

Staggered Projections for Frictional Contact in Multibody Systems

Danny M. Kaufman^{*,‡}

Shinjiro Sueda^{*}

Doug L. James[†]

Dinesh K. Pai^{*}

^{*}University of British Columbia

[†]Cornell University

[‡]Rutgers University




Figure 1: *Staggered Projections* resolves frictional contact between a wide range of rigid and deformable models at rates suitable for (a) interactive haptic simulations, as well as accurate animations that capture important frictional contact phenomena such as: (b) large-deformation frictional contact with reduced SrVK models, (c) large-scale, frictional stacking and jamming of large numbers of objects without constraint drift, (d) and long term stable simulation of difficult, frictionally dependent structures.

Abstract

We present a new discrete velocity-level formulation of frictional contact dynamics that reduces to a pair of coupled projections and introduce a simple fixed-point property of this coupled system. This allows us to construct a novel algorithm for accurate frictional contact resolution based on a simple staggered sequence of projections. The algorithm accelerates performance using warm starts to leverage the potentially high temporal coherence between contact states and provides users with direct control over frictional accuracy. Applying this algorithm to rigid and deformable systems, we obtain robust and accurate simulations of frictional contact behavior not previously possible, at rates suitable for interactive haptic simulations, as well as large-scale animations. By construction, the proposed algorithm guarantees exact, velocity-level contact constraint enforcement and obtains long-term stable and robust integration. Examples are given to illustrate the performance, plausibility and accuracy of the obtained solutions.

CR Categories: I.6.8 [Simulation and Modeling]: Types of Simulation—Animation, I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Physically based modeling, I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Virtual Reality

Keywords: friction, contact, deformation, multibody dynamics

1 Introduction

Frictional contact and deformation are two of the most challenging physical phenomena to simulate in a computer animation. They are also essential for physical realism, and therefore both phenomena have been extensively investigated in computer graphics. While deformation modeling is now well understood and friction modeling is becoming more so, existing methods for simulating realistic frictional contact still often fail for many interesting and important examples.

No previous system can accurately, robustly and efficiently simulate the complex behaviors created by frictional contact between large deformation models at rates suitable for both haptic simulation and computer animation. Likewise, it has not been previously possible to stably and accurately simulate large-scale frictionally dependent structures and composites.

Floating frame deformable models [Terzopoulos and Witkin 1988] and rigid body models explicitly include rigid degrees of freedom (DoFs) and are popular in graphics because of the compact representation and resulting speedup they offer. However, these models present extra difficulties for contact algorithms because of the large potential interaction between friction and contact impulses [Erdmann 1994; Kaufman et al. 2005] (see §4). Because of this, many algorithms customized for resolving frictional contact between finite element models (FEM) and/or other mesh and node based models are not applicable for systems with rigid DoFs.

In practical applications, however, there is a critical need to simulate hybrid models that combine many different models, including rigid and deformable objects, in frictional contact, without resorting to complicated software glue.

Finally, we note that friction is often added on to contact algorithms as an afterthought. Generally, little or no facility is provided to analyze and/or control the accuracy of the friction model applied. It would be useful if instead the cost of frictional accuracy could easily be adjusted according to the needs of a particular simulation.

Contributions In this paper we address these issues with an efficient and robust frictional contact algorithm (§7.1), suitable for the accurate simulation of complex contacting systems, composed of many different models, including reduced deformable bodies, rigid bodies, and jointed mechanical systems.

ACM Reference Format

Kaufman, D., Sueda, S., James, D., Pai, D. 2008. Staggered Projections for Frictional Contact in Multibody Systems. *ACM Trans. Graph.* 27, 5, Article 164 (December 2008), 11 pages.
DOI = 10.1145/1409060.1409117 <http://doi.acm.org/10.1145/1409060.1409117>.

Copyright Notice

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701, fax +1 (212) 869-0481, or permissions@acm.org.
© 2008 ACM 0730-0301/2008/05-ART164 \$5.00 DOI 10.1145/1409060.1409117
<http://doi.acm.org/10.1145/1409060.1409117>

We propose a new discrete velocity-level formulation of frictional contact (§5) that reduces to a pair of coupled projections (§6), and show how this essential coupling makes most non-trivial frictional contact problems NP-hard (§5.3). We introduce a simple fixed-point property of the projections (§6.1) and an associated geometric interpretation that allows us to, despite the NP-hard classification, construct a novel algorithm for accurate contact resolution based on a simple staggered sequence of projections. A related approach [Anitescu and Hart 2004] is discussed in §2.

The resulting algorithm guarantees velocity-level, contact constraint satisfaction (§7). This guarantee is obtained immediately, from the first applied projection, while further projections increase the accuracy of the frictional behavior. Frictional accuracy can then be controlled by a user supplied tolerance while, for time-critical applications, the algorithm supports graceful degradation (§7). Finally, the algorithm utilizes warm starts to leverage the potentially high temporal coherence between time steps in contacting systems (§6.4).

While we do not guarantee convergence, we show that the algorithm obtains robust, stable and accurate solutions for many difficult contacting systems that were not previously possible to solve. In particular, in our validation examples (§8), we focus on rigid and reduced deformable bodies with rigid modes since, as discussed above and also covered in §4, systems with rigid DoFs present unique difficulties for frictional contact algorithms.

2 Related Work

Rigid bodies are one of the most common physical models used in graphics and engineering, while the simulation of deformable systems is also a well studied area. We refer the reader to the surveys of Nealen et al. [2005] and Gibson and Mirtich [1997]. In our examples we leverage linear modal analysis and modal derivative methods [Barbič and James 2005] to obtain small but representative deformation bases.

The accurate modeling of contact for continuum mechanics has long been an active area of research [Hertz 1882]. Contact mechanics has been broadly covered, with varying degrees of formalism [Johnson 1985; Kikuchi and Oden 1988; Brogliato 1999; Wriggers 2002]. Yet resolving multiple contacts with friction for both rigid and deformable bodies is still challenging [Stewart 2000]. In particular, flexible multibody systems have received increased attention in the last decade [Wasfy and Noor 2003], but fast and reliable algorithms for impact and friction remain open problems. Also, existing algorithms tend to support either rigid *or* flexible multibody systems, but rarely work for the hybrid cases desired in practice. Finally, we note that, although contacting configurations often maintain high temporal coherence over time, to the best of our knowledge, current contact resolution methods do not support the reuse of earlier solutions via warm starts.

Penalty-based methods, well known in mechanics [Kikuchi and Oden 1988], first introduced contact resolution problems to the graphics community [Hahn 1988; Bridson et al. 2002; Hauser et al. 2003], but suffer from stiffness and stability issues despite recent advances [Spillmann et al. 2007]. Optimization-based approaches in graphics [Baraff 1989; Baraff 1991; Baraff 1994; Redon et al. 2002; Raghupathi and Faure 2006], which are generally equivalent to Linear Complementarity Programming (LCP) formulations, introduced improved accuracy in exchange for greater computational cost and motivated several extensions using approximations of friction [Milenkovic and Schmidl 2001; Kaufman et al. 2005]. In some cases the frictional contact problem has been simplified by reducing the contact inequality constraints to equalities [Irving et al. 2007; Harmon et al. 2008]; in others it is mollified by treating individual constraints sequentially rather than simultaneously [Moore and Wilhelm 1988; Mirtich and Canny 1995; Cirak and West 2005].

In general, with the exception of LCP-based approaches, which

we will discuss further below, existing methods for frictional contact resolution generally suffer from instabilities, often exhibited by unnatural vibrations at contacts, constraint drift (i.e., elements sinking into one another over time), popping artifacts (created by stabilization methods applied to combat drift), energy gain (although contact resolution should be dissipative) and other issues which have so far made the accurate simulation of many frictional contact phenomena out of reach. In the following discussion, as well as above, we use the term “robustness” to denote that these problems do not manifest, regardless of the accuracy of the obtained solution. In §4 we will discuss and motivate reasons why accurate frictional contact modeling has proven so difficult, while here, in the remaining part of this section, we will focus on a brief discussion of LCP methods and other approaches related to our proposed method.

Restricting the frictional contact problem to rigid body models opened the development of acceleration-level LCP solutions [Lötstedt 1984; Baraff 1991; Baraff 1994; Trinkle et al. 1995]. Baraff [1993] noted that these acceleration formulations were not guaranteed to always have a solution and used these inconsistencies to show that the acceleration LCP is NP-hard. This observation led to more recent velocity-level LCP methods [Stewart and Trinkle 1996; Anitescu and Potra 1997]. While velocity-level LCPs can be shown to always have a solution, Anitescu and Hart [2004] demonstrate examples where the solution set is non-convex and, based on this observation, suggest that velocity-level LCPs may be difficult, in general, to solve. LCP approaches have also been extended to quasi-rigid [Song and Kumar 2003; Pauly et al. 2004], linear modal [Stewart 2001; Hauser et al. 2003] and FEM models [Klarbring 1986; Baraff and Witkin 1992; Jourdan et al. 1998; Duriez et al. 2006; Otaduy et al. 2007].

Recent attention in graphics has been focused on iterative LCP solution methods [Murty 1988] customized for contacting systems [Guendelman et al. 2003; Duriez et al. 2006; Erleben 2007]. Although potentially fast, these approaches require, in practice, an exponentially large number of iterations to converge to *both* accurate contact constraint enforcement *and* accurate friction response [Erleben 2007]. The inherent lack of accuracy, stability and robustness this causes generally requires hand tuning and large amounts of non-physical constraint stabilization [Erleben 2007] to make multibody simulations work. Effectively, iterative LCP methods often reintroduce many of the same errors and artifacts, discussed above, that LCP methods were designed to avoid in the first place.

At the other end of the spectrum, it has been widely assumed that, when fully accurate and robust simulations were desired, direct LCP solvers could be applied, provided enough computation time was available. It has been recently noted, however, that direct LCP solvers do not, in practice, scale. They are not, in fact, currently able to return solutions for non-trivial contacting problems beyond relatively small-scale examples [Anitescu and Hart 2004; Erleben 2007].

Our proposed approach provides robust and accurate solutions for complex rigid and deformable contacting systems where both direct and iterative LCP solvers fail. It is most closely related to the velocity-based, fixed-point method for rigid body frictional contact introduced by Anitescu and Hart [2004]. Their method, however, is constructed by symmetrizing the LCP formulation. Contact and friction constraints are folded together, which effectively moves the contact constraint directions away from the surface normals, as the coefficients of friction increase. This results in a pair of minimizations quite different from the ones we propose. While their approach is convergent for very small coefficients of friction ($\mu \simeq 0.05$), the authors report that it fails for coefficient values where frictional behavior becomes pronounced, e.g., $\mu = 0.2$.

3 Background on Contacting Systems

In this section we provide a general framework for handling frictional contact between many classes of models. We define the assumptions we make about these models, construct generalized constraints, and discuss other associated machinery needed to implement the Staggered Projections frictional contact algorithm given in §7.1.

In the following we consider systems of bodies with a generalized coordinate representation given by \mathbf{q} , a corresponding generalized velocity¹, $\dot{\mathbf{q}}$, and a mass matrix denoted by \mathbf{M} . The world-space position of a point i is then presumed to be given by a function² $\mathbf{x}_i(\mathbf{q})$. Finally, we obtain point velocities, $\dot{\mathbf{x}}_i$, using the Jacobian, $\nabla \mathbf{x}_i(\mathbf{q})$, by left multiplying with the generalized velocity, $\dot{\mathbf{x}}_i = \nabla \mathbf{x}_i \dot{\mathbf{q}}$.


3.1 Multibody Systems

For the accurate simulation of contact *between* multiple bodies we adopt a multibody formulation in which we use generalized coordinates to represent the *entire* system's state as a single vector. Generalized terms are formed by the concatenation of all bodies' corresponding coordinate vectors. Thus the generalized velocity is a large vector constructed by the concatenation of all velocities. Generalized forces and configurations follow in the same manner, while the full system mass matrix is just the large sparse block diagonal matrix formed by the individual mass matrices that correspond to their respective entries in the subcomponents of the generalized velocity. Finally, the generalized Jacobians, $\nabla \mathbf{x}_i$, should perform the same task as the single body version described above. Now, however, they will be large and sparse with just a 3 by $|\dot{\mathbf{q}}_{body_i}|$ nonzero block submatrix that functions as a stencil to extract the velocity of a point i from the correct subcomponent of the generalized velocity. In an actual implementation, simple optimizations can be made to account for the sparsity of these terms. Here, however, it's helpful to use this structure so that in the following sections our results apply equally well to a single body experiencing self-contact or to a large system composed of many bodies that are in contact with one another.

3.2 Frictional Contact Constraints and Impulses

Here we'll briefly describe the velocity-level constraints that are imposed on contacting systems and the impulses they can induce. Duality between these constraints and their associated impulses allow us, in §5, to derive discrete minimization principles for the integration of contacting systems.

Pairwise Point Contacts A discrete system of contacts can be described by a set, \mathcal{C} , of point contacts. Here, each contact $k \in \mathcal{C}$ is between two points, indexed by i and j , and occurs at a location, $\mathbf{x}_k \in \mathbb{R}^3$. Contacts are also presumed to be between two surfaces having a common tangent plane at \mathbf{x}_k with an associated unit length normal given by $\mathbf{n}_k \in \mathbb{R}^3$. Finally, in the following, we define each \mathbf{n}_k so that it points outward from i towards j .



To simplify the following, we define the relative velocity Jacobian for contact k as

$$\Gamma_k \stackrel{\text{def}}{=} (\nabla \mathbf{x}_i - \nabla \mathbf{x}_j). \quad (1)$$

Then, if $\mathbf{y} \in \mathbb{R}^3$ is a unit length vector, $\mathbf{y}^T \Gamma_k \dot{\mathbf{q}}$ gives the magnitude, in the direction of \mathbf{y} , of the relative velocity between i and j . Similarly, if \mathbf{y} is now an impulse applied to point i and an equal but opposite impulse is applied to point j , then $\Gamma_k^T \mathbf{y}$ is the resulting generalized impulse.

Contact Constraints If we ignore friction for a moment, the chief concern for a contacting system is the constraint that all points

in contact should not interpenetrate [Baraff 1989]. To prevent interpenetration at a contact k , the relative velocity between the two contacting points along \mathbf{n}_k must be non-negative. This is equivalent to enforcing the contact constraint $\mathbf{n}_k^T \Gamma_k \dot{\mathbf{q}} \geq 0$.

Defining generalized normals as

$$\mathbf{n}_k \stackrel{\text{def}}{=} \Gamma_k^T \mathbf{n}_k, \quad (2)$$

we denote the subspace of generalized normal directions by

$$\mathbf{N} \stackrel{\text{def}}{=} (\mathbf{n}_1 \dots \mathbf{n}_{|\mathcal{C}|}). \quad (3)$$

The nonpenetration inequality constraint, for all points of contact simultaneously, is then given by the global contact constraint

$$\mathbf{N}^T \dot{\mathbf{q}} \geq 0. \quad (4)$$


Contact Impulses If a contacting system has a negative relative velocity along contact k 's normal, we can prevent interpenetration at k by applying equal and opposite impulses at \mathbf{x}_k , along the contact normal. Thus generalized normals have a useful dual interpretation, serving both as potential generalized contact impulse directions *and* as generalized coordinate representations of the velocity constraints imposed by contacts k between points i and j .

In the following, the magnitude of each contact impulse, applied along a corresponding normal, \mathbf{n}_k , is given by α_k , while the set of all contact impulse magnitudes, corresponding to the contact normals in \mathbf{N} , is given by the vector

$$\alpha = (\alpha_1 \dots \alpha_{|\mathcal{C}|})^T. \quad (5)$$

Friction Impulses A friction impulse, applied at a contact point, lies in the tangent plane orthogonal to the contact normal.

At each contact k we uniformly sample a symmetric set of unit length vectors from the tangent plane. The matrix composed column-wise of these samples is given by \mathbf{T}_k so that a friction impulse, $\mathbf{f}_k \in \mathbb{R}^3$, applied at a contact k , will lie in the span of \mathbf{T}_k .



Coulomb Constraints An isotropic Coulomb constraint then restricts the magnitude of this friction impulse by the inequality

$$\|\mathbf{f}_k\| \leq \mu_k \alpha_k, \quad (6)$$

where μ_k is the coefficient of friction and α_k is the normal (contact) impulse magnitude at k . This restricts the friction impulse to lie within a disk in the tangent plane with a radius of $\mu_k \alpha_k$.

To linearize the Coulomb constraint, we can set $\mathbf{f}_k = \mathbf{T}_k \beta_k$, where β_k is the vector of impulse magnitudes along each direction spanned by \mathbf{T}_k , and then approximate the friction disk with an inscribed polygon [Stewart 2000] using the constraints

$$\mathbf{e}^T \beta_k \leq \mu_k \alpha_k, \quad \beta_k \geq 0, \quad (7)$$

where $\mathbf{e} = (1 \dots 1)^T$.

For multiple points of contact we can concatenate all of the respective μ_k values into a single vector μ and correspondingly define a vector that concatenates the friction impulse magnitudes for all of the contacts

$$\beta = (\beta_1^T \dots \beta_{|\mathcal{C}|}^T)^T. \quad (8)$$

Then the linearized Coulomb constraint, for all contacts, is

$$\mathbf{E}^T \beta \leq \text{diag}(\mu) \alpha, \quad \beta \geq 0. \quad (9)$$

Here $\text{diag}(\mu)$ constructs the diagonal matrix composed of the μ_k values along the diagonal and zeros elsewhere. The matrix \mathbf{E} corresponds to the vector β such that each column k of \mathbf{E} has ones in rows that correspond to entries in the subvector $\beta_k \in \beta$ and zeros in all other rows.

¹Note that \mathbf{q} and $\dot{\mathbf{q}}$ do not necessarily have the same dimensions here, and different rotation representations are commonly used for pedagogical and implementation convenience [Shabana 2005].

²The form of this function depends on the type of physical model.

4 Frictional Contact Challenges

Much of the difficulty in frictional contact algorithms comes from the feedback between friction impulses and contact (nonpenetration) impulses. Many contact resolution algorithms in graphics propose a two pass strategy in which contact constraints are processed in a first pass, followed by a second pass in which friction impulses are applied. For deformable models, with low stiffness, this can often work reasonably well. As stiffness increases, however, friction impulses can begin to apply significant global changes in the system's velocity that, in turn, can create new contact constraint violations. In a similar way, contact correction impulses can effect global changes in the system's sliding velocities and the Coulomb constraints, which must affect friction impulse calculations.




Figure 2: **Rigid Card House:** Constraints and friction impulse coupling.

Rigid body systems, which are essentially deformable bodies with infinite stiffness, pose an extreme version of this problem. Consider the card house example where the structure is composed of thin, rigid planks. A tangential friction impulse applied at the base of a card, to oppose outward sliding, will also induce a torque on the card. The torque, in turn, will cause a negative velocity along the normal at the same point where the friction impulse was originally applied. Similarly, a normal impulse, applied at the same contact point, would generate a faster outward sliding velocity (see Figure 2). These effects are related to the Painlevé Paradox [Painlevé 1895], and the observation by Erdmann [1994] that generalized normal and friction impulses, for rigid bodies, are not, in general, orthogonal. They can, in fact, oppose or even reinforce one another. We also note that, while this effect is most marked for rigid bodies, we can extend the same observations to deformable models.

There has recently been a focus on stable structured stacking in gaming and graphics [Erleben 2007]. While these examples require accurate contact constraint enforcement they require very little accuracy from friction. Structures like the rigid card house in Figure 1, masonry arches *without* adhesives (see supplemental video), or woven elastic-frictional composites (see Figure 5), on the other hand, depend on very accurate coupling between friction and contact impulse calculations. Non-orthogonal friction and contact impulses must be accurately balanced, at each active contact point, to prevent incremental errors from crashing or breaking apart the structures. Because of this, inaccuracies imposed by two-pass, penalty, iterative LCP, and other methods have previously made the stable simulation of such frictionally dependent structures impractical.

Accurate coupling between friction, contact impulses and deformation is also extremely important when taking large step sizes. Contact algorithms generally impose large impulses to contacting bodies. When inaccurate, they can cause system instabilities.

5 A Coupled Potential Formulation

In this section we show that discrete frictional contact dynamics can be described by a pair of discrete, coupled variational principles. These coupled principles are obtained by a numerical discretization of contacting systems that employs a novel generalization of the

frictional maximal dissipation principle.

This results in a predictor-corrector method that first generates a predictor velocity, \dot{q}^p , by integrating the conservative portion of the contacting system forward using an arbitrary numerical integration method. A corrector step then handles the dissipative portion of the system, implicitly, by minimizing two separate, but dependent, functions corresponding to the optimality conditions of the two associated variational principles.

5.1 Discrete Euler-Lagrange Equations for Contact

We first define a velocity-level, predictor-corrector form of the discrete Euler-Lagrange equations (DEL) of motion for integrating the contacting system from time t to $t + 1$, with a step-size of h ,

$$M(\dot{q}^{t+1} - \dot{q}^t) = hg(q^p, \dot{q}^p) + hf_{ext}^p + r^{t+1} \quad (10)$$

Here f_{ext} is a specified external force, r^{t+1} is the frictional contact response impulse, and the generalized force is

$$g(q, \dot{q}) = -\nabla V(q) + \left(-\dot{M}\dot{q} + \nabla_q K(q, \dot{q}) \right) = -\nabla V(q) + g_{qv}, \quad (11)$$

where V and K are the potential³ and kinetic energies of the system respectively, and g_{qv} is the so-called *quadratic velocity vector* that provides Coriolis and centrifugal forces in multibody systems [Shabana 2005].

Predictor Depending on specific animation needs and the physical models in use, we select an appropriate integration method and use it to integrate *only* the terms indexed by p in (10). This amounts to solving the system

$$M(\dot{q}^p - \dot{q}^t) = hg(q^p, \dot{q}^p) + hf_{ext}^p \quad (12)$$

and generates the predictor velocity, \dot{q}^p .

Corrector Then, by subtracting (12) from (10), the remaining unknown, r^{t+1} , is determined by

$$M(\dot{q}^{t+1} - \dot{q}^p) = r^{t+1}. \quad (13)$$

5.2 Coupled Potentials

While r^{t+1} is a purely dissipative term, similar to a damping force, it can not, unlike many damping forces, be obtained from the minimization of a single potential function. Instead, r^{t+1} must be partitioned into two components, a friction impulse, f^{t+1} , and a contact impulse, c^{t+1} , each of which is given by the minimization of a separate but *coupled* function. In this context, we can treat each of these functions as a discrete dissipative potential. Because of the coupling between them, however, the minimizations of each of these potentials can not, in general, be obtained independently. As we discussed in §4 and will show in §5.3, this is the root cause of difficulty in the simulation of contacting systems with friction.

In the following discussion the two discrete contact potentials could be introduced explicitly using a differential inclusion formulation or variational inequalities. In this setting, however, we will introduce the contact potentials implicitly, using an equivalent formulation that uses complementarity and minimization forms that will be more familiar to the graphics community.

³This could simply describe a gravity potential, while for a hyperelastic deformable body this should also include internal strain energy.




Figure 3: **Staggered Projection Sequence:** A simplified example of a few steps in a possible staggered sequence is shown above. (a) We start the staggered sequence with the negative momentum predicted by (12) and depicted above by the point x . We also start with the set of all possible contact normal directions, in this case composed of only a single vector, and depicted above as the green arrow C . Finally, we also initialize the unscaled friction set, F , depicted above as a red line segment. Note that, in this example, we show the likely case where the friction set and the contact normal directions are not orthogonal. (b) Starting with $f^0 = 0$, for the initial (warm start) estimate of the friction impulse, we project $(x - f^0)$ onto C to obtain the first contact impulse estimate, c^1 . (c) We then scale F by $\alpha^1 = |c^1|$ and project $(x - c^1)$ onto it, to get the first friction impulse estimate f^1 . This completes the first subsequence of the staggered projection sequence. Steps (d) and (e) then repeat the process for the second subsequence, which will converge, in this simple example, after two more subsequences are completed, to the optimal solution shown in Figure 4.

Signorini-Fichera Condition In order to formulate the contact potential, we first apply an impulse-level discretization of the *Signorini-Fichera Condition* (SFC) [Signorini 1933; Fichera 1964]. This approach directly enforces nonlinear boundary conditions between final velocities at contact points and the corrective contact impulses applied at those points,

$$0 \leq N^T \dot{q}^{t+1} \perp \alpha^{t+1} \geq 0. \quad (14)$$

The left-hand side of this complementarity condition is simply the contact constraint inequality⁴ discussed above in §3.2, while the right-hand side requires all contact impulses to be non-negative along the normals; i.e., no adhesive impulses are allowed. Note that, as discussed above, the full contact impulse is given by $c^{t+1} = N\alpha^{t+1}$. Finally, the complementarity condition itself ensures that contact impulses are only applied at points where the body *is not* moving away from the contacting surface.

Maximal Dissipation Principle Moreau [1973] introduced the application of a *Maximal Dissipation Principle* (MDP) for single-point frictional contact conditions. Letting f_k denote a frictional force applied at a sliding contact point, x_k , MDP requires that friction maximize the rate of negative work at the contact, given by $-f_k^T \dot{x}_k$. Combined with a suitable friction constraint, as described in §3.2, this simply provides a variational interpretation of the more familiar Coulomb friction.

To extend the MDP to generalized coordinates we enforce it simultaneously at *all* contact points to obtain

$$\max_{f_k} \sum_{k \in \mathcal{C}} (-f_k^T \dot{x}_k) = \min_{f_k} \left[\sum_{k \in \mathcal{C}} f_k^T \Gamma_k \right] \dot{q} = \min_f f^T \dot{q}. \quad (15)$$

This, in turn, requires that the total generalized friction impulse is given by $f = \sum_{k \in \mathcal{C}} \Gamma_k^T f_k$.

Then, using the above identity and remembering our earlier definition of the friction impulse magnitude vector, β , from §3.2, the generalized friction impulse at each contact is $f_k = \Gamma_k^T T_k \beta_k$. Note that, by construction, friction now appropriately applies an equal and opposite impulse at each point of contact.

For convenience we can then define a generalized basis for friction impulses at contact k ,

$$D_k \stackrel{\text{def}}{=} \Gamma_k^T T_k, \quad (16)$$

⁴For rigid body systems we can include a Newtonian restitution model by modifying the contact constraints to $n_k^T (\dot{q}^{t+1} + e \dot{q}^t) \geq 0$. Here the coefficient of restitution is given by $e \in [0, 1]$.

and the corresponding subspace of all generalized friction impulses,

$$D = (D_1 \dots D_{|\mathcal{C}|}). \quad (17)$$

Now, given this basis, at each time step, the final friction impulse is $f^{t+1} = D\beta^{t+1}$.

Finally, by applying an implicit discretization to (15), and using the linearized Coulomb constraint from §3.2, we obtain a discrete generalized maximal dissipation principle,

$$\beta^{t+1} = \underset{\beta}{\operatorname{argmin}} \left(\beta^T D^T \dot{q}^{t+1} : E^T \beta \leq \operatorname{diag}(\mu) \alpha^{t+1}, \beta \geq 0 \right). \quad (18)$$

Discrete Hamilton's Principle The full frictional contact response impulse is now given by⁵

$$r^{t+1} = c^{t+1} + f^{t+1} = N\alpha^{t+1} + D\beta^{t+1}. \quad (19)$$

Equation (19) together with (13) and (14) now form the KKT conditions for a second minimization [Moreau 1966; Boyd and Vandenberghe 2004]:

$$\dot{q}^{t+1} = \underset{v}{\operatorname{argmin}} \left(\frac{1}{2} v^T M v - v^T (M \dot{q}^t + D\beta^{t+1}) : N^T v \geq 0 \right). \quad (20)$$

This minimization implicitly solves for the contact impulse, c^{t+1} , and, when combined with (10), can be interpreted as a discrete analog of Hamilton's Stationarity Principle.

5.3 Discussion: NP-Hardness

While each of the above minimizations, (18) and (20), is separately convex, and thus individually solvable utilizing available robust algorithms and solvers [Boyd and Vandenberghe 2004], the solution to the frictional contact problem is given by obtaining a solution that satisfies both minimizations simultaneously. At each time step, to solve the velocity minimization (20), we need the friction impulse given by the solution to the discrete MDP minimization (18). To solve (18), however, we conversely require the optimality conditions given by (20). Thus, as discussed in §5.2, to resolve the frictional contact problem we need to solve a system composed of two minimizations that are coupled, and so can not be solved independently. This coupled minimization problem is much harder.

⁵Note that because r^{t+1} , by construction, only opposes relative sliding and relative normal velocities, it will conserve the linear and angular momentum of the contacting system it is applied to.

In particular, the problem of finding the solution of the coupled problem is equivalent to finding the *global* minimum of a non-convex QP that can exhibit multiple distinct global minima for its solutions. In general, the solution of non-convex QPs is an NP-hard problem in global optimization [Murty and Kabadi 1987].

A natural approach to try is to instead find a *local* minimum of the non-convex QP, using Sequential QP or other efficient local minimization methods. Unfortunately, local minima will not, in general, satisfy SFC and/or MDP. We have experimentally verified this with an implementation of Sequential QP. In the following sections, we'll introduce an algorithm for obtaining accurate approximations of the *global* minimum to this problem.

6 Staggered Projections

The solution to the corrector step is then given by an optimal pair, \dot{q}^{t+1} and β^{t+1} , that satisfy *both* (18) and (20); finding it is, as discussed above, an NP-hard problem. However, a simple geometric interpretation, and transformation, of the minimizations introduced in §5, leads to a projective fixed-point property that allows us to construct an algorithm for accurately resolving the coupled problem.

The resulting algorithm leverages the particular structure and properties of the discrete formulation to quickly converge to numerically close approximations of true coupled solutions. Additionally, we are able to do this in a way that strongly facilitates warm starts to reuse computations from earlier steps.

In the following we will reinterpret both of the above minimizations as projections in impulse space, using the kinetic metric. We define these projections, so that, given an arbitrary convex set A , the projection of some z onto A is given by

$$P_A(z) \stackrel{\text{def}}{=} \operatorname{argmin}_{y \in A} (y - z)^T M^{-1} (y - z). \quad (21)$$

6.1 A Fixed-Point Projective Property

If the generalized velocity, \dot{q} , has n DoFs, the set of possible normal impulse directions, given by N , forms a polyhedral cone in \mathbb{R}^n , while the Coulomb constraint, (7), similarly defines a more complex scaled, convex subset of \mathbb{R}^n .

More rigorously we can define the polyhedral cone of all possible contact impulses as

$$C \stackrel{\text{def}}{=} \{N\alpha : \alpha \geq 0\}, \quad (22)$$

the polyhedral, scaled, convex set of possible friction impulses as

$$F(\alpha) \stackrel{\text{def}}{=} \{D\beta : E^T \beta \leq \operatorname{diag}(\mu)\alpha, \beta \geq 0\}, \quad (23)$$

and then, to simplify the discussion, we also let $x = -M\dot{q}^P$ denote the negative momentum of the predictor velocity.

Using these definitions, the coupled solution of the minimizations, (18) and (20), is equivalent to the two coupled projections

$$f^{t+1} = P_{F(\alpha^{t+1})}(x - c^{t+1}), \quad (24)$$

$$c^{t+1} = P_C(x - f^{t+1}). \quad (25)$$

See Figure 4 (left) for a simplified geometric example of the coupled relationship. Then, by substituting (24) into (25), we obtain the fixed-point property that characterizes all solutions of the coupled contact problem:

$$N\alpha^{t+1} = P_C(x - P_{F(\alpha^{t+1})}(x - N\alpha^{t+1})). \quad (26)$$

By construction, any fixed-point solution of (26), together with the resulting friction response given by (24), then satisfies the optimality conditions of both (18) and (20), and thus, correspondingly,

defines a global minimum for the non-convex optimization problem discussed in §5.3.

Because the right-hand side of the above mapping is composed of a concatenation of projections, which are generally non-expansive, and because, for many cases, this map generates a contraction near solutions, we propose applying a staggered sequence of the two projections (see Figure 3), to obtain the solution to the predictor step. Each subsequence is then defined so that we obtain the $i + 1$ solution estimate by applying

$$\begin{aligned} f^{i+1} &\leftarrow P_{F(\alpha^i)}(x - N\alpha^i), \\ N\alpha^{i+1} &\leftarrow P_C(x - f^{i+1}). \end{aligned} \quad (27)$$

Below we further motivate the Staggered Projection method and then explain additional properties and structures that this approach allows us to take advantage of. This will then lead to our proposed algorithm (presented in §7.1).

6.2 Closest Point Geometry

We describe the proposed approach as a staggered projection method to differentiate it from alternating projection methods [Bauschke 2000] which superficially resemble our approach, but have very different properties. Effectively, the staggered projection sequence maintains the unchanging point x which is repeatedly projected, in a staggered order, onto a pair of converging sets, $\{F(\alpha) + N\alpha\}$ and $\{C + f\}$. Geometrically, convergence is satisfied when the closest points to x , in both sets, are the same (see Figure 4, right). This type of minimum property can not be satisfied by alternating projection methods which can only apply to finding *arbitrary* feasible points in set intersections.

The staggered sequence of projections is thus essentially a self-correcting process, alternating between two changing sets. The output of each successive projection onto one set corrects for any change it might impose on the current predicted solution. From a physical standpoint these modifications account for the effect of changes in a friction (or contact) impulse, just obtained from the last projection, may have on the next projection. For instance, after subsequence i , a projection onto C , offset by the current friction impulse estimate, f^i , will generate a new contact impulse that takes into account any possible changes in constraint status that could be caused by applying f^i .

6.3 Subproblem Dimensions and Complexity

Both of the projections in (27) can be implemented as convex QP solves. Roughly speaking the cost of each successive QP solve will be polynomial in its dimension, and linear in the number of constraints. In practice, in our implementations, for QPs with dimension n we've found this to be approximately $O(n^3)$.

For systems and complex contacting scenarios where we expect the system's DoFs to be much smaller than the number of contacts, i.e., $|\dot{q}| \ll |\mathcal{C}|$, both of the minimizations can be formulated as QP solves with a dimension proportional to the system's DoFs rather than the number of contacts. Alternately for systems where we expect $|\dot{q}| \gg |\mathcal{C}|$ we can instead apply the dual formulation of the minimizations using QP solves with a dimension proportional to the number of contacts.

6.4 Warm Start

The structure of the staggered projection formulation also encourages the reuse of computations. Because of the fixed-point property in (26), initializing a staggered projection algorithm with *either* an optimal contact impulse *or* an optimal friction impulse brings the system to convergence after a single projection. Similarly, we find that initializing the algorithm with a *close to optimal* impulse dramatically speeds up convergence. Because most steps of contacting simulations maintain a reasonable degree of coherence, we can




Figure 4: **Geometric Optimality:** The optimal solution at time $t + 1$ is given by the pair c^{t+1} and f^{t+1} . A simplified example of an optimal geometry characterizing this solution is depicted above. **Left:** A pair, c and f , are optimal if and only if the projection of $(x - f)$ onto the contact cone, here depicted by the larger green arrow, C , returns c , and if the projection of $(x - c)$ onto the scaled friction set, here depicted by the red line segment, $F(\alpha)$, similarly returns f . If this is true, then the sum of c and f gives r^{t+1} , an optimal frictional-contact response. **Right:** Equivalently, c and f are optimal when the projection of x onto both of the sets, $\{F(\alpha) + c\}$ (with $|c| = \alpha$) and $\{C + f\}$, return the same point. If this is true, then this projection point is r^{t+1} , an optimal frictional-contact response.

leverage this warm start property to accelerate contact resolution solves. In particular, many of the most difficult frictional-contact configurations, such as those maintaining complex but stable sticking behavior, exhibit very high temporal coherence that allows for the reuse of solutions over many time steps. Overall we typically observe up to two orders of magnitude speed-up when applying warm starts (see §8).

7 Staggered Projections in Practice

Sequence Ordering Because storing and reusing prior contact correction impulses requires more overhead, it is simpler, and more efficient, to warm start using friction impulses. We thus begin each staggered projection sequence with a contact projection that applies the friction impulse obtained from the prior step, in order to encourage rapid convergence.

Graceful Degradation Each staggered projection sequence finishes its sequence with a final contact projection. Ordering the sequences in this way guarantees that the final solution obtained will satisfy both the DEL and the Signorini-Fichera boundary conditions and so will generate robust behavior, regardless of whether convergence has been obtained. For time-critical applications, where hard time limits may require setting a fixed upper bound on the maximum number of iterations, this allows for graceful degradation in case of early exits from the projective sequence.

Convergence Criteria Given the staggered sequence ordering, we determine numerical convergence at the end of subsequence i , to a fixed-point solution of (26), using the relative, kinetic metric error,

$$\text{rel_err} = \frac{(f^i - f^{i-1})^T M^{-1} (f^i - f^{i-1})}{f^{i-1}{}^T M^{-1} f^{i-1}}, \quad (28)$$

and a user supplied friction error threshold.

Non-Monotonicity While generally monotone, the staggered projection sequences often exhibit local non-monotone behavior. Thus, for time-critical applications we can cache and update current best solutions, based on residual values, for applications where

early exiting may be necessary. Here a useful residual, at the end of subsequence i , is given by the sum,

$$\text{residual} = \sum_{k \in \mathcal{C}} |\alpha_k^i n_k^T (\dot{q}^p + N\alpha^i + D\beta^i)|. \quad (29)$$

7.1 Algorithm

Once an acceptable friction error tolerance, ϵ , is selected, the Staggered Projection algorithm for stepping a contacting system from time t to $t + 1$ is given in Algorithm 1. As in the preceding sections, we continue to use the superscripts t and $t + 1$ to indicate discrete time increments and reserve the superscripts i and $i + 1$ to denote subsequence numbers within a single time step.

Algorithm 1 Staggered Projections.

```

1:  $\dot{q}^p \leftarrow \text{solve: } M\dot{q}^p = g(q^p, \dot{q}^p) + f_{ext}^p$ 
2:  $f^0 \leftarrow f^t$  // warm start
3:  $i \leftarrow 1, \text{min\_res} \leftarrow \infty, \text{rel\_err} \leftarrow \infty$ 
4: while  $\text{rel\_err} > \epsilon$  and  $i < \text{max\_iters}$  do
5:    $N\alpha^i \leftarrow P_C(-M\dot{q}^p - f^{i-1})$ 
6:    $f^i \leftarrow P_{F(\alpha^i)}(-M\dot{q}^p - N\alpha^i)$ 
7:   compute  $\text{rel\_err}$  // Equation (28)
8:   compute  $\text{residual}$  // Equation (29)
9:   if  $\text{residual} < \text{min\_res}$  then
10:      $\text{min\_res} \leftarrow \text{residual}$ 
11:      $f^* \leftarrow f^i$  // cache best solution
12:   end if
13:    $i \leftarrow i + 1$ 
14: end while
15:  $f^{t+1} \leftarrow f^*$ 
16:  $c^{t+1} \leftarrow N\alpha^{t+1} \leftarrow P_C(-M\dot{q}^p - f^{t+1})$ 
17:  $\dot{q}^{t+1} \leftarrow \dot{q}^p + M^{-1}(c^{t+1} + f^{t+1})$ 
    
```

8 Results

We instrumented several examples to examine both the behavior of our algorithm and the performance of our implementation. Please see the accompanying video and Table 1. All timings were measured on a PC with a 3 GHz Intel Pentium D processor, 2G RAM. We used Java (JDK 1.6) for our implementation. Note that while faster timings could be obtained via parallelization (§8.2) all simulations were run single threaded to get accurate timings.

In the following examples we simulate contacting systems composed of large-deformation reduced St.Venant-Kirchhoff (StVK) models [Barbič and James 2005] and/or linear modal models, both augmented with a floating frame [Shabana 2005]. We also include rigid bodies. In our implementation, narrow-phase collision detection is performed using a sphere-based Bounded Deformation Tree (BD-Tree) [James and Pai 2004] for reduced deformable bodies, a customized bounding-sphere hierarchy for rigid bodies, and a uniform subdivision for broad-phase culling. No stabilization was employed. Contact samples are generated by sampling the base geometries of deformable bodies, and the zero level set of pre-computed distance fields of rigid bodies. To facilitate comparison with existing packages, we wrote a Java wrapper for ODE's [Smith 2006] box/box implementation. We use this to compute narrow-phase collision detection and contact sampling between box-based rigid body geometries.

A variety of numerical integrators can be used to both advance the unconstrained state and to generate predictor velocities during contact. In our implementation, we use a hybrid implicit-explicit (IMEX) integration scheme that is explicit in rigid motion and the quadratic velocity forces g_{qv} , but implicit in the internal deformation forces (and damping). For the latter, we use the semi-implicit Newmark subspace integrator of Barbič and James [2005].

All QP solves were implemented using QL [Schittkowski 2005], a robust implementation of the dual active set algorithm of Goldfarb and Idrani [1983]. QL is a dense solver so speedups could potentially be obtained with sparse solvers.

8.1 Validation Tests

We tested our algorithm on a wide range of difficult contact scenarios to verify that we obtain plausible behavior.




Figure 5: **Hanging woven elastic-frictional composites (left):** We increase stiffness (bottom to top) and the coefficient of friction (right to left). The woven composites come apart for all cases where the combined friction and stiffness are insufficient to resist gravity. **Stick-slip instability (right):** Plastic chair legs chatter while sliding on an inclined surface due to fictionally induced oscillations. This behavior is obtained from accurate solutions of the coupled friction, contact and deformation modes.

Elastic-Frictional Composites We wove flattened strands of a stiff rubbery material into frictional composites. These structures demonstrate the importance of tight coupling between friction and deformation. In our first experiment we hang a grid of composites on pegs. We vary μ , the coefficient of friction, right to left, with the increasing values 0.1, 0.3 and 0.5, and increase stiffness slightly for the upper row. On the top row the two leftmost composites stay cohered, due to a combined high friction and stiffness, while the right hand composite comes apart due to low friction. On the bottom row the two rightmost composites come apart quickly due to a combined low friction and stiffness, while the left composite slowly creeps apart due to low stiffness but high friction. See Figure 5, left. In our second example we throw stiff composites, with $\mu = 0.5$, at a peg board. The composite’s cohesive behavior varies with the type of the impact, in some cases deforming as a whole, while for other impacts coming apart quickly.

Stick-Slip Instability *Stick-slip oscillation* is an important frictionally induced instability in deformation dynamics. High friction during sliding creates a buildup of elastic energy in contacting systems. This energy is partially stored in the global deformation of the system, but also builds up at contacting interfaces. Elastic energy is then released suddenly when the magnitude of the friction force, opposing sliding contact, is exceeded by the tangential stiffness at the contact. This interaction between the friction forces and the sliding velocities is generally periodic, and so can induce self-excited oscillations that involve the buildup and dissipation of significant amounts of energy. In this example we show that accurate solutions of the coupled friction, contact and deformation modes in our algorithm capture the stick slip induced instability behavior in the chattering legs of plastic chairs sliding on an inclining surface (see Figure 5, right).

Jamming Another behavior that arises due to deformation and friction is *jamming*. Stable deformable jamming is another difficult frictional phenomenon to simulate. We dropped deformable and rigid objects into tight tubes with varying coefficients of friction, obtaining stable and plausible changes in jamming behavior as μ decreases. We also demonstrate robust and stable deformable frictional jamming at real-time rates in our haptics example.

scene	solve (s)	contacts	iters.	DoFs	time-step (s)	models
composites	0.56	921	6.4	192	3e-3	modal
chairs	0.72	9,025	2.9	1,766	1e-2	modal
catenary arch	1.2	2,042	5.6	162	1e-2	rigid
card house	4.7	528	3.2	348	1e-3	rigid
dinosaurs	0.61	4,690	2.7	315	1e-2	StVK
haptic	0.0032	117	2.4	72	1e-2	modal & rigid

Table 1: **Performance Evaluation:** This table summarizes the average solve time, number of contacts and iterations, per simulation step, for the validation examples. The tangent sample size was 8 for all friction solves.

Friction-Dependent Masonry For suitably high coefficients of friction, complex masonry structures, such as arches, can be constructed *without* adhesives. Difficulties in accurately simulating such structures are discussed in §4. We stably simulate a catenary arch using rigid blocks with $\mu = 0.6$. We first drop blocks onto the arch’s keystone (the most stable point in the arch) and note that the structure deforms globally but does not fall apart; instead it finds a new stable equilibrium. We then throw blocks at one of the arch’s legs (a less stable point) and the arch falls apart.

Rigid Card House In this example, a stable house of cards is made out of rigid cards with high friction. We use $e = 0.1$ and $\mu = 0.8$ for all cards. The card house stands stably long term and then, under successive impacts from small blocks, it repeatedly has sections fall down and then regains equilibrium (see Figure 7). Difficulties in accurately simulating these behaviors are also discussed in §4.

Large Deformation Frictional Contact In these examples we simulate large deformation frictional contact using reduced StVK models. In the first example dinosaurs are thrown together undergoing fast impacts with $\mu = 0.5$ (see Figure 1 (b)). In a second example we drop two groups of bunnies, with differing coefficients of friction, onto a ramp and then into collision with semi-circular arches composed of rigid blocks. The red bunnies use $\mu = 0.4$, blue bunnies, $\mu = 0.1$ and the rigid arch blocks, $\mu = 0.5$ (see Figure 6). The bunnies exhibit differing sticking, and stick-slip sliding behaviors depending on the incidence of impact and μ . In both examples the models end the simulation in stable piles.

8.2 Algorithm Behavior and Comparison

Stability and Long Term Integration To test long term integration and stability properties, all of the above validation examples were stepped for 10 minutes of simulation time. All simulations ran without blow-up, or constraint drift other than possibly initial (small) penetrations sometimes caused by the collision of sharp asperities (bunny ears, cow horns, etc.) during the fast phases of impacts.

Comparison with LCP Solvers To compare our algorithm with direct LCP solvers we also implemented our base collision and contact sampling code with the Stewart-Trinkle [Stewart and Trinkle 1996] velocity-level LCP using the PATH solver [Ferris and Munson 1998]. PATH supports sparse LCP solves and is the most broadly referenced direct LCP solver in the frictional contact literature. We also tested examples with the Gauss-Seidel iterative LCP solvers in the ODE [Smith 2006] and the Open-Tissue [Erleben 2007] packages. In these examples, all methods employ semi-implicit Euler for non-contacting integration. For collision detection and contact sampling direct LCP, Staggered Projections and ODE apply ODE’s box/box implementation; Open-Tissue applies a box/box variant.

Stacking: Both the ODE and Open-Tissue iterative solvers generate sidewise sliding errors for simple rigid stacking block examples. While these errors build more slowly in Open-Tissue, for both systems this results in stacks falling over after a few seconds. We

find that this occurs even when large iteration counts (100) and very small step sizes (10^{-4}) are applied. If we make the influence of dynamics negligible in Open-Tissue, however, by setting Open-Tissue's fraction parameter to ~ 0.002 , and thus make Open-Tissue's stabilization method, shock-propagation, almost entirely dominant, we can obtain stable stacks of blocks. The direct LCP slowly generates stable stacking for small numbers of blocks (≤ 15) but can not scale, and begins failing for larger problems. This is consistent with the results reported in Anitescu and Hart [2004]. Finally, the Staggered Projections implementation generates long term stable stacking⁶ for both rigid and deformable bodies.

Rigid Card House: We also attempted to duplicate the card house example from the validation tests using both direct and iterative LCP methods. Direct LCP generates, very slowly, a stable solution for a small, two-level card house, but again, does not scale and fails on card house examples of larger sizes. ODE's iterative solver fails for card houses at all sizes. Again this is caused by a sidewise sliding error, no matter how high we set friction. Open-Tissue's iterative solver fails similarly. If we, in addition, also apply Open-Tissue's shock-propagation based stabilization, some of the sliding error is reduced. However, vibratory artifacts induced by the stabilization method still cause the Open-Tissue solver to fail on card houses of all sizes (see the accompanying video).

Convergence and Warm Starts While we do not guarantee convergence, all validation examples obtained convergent solutions. All of the above validation examples were run to convergence with a tolerance of $\epsilon = 10^{-4}$. A tolerance of $\epsilon = 10^{-1}$, however, was generally sufficient to generate convincing frictional behavior, while convergence was also tested for tolerances down to $\epsilon = 10^{-8}$. In general, setting ϵ in this range allows the user detailed control over the accuracy of the frictional response obtained by the algorithm.

We note an appreciable acceleration in convergence when warm starting is applied in these examples. In Table 2 (left) we summarize the number of iterations to convergence for the validation examples *with* and then *without* warm starts. In many cases a speed-up of up to two orders of magnitude was obtained. Warm starting also generated faster simulations in examples where a hard maximum iteration limit was enforced. In these simulations convergence may not be obtained for individual steps. When temporal coherence is sufficiently high, however, warm starting appears to help convergent behavior emerge over multiple steps.

scene	warm	w/o warm	scene	solve (s)	DoFs
composites	6.4	25.8	log house	4.9	414
chairs	2.9	30.9	log house dec.	0.061	414
catenary arch	5.6	133.7	bunny drop	37.5	12,276
card house	3.2	57.1			
dinosaurs	2.7	18.3			

Table 2: **Convergence Behavior (left):** The average number of iterations for the validation examples. **Decomposition Performance (right):** The average solve time for the decomposition examples.

Haptic Rate Interaction We also implemented a highly unoptimized prototype force-feedback haptic rendering of deformable and rigid frictional contact interactions by plugging a Phantom Premium 1.5 haptic device directly into our existing integration cycle. A timing and convergence summary for the interaction session shown in the accompanying video and Figure 1 (a) is given in Table 1, bottom. Note that, here, hard time limits required us to implement graceful degradation. Even so we found that most ($\sim 80\%$) of the integration steps were still convergent and that complex frictional behaviors such as jamming, sticking and stick-slip were obtained.

⁶These S.P. simulations were also run out to 10 minutes of virtual time.

Decomposition Without constraints, each body in a multibody system can be integrated independently. When contacts are imposed, however, the DoFs of individual bodies are effectively “glued” together. Generally, though, multiple *independent* connected sets of bodies are formed by contact constraints, each of which can be solved separately, and, if parallel processing is available, each such set can be solved in parallel. This is a standard feature in multibody packages (e.g., “islands” in ODE [Smith 2006]).

Additional decompositions can be applied to Staggered Projections steps to extract further sparsity. This allows us to effectively decompose the contacting system down to even smaller independent sets. Initial investigations reveal the potential for large speed-ups. While *all* of the above examples and timings provided were obtained *without* decomposition, we implemented two examples with a preliminary decomposition approach. We constructed a “Lincoln Log” scene composed of a rigid log house frictionally sliding, jumping a ramp and then colliding. The scene was simulated twice, once with decomposition and the second time without. We provide a side-by-side comparison of the results in the accompanying video. We also simulated a larger drop of rigid bunnies using decomposition (see Figure 1 (c)). The average solve time, per step, and the number of DoFs per scene for these three simulations is given in Table 2 (right).

9 Conclusion

We have proposed a discrete formulation of frictional contact suitable for both rigid and deformable systems, with a robust and efficient algorithm for accurately capturing the complex interplay between contact, friction and deformation. The algorithm behaves well in practice, as demonstrated by our instrumented examples and video.

Limitations and Future Work While the timings for the haptic demo (~ 100 Hz.) are sufficient for rendering soft contact, a separate code base specifically optimized for haptic rendering should be able to scale the timings down to enable haptic update rates of 1000 Hz.

We are also interested in the future inclusion of additional deformation models such as FEM, discrete shells, and other mesh and node based formulations. We expect that Staggered Projections should extend well to many of these types of generic deformable models. For some cases, however, where the interaction between contacts is almost entirely mediated by internal energy, such as models with low stiffness and no explicit rigid DoFs, internal energy terms will likely need to be included in the corrector step.

The Staggered Projection algorithm specifically addresses the *resolution* of frictional contact and so does not address *collision detection* and *contact sampling*. Step sizes are therefore somewhat dependent on the accuracy and quality of collision detection code. Continuous collision detection methods and other associated approaches should allow for larger step sizes in examples whose geometries and dynamics include high speed impacts and/or thin materials.

We are also focusing on clarifying conditions under which convergence can be guaranteed for Staggered Projections. Given the robust, accurate and convergent behavior exhibited so far in testing it is certainly possible that a broad convergence guarantee can be obtained. Currently, however, this is a work in progress. Research in this area is also exciting because it may potentially lead to further convergence accelerations.




Figure 6: **Hybrid scene:** Reduced deformable StVK bunnies collide with each other and with initially stable semi-circular arches composed of rigid blocks.




Figure 7: **Card house:** A rigid card house is initially stable and then partially knocked down by dropping blocks.

Acknowledgements: We would like to thank Timothy Edmunds, Mclean Edwards, Chris Cameron and David Levin for helpful discussions; the anonymous reviewers for their comments and suggestions; Paul Kry for sharing his FEM modeling code and Jernej Barbič for sharing his reduced StVK models and code. D.M.K., S.S. and D.K.P. were supported in part by the Canada Research Chairs Program, MITACS, NSERC, the Peter Wall Institute for Advanced Studies, and Autodesk. D.L.J. was supported in part by NSF Grant CAREER-0430528, Pixar, the Alfred P. Sloan Foundation, and Intel.

References

- ANITESCU, M., AND HART, G. D. 2004. A fixed-point iteration approach for multibody dynamics with contact and small friction. *Mathematical Programming* 101, 1, 3–32.
- ANITESCU, M., AND POTRA, F. R. 1997. Formulating dynamic multirigid-body contact problems with friction as solvable linear complementarity problems. *ASME Nonlinear Dynamics* 14, 231–247.
- BARAFF, D., AND WITKIN, A. P. 1992. Dynamics simulation of non-penetrating flexible bodies. In *Computer Graphics (SIGGRAPH 92)*, 303–308.
- BARAFF, D. 1989. Analytical methods for dynamic simulation of non-penetrating rigid bodies. In *Computer Graphics (SIGGRAPH 89)*, 223–232.
- BARAFF, D. 1991. Coping with friction for non-penetrating rigid body simulation. In *Computer Graphics (SIGGRAPH 91)*, 31–41.
- BARAFF, D. 1993. Issues in computing contact forces for non-penetrating rigid bodies. *Algorithmica* 10, 2–4, 292–352.
- BARAFF, D. 1994. Fast contact force computation for nonpenetrating rigid bodies. In *Proc. ACM SIGGRAPH 94*, 23–34.
- BARBIČ, J., AND JAMES, D. L. 2005. Real-time subspace integration for St. Venant-Kirchhoff deformable models. *ACM Trans. Graph. (SIGGRAPH 05)* 24, 3, 982–990.
- BAUSCHKE, H. 2000. Projection algorithms: results and open problems. In *Inherently Parallel Algorithms in Feasibility and Optimization and their Applications*, 11–22.
- BOYD, S., AND VANDENBERGHE, L. 2004. *Convex Optimization*. Cambridge.
- BRIDSON, R., FEDKIW, R. P., AND ANDERSON, J. 2002. Robust treatment of collisions, contact, and friction for cloth animation. *ACM Trans. Graph. (SIGGRAPH 02)* 21, 3, 594–603.
- BROGLIATO, B. 1999. *Nonsmooth Mechanics*. Springer-Verlag.
- CIRAK, F., AND WEST, M. 2005. Decomposition contact response (DCR) for explicit finite element dynamics. *International Journal for Numerical Methods in Engineering* 64, 8, 1078–1110.
- DURIEZ, C., DUBOIS, F., ANDRIOT, C., AND KHEDDAR, A. 2006. Realistic haptic rendering of interacting deformable objects in virtual environments. *IEEE Transactions on Visualization and Computer Graphics*.
- ERDMANN, M. E. 1994. On a representation of friction in configuration space. *The International Journal of Robotics Research* 13, 3, 240–271.
- ERLEBEN, K. 2007. Velocity-based shock propagation for multibody dynamics animation. *ACM Trans. Graph.* 26, 2.
- FERRIS, M. C., AND MUNSON, T. S. 1998. Complementarity problems in GAMS and the PATH solver. *Mathematical Programming Technical Report* 98-12.
- FICHERA, G. 1964. Problemi elastostatici con vincoli unilaterali: il problema di Signorini con ambigue condizioni al contorno. *Mem. Acc. Naz. Lincei* 8, 7, 91–14.
- GIBSON, S. F., AND MIRTICH, B. 1997. A survey of deformable models in computer graphics. *Technical Report TR-97-19, MERL*.
- GOLDFARB, D., AND IDNANI, G. 1983. A numerically stable dual method for solving strictly convex quadratic programs. *Mathematical Programming* 27, 1–33.

- GUENDELMAN, E., BRIDSON, R., AND FEDKIW, R. 2003. Non-convex rigid bodies with stacking. *ACM Trans. Graph. (SIGGRAPH 03)* 22, 3, 871–878.
- HAHN, J. K. 1988. Realistic animation of rigid bodies. In *Computer Graphics (SIGGRAPH 88)*, 299–308.
- HARMON, D., VOUGA, E., TAMSTORF, R., AND GRINSPUN, E. 2008. Robust treatment of simultaneous collisions. *ACM Trans. Graph. (SIGGRAPH 08)* 27, 3.
- HAUSER, K. K., SHEN, C., AND O'BRIEN, J. F. 2003. Interactive deformation using modal analysis with constraints. In *Graphics Interface*, 247–255.
- HERTZ, H. 1882. On the contact of elastic solids. *J. Reine und Angewandte Mathematik* 92, 156, 156–171.
- IRVING, G., SCHROEDER, C., AND FEDKIW, R. 2007. Volume conserving finite element simulations of deformable models. *ACM Trans. Graph. (SIGGRAPH 07)* 26, 3.
- JAMES, D. L., AND PAI, D. K. 2004. BD-Tree: Output-sensitive collision detection for reduced deformable models. *ACM Trans. Graph. (SIGGRAPH 04)* 23, 3, 393–398.
- JOHNSON, K. L. 1985. *Contact Mechanics*. Cambridge.
- JOURDAN, F., ALART, P., AND JEAN, M. 1998. A Gauss-Seidel like algorithm to solve frictional contact problems. *Comput. Methods Appl. Mech. Eng.* 155, 31–47.
- KAUFMAN, D. M., EDMUNDS, T., AND PAI, D. K. 2005. Fast frictional dynamics for rigid bodies. *ACM Trans. Graph. (SIGGRAPH 05)* 24, 3, 946–956.
- KIKUCHI, N., AND ODEN, J. 1988. *Contact Problems in Elasticity: A study of variational inequalities and finite element methods*. SIAM Studies in Applied and Numerical Mathematics.
- KLARBRING, A. 1986. A mathematical programming approach to three-dimensional contact problems with friction. *Comput. Methods Appl. Mech. Eng.* 58, 175–200.
- LÖTSTEDT, P. 1984. Numerical simulation of time-dependent contact friction problems in rigid body mechanics. *SIAM Journal of Scientific Statistical Computing* 5, 2, 370–393.
- MILENKOVIC, V. J., AND SCHMIDL, H. 2001. Optimization-based animation. In *Proc. ACM SIGGRAPH 01*, 37–46.
- MIRTICH, B., AND CANNY, J. F. 1995. Impulse-based dynamic simulation of rigid bