it should be pointed out that the algorithm in [18] forces in some sense all collisions to occur at the end of the time step which is the main reason why it can proceed with fixed time step without the need of stopping at collisions and why the effort per time step is predictable. While it has been proven that the resulting algorithm is stable [18, 20], there is the conceptual risk that the convenient time step is too large for accurate prediction of the granular material behavior.

We therefore turn our sights to validation, that is, assessing whether CCP methods produce good agreement with the experiment and DEM method in terms of the macroscopic behavior of dense granular flow. We choose macroscopic behavior metrics – such as aggregate volume fraction or mean velocity – since the behavior of individual particles in vertical granular channels is known to be chaotic [21] and thus agreement among the various methods in particle-by-particle prediction is unlikely to occur.

To validate the approach, in this paper we compare our results to data from experiments and from other authors [2] [22]. We also discuss the application of our method to a study of the fuel flow in a fourth-generation pebble-bed nuclear reactor (PBR). Our CCP approach compares favourably to the DEM method in terms of increased computational efficiency [7].

2 The Pebble-Bed Reactor

The Pebble-Bed nuclear reactor is one of the most promising designs of the fourth generation reactors [23]. It is a high-temperature reactor, with an output temperature exceeding 1000 °C, which makes it highly suitable for a hydrogen economy. In addition, it can run continuously for years, without needing to be stopped for refueling, which is the case in the currently used light water reactors. By continuously recirculating the fuel through the reactor core (Figure 1), the PBR dramatically reduces operation costs [24].

The fuel is encased in tennis-ball-size graphite spheres, each filled with TRISO nuclear fuel (tristructural-isotropic particles of coated UO_2 , with sub-millimeter diameter). The $\approx 400,000$ pebbles are continuously recirculated or refreshed at a rate of about 2 per minute [22]. They are densely packed, at volume fractions approaching 0.6, and thus constitute a dense granular flow [25]. The center pebbles are moderator pebbles with comparable weight to the fuel pebbles, even if they do not contain UO_2 . The reactor is cooled with a fast helium flow that is blown top-down, which has negligible drag effects on the spheres when compared to gravitational forces [25]. Because of its use of graphite, the reactor is passively safe, being far from meltdown even in the case of total loss of coolant [25].

Predicting the distribution of the fuel pebbles in the pebble-bed reactor is crucial for its safety and performance properties [26], hence motivating the current research.

Specifically, our focus is on the efficient prediction of the ensemble behavior of the pebbles, under normal gravitational field.

3 Granular Flow Model

We assume that the granular material is composed of a set of rigid particles, whose relative motion is constrained by rigid contacts with friction. Therefore, the behavior of the material depends on a few simple parameters such as the geometric shapes of the particles and their friction coefficients. In many cases (simulation of stiff particles such as sand or pebbles, for instance) such a rigid-contact assumption represents a welcome simplification with respect to the DEM approach, because we do not need to introduce complex constitutive laws or penalty functions that model the contact, and because there is no unwanted compliance in the contacts.

In the following we adopt the mathematical formulation that we presented in [16]. Assume n_b is the number of particles. For the j -th particle, \mathbf{r}_j is the absolute position of the center of mass, and the unimodular quaternion $\epsilon_j \in \mathbf{S}^3$ represents the rotation. The state of the granular material is represented by the large vector $\mathbf{q} = [\mathbf{r}_1^T, \epsilon_1^T, \dots, \mathbf{r}_{n_b}^T, \epsilon_{n_b}^T]^T \in \mathbb{R}^{7n_b}$ along with the time derivative $\dot{\mathbf{q}} = [\dot{\mathbf{r}}_1^T, \dot{\epsilon}_1^T, \dots, \dot{\mathbf{r}}_{n_b}^T, \dot{\epsilon}_{n_b}^T]^T \in \mathbb{R}^{7n_b}$. Optionally one can avoid using quaternion derivatives in $\dot{\mathbf{q}}$ because, remembering $\dot{\epsilon}_i = \frac{1}{2}\mathbf{G}^T(\mathbf{q})\bar{\omega}_i$ with 3×4 matrix $\mathbf{G}(\mathbf{q})$ as defined in [27], it is possible to introduce the generalized velocities $\mathbf{v} = [\dot{\mathbf{r}}_1^T, \bar{\omega}_1^T, \dots, \dot{\mathbf{r}}_{n_b}^T, \bar{\omega}_{n_b}^T]^T \in \mathbb{R}^{6n_b}$, with linear mapping $\dot{\mathbf{q}} = \mathbf{L}(\mathbf{q})\mathbf{v}$.

We denote by $\mathbf{f}^A(t, \mathbf{q}, \mathbf{v})$ the set of applied forces, and we assume that mass matrices of each particle remain constant.

We assume that the j -th particle is described by a convex shape $\Omega_j \in \mathbb{R}^3$, so that for each pair of particles j, k a nonpenetration condition $\Phi_{j,k}(\mathbf{q}_j, \Omega_j, \mathbf{q}_k, \Omega_k) \geq 0$ can be defined.

Although for simple shapes this condition corresponds to a scalar distance inequality between the nearest points of the two surfaces, there are cases (as in contact between two flat faces of polytopes) where multiple contact points could be considered, hence leading to multiple contacts and multiple inequalities even for a single particle pair (see Section 6). Furthermore, given positions \mathbf{q} and shapes $\Omega = \bigcup_{j=1..n_b} \Omega_j$ of all particles, one can find a set of n_c contacts (couples of contact points on two surfaces, aligned to the distance between the two surfaces), each corresponding to a nonpenetration constraint $\Phi_i(\mathbf{q}, \Omega) \geq 0$.

By means of a proper collision-detection method, we introduce contact constraint only for surfaces that are separated by a distance smaller than a threshold $\delta > 0$, so that among all n_c potential contacts, a smaller set \mathcal{A} of more relevant contact constraints can be selected:

$$\mathcal{A}(\mathbf{q}, \Omega, \delta) = \{i \mid i \in \{1, 2, \dots, n_c\}, \Phi_i(\mathbf{q}, \Omega) \leq \delta\},$$

with cardinality $n_{\mathcal{A}}$. From now on, with $i \in \mathcal{A}$ we will denote the i -th relevant contact.

We found that, in many scenarios involving dense granular flows, this approach leads to a number of contact constraints $n_{\mathcal{A}}$ that is only $O(n_b)$. For instance, with equal-sized spheres, we had on average $n_{\mathcal{A}} \propto C n_b$ with C in the $4 \div 6$ range.

Given the i -th contact, among two bodies A and B , let \mathbf{n}_i be the normal at the contact point, directed toward the exterior of the A body. Let \mathbf{u}_i and \mathbf{w}_i be two vectors in the contact plane such that $\mathbf{n}_i, \mathbf{u}_i, \mathbf{w}_i \in \mathbb{R}^3$ are mutually orthogonal vectors. Also, the frictional contact force is represented by multipliers $\hat{\gamma}_{i,n} \geq 0$, $\hat{\gamma}_{i,u}$, and $\hat{\gamma}_{i,w}$, which lead to the normal component of the force $\mathbf{F}_{i,N} = \hat{\gamma}_{i,n} \mathbf{n}_i$ and the tangential component of the force $\mathbf{F}_{i,T} = \hat{\gamma}_{i,u} \mathbf{u}_i + \hat{\gamma}_{i,w} \mathbf{w}_i$.

For each contact in \mathcal{A} that is active, that is, $\Phi_i(\cdot) = 0$ because bodies are touching, we introduce the contact forces; inactive contacts, $\Phi_i(\cdot) > 0$, do not enforce any reaction force. Hence the mathematical description of this unilateral model leads to a complementarity problem [28]:

$$\hat{\gamma}_{i,n} \geq 0 \quad \perp \quad \Phi_i(\cdot) \geq 0. \quad (1)$$

Moreover, the friction coefficient μ_i limits the ratio between the normal and the tangential force, and the tangential force must have a direction that is opposite to the tangential speed $\mathbf{v}_{i,T}$. This results in the following relationships:

$$\begin{aligned} \mu_i \hat{\gamma}_{i,n} &\geq \sqrt{\hat{\gamma}_{i,u}^2 + \hat{\gamma}_{i,w}^2}, \quad \|\mathbf{v}_{i,T}\| \left(\mu_i \hat{\gamma}_{i,n} - \sqrt{\hat{\gamma}_{i,u}^2 + \hat{\gamma}_{i,w}^2} \right) = 0, \\ \langle \mathbf{F}_{i,T}, \mathbf{v}_{i,T} \rangle &= - \|\mathbf{F}_{i,T}\| \|\mathbf{v}_{i,T}\|. \end{aligned} \quad (2)$$

We note that equations (1) and (2) are the first-order necessary Karush-Kuhn-Tucker conditions [29] for a minimization problem:

$$(\hat{\gamma}_{i,u}, \hat{\gamma}_{i,w}) = \underset{\sqrt{\hat{\gamma}_{i,u}^2 + \hat{\gamma}_{i,w}^2} \leq \mu_i \hat{\gamma}_{i,n}}{\operatorname{argmin}} \quad \mathbf{v}_{i,T}^T (\hat{\gamma}_{i,u} \mathbf{u}_i + \hat{\gamma}_{i,w} \mathbf{w}_i). \quad (3)$$

Let us introduce $\mathbf{A}_{i,p} = [\mathbf{n}_i, \mathbf{u}_i, \mathbf{w}_i]$ as the $\mathbb{R}^{3 \times 3}$ matrix representing the alignment of the i th contact, \mathbf{A}_A and \mathbf{A}_B as the rotation matrices $\mathbf{A}_i = \mathbf{A}(\epsilon_i)$ of the two bodies, and the vectors $\bar{\mathbf{s}}_{i,A}$ and $\bar{\mathbf{s}}_{i,B}$ as the contact point positions in body-relative coordinates. Given a triplet of multipliers $\hat{\gamma}_i = \{\hat{\gamma}_{i,n}, \hat{\gamma}_{i,u}, \hat{\gamma}_{i,w}\}^T$, the contact exerts two forces on the center of mass of the two bodies A and B ,

$$\mathbf{F}_{i,A} = -\mathbf{A}_{i,p}^T \hat{\gamma}_i, \quad \mathbf{F}_{i,B} = \mathbf{A}_{i,p}^T \hat{\gamma}_i, \quad (4)$$

and two torques, expressed in body-relative coordinates:

$$\mathbf{T}_{i,A} = \mathbf{A}_{i,p}^T \mathbf{A}_A \tilde{\mathbf{s}}_{i,A} \hat{\gamma}_i, \quad \mathbf{T}_{i,B} = -\mathbf{A}_{i,p}^T \mathbf{A}_B \tilde{\mathbf{s}}_{i,B} \hat{\gamma}_i, \quad (5)$$

where we used the skew-symmetric matrices $\tilde{\mathbf{s}}$ defined as

$$\tilde{\mathbf{s}} = \begin{bmatrix} 0 & -\bar{s}_z & \bar{s}_y \\ \bar{s}_z & 0 & -\bar{s}_x \\ -\bar{s}_y & \bar{s}_x & 0 \end{bmatrix}.$$

This allows us to introduce the tangent space generators $\mathbf{D}_i = [\mathbf{D}_{i,n}, \mathbf{D}_{i,u}, \mathbf{D}_{i,w}] \in \mathbb{R}^{6n_b \times 3}$, defined as

$$\mathbf{D}_i^T = \begin{bmatrix} \mathbf{0} & \dots & -\mathbf{A}_{i,p}^T & \mathbf{A}_{i,p}^T \mathbf{A}_A \tilde{\mathbf{s}}_{i,A} & \mathbf{0} & \dots & \mathbf{0} & \mathbf{A}_{i,p}^T & -\mathbf{A}_{i,p}^T \mathbf{A}_B \tilde{\mathbf{s}}_{i,B} & \dots & \mathbf{0} \end{bmatrix}, \quad (6)$$

and to write the complete expression of the dynamical model. The time evolution of the system is governed by the following differential problem with equilibrium constraints, which is equivalent to a differential variational inequality (see [30]):

$$\begin{aligned} \dot{\mathbf{q}} &= \mathbf{L}(\mathbf{q}) \mathbf{v} \\ \mathbf{M} \dot{\mathbf{v}} &= \mathbf{f}(t, \mathbf{q}, \mathbf{v}) + \sum_{i \in \mathcal{A}(\mathbf{q}, \Omega, \delta)} (\hat{\gamma}_{i,n} \mathbf{D}_{i,n} + \hat{\gamma}_{i,u} \mathbf{D}_{i,u} + \hat{\gamma}_{i,w} \mathbf{D}_{i,w}) \\ i \in \mathcal{A}(\mathbf{q}, \Omega, \delta) &: \hat{\gamma}_{i,n} \geq 0 \perp \Phi_i(\mathbf{q}) \geq 0, \\ (\hat{\gamma}_{i,u}, \hat{\gamma}_{i,w}) &= \underset{\mu_i \hat{\gamma}_{i,n} \geq \sqrt{(\hat{\gamma}_{i,u})^2 + (\hat{\gamma}_{i,w})^2}}{\operatorname{argmin}} \mathbf{v}^T (\hat{\gamma}_{i,u} \mathbf{D}_{i,u} + \hat{\gamma}_{i,w} \mathbf{D}_{i,w}). \end{aligned} \quad (7)$$

This DVI also includes the vector of forces $\mathbf{f}(t, \mathbf{q}, \mathbf{v})$, in Lagrangian coordinates, caused by gyroscopic effects and other applied forces such as gravity and aerodynamical drag.

Note that this model does not take into account elastic rebound effects. However, for the type of application targeted — dense granular flow that appears, for example, in pebble-bed nuclear reactor simulation — all collisions occurring during the simulation can be assumed to be of the inelastic type.

There exist configurations for which this type of DVI might not have a solution [31, 32]. Nonetheless, as shown in [33], it is still possible to obtain weaker solutions in a vector-measure differential-inclusion sense using time stepping schemes that do not solve in terms of accelerations and forces, but rather use speed increments and impulses as unknowns.

4 Solving the DVI

Given a position $\mathbf{q}^{(l)}$ and velocity $\mathbf{v}^{(l)}$ at the time-step $t^{(l)}$, the numerical solution is found at the new time-step $t^{(l+1)} = t^{(l)} + h$ by solving the following optimization problem with equilibrium constraints [34]:

$$\mathbf{M}(\mathbf{v}^{(l+1)} - \mathbf{v}^{(l)}) = h \mathbf{f}(t^{(l)}, \mathbf{q}^{(l)}, \mathbf{v}^{(l)}) + \sum_{i \in \mathcal{A}(q^{(l)}, \Omega, \delta)} (\gamma_{i,n} \mathbf{D}_{i,n} + \gamma_{i,u} \mathbf{D}_{i,u} + \gamma_{i,w} \mathbf{D}_{i,w}) \quad (8)$$

$$i \in \mathcal{A}(q^{(l)}, \Omega, \delta) : 0 \leq \frac{1}{h} \Phi_i(\mathbf{q}^{(l)}) + \mathbf{D}_{i,n}^T \mathbf{v}^{(l+1)} \perp \gamma_n^i \geq 0, \text{ and} \quad (9)$$

$$(\gamma_{i,u}, \gamma_{i,w}) = \underset{\mu_i \gamma_{i,n} \geq \sqrt{\gamma_{i,u}^2 + \gamma_{i,w}^2}}{\operatorname{argmin}} \mathbf{v}^T (\gamma_{i,u} \mathbf{D}_{i,u} + \gamma_{i,w} \mathbf{D}_{i,w}), \quad (10)$$

$$\mathbf{q}^{(l+1)} = \mathbf{q}^{(l)} + h \mathbf{L}(\mathbf{q}^{(l)}) \mathbf{v}^{(l+1)}. \quad (11)$$

Here, γ_s represents the constraint impulse of a contact constraint, that is, $\gamma_s = h\widehat{\gamma}_s$, for $s = n, u, w$. The $\frac{1}{h}\Phi_i(\mathbf{q}^{(l)})$ term achieves constraint stabilization; its effect is discussed in [20]. The scheme converges to the solution of a measure differential inclusion [18] when the step size $h \rightarrow 0$.

We observe that such a stepping method, in the case of no contacts, behaves like a semi-implicit Euler integrator, a first-order integration scheme that features the property of being symplectic [35].

Unknown impulses and speeds in (8)-(10) can be found by solving a nonlinear complementarity problem (NCP), that represents the most complex part of the entire approach. [36, 37].

Among the different approaches that can be used to solve (8)-(10), [12, 37], we use a fixed-point iterative scheme that is completely matrixless and fits well in scenarios with thousands of contacts and particles. The theoretical properties of the method (convergence, existence of the solution) are detailed in [16]; in this paper we focus on its efficient implementation.

Using the approach proposed in [16], one solves the problem as a monotone optimization by introducing a relaxation over the complementarity constraints: Eq. (9) is replaced with

$$i \in \mathcal{A}(q^{(l)}, \Omega, \delta) : 0 \leq \frac{1}{h}\Phi_i(\mathbf{q}^{(l)}) + \mathbf{D}_{i,n}^T \mathbf{v}^{(l+1)} - \mu_i \sqrt{(\mathbf{v}^T \mathbf{D}_{i,u})^2 + (\mathbf{v}^T \mathbf{D}_{i,w})^2} \perp \gamma_n^i \geq 0. \quad (12)$$

In [18] it was demonstrated that, for $h \rightarrow 0$, the solution of the modified time-stepping scheme will approach the solution of the same measure differential inclusion as the scheme, presented in [36, 37], that uses unrelaxed constraints. For large time steps, the above relaxation causes the objects in sliding contact to float apart by a quantity proportional to $h\mu\|\mathbf{v}_T\|$; this artifact introduces an unnatural “dilation layer” around sliding objects, but it can be tolerated in case of small friction and/or small tangential speed (that is often the case in dense granular flows) and tends to zero thickness anyway, for $h \rightarrow 0$. In addition, the equation (12) is the key to being able to integrate with fixed time step and not have to stop at collisions, at least for totally plastic flows [20]. This is best seen if we assume that $\mu_{(i)} = 0$ and Φ is linear. In this case, if two objects in proximity enter contact during a time step, (12) ensures that exact contact occurs precisely at the end of the time step). Therefore all collisions during a time step are now simultaneous, and there is no need to stop and restart at each one of them. The convergence result in [18] ensures that this strategy is consistent – converges to the correct limit behavior – in the limit of $h \rightarrow 0$.

With the above assumptions, the problem is a special class of convex complementarity called second-order cone complementarity problem.

The algorithm we proposed in [16] is an evolution of the method introduced by Mangasarian for LCPs [17], that is, a fixed-point map (a smoother like Gauss-Seidel or Jacobi) interleaved by projections onto the feasible sets.

We introduce the following quantities:

$$\tilde{\mathbf{k}} \equiv \mathbf{M}\mathbf{v}^{(l)} + h\mathbf{f}(t^{(l)}, \mathbf{q}^{(l)}, \mathbf{v}^{(l)}) \quad (13)$$

$$\mathbf{b}_i \equiv \left\{ \frac{1}{h}\Phi_i(\mathbf{q}^{(l)}), 0, 0 \right\}^T \quad i \in \mathcal{A}(q^{(l)}, \Omega, \delta). \quad (14)$$

The method iterates over index r until convergence is reached:

$$\forall i \in \mathcal{A}(q^{(l)}, \Omega, \delta) :$$

$$\delta_i^{r+1} = \gamma_i^r - \omega \eta_i \left[D_i^T \mathbf{M}^{-1} \left(\sum_{z \in \mathcal{A}(q^{(l)}, \Omega, \delta)} D_z \gamma_z^r + \tilde{\mathbf{k}} \right) + \mathbf{b}_i \right] \quad (15)$$

$$\gamma_i^{r+1} = \lambda \Pi_{\Upsilon_i}(\delta_i^{r+1}) + (1 - \lambda) \gamma_i^r. \quad (16)$$

The results of that iterative process are the contact multipliers γ_i (the dual variables); in a following step one can also easily compute the unknown speeds $\mathbf{v}^{(l+1)}$ (the primal variables) using Eq. (8), as

$$\mathbf{v}^{(l+1)} = \mathbf{M}^{-1} \sum_{z \in \mathcal{A}(q^{(l)}, \Omega, \delta)} D_z \gamma_z^{r+1} + h \mathbf{M}^{-1} \mathbf{f}(t^{(l)}, \mathbf{q}^{(l)}, \mathbf{v}^{(l)}) + \mathbf{v}^{(l)}. \quad (17)$$

Note that, right after the computation of δ_i^{r+1} , the iterative process uses the projection operator $\Pi_{\Upsilon_i}(\cdot)$, a nonexpansive map $\Pi_{\Upsilon_i} : \mathbb{R}^3 \rightarrow \mathbb{R}^3$. If the multipliers fall into the friction cone Υ_i , they are not modified; if they are in the polar cone Υ_i^o , they are set to zero; in the remaining cases they are projected orthogonally onto the surface of the friction cone. For more details on $\Pi_{\Upsilon_i}(\cdot)$ and convergence, see [16]. The overrelaxation factor ω and the parameters λ and η_i are adjusted to control the convergence [16].

5 Implementation of the Complementarity Solver

The iteration (15)-(16) can be implemented in different flavors; in this section we discuss the most relevant.

5.1 Projected Jacobi

We aim at avoiding the summation in the right-hand side of Eq. (15). Using Eqs. (8) and (13), we can rewrite the iteration in a more compact form:

$$\forall i \in \mathcal{A}(q^{(l)}, \Omega, \delta) : \quad \gamma_i^{r+1} = \lambda \Pi_{\Upsilon_i} \left[\gamma_i^r - \omega \eta_i (D_i^T \mathbf{v}^{(l+1),r} + \mathbf{b}_i) \right] + (1 - \lambda) \gamma_i^r \quad (18)$$

In this case, for each iteration, before repeating (18), also primal variables (the velocities $\mathbf{v}^{(l+1)}$) are updated as

$$\mathbf{v}^{(l+1),r+1} = \mathbf{M}^{-1} \sum_{z \in \mathcal{A}(q^{(l)}, \Omega, \delta)} D_z \gamma_z^{r+1} + h \mathbf{M}^{-1} \mathbf{f}(t^{(l)}, \mathbf{q}^{(l)}, \mathbf{v}^{(l)}) + \mathbf{v}^{(l)}. \quad (19)$$

Note that Eq. (19) can be implemented in an incremental form for even better performance, because at each r -th iteration the correction $h \mathbf{M}^{-1} \mathbf{f}(t^{(l)}, \mathbf{q}^{(l)}, \mathbf{v}^{(l)}) + \mathbf{v}^{(l)}$ does not change: it is enough to start the process with $\mathbf{v}^{(l+1),0} = h \mathbf{M}^{-1} \mathbf{f}(t^{(l)}, \mathbf{q}^{(l)}, \mathbf{v}^{(l)}) + \mathbf{v}^{(l)}$ and add $\mathbf{M}^{-1} \sum_{z \in \mathcal{A}} D_z \Delta \gamma_z^{r+1}$ at each iteration.

Note that, except for the projection, this method has the same properties of a Jacobi method with overrelaxation. As for the ω and η parameters we found that good default values are $\eta_i = 3/\text{Trace}(D_i^T \mathbf{M}^{-1} D_i)$ and $\omega = 0.2$.

5.2 Projected Gauss-Seidel

If we perform Eq. (19) right after Eq. (18) instead of waiting for all the impulses to be computed in an iteration, we obtain a scheme that is similar to a successive overrelaxation method or a Gauss-Seidel, because it exploits the immediate update of multiplier in the computation of the following multipliers.

This method can be summarized as follows. First the primal variables are initialized as

$$\mathbf{v}^{(l+1),0} = h\mathbf{M}^{-1}\mathbf{f}(t^{(l)}, \mathbf{q}^{(l)}, \mathbf{v}^{(l)}) + \mathbf{v}^{(l)}. \quad (20)$$

Then one must iterate the following formulas over r until convergence:

$$\forall i \in \mathcal{A}(q^{(l)}, \Omega, \delta) \left\{ \begin{array}{l} \gamma_i^{r+1} = \lambda \Pi_{\Gamma_i} [\gamma_i^r - \omega \eta_i (D_i^T \mathbf{v}^{(l+1),r} + \mathbf{b}_i)] + (1 - \lambda) \gamma_i^r \\ \Delta \gamma_z^{r+1} = \gamma_z^{r+1} - \gamma_z^r \\ \mathbf{v}^{(l+1),r+1} + = \mathbf{M}^{-1} \sum_{z \in \mathcal{A}(q^{(l)}, \Omega, \delta)} D_z \Delta \gamma_z^{r+1}. \end{array} \right. \quad (21)$$

Convergence can be monitored looking at the residual terms $D_i^T \mathbf{v}^{(l+1),r} + \mathbf{b}_i$: the iteration is stopped when they fall below a tolerance threshold.

The convergence properties are better than the previous projected Jacobi, especially in case of redundant contacts. Moreover, larger values for ω can be used without risking divergence. For instance we successfully tested values up to $\omega = 1$. This method behaves better than the projected Jacobi especially in case of redundant constraints, as occur in the case of many contact points between two flat surfaces.

5.3 Multithreaded Parallel Solver

A multithreaded version of the solver can leverage modern multicore processors. In our implementation, given n_μ cores, the solver partitions the set of constraints $\mathcal{A}(q^{(l)}, \Omega, \delta)$ into n_{th} subsets $\mathcal{A}_s(q^{(l)}, \Omega, \delta)$, with $n_{th} \geq n_\mu$; $\mathcal{A}(q^{(l)}, \Omega, \delta) = \bigcup_{s=1..n_{th}} \mathcal{A}_s(q^{(l)}, \Omega, \delta)$.

As in block-Gauss-Seidel solvers, after initialization of primal variables with (20), each $\mathcal{A}_s(q^{(l)}, \Omega, \delta)$ subset is assigned to a thread, each running the iteration (21) where $\mathcal{A}(q^{(l)}, \Omega, \delta)$ has been replaced by $\mathcal{A}_s(q^{(l)}, \Omega, \delta)$. Then the threads are synchronized for completion, and the iteration (21) is started again. This process continues until convergence.

We also developed a hyperthreaded version of the projected Jacobi scheme that assigns one contact per thread on massively parallel architectures of CUDA type [38]. With such an implementation we obtained a speed improvement of more than one order of magnitude on a NVIDIA 9800GX2 GPU board.

6 Collision Detection

Since granular flows usually involve thousands, if not millions, of particles, it is not practical to search for contact points for all the particle pairs because we would obtain a method with superlinear complexity $O(n_b^2)$ both in algorithmic time and in memory space.

Hence a preliminary stage, called *broad-phase collision detection*, is used to detect pairs of particles whose bounding boxes are near enough. We explored two approaches to this problem: the highly-optimized SAP (sweep and prune) algorithm available in the Bullet

open-source library [39], and a GPU hyperthreaded implementation that partitions the space into a 3D grid of particle bins and creates collision pairs for the particles in a single bin. For a large amount of particles with the same size, densely packed in a limited volume as in our tests, the second approach showed superior performance.

After the broad-phase stage there is the *narrow-phase* process, where we create the contacts—if any—for each pair of particles that resulted near enough. For convex geometries, such as spheres, ellipsoids, and polytopes, we use the Gilbert-Johnson-Keerthi (GJK) algorithm [40]; concave shapes, on the other hand, are rather decomposed into clusters of concave polytopes. The GJK algorithm easily obtains the two points with minimum distance between two convex shapes, so it can be used to create contact points when shapes are not already in contact. However it cannot deal with interpenetrations; we solve this problem by performing GJK on shapes that are smaller than the original ones, removing a layer with δ_{in} thickness. After the nearest contact points have been obtained, they are offset by amounts δ_{in} toward the exterior. This efficient workaround allows interpenetrations up to $2\delta_{in}$ between the shapes, although it has the side effect that sharp corners of polyhedrons are smoothed, as they have fillets of radius δ_{in} (the collision shape is the Minkowski sum of a smaller shape and a ball with $r = \delta_{in}$). This side effect is not an issue when working with spherical particles, since the Minkowski sum of two spheres is still a sphere.

Penetrations larger than $2\delta_{in}$ —which should be rare—can be handled by falling back to an EPA (expanding polytope) algorithm. Note that, in the case of a sphere, choosing δ_{in} as the radius of the sphere means that EPA is not even needed.

The case of the collisions between the particles and the interior of the PBR reactor vessel is handled by a custom algorithm. This is needed because the inner wall of the vessel is made with a hollow cylinder and a funnel, both concave shapes that we do not approximate with convex decomposition because the resulting cluster of polyhedra would introduce small artifacts (the curved surface would be approximated by faceted shapes).

Contacts whose distance is larger than δ do not enter the set $\mathcal{A}(q^{(l)}, \Omega, \delta)$. Choosing the proper δ and δ_{in} parameters is not trivial, because too large a δ means that too many unneeded potential contacts will enter into the complementary solver, thus slowing its performance, while too low a δ with fast-moving particles could result in some contacts being created when it is too late and there’s already a small interpenetration. Similarly, a large δ_{in} should be avoided in the case of sharp shapes because of the abovementioned “smoothing” effect, while too small δ_{in} means that even the least interpenetration may break the GJK algorithm and require the less robust EPA fallback. In the tests discussed here, using only spheres of radius r_b , we found a good trade off with $\delta = 0.2r_b$ and $\delta_{in} = r_b$.

We remark that the above mentioned software implementation allows simulations of granular materials with particles of arbitrary shapes, such as fragments of rocks. However problems involving spherical bodies pose less difficulties than in the case of complex shapes: particles with irregular surfaces may lead to multiple contact points and concave shapes can lead to locking phenomena or ill-posed contact constraints that negatively affect the performance of the complementarity solver.

7 Experimental Setup

We built two test beds for validating our simulation method; see Fig.(2). Both are filled with plastic spheres with diameter $6 \text{ mm} \pm 0.05$; the density of the plastic is 1060 kg/m^3 .

The first test bed is a flat hopper, with 30 mm thickness and 440 mm width. The bottom funnel can be adjusted either in terms of slope and width of the outlet. The front wall is made with transparent plexiglas. It contains a maximum of 50,000 spheres.

The second test bed is a half small-scale reproduction of the PBR reactor. The diameter of the cylinder is 348 mm; the slope of the conical funnel can be selected between 30° and 60° . The outlet diameter is 60 mm. All the walls are transparent. The test bed contains a maximum of 200'000 spheres (we note that a similar small-scale PBR has been presented in [22], obtaining similar results).

In both cases the spheres can be followed from the outside by using a high resolution camera, shooting at fixed time intervals, see Fig.(3). Digital postprocessing of pictures using the MatlabTM Image Processing Toolbox allows us to obtain mean trajectories and profiles of mean speeds at different heights. All spheres are initially yellow, but to ease the image processing, we dip half amount of the spheres into a special ink, so that the spheres can be layered into the test beds as bands of different colors. After each test all the spheres are washed with alcohol and are ready for another test.

In controlled air conditions, sphere-wall friction is $\mu = 0.43$ and friction between the spheres is $\mu = 0.46$. Air conditioning is important because we experienced that temperature, humidity and static electricity can affect these values.

8 Results

Both the flat hopper with 50,000 spheres and the half reactor with 165,000 spheres ¹ have been simulated by using the convex cone complementarity approach, implemented in our C++ software Chrono::Engine. Two types of reactor models have been simulated: the small-scale reactor, and a full-scale reactor. In the small-scale case we take into account the friction between the spheres and the flat wall because it is meant as a comparison with experimental data, where friction against the plexiglas window cannot be avoided. Instead, the full-scale is used for comparison with a half-reactor DEM model: in this case there is no friction on the vertical plane that splits the vessel, because it aims at reproducing the entire reactor without artifacts caused by the vertical wall.

After each simulation we obtained many gigabytes of data, with the position of the spheres and the application point, alignment, and strength of the contact forces saved at each time step. We rendered the results of the simulations using a custom script for the POV raytracing software, as shown in Fig.(4,5), to make pictures and animations.

In Fig.(6) one can appreciate the matching between a filmed experiment and the corresponding simulated flow. We noticed that the initial filling process can lead to different initial

¹The simulation is limited by RAM constraints on 32 bit systems. In fact, we simulated also a reactor with 400'000 spheres, using a GPU-based parallel version of the solver on a 64 bit system, but results are not discussed here.

conditions. For instance Fig.(6) shows that the simulated material starts with less crystallization (i.e. arrangement in layered regular patterns), because initial packing was created from almost a random position followed by an iterative adjustment, while the experimental hopper was slowly filled by pouring spheres from above.²

We also developed various MatlabTM postprocessing programs to load the many output files and produce statistical results. As demonstrated in [7], the granular flow of pebbles results in the same statistics if the pebbles are continuously extracted, as opposed to a slow schedule, for instance, two spheres per minute in a real PBR reactor.

One of the most interesting results is the mean speed of the granular flow. Given a couple of position files, separated by a time interval, the speed of each sphere can be computed; the space is then divided in rectangular bins, and all the spheres that fall within a certain bin will contribute to the average speed of that bin. For the flat hopper we use a 2D grid of 18x20 bins (x,y), while for the PBR reactor we use a 3D Cartesian grid of 28x15x15 bins (x,y,z) or a polar grid of 14x15 (r,y) circumferential bins to plot data for a vertical slice of the reactor, averaged along the circumference.

Although many spheres contribute to the average of a single bin, we improved the smoothing of the plotted data by averaging the data of those grids over few seconds of simulated time, thus repeating the process for multiple couples of position states. This produced the results of Fig.(7, 8 and 9). For the PBR flow, these plots match the results reported in [22, 7]; in particular one can note the typical transition to a plug-like motion in the upper part of the vessel [41], with almost no radial dispersion except near the outlet (Figure 10). Also contacts can be plotted, either averaged in bins (and matching other results from [42]) or drawn with contact nets as in Fig.(11). One should note that, in addition to this qualitative agreement, Fig.(7) shows that our simulation approach matches *the experimental values* of the velocity statistics within 5% – no small accomplishment if one considers that friction coefficients can easily vary by up to 10% due to humidity and other environmental factors.

Other interesting results are the volume fraction v_m (material volume per unit of volume) and the porosity, $\phi_m = 1 - v_m$. The value of porosity has been compared to results coming from a DEM simulation of the PBR reactor [7], obtaining good agreement: no more than 5 % discrepancy in all areas except at the wall, where it is 10 %; see Fig.(12).

For the PBR reactor we used 240 iterations per time step, whereas 150 iterations per time step were enough in the hopper simulation to guarantee good precision, with a time step of 0.0025 s. The plot of Fig.(13) shows acceptable constraint violation as an intuitive meaning of the precision of the solver. We note that particles that are at the bottom of the silo may experience larger interpenetration errors; hence the convergence of the solver is more critical with high stacks. This suggests that future improvements could exploit multigrid approaches to alleviate this problem and require fewer iterations. We also note that larger time steps, up to 0.01 s, fewer iterations, and consequent lower precision may still result in stable simulations, where the overall motion of the granular material is still meaningful, although low precision tends to “smooth” speed profiles, to miss part of the crystallization effects, and to distort volume fraction.

Here, our main objective is the validation of the predictive power of our DVI-based

²From repeated tests we noticed that the height of the stacks can differ up to 10% if the simulated and experimental filling process are quite different.

approach. Given our ultimate objective of developing DVI approaches that are both accurate and much faster than DEM, however, we briefly discuss the current computational effort situation. For the PBR simulations with 165,000 spheres, each 2.5 millisecond time step required about 2 minutes of CPU time on a 2GHz Pentium computer, whereas the hopper simulation required 30 seconds per step; we remark that these time steps are larger than those required by DEM methods by a factor of more than 300. Assuming that DEM effort behaves linearly with the problem, the 780 CPU hours needed for 7.8 s of simulation of 440'000 spheres as in [7], suggests that DEM would need 5.6 minutes on a 3GHz Intel Xeon for a 2.5 millisecond simulation, several times larger than our method.

We note, moreover, that the computational performance potential of our method is much larger. Indeed, our calculation remains stable even if the time step is increased by a factor of 10 and the number of iterations is decreased by a factor of 3, at the cost of some degradation of the accuracy. In addition, we point out that the DEM parameters in [7] were larger than the one recommended by Hertzian contact theory by factors that resulted in time steps that were more than 300 times larger. For the PBR experiments, DEM still accurately predicted the macroscopic behavior; but that result is far from guaranteed if the configuration is changed (say, by increasing the load on the pebbles).

We conclude that, for the configurations studied, the CCP method is at least as fast or faster than DEM, and that it produces simulations whose statistics match well physical experiments. To be sure, CCP is more complex to implement than DEM. In addition, certain physical features, such as restitution, while implementable in our approach [36], need more work for validation-quality simulation compared to DEM.

Finally, more research is under way to accelerate the resolution of CCP problems, the linchpin of CCP being able to outperform DEM. Given the many options in fine tuning of our algorithm, some experimentation is needed before we determine the optimal accuracy–performance trade-off. But the main reasons why these options are available to us is the stability of the method for a broad range of the parameter choices, which is due to the absence of the artificial stiffness, an option not available to DEM. In future work we plan to carry out comprehensive performance profiling tests for both our method and DEM.

9 Conclusions

We adopted a numerical method based on convex complementarity problems as an efficient alternative to DEM methods for the simulation of granular dynamics. Two test beds have been built for validation; one of them represents a small scale version of a PBR nuclear reactor, that is the application that motivates our research efforts.

The agreement between the numerical and the experimental results demonstrate that the CCP approach is a reliable method for simulating large granular flows, suggesting its adoption also in fields such as rock-soil dynamics, pharmaceutical engineering, and powder mechanics as a faster alternative to the DEM approach even in case of hundreds of thousands of particles.

Future efforts will address the improvement of the convergence properties of the method by endorsing multigrid approaches and domain decomposition on parallel architectures.

Acknowledgments

Financial support for A.Tasora is provided in part by the Italian Ministry of University under PRIN grant 2007Z7K4ZB. Mihai Anitescu was supported by the Office of Advanced Scientific Computing Research, Office of Science, U.S. Department of Energy, under Contract DE-AC02-06CH11357.

References

- [1] Radjai, F., Jean, M., Moreau, J.-J., and Roux, S., 1996, “Force distributions in dense two-dimensional granular systems,” *Physical Review Letters*, **77(2)**, pp. 274–277.
- [2] J. Choi, A. K. and Bazant, M. Z., 2005, “Velocity profile of granular flows in silos and hoppers,” *J. Phys.: Condens. Matter*, **17**, pp. S2533–S2548.
- [3] Kamrin, K. and Bazant, M. Z., 2007, “Stochastic flow rule for granular materials,” *Physical Review E*, **75**.
- [4] K. Kamrin, C. H. R. and Bazant, M. Z., 2007, “The stochastic flow rule: A multi-scale model for granular plasticity,” *Modelling Simul. Mater. Sci. Eng.*, **15**, pp. S449–S464.
- [5] Cundall, P. and Strack, O., 1979, “A discrete numerical model for granular assemblies,” *Geotechnique*, **29(1)**, pp. 47–65.
- [6] Hirshfeld, D. and Rapaporta, D., 2001, “Granular flow from a silo: Discrete-particle simulations in three dimensions,” *Eur. Phys. J. E*, **4**, pp. 193–199.
- [7] Rycroft, C. H., Grest, G. S., Landry, J. W., and Bazant, M. Z., 2006, “Analysis of granular flow in a pebble-bed nuclear reactor,” *Physical Review E*, **74**, **021306**.
- [8] Madsen, J., Pechdimaljian, N., and Negrut, D., 2007, “Penalty versus complementarity-based frictional contact of rigid bodies: A CPU time comparison,” Tech. Rep. TR-2007-05, Simulation-Based Engineering Lab, University of Wisconsin, Madison.
- [9] Pang, J.-S. and Stewart, D., 2008, “Differential variational inequalities,” *Mathematical Programming*, **113(2)**, pp. 345–424.
- [10] Moreau, J. J., 1988, “Unilateral contact and dry friction in finite freedom dynamics,” *Nonsmooth Mechanics and Applications*, J. J. Moreau and P. D. Panagiotopoulos, eds., Springer-Verlag, Berlin, pp. 1–82.
- [11] Moreau, J. J. and Jean, M., 1996, “Numerical treatment of contact and friction: The contact dynamics method,” *Proceedings of the Third Biennial Joint Conference on Engineering Systems and Analysis*, Montpellier, France, pp. 201–208.
- [12] Jourdan, F., Alart, P., and Jean, M., 1998, “A Gauss Seidel like algorithm to solve frictional contract problems,” *Computer Methods in Applied Mechanics and Engineering*, **155**, pp. 31–47.

- [13] Cottle, R. W., Pang, J.-S., and Stone, R. E., 1992, *The Linear Complementarity Problem*, Academic Press, Boston.
- [14] Murty, K. G., 1988, *Linear Complementarity, Linear and Nonlinear Programming*, Helderman Verlag, Berlin.
- [15] Tasora, A. and Anitescu, M., 2007, “A fast contraction mapping for solving multibody systems,” *PAMM*, **7**, pp. 1062401–1062402.
- [16] Anitescu, M. and Tasora, A., 2008, in press, “An iterative approach for cone complementarity problems for nonsmooth dynamics,” *Computational Optimization and Applications*.
- [17] Mangasarian, O., 1977, “Solution of symmetric linear complementarity problems by iterative methods,” *Journal of Optimization Theory and Applications*, **22**(4), pp. 465–485.
- [18] Anitescu, M., 2006, “Optimization-based simulation of nonsmooth rigid multibody dynamics,” *Mathematical Programming*, **105**(1), pp. 113–143.
- [19] Stewart, D., 1997, “Existence of solutions to rigid body dynamics and the Painlevé paradoxes,” *C. R. Acad. Sci. Paris*, **325**, pp. 689–693.
- [20] Anitescu, M. and Hart, G. D., 2004, “A constraint-stabilized time-stepping approach for rigid multibody dynamics with joints, contact and friction,” *International Journal for Numerical Methods in Engineering*, **60**(14), pp. 2335–2371.
- [21] Savage, S. B., 1979, “Gravity flow of cohesionless granular materials in chutes and channels,” *Journal of Fluid Mechanics Digital Archive*, **92**(01), pp. 53–96.
- [22] Kadak, A. and Bazant, Z., 2004, “Pebble flow experiments for pebble bed reactors,” *Proceedings of 2nd International Topical Meeting on High Temperature Reactor Technology, Beijing, China*.
- [23] Schultz, R., Ougouag, A., Nigg, D., Gougar, H., Johnson, R., Terry, W., Oh, C., McEligot, D., Johnsen, G., McCreery, G., et al., 2007, “Next generation nuclear plant methods technical program plan,” Tech. Rep. INL/EXT-06-11804, Idaho National Laboratory.
- [24] Rothwell, G. and Rust, J., 1997, “On the optimal lifetime of nuclear power plants,” *Journal of Business and Economic Statistics*, **15**, pp. 195–208.
- [25] Gougar, H. D., 2004, *Advanced core design and fuel management for pebble-bed reactors*, Ph.D. thesis, Department of Nuclear Engineering, Penn State University.
- [26] Ougouag, A., Ortensi, J., and Hiruta, H., 2009, “Analysis of an earthquake-initiated-transient in a PBR,” Tech. Rep. INL/CON-08-14876,, Idaho National Laboratory (INL).
- [27] Shabana, A. A., 2005, *Dynamics of Multibody Systems*, Cambridge University Press, third ed.

- [28] Stewart, D. E. and Trinkle, J. C., 1996, “An implicit time-stepping scheme for rigid-body dynamics with inelastic collisions and Coulomb friction,” *International Journal for Numerical Methods in Engineering*, **39**, pp. 2673–2691.
- [29] Bertsekas, D. P., 1995, *Nonlinear Programming*, Athena Scientific, Belmont, MA.
- [30] Pang, J. S. and Stewart, D. E., 2008, “Differential variational inequalities,” *Mathematical Programming*, **113**, pp. 1–80.
- [31] Baraff, D., 1993, “Issues in computing contact forces for non-penetrating rigid bodies,” *Algorithmica*, **10**, pp. 292–352.
- [32] Stewart, D. E., 1998, “Convergence of a time-stepping scheme for rigid body dynamics and resolution of Painleve’s problems,” *Archive Rational Mechanics and Analysis*, **145(3)**, pp. 215–260.
- [33] Stewart, D. E., 2000, “Rigid-body dynamics with friction and impact,” *SIAM Review*, **42(1)**, pp. 3–39.
- [34] Tasora, A., 2007, “High performance complementarity solver for non-smooth dynamics,” *Proceedings of the ECCOMAS Multibody Dynamics Conference*, C. L. Bottasso, P. Masarati, and L. Trainelli, eds., Milan, Italy.
- [35] Hairer, E., Lubich, C., and G.Wanner, 2004, *Geometric numerical integration*, Springer, Berlin.
- [36] Anitescu, M. and Potra, F. A., 1997, “Formulating dynamic multi-rigid-body contact problems with friction as solvable linear complementarity problems,” *Nonlinear Dynamics*, **14**, pp. 231–247.
- [37] Anitescu, M. and Hart, G. D., 2004, “A fixed-point iteration approach for multibody dynamics with contact and friction,” *Mathematical Programming, Series B*, **101(1)**, pp. 3–32.
- [38] A. Tasora, D. N. and Anitescu, M., 2008, “Large-scale parallel multi-body dynamics with frictional contact on the graphical processing unit,” *Journal of Multi-body Dynamics*, **222(4)**, pp. 315–326.
- [39] Forum, 2008, “Bullet Physics Library,” Available online at <http://www.bulletphysics.com/Bullet/wordpress/bullet>.
- [40] E.G. Gilbert, S. K., D.W. Johnson, 1988, “A fast procedure for computing the distance between complex objects in three-dimensional space,” *Robotics and Automation*, **4(2)**, pp. 193–203.
- [41] Schulze, D., 2007, *Powders and Bulk Solids*, Springer, Berlin.
- [42] J. Sykut, M. M. and Horabik, J., 2008, “DEM simulation of the packing structure and wall load in a 2-dimensional silo,” *Granular Matter*, **10**, pp. 273–278.

Figures

Note: all pictures are available in EPS format. In this PDF we resized all the pictures as they should appear in an ASME paper: their width is half page width (one column of text in a double-column layout), except the three pictures of Figure 6, that could run through the top of a page, sharing two columns. Also, some photos as Figure 3 could be further reduced in size, if you need space.

Although figures are presented in color in this PDF for better readability, they are meant to be printed in B/W without major loss of quality.




Figure 1: Schematic representation of the pebble bed reactor

The submitted manuscript has been created by the University of Chicago as Operator of Argonne National Laboratory ("Argonne") under Contract No. DE-AC02-06CH11357 with the U.S. Department of Energy. The U.S. Government retains for itself, and others acting on its behalf, a paid-up, nonexclusive, irrevocable worldwide license in said article to reproduce, prepare derivative works, distribute copies to the public, and perform publicly and display publicly, by or on behalf of the Government.




Figure 2: The experimental testbeds: the half reactor and the flat hopper




Figure 3: Snapshot from one of the experiments of flow through the PBR half reactor model, with the 60° funnel case




Figure 4: Simulation of the 2D hopper model




Figure 5: Simulation of the granular flow in the PBR nuclear reactor (half-reactor model, 165,000 pebbles)



Figure 6: Comparison between experimental flow (half left of figures) and simulated flow (half right of figures) in the hopper, at time $t=0$ s, $t=0.6$ s, $t=1.2$ s

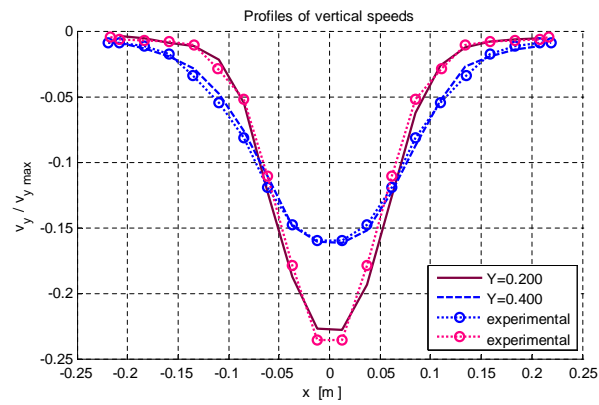


Figure 7: Comparison between numerical and experimental speed profiles in the flat hopper, at $t=1$ s

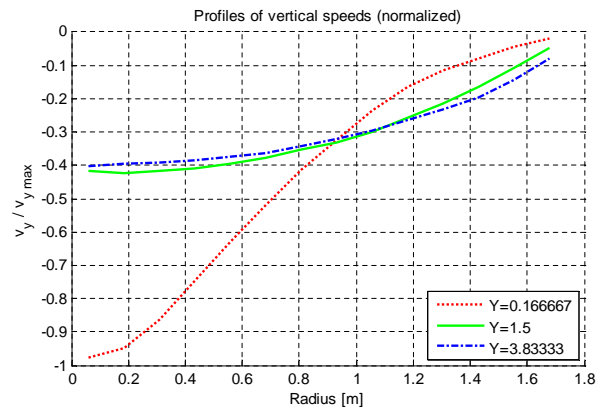


Figure 8: Speed profiles at different heights in the PBR reactor vessel




Figure 9: Reactor flow speed




Figure 10: Vertical speeds near the outlet of the reactor




Figure 11: Contacts near the orifice of the 2D hopper, during steady flow




Figure 12: Porosity in the reactor: comparison with the DEM method

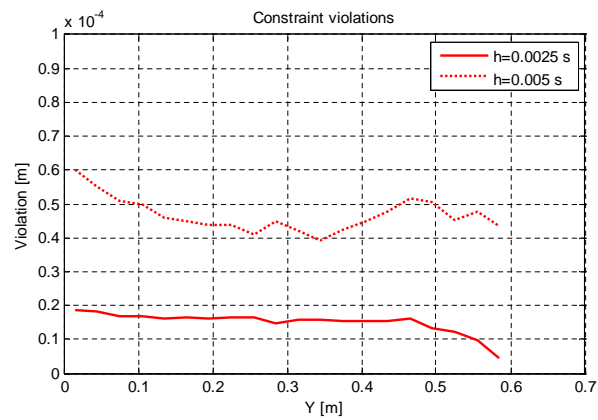


Figure 13: Maximum violations in contact constraints at different heights in the hopper, for different time steps