

LARGE-SCALE PARALLEL MULTIBODY DYNAMICS WITH FRICTIONAL CONTACT ON THE GPU

Dan Negrut*

Dept. of Mechanical Engineering
University of Wisconsin–Madison
Madison, WI 53706
Email: negrut@wisc.edu

Alessandro Tasora

Dept. of Mechanical Engineering
University of Parma
Parma, Italy
Email: tasora@ied.unipr.it

Mihai Anitescu

Math. and Comp. Sci. Division
Argonne National Laboratory
Argonne, IL 60439
Email: anitescu@mcs.anl.gov

ABSTRACT

In the context of simulating the frictional contact dynamics of large systems of rigid bodies, this paper reviews a novel method for solving large cone complementarity problems by means of a fixed-point iteration algorithm. The method is an extension of the Gauss-Seidel and Gauss-Jacobi methods with overrelaxation for symmetric convex linear complementarity problems. Convergent under fairly standard assumptions, the method is implemented in a parallel framework by using a single instruction multiple data (SIMD) computation paradigm promoted by the Compute Unified Device Architecture (CUDA) library for graphical processing unit (GPU) programming. The framework supports the simulation of problems with more than one million bodies in contact. Simulation thus becomes a viable tool for investigating the dynamics of complex systems such as ground vehicles running on sand, powder composites, and granular material flow.

INTRODUCTION

Approximating through numerical simulation the time evolution of a multibody system in the presence of friction and contact/impact phenomena continues to be a very challenging task. For instance, results reported in [1] indicate that the most widely used commercial software package for multibody dynamics simulation has significant difficulties in handling a simple problem such as a collection of balls falling in a box, whenever the number of balls becomes larger than 50. In fact, the problem becomes practically intractable when the number of bodies becomes larger than 100. Presented here is an algorithm that can robustly and efficiently approximate the dynamics of rigid bodies undergoing frictional contact [2]. Posing challenges of its own, the case

of *deformable* frictional contact is extensively discussed in [3,4] and falls outside the scope of this work.

Two approaches are most often considered when simulating the dynamics of a multibody system with frictional contact. First is the class of so-called penalty methods, where it is assumed that every time two rigid bodies come in frictional contact, the interaction can be represented by a collection of stiff springs along with damping elements that act at the interface of the two bodies [5–8]. Implementing these regularization approaches requires little effort beyond that usually associated with developing a multibody dynamics simulation code. Furthermore, this methodology can easily accommodate complex frictional contact mechanisms, as it allows for a large number of “tuning” parameters that, in general, can be adjusted to control the dynamics of the frictional contact interaction. What has prevented the widespread use of this solution is the small step-size at which the numerical integration formula, because of stability limitations, is able to advance the simulation, a drawback related to the stiff spring elements artificially included in the model. Most of the time, this step-size limitation is counterbalanced by the use of implicit integration formulas, a proposition that typically comes at a price as it requires the solution of a discretization nonlinear system at each integration time-step. This in turn leads to a heavy computational burden for scenarios with a large number of active frictional contact events.

A second approach, and the one pursued in this work, relies on an different mathematical framework capable of handling applications with hundreds of thousands of frictional contact events. The algorithms in this class draw on time-stepping procedures that produce weak solutions of the differential variational inequality (DVI) that describes the time evolution of rigid bodies with collision, contact, and friction. The DVI as a problem formulation was recently introduced in full generality and classified

*Address all correspondence to this author.

by differential index [9, 10], though earlier numerical approaches based on DVI formulations do exist [11–13]. Recent work on time-stepping schemes has included both acceleration-force linear complementarity problem (LCP) approaches [14–16] and velocity-impulse LCP-based time-stepping methods [17–20]. The LCP problems, obtained as a result of the introduction of inequalities in time-stepping schemes for DVI, coupled with a polyhedral approximation of the friction cone, must be solved at each time step in order to determine the system state configuration as well as the Lagrange multipliers associated with the frictional contact problem [12, 17]. If the simulation entails a large number of contacts and rigid bodies, as is the case of part feeders, packaging machines, and granular flows, the computational burden of classical LCP solvers can become significant. Indeed, a well-known class of numerical solutions for LCP problems is based on *simplex methods*, also known as *direct* or *pivoting* methods [21]. However, these methods may exhibit exponential worst-case complexity [22]. Further complicating the numerical solution, since the three-dimensional Coulomb friction case leads to a nonlinear complementarity problem (NCP), the use of a polyhedral approximation to morph the NCP into a LCP problem introduces artificial anisotropy, which affects friction because friction cones become faceted friction pyramids [16–18]. In fact, this discrete and finite approximation of friction cones is one of the reasons for the large dimension of the problem that needs to be solved in multibody dynamics with frictional contact.

In order to circumvent the limitations imposed by the use of classical LCP solvers and the limited accuracy associated with polyhedral approximations of the friction cone, a parallel fixed-point iteration method with projection on a convex set is proposed, which can directly solve large cone complementarity problems with low computational overhead. The method is based on a time-stepping formulation that solves at every step a cone constrained optimization problem [23]. The time-stepping scheme has been proved to converge in a measure differential inclusion sense to the solution of the original continuous-time DVI. For the proposed approach, about 80% of the computational effort in simulating frictional contact dynamics is spent solving the cone complementarity problem (CCP). The goal of this work is to solve the CCP in parallel using commodity high-performance computing hardware. Specifically, a methodology is proposed that hinges on the use of parallel computational resources available on NVIDIA's GPU cards, which can currently handle 12,288 live computational threads simultaneously on the GeForce 8800 series. Tapping into this massively parallel computational resource has been facilitated by NVIDIA's sharing of a well-integrated application programming interface supported by the CUDA library [24].

FORMULATION OF THE MULTIBODY DYNAMICS WITH FRICTIONAL CONTACT PROBLEM

The equations that govern the time evolution of a multibody system can be expressed in the form (see, for instance, [25])

$$\begin{aligned} \dot{\mathbf{q}} &= \mathbf{L}(\mathbf{q})\mathbf{v} \\ \mathbf{M}\dot{\mathbf{v}} &= \mathbf{f}^A(t, \mathbf{q}, \mathbf{v}), \end{aligned} \quad (1)$$

where $\mathbf{q} = [\mathbf{r}_1^T, \boldsymbol{\varepsilon}_1^T, \dots, \mathbf{r}_{n_b}^T, \boldsymbol{\varepsilon}_{n_b}^T]^T \in \mathbb{R}^{6n_b}$ are generalized positions, $\mathbf{v} = [\dot{\mathbf{r}}_1^T, \dot{\boldsymbol{\omega}}_1^T, \dots, \dot{\mathbf{r}}_{n_b}^T, \dot{\boldsymbol{\omega}}_{n_b}^T]^T \in \mathbb{R}^{6n_b}$ are generalized velocities, and n_b represents the number of bodies in the system. The matrix \mathbf{M} is the generalized mass matrix, and $\mathbf{f}^A(t, \mathbf{q}, \mathbf{v})$ represents the vector of generalized applied forces. The convention used here is that any symbol in bold represents a vector or matrix quantity, while an overbar represents a vector quantity represented in the local, body-fixed reference frame associated with a body that is inferred from the context.

The formulation of the equations of motion draws on the so-called absolute, or Cartesian, representation of the attitude of each rigid body in the system. For each body i , its orientation is described by a set of three Euler angles, $\boldsymbol{\varepsilon}_i \in \mathbb{R}^3$, following the 3-1-3 local rotation sequence (see, for instance, [25]). The rate at which each body changes its orientation is captured by the local angular velocity $\dot{\boldsymbol{\omega}}_i \in \mathbb{R}^3$. The location of each body is uniquely determined by a position vector $\mathbf{r}_i = [x_i, y_i, z_i]^T$ that specifies where the body-fixed centroidal reference frame is located. The translational velocity of the body is simply $\dot{\mathbf{r}}_i$, where an overdot represents time differentiation. With this set of generalized coordinates, the mass matrix \mathbf{M} remains constant and diagonal between any realigning of a body-fixed centroidal reference frame, which can potentially be employed to avoid Euler angles singularities. Also note that since for each body i there is a locally nonsingular matrix $\mathbf{B}(\boldsymbol{\varepsilon}_i)$ such that $\dot{\boldsymbol{\omega}}_i = \mathbf{B}(\boldsymbol{\varepsilon}_i)\dot{\boldsymbol{\varepsilon}}_i$, the operator $\mathbf{L}(\mathbf{q})$ that relates the time derivative of the level-zero generalized coordinates to the level-one generalized coordinates is generally not the identity matrix. Note that there are no bilateral constraints present in the current formulation. This case is discussed in [2, 26], and a paper presenting a parallel methodology for the general case of bilateral and unilateral constraints is forthcoming.

Two rigid bodies should not penetrate, and, if they are in contact, there should be friction acting at the interface. In order to enforce the nonpenetration constraint, a gap function $\Phi(\mathbf{q}) \in \mathbb{R}$ is assumed to exist and satisfy

$$\Phi(\mathbf{q}) = \begin{cases} > 0 & \text{if the bodies are separated,} \\ = 0 & \text{if the bodies touch each other,} \\ < 0 & \text{if the bodies are interpenetrating.} \end{cases} \quad (2)$$

For such a function, the nonpenetration constraint becomes $\Phi(\mathbf{q}) \geq 0$. An example of such a mapping is the signed distance function [27], which, when the bodies are smooth and convex, is differentiable at least up to some value of the interpenetration [28]. For most cases, even simple ones involving the relative

position of two spheres, a differentiable signed distance function cannot be defined for certain configurations \mathbf{q} . The fact that $\Phi(\mathbf{q})$ can be differentially defined only on a neighborhood of the set $\Phi(\mathbf{q}) \geq 0$ can be accommodated at the cost of making the analysis substantially more involved [29]. This will not be done here. Indeed, in order to simplify the discussion, a *differentiability of geometrical constraint data assumption* is made, in that any contact is assumed to be described by a gap function $\Phi(\mathbf{q})$ that is twice continuously differentiable. Note that for an overwhelming majority of applications, when one deals with convex geometries and with suitably small numerical integration step-sizes, this assumption is easily verified.

The friction model used here is the Coulomb model, which leads to frictional conic constraints regarded as an extension of complementarity models discussed in [17, 18]. If the configuration of the system \mathbf{q} is such that a contact i is active, that is, $\Phi_i(\mathbf{q}) = 0$, then a normal force and a tangential force are going to act on each of the two bodies at the contact point. Let \mathbf{n}_i be the normal at the contact pointing toward the exterior of the first contact body, where of the two bodies in contact j and k , the first contact body is always the body with lower index. 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 orthonormal vectors. Although they typically depend on \mathbf{q} , this dependency is not explicitly indicated, in order to keep the notation simple.

The frictional contact force is impressed on the system by means of 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 $\mathbf{F}_{i,T} = \hat{\gamma}_{i,u} \mathbf{u}_i + \hat{\gamma}_{i,w} \mathbf{w}_i$.

The Coulomb model consists of the following constraints:

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

where $\mathbf{v}_{i,T}$ is the relative tangential velocity at contact i . The magnitude of the friction force depends on the friction coefficient $\mu_i \in \mathbb{R}^+$, which typically has a value between 0 and 1 for most materials, and is instrumental in linking the magnitude of the tangential and normal forces through a constitutive type equation¹.

The first part of the constraint can be restated as

$$\mathbf{F}_i = \mathbf{F}_{i,N} + \mathbf{F}_{i,T} = \hat{\gamma}_{i,n} \mathbf{n}_i + \hat{\gamma}_{i,u} \mathbf{u}_i + \hat{\gamma}_{i,w} \mathbf{w}_i \in \mathcal{C}, \quad (3)$$

where \mathcal{C} is a cone in three dimensions, whose slope is $\tan^{-1} \mu_i$. Defining by $\langle \cdot, \cdot \rangle$ the inner product of two vectors, the constraint $\langle \mathbf{F}_{i,T}, \mathbf{v}_{i,T} \rangle = -\|\mathbf{F}_{i,T}\| \|\mathbf{v}_{i,T}\|$ requires that the tangential force

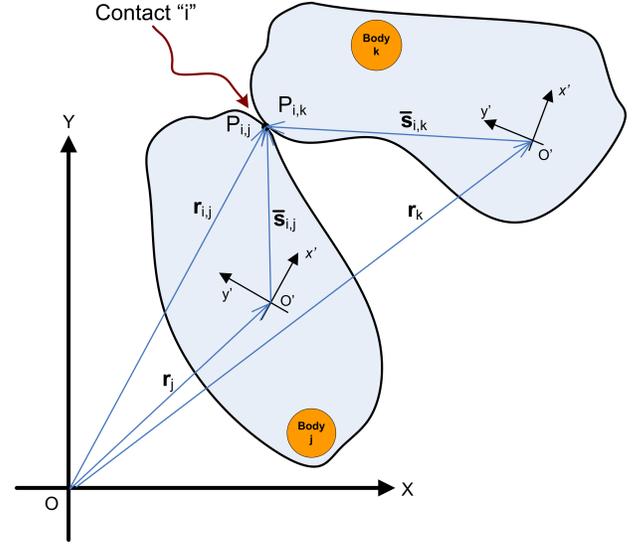


Figure 1. Contact i active between two bodies $j, k \in \{1, 2, \dots, n_b\}$

be opposite to the tangential velocity. This results in the friction force being dissipative. In fact, an equivalent convenient way of expressing this constraint is by using the maximum dissipation principle $(\hat{\gamma}_{i,u}, \hat{\gamma}_{i,w}) = \underset{\sqrt{\hat{\gamma}_{i,u}^2 + \hat{\gamma}_{i,w}^2} \leq \mu_i \hat{\gamma}_{i,n}}{\operatorname{argmin}} \mathbf{v}_{i,T}^T (\hat{\gamma}_{i,u} \mathbf{u}_i + \hat{\gamma}_{i,w} \mathbf{w}_i)$ [13, 30].

For this minimization problem, it is relatively straightforward to establish a connection between the first-order necessary KKT conditions [31] and the Coulomb model above. Effectively, the condition in this equation states that the friction force is such that, given a tangential velocity and a normal force, the power dissipated is maximized.

The contribution of the frictional contact forces in the equations of motion, Eq. (1), is through a set of generalized forces associated with each active contact in the model. Based on Newton's third law, each body experiences a force of the same magnitude but opposite direction at the point of contact. Therefore, the virtual work associated with the frictional contact force \mathbf{F}_i between bodies j and k becomes $\delta W_i = \delta \mathbf{r}_{i,j}^T \mathbf{F}_i - \delta \mathbf{r}_{i,k}^T \mathbf{F}_i$. As illustrated in Fig. (1), $\mathbf{r}_{i,j} = \mathbf{r}_j + \mathbf{A}_j \tilde{\mathbf{s}}_{i,j}$ gives the position, expressed in the global inertial reference frame, of the contact point $P_{i,j}$ on body j , and $\delta \mathbf{r}_{i,j} = \delta \mathbf{r}_j + \mathbf{A}_j \delta \tilde{\pi}_j \tilde{\mathbf{s}}_{i,j} = \delta \mathbf{r}_j - \mathbf{A}_j \tilde{\mathbf{s}}_{i,j} \delta \tilde{\pi}_j$ represents a virtual displacement of body j , which is due to a virtual translational displacement of the body center of mass, $\delta \mathbf{r}_j$, and a virtual rotation $\delta \tilde{\pi}_j$, expressed in the local body j reference frame. Similar quantities are defined in conjunction with body k . Note that the operator $\tilde{\cdot}$ acting on a vector $\mathbf{h} = [h_1, h_2, h_3]^T$ produces a skew symmetric matrix $\tilde{\mathbf{h}} \equiv \mathbf{H} \in \mathbb{R}^{3 \times 3}$ with $\mathbf{H}(1,2) = -h_3$, $\mathbf{H}(1,3) = h_2$, and $\mathbf{H}(2,3) = -h_1$. From Eq. (3),

$$\begin{aligned} \delta W_i &= (\delta \mathbf{r}_j^T + \delta \tilde{\pi}_j^T \tilde{\mathbf{s}}_{i,j} \mathbf{A}_j^T) (\hat{\gamma}_{i,n} \mathbf{n}_i + \hat{\gamma}_{i,u} \mathbf{u}_i + \hat{\gamma}_{i,w} \mathbf{w}_i) \\ &\quad - (\delta \mathbf{r}_k^T + \delta \tilde{\pi}_k^T \tilde{\mathbf{s}}_{i,k} \mathbf{A}_k^T) (\hat{\gamma}_{i,n} \mathbf{n}_i + \hat{\gamma}_{i,u} \mathbf{u}_i + \hat{\gamma}_{i,w} \mathbf{w}_i) \\ &= \delta \mathbf{q}^T \mathbf{D}_i (\hat{\gamma}_{i,n} \mathbf{n}_i + \hat{\gamma}_{i,u} \mathbf{u}_i + \hat{\gamma}_{i,w} \mathbf{w}_i) \\ &= \delta \mathbf{q}^T (\hat{\gamma}_{i,n} \mathbf{D}_{i,n} + \hat{\gamma}_{i,u} \mathbf{D}_{i,u} + \hat{\gamma}_{i,w} \mathbf{D}_{i,w}), \end{aligned}$$

¹Though the original Coulomb model distinguishes between static μ_s and kinetic μ_k friction coefficients, where usually the kinetic coefficient is slightly lower than its static counterpart, in this work both are considered to have the same value μ . The difference is not relevant for the discussion; it suffices to say that to correct this approach would require one to adjust the friction coefficient adaptively during the simulation depending on the slipping speed, so as to express complex nonlinearities in μ as a function of speed.

where, with \mathbf{I}_3 the 3×3 identity matrix, the projection matrix $\mathbf{D}_i \in \mathbb{R}^{6n_b \times 3}$ is defined for contact i as

$$\mathbf{D}_i^T = [\mathbf{0} \dots \mathbf{I}_3 (\tilde{\mathbf{s}}_{i,j} \mathbf{A}_j^T)^T \mathbf{0} \dots \mathbf{0} -\mathbf{I}_3 -(\tilde{\mathbf{s}}_{i,k} \mathbf{A}_k^T)^T \dots \mathbf{0}],$$

and $\mathbf{D}_{i,n} \equiv \mathbf{D}_i \mathbf{n}_i$, $\mathbf{D}_{i,u} \equiv \mathbf{D}_i \mathbf{u}_i$, and $\mathbf{D}_{i,w} \equiv \mathbf{D}_i \mathbf{w}_i$. Note that the velocity at the point of contact can also be expressed in terms of $\mathbf{D}_{i,u}$ and $\mathbf{D}_{i,w}$. To this end, the tangential velocity can be expressed as $\mathbf{v}_{i,T}^T = \mathbf{v}^T \mathbf{D}_i$, and therefore the power dissipated can be equivalently expressed as $\mathbf{v}_{i,T}^T (\hat{\gamma}_{i,u} \mathbf{u}_i + \hat{\gamma}_{i,w} \mathbf{w}_i) = \mathbf{v}^T \mathbf{D}_i (\hat{\gamma}_{i,u} \mathbf{u}_i + \hat{\gamma}_{i,w} \mathbf{w}_i) = \mathbf{v}^T (\hat{\gamma}_{i,u} \mathbf{D}_{i,u} + \hat{\gamma}_{i,w} \mathbf{D}_{i,w})$

When one revisits Eq. (1) and assumes a set of p active constraints at time t , a more specific expression can be provided for the differential equations governing the time evolution of the multibody system by singling out the contribution of the frictional contact force. Drawing on the Coulomb model discussed, the following differential variational inequality is associated with the time evolution of the multibody system [32]:

$$\begin{aligned} \dot{\mathbf{q}} &= \mathbf{L}(\mathbf{q}) \mathbf{v} \\ \mathbf{M} \dot{\mathbf{v}} &= \mathbf{f}(t, \mathbf{q}, \mathbf{v}) + \sum_{i=1}^p (\hat{\gamma}_{i,n} \mathbf{D}_{i,n} + \hat{\gamma}_{i,u} \mathbf{D}_{i,u} + \hat{\gamma}_{i,w} \mathbf{D}_{i,w}) \\ 1 \leq i \leq p : \hat{\gamma}_{i,n} &\geq 0 \perp \Phi_i(\mathbf{q}) \geq 0, \quad \text{and} \\ (\hat{\gamma}_{i,u}, \hat{\gamma}_{i,w}) &= \underset{\mu \hat{\gamma}_{i,n} \geq \sqrt{(\hat{\gamma}_{i,u})^2 + (\hat{\gamma}_{i,w})^2}}{\operatorname{argmin}} \mathbf{v}^T (\hat{\gamma}_{i,u} \mathbf{D}_{i,u} + \hat{\gamma}_{i,w} \mathbf{D}_{i,w}). \end{aligned} \quad (4)$$

The Coulomb model used in this work is the predominant model used in the engineering literature to describe dry friction. Unfortunately, the model may be inconsistent: there exist configurations for which the model does not have a solution [14, 20]. This situation has led to the need to explore weaker formulations where the forces are measures and Newton's law is satisfied in a measure differential inclusion sense [20]. It has been shown that solutions in that sense do exist and can be found by time-stepping schemes [33]. Note that the type of applications targeted — pebble-bed nuclear reactor simulation, granular flow dynamics, etc. — all collisions that appear during the simulation can be assumed of the inelastic type.

PROPOSED SOLUTION TO THE PROBLEM OF MULTIBODY DYNAMICS WITH FRICTIONAL CONTACT

Starting from a time-step $t^{(l)}$ with position $\mathbf{q}^{(l)}$ and velocity $\mathbf{v}^{(l)}$, the solution is found at the new time-step $t^{(l+1)} = t^{(l)} + h$ as the solution of the following optimization problem with equilib-

rium constraints:

$$\begin{aligned} \mathbf{M}(\mathbf{v}^{(l+1)} - \mathbf{v}^l) &= h \mathbf{f}(t^{(l)}, \mathbf{q}^{(l)}, \mathbf{v}^{(l)}) \\ &+ \sum_{i \in \mathcal{A}(q^{(l)}, \delta)} (\gamma_{i,n} \mathbf{D}_{i,n} + \gamma_{i,u} \mathbf{D}_{i,u} + \gamma_{i,w} \mathbf{D}_{i,w}), \end{aligned} \quad (5)$$

$$i \in \mathcal{A}(q^{(l)}, \delta) : 0 \leq \frac{1}{h} \Phi_i(\mathbf{q}^{(l)}) + \mathbf{D}_{i,n}^T \mathbf{v}^{(l+1)} \perp \gamma_n^i \geq 0, \text{ and} \quad (6)$$

$$(\gamma_{i,u}, \gamma_{i,w}) = \underset{\mu \gamma_{i,n} \geq \sqrt{\gamma_{i,u}^2 + \gamma_{i,w}^2}}{\operatorname{argmin}} \mathbf{v}^T (\gamma_{i,u} \mathbf{D}_{i,u} + \gamma_{i,w} \mathbf{D}_{i,w}), \quad (7)$$

$$\mathbf{q}^{(l+1)} = \mathbf{q}^{(l)} + h \mathbf{L}(\mathbf{q}^{(l)}) \mathbf{v}^{(l+1)}. \quad (8)$$

Here, for a conveniently chosen small value of $\delta > 0$,

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

γ_s represents the constraint impulse of a contact constraint, that is, $\gamma_s = h \hat{\gamma}_s$, for $s = n, u, w$. The $\frac{1}{h} \Phi_i(\mathbf{q}^{(l)})$ term achieves constraint stabilization; its effect is discussed in [29]. It has been shown that as the step size $h \rightarrow 0$, the scheme converges to the solution of a measure differential inclusion [23]. Numerical solutions for the case when the nonlinear frictional contact constraint is approximated by a piecewise linear cone can be found by Lemke's algorithm [18]. Nonetheless, it has been shown in [34] that as the number of constraints in the problem increases, the computational cost of Lemke's method increases far faster than linearly with the size of the problem. As an alternative, the problem is cast as a monotone optimization problem by introducing a relaxation over the complementarity constraints; that is, the time-stepping scheme is modified by replacing Eq. (6) with

$$i \in \mathcal{A}(q^{(l)}, \delta) : 0 \leq \frac{1}{h} \Phi_i(\mathbf{q}^{(l)}) + \mathbf{D}_{i,n}^T \mathbf{v}^{(l+1)} - \mu_i \sqrt{(\mathbf{v}^T \mathbf{D}_{i,u})^2 + (\mathbf{v}^T \mathbf{D}_{i,w})^2} \perp \gamma_n^i \geq 0. \quad (9)$$

It has been shown in [23] that, as $h \rightarrow 0$, the solution of the modified time-stepping scheme will approach the solution of the same measure differential inclusion as the original scheme. Additionally, the iterates produced by the modified scheme also approach the ones of the original scheme provided that $\mu_i \gamma_{i,n} \sqrt{(\mathbf{v}^T \mathbf{D}_{i,u})^2 + (\mathbf{v}^T \mathbf{D}_{i,w})^2} \ll 1$ [34]. This assumption is satisfied in the regime in which pebble-bed nuclear reactors operate [35], as well as for granular flow applications, two classes of applications targeted by the proposed approach. Generally, the above assumption will be satisfied whenever there is little friction between the two bodies in contact (low μ) or when the relative velocity at the contact point is small.

The KKT optimality conditions for the equilibrium constraint in Eq. (7) state that, for any $i \in \mathcal{A}(q^{(l)}, \delta)$, there exists a Lagrange multiplier λ_i such that

$$\begin{aligned} \lambda_i \gamma_{i,u} &= -\mathbf{v}^T \mathbf{D}_{i,u}, \quad \lambda_i \gamma_{i,w} = -\mathbf{v}^T \mathbf{D}_{i,w}, \\ \lambda_i &\geq 0 \perp \mu_i \gamma_{i,n} - \sqrt{\gamma_{i,u}^2 + \gamma_{i,w}^2} \geq 0. \end{aligned} \quad (10)$$

If for $i \in \mathcal{A}(\mathbf{q}^{(l)}, \delta)$, $\mathbf{c}_i^T \equiv [\gamma_{i,n}, \gamma_{i,u}, \gamma_{i,w}]$, and $\mathbf{g}_i^T \equiv \left[\frac{1}{h} \Phi_i(\mathbf{q}^{(l)}) + \mathbf{D}_{i,n}^T \mathbf{v}^{(l+1)}, \mathbf{D}_{i,u}^T \mathbf{v}^{(l+1)}, \mathbf{D}_{i,w}^T \mathbf{v}^{(l+1)} \right]$, then drawing on Eqs. (9) and (10), one can show that $\mathbf{g}_i^T \mathbf{c}_i = 0$, and thus $\mathbf{g}_i \perp \mathbf{c}_i$. Defining the cones

$$\Lambda_i = \left\{ \mathbf{g} = [g_1, g_2, g_3]^T \in \mathbb{R}^3 \mid g_1 \geq \mu_i \sqrt{g_2^2 + g_3^2} \right\},$$

$$C_i = \left\{ \mathbf{c} = [c_1, c_2, c_3]^T \in \mathbb{R}^3 \mid \mu_i c_1 \geq \sqrt{c_2^2 + c_3^2} \right\},$$

then Λ_i is the negative polar cone of C_i ; that is, $\mathbf{g} \in \Lambda_i$ and $\mathbf{c} \in C_i$ implies that $\mathbf{g}^T \mathbf{c} \geq 0$. Here, C° , the polar cone of a given cone $C \subset \mathbb{R}^m$, is defined as $C^\circ = \{ \mathbf{x} \in \mathbb{R}^m \mid \langle \mathbf{x}, \mathbf{y} \rangle \leq 0, \forall \mathbf{y} \in C \}$. Then based on Eqs. (7) and (9) the following set of cone complementarity constraints holds:

$$-\mathbf{g}_i \in C_i^\circ \perp \mathbf{c}_i \in C_i, \quad \forall i \in \mathcal{A}(\mathbf{q}^{(l)}, \delta). \quad (11)$$

In what follows, the focus shifts back to reformulating the optimization problem with equilibrium constraints of Eqs. (5) through (8) to account for the above cone complementarity constraints. To this end, the notation $\mathbf{Q} \equiv \mathbf{M}\mathbf{v}^{(l)} + h\mathbf{f}(t^{(l)}, \mathbf{q}^{(l)}, \mathbf{v}^{(l)})$, is introduced. Then, Eq. (5) is reformulated as $\mathbf{M}\mathbf{v}^{(l+1)} = \mathbf{Q} + \mathbf{D}\boldsymbol{\gamma}$, where with $n_{\mathcal{A}}$ being the cardinality of $\mathcal{A}(\mathbf{q}^{(l)}, \delta)$, $\mathbf{D} \equiv [\mathbf{D}_1 \dots \mathbf{D}_{n_{\mathcal{A}}}] \in \mathbb{R}^{6n_b \times 3n_{\mathcal{A}}}$, $\boldsymbol{\gamma} = [\gamma_1^T \dots \gamma_{n_{\mathcal{A}}}^T]^T \in \mathbb{R}^{3n_{\mathcal{A}}}$, and, for $i \in \mathcal{A}(\mathbf{q}^{(l)}, \delta)$, $\mathbf{D}_i \equiv [\mathbf{D}_{i,n} \ \mathbf{D}_{i,u} \ \mathbf{D}_{i,w}] \in \mathbb{R}^{6n_b \times 3}$ and $\gamma_i^T \equiv [\gamma_{i,n} \ \gamma_{i,u} \ \gamma_{i,w}]$. Next, the vector $\mathbf{b} \in \mathbb{R}^{3n_{\mathcal{A}}}$ is defined as $\mathbf{b}^T \equiv \left[\frac{1}{h} \Phi_1(\mathbf{q}^{(l)}) \ 0 \ 0 \dots \frac{1}{h} \Phi_{n_{\mathcal{A}}}(\mathbf{q}^{(l)}) \ 0 \ 0 \right]$. Here, for an arbitrary vector $\mathbf{h} \in \mathbb{R}^{3n_{\mathcal{A}}}$, the notation \mathbf{h}_i , $1 \leq i \leq n_{\mathcal{A}}$, is used to represent the entries in rows $3(i-1)+1$, $3(i-1)+2$, $3i$, that is, the entries associated with contact i . Finally, with the notation $\mathbf{N} \equiv \mathbf{D}^T \mathbf{M}^{-1} \mathbf{D}$, and $\mathbf{d} \equiv \mathbf{b} + \mathbf{D}^T \mathbf{M}^{-1} \mathbf{Q}$, the configuration of the multibody system at $t^{(l+1)}$ is obtained as the solution of the following:

$$-\left(\mathbf{N}\boldsymbol{\gamma}^{(l+1)} + \mathbf{d} \right)_i \in C_i^\circ \perp \gamma_i^{(l+1)} \in C_i, \quad 1 \leq i \leq n_{\mathcal{A}}, \quad (12)$$

$$\mathbf{M}\mathbf{v}^{(l+1)} = \mathbf{Q} + \mathbf{D}\boldsymbol{\gamma}, \quad (13)$$

$$\mathbf{q}^{(l+1)} = \mathbf{q}^{(l)} + h\mathbf{L}(\mathbf{q}^{(l)})\mathbf{v}^{(l+1)}. \quad (14)$$

The frictional contact forces $\boldsymbol{\gamma}^{(l+1)}$ are obtained by solving the cone complementarity problem in Eqs. (12) and (13), where, at each iteration, the velocity is evaluated as $\mathbf{v}^{(l+1)} = \mathbf{M}^{-1} \left(\mathbf{Q} + \mathbf{D}\boldsymbol{\gamma}^{(l+1)} \right)$. The attitude of each body in the system is obtained by using Eq. (14). Note that \mathbf{N} is positive semidefinite, while \mathbf{M} is diagonal, constant, and positive definite. Equation (12) in fact represents the optimality condition of the following cone complementarity problem (CCP):

$$\min f(\boldsymbol{\gamma}) = \frac{1}{2} \boldsymbol{\gamma}^T \mathbf{N}\boldsymbol{\gamma} + \mathbf{d}^T \boldsymbol{\gamma} \quad \text{s.t.} \quad \gamma_i \in C_i, \quad \forall 1 \leq i \leq n_{\mathcal{A}},$$

whose solution is found by using an iterative algorithm that, start-

ing with an arbitrarily chosen $\boldsymbol{\gamma}^{(0)}$, computes the iteration $r+1$, $r \geq 0$,

$$\boldsymbol{\gamma}_i^{r+1} = \rho \Pi_{C_i} \left[\boldsymbol{\gamma}_i^r - \omega \eta_i (\mathbf{N}\boldsymbol{\gamma} + \mathbf{d}) + \sum_{m \in \mathcal{D}(i,r)} \mathbf{K}_{i,m}^r (\boldsymbol{\gamma}_m^{r+1} - \boldsymbol{\gamma}_m^r) \right] + (1-\rho) \boldsymbol{\gamma}_i^r, \quad (15)$$

where η_i will be defined in Algorithm 1 below, $0 < \rho \leq 1$, and $\omega > 0$ are two parameters, and, for each r , $\mathbf{K}_{i,m}^r$ is a coefficient matrix that indicates how the frictional contact force associated with contact m gets reflected in the computation of the frictional contact force associated with contact i . Here $\mathcal{D}(i,r)$ represents the set of contacts considered when updating the frictional contact forces associated with contact i . Note that for a Gauss-Jacobi-type iteration $\mathcal{D}(i,r) = \emptyset$ and for a Gauss-Seidel type iteration $\mathcal{D}(i,r) = \{1, 2, \dots, i-1\}$, but other update strategies, which might depend on the iteration index r , can and will be pursued. Finally, the operator Π_{C_i} is the cone i projection operator; see, for instance, [26].

The iterative scheme in Eq. (15) was proved to converge under mild assumptions that can be met by a suitable choice of relaxation parameter ω [2]. Therein, it was also pointed out that a Gauss-Seidel update sequence in the iterative process led to a robust algorithm. Although convenient for the convergence analysis, Eq. (15) is not the form that is considered for software implementation. Rather, an inner loop iteration algorithm, that also updates the speed $\mathbf{v}^{(l+1)}$, is provided in pseudocode below:

Algorithm 1: The Inner Iteration Loop.

1. For $1 \leq i \leq n_{\mathcal{A}}$, evaluate $\mathbf{S}^i = \mathbf{M}^{-1} \mathbf{D}_i \in \mathbb{R}^{6n_b \times 3}$, and $\eta_i = 3/\text{Trace}(\mathbf{D}_i^T \mathbf{S}^i)$.
 2. If warm start with initial guess $\boldsymbol{\gamma}^*$, then set $\boldsymbol{\gamma}^0 = \boldsymbol{\gamma}^*$, otherwise $\boldsymbol{\gamma}^0 = \mathbf{0}$.
 3. Initialize speeds: $\mathbf{v}^{(l+1)} = \sum_{i=1}^{n_{\mathcal{A}}} \mathbf{S}^i \boldsymbol{\gamma}^0 + \mathbf{M}^{-1} \mathbf{Q}$.
 4. For $i = 1, \dots, n_{\mathcal{A}}$, perform the updates:
$$\boldsymbol{\gamma}_i^{\text{prelim}} = \boldsymbol{\gamma}_i^r - \omega \eta_i \left(\mathbf{D}_i^T \mathbf{v}^{(l+1)} + \mathbf{b}_i \right);$$

$$\boldsymbol{\gamma}_i^{r+1} = \rho \Pi_{C_i} \left(\boldsymbol{\gamma}_i^{\text{prelim}} \right) + (1-\rho) \boldsymbol{\gamma}_i^r;$$

$$\Delta \boldsymbol{\gamma}_i^{r+1} = \boldsymbol{\gamma}_i^{r+1} - \boldsymbol{\gamma}_i^r;$$

$$\mathbf{v}^{(l+1)} := \mathbf{v}^{(l+1)} + \mathbf{S}_i \Delta \boldsymbol{\gamma}_i^{r+1}.$$
 5. Repeat step 4 by looping on the list of contacts in reverse order, if symmetric updates are desired.
 6. $r := r+1$. Repeat from 4 until convergence, or until $r > r_{\max}$.
-

The stopping criteria is based on the value of the velocity update. The overall algorithm that provides an approximation to the solution of Eqs. (12) through (14) relies on Algorithm 1 above and requires the following steps:

Algorithm 2: The Outer, Time-Stepping, Loop.

1. Set $t = 0$, step counter $l = 0$, provide initial values for $\mathbf{q}^{(l)}$ and $\mathbf{v}^{(l)}$.
2. Perform collision detection between bodies, obtaining $n_{\mathcal{A}}$ possible contact points within a distance δ . For each con-

tact i , compute $\mathbf{D}_{i,n}$, $\mathbf{D}_{i,u}$, $\mathbf{D}_{i,w}$, and residual $\Phi_i(\mathbf{q})$, which also provides \mathbf{b}_i .

3. For each body, compute forces $\mathbf{f}(t^{(l)}, \mathbf{q}^{(l)}, \mathbf{v}^{(l)})$ and then \mathbf{Q} .
4. Use Algorithm 1 to solve the cone complementarity problem and obtain unknown impulse γ and velocity $\mathbf{v}^{(l+1)}$.
5. Update positions using $\mathbf{q}^{(l+1)} = \mathbf{q}^{(l)} + h\mathbf{L}(\mathbf{q}^{(l)})\mathbf{v}^{(l+1)}$.
6. Increment $t := t + h$, $l := l + 1$, and repeat from step 2 until $t > t_{\text{end}}$

Note that choosing a proper value δ for the collision envelope is not trivial. On the one hand, if a very small or zero value is used, contacts will enter the CCP solver only when it is too late and some amount of interpenetration will be unavoidable, which in turn adversely impacts the stability of the method. On the other hand, if too large a value is used, the collision detection algorithm will return too many potential contacts that waste computational resources and could occasionally create troubles with convex shapes, which decreases the efficiency of the method. Knowing \mathbf{v} , a simple yet efficient heuristic is to choose δ as a rough approximation of the maximum distance that can be traveled within a step-size by an arbitrary point of a moving body.

PARALLEL IMPLEMENTATION DETAILS

The parallel implementation for the multibody dynamics frictional contact problem is heavily dependent on the underlying hardware used to run the simulation. Specifically, the algorithm proposed is well suited for running on parallel platforms that support the single instruction multiple data (SIMD) computational paradigm, which is ideally suited for handling problems with contacts in excess of hundreds of thousand. NVIDIA's GeForce 80 family of Graphics Processing Units (GPUs) has been adopted as the implementation platform. Priced at \$460 per card, the GeForce 8800 GTX has been clocked at 320 GFlop, about seven times faster than an Intel Core 2 Duo running at 3 GHz, and this gap is bound to grow in the immediate future². A description of the hardware behind the GeForce 8800 GTX card used is beyond the scope of this paper. It suffices to indicate that for the frictional contact problem discussed here, contacts are processed by each of the 16 streaming multiprocessors in warps of 16 contacts per streaming multiprocessor. As far as memory allocation is concerned, each streaming multiprocessor has 8192 registers that each can hold a float or an integer, and each has 16 KB of shared memory that can be shared by all live threads associated with the respective streaming multiprocessor. In this context, the maximum number of live threads is at 768 much larger than the batch of 16 threads that are executed at each time, since some of threads are active but parked while waiting for global memory fetches, a strategy that hides the memory latency and improves the overall efficiency. The recommended strategy for GPU computing on GeForce 8800 GTX is to try to maximize the number

of parallel threads active at any given time, which in an ideal situation would be at 12,288 threads. Since register and shared memory resources are fixed, however, the more threads active, the less the resources per thread. In the end, an optimal point is reached, where, while the number of threads per streaming processor is still large, the memory allocated to each thread is enough to allow it to run a batch of commands associated with Algorithm 1. In the current parallel implementation, Algorithm 1 runs with 6,144 active threads at any given time, or in other words, the GPU is working with 6,144 contacts at each time. The number of active threads can be further increased but at the price of an increased number of global memory accesses, where "global" here represents the 768 MB GDDR3 on-chip GPU memory. While global memory fetches are still fast due to low latency and high memory bandwidth (86.4 GB/sec), they are still two orders of magnitude slower than register or shared memory access and should be avoided whenever possible.

In the context of this work, the most relevant consequence of using a parallel execution approach is that this execution model can lead to a random velocity update sequence. In a *sequential* execution mode, the loop over the active contact set in Algorithm 1 is carried out in an orderly fashion. If symmetric updates are desired, then the order is reversed, but it is yet predefined and deterministic. When implementing a *parallel* version of Algorithm 1, two different approaches can be followed. The first falls back on a Gauss-Jacobi approach and, referring back to Eq. (15), corresponds to $\mathcal{D}(i, r) = \emptyset$. Previous numerical experiments suggest that the convergence rate decreases and, for certain models, leads to a large number of iterations r or even lack of convergence. The second approach investigated in this work allows for a random update sequence, to the extent that in general $\mathcal{D}(i, r) \neq \mathcal{D}(i, p)$, for $r \neq p$. In this context, note that for the Gauss-Seidel there is no dependency in Algorithm 1 of $\mathcal{D}(i, r) = \{1, 2, \dots, i-1\}$ on r , and in fact each iteration follows the same update sequence. Enforcing a similar update sequence in a parallel execution scenario would unacceptably compromise performance, as this has to rely on a system of semaphores similar to the `mutex` support in the POSIX standard threading library API, which in fact is not supported by CUDA library support [24]. Note that there is a remote chance that during some iteration r a certain update produced by contact i would fail altogether because of a race condition and lack of `mutex`-type support. This situation is not expected to impact the convergence of the algorithm, particularly for large frictional contact problems with hundreds of thousands of contacts where the probability of two out of 256 contacts to be associated with the same body and try to update its velocity at the same time is very small, but not zero. Note also that the 256 value corresponds on GeForce 8800 GTX to the number of streaming multiprocessors (16), each processing simultaneously a warp of 16 threads (frictional contacts).

²Note that through a new product line called Larrabee, Intel is integrating the GPU and CPU on the same chip and expected to support in early 2010 the same parallel computation model promoted by NVIDIA, only with larger shared memory and L1 cache.

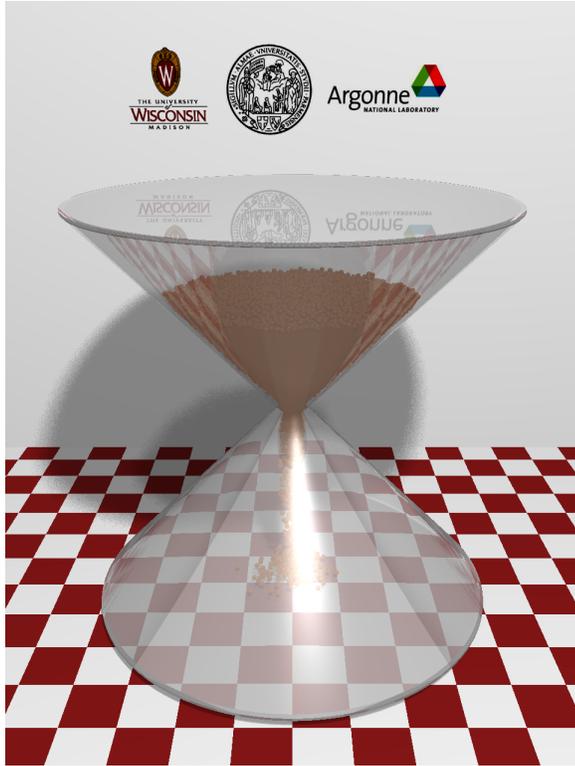


Figure 2. A frame from the simulation of 25,000 rigid bodies with frictional contacts flowing in an hourglass.

NUMERICAL RESULTS

(TO BE ADDED BEFORE FINAL DEADLINE)

Hourglass simulation

The hourglass simulation of 25,000 spheres (Fig.2) captures 3 minutes of physical time. The spheres' diameter is $d = 0.2\text{m}$ and friction coefficient is $\mu = 0.2$. The simulation was carried out with a numerical integration step-size $h = 0.01\text{s}$, which is three orders of magnitude larger than typically required by the DEM method in [35]. The simulation took 70 hours to complete. Timing, since it depends on issues such as cache management that can vastly change with different compilers and optimization levels, is perhaps less relevant than the fact that the simulation completed with low penetration error for a fixed and relatively small number of iterations, 140, for a very high density configuration.

Conclusions

The paper proposes a theoretically rigorous approach to simulating multibody dynamics problems with frictional contact. The algorithm proposed is backed by convergence results derived under the assumption of either small friction or small relative sliding velocity in [2]. The methodology leverages commodity high-performance parallel computing on the GPU. Preliminary results obtained with the proposed parallel algorithm demonstrate that for very large problems the computational bottleneck associated with the sequential algorithm, that is, the solution of

the cone complementarity problem has been eliminated. The computation is now dominated by the collision detection stage, which at this time runs sequentially. Four issues remain to be addressed. First, a rigorous convergence analysis for the case of random velocity updates is needed. Although preliminary results show that this update strategy works in conjunction with large frictional contact models, it is important to understand and possibly address some of the limitations associated with this approach. Second, the methodology should be expanded to include the case of bilateral and unilateral constraints present in a multitude of mechanical system models. It is expected that the latter will positively impact the RATTLE and SHAKE algorithms in molecular dynamics simulation. Third, the cone complementarity problem approach needs to be extended to deformable multibody dynamics problems. Fourth, a parallel collision detection engine would allow for an entirely parallel approach to multibody dynamics with frictional contact that, given the recent advances in commodity high-performance parallel computing hardware, opens up new fields of simulation-based engineering in materials science, pharmaceutical industry, granular flow dynamics. The first three open issues will be addressed in future work; the fourth one is essentially a challenging computer science problem that remains to be addressed by that community.

ACKNOWLEDGMENT

Financial support for the first author is provided in part by National Science Foundation under Grant No. CMMI-0700191. Dan Negrut and Alessandro Tasora thank the NVIDIA corporation for sponsoring their research program in the area of high performance multibody dynamics simulation. Toby Heyn is also acknowledged for generating the hourglass and Mars rover examples model that serve as the benchmark tests for this paper. Mihai Anitescu was supported by Contract No. DE-AC02-06CH11357 of the U.S. Department of Energy.

REFERENCES

- [1] 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.
- [2] Anitescu, M., and Tasora, A., 2008, in press. "An iterative approach for cone complementarity problems for non-smooth dynamics". *Computational Optimization and Applications*.
- [3] Kikuchi, N., and Oden, J., 1998. *Contact Problems in Elasticity: A Study of Variational Inequalities and Finite Element Methods*. SIAM, Philadelphia, PA.
- [4] Wriggers, P., 2006. *Computational Contact Mechanics*, second ed. Springer-Verlag, Berlin Heidelberg.
- [5] Donald, B. R., and Pai, D. K., 1990. "On the motion of compliantly connected rigid bodies in contact: a system for analyzing designs for assembly". In Proceedings of the Conf. on Robotics and Automation, IEEE, pp. 1756–1762.

- [6] Song, P., Kraus, P., Kumar, V., and Dupont, P., 2001. “Analysis of rigid-body dynamic models for simulation of systems with frictional contacts”. *Journal of Applied Mechanics*, **68**(1), pp. 118–128.
- [7] Song, P., Pang, J.-S., and Kumar, V., 2004. “A semi-implicit time-stepping model for frictional compliant contact problems”. *International Journal of Numerical Methods in Engineering*, **60**(13), pp. 267–279.
- [8] Pang, J.-S., Kumar, V., and Song, P., 2005. “Convergence of time-stepping method for initial and boundary-value frictional compliant contact problems”. *SIAM J. Numer. Anal.*, **43**(5), pp. 2200–2226.
- [9] Pang, J.-S., and Stewart, D., 2008. “Differential variational inequalities”. *Mathematical Programming*, **113**(2), pp. 345–424.
- [10] Pang, J.-S., and Stewart, D., 2004. “Solution dependence on initial conditions in differential variational inequalities”. *Set-Valued Analysis*. To Appear.
- [11] Moreau, J. J., 1983. “Standard inelastic shocks and the dynamics of unilateral constraints”. In *Unilateral Problems in Structural Analysis*, G. D. Piero and F. Maciari, eds., CISM Courses and Lectures no. 288, Springer-Verlag, pp. 173–221.
- [12] Lotstedt, P., 1982. “Mechanical systems of rigid bodies subject to unilateral constraints”. *SIAM Journal of Applied Mathematics*, **42**(2), pp. 281–296.
- [13] Marques, M. D. P., 1993. *Differential Inclusions in Non-smooth Mechanical Problems: Shocks and Dry Friction*, Vol. 9 of *Progress in Nonlinear Differential Equations and Their Applications*. Birkhäuser Verlag, Basel, Boston, Berlin.
- [14] Baraff, D., 1993. “Issues in computing contact forces for non-penetrating rigid bodies”. *Algorithmica*, **10**, pp. 292–352.
- [15] Pang, J.-S., and Trinkle, J. C., 1996. “Complementarity formulations and existence of solutions of dynamic multi-rigid-body contact problems with coulomb friction”. *Math. Program.*, **73**(2), pp. 199–226.
- [16] Trinkle, J., Pang, J.-S., Sudarsky, S., and Lo, G., 1997. “On dynamic multi-rigid-body contact problems with coulomb friction”. *Zeitschrift für Angewandte Mathematik und Mechanik*, **77**, pp. 267–279.
- [17] 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.
- [18] 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.
- [19] Anitescu, M., Potra, F. A., and Stewart, D., 1999. “Time-stepping for three-dimensional rigid-body dynamics”. *Computer Methods in Applied Mechanics and Engineering*, **177**, pp. 183–197.
- [20] Stewart, D. E., 2000. “Rigid-body dynamics with friction and impact”. *SIAM Review*, **42**(1), pp. 3–39.
- [21] Cottle, R., and Dantzig, G., 1968. “Complementary pivot theory of mathematical programming”. *Linear Algebra and Its Applications*, **1**, pp. 103–125.
- [22] Baraff, D., 1994. “Fast contact force computation for non-penetrating rigid bodies”. In *Computer Graphics (Proceedings of SIGGRAPH)*, pp. 23–34.
- [23] Anitescu, M., 2006. “Optimization-based simulation of nonsmooth rigid multibody dynamics”. *Math. Program.*, **105**(1), pp. 113–143.
- [24] NVIDIA, 2007. *CUDA Programming Guide*. Available online at http://developer.download.nvidia.com/compute/cuda/1_1/NVIDIA_CUDA_Programming_Guide_1.1.pdf.
- [25] Shabana, A. A., 2005. *Dynamics of Multibody Systems*, third ed. Cambridge University Press.
- [26] Tasora, A., 2007. “High performance complementarity solver for non-smooth dynamics”. In *Proceedings of the ECCOMAS Multibody Dynamics Conference*, C. L. Bottasso, P. Masarati, and L. Trainelli, eds.
- [27] Kim, Y. J., Lin, M. C., and Manocha, D., 2002. “Deep: Dual-space expansion for estimating penetration depth between convex polytopes”. In *Proceedings of the 2002 International Conference on Robotics and Automation*, Vol. 1, Institute for Electrical and Electronics Engineering, pp. 921–926.
- [28] Anitescu, M., Cremer, J. F., and Potra, F. A., 1996. “Formulating 3d contact dynamics problems”. *Mechanics of Structures and Machines*, **24**(4), pp. 405–437.
- [29] 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.
- [30] Moreau, J. J., 1988. “Unilateral contact and dry friction in finite freedom dynamics”. In *Nonsmooth Mechanics and Applications*, J. J. Moreau and P. D. Panagiotopoulos, eds., Springer-Verlag, pp. 1–82.
- [31] Bertsekas, D., 1995. *Nonlinear Programming*, Belmont, MA: Athena Scientific.
- [32] Pang, J. S., and Stewart, D. E., 2008. “Differential variational inequalities”. *Mathematical Programming*, **113**, pp. 1–80.
- [33] 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.
- [34] 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.
- [35] Rycroft, C., Grest, G., Landry, J., and Bazant, M., 2006. “Analysis of granular flow in a pebble-bed nuclear reactor”. *Physical Review E*, **74**, 021306.

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.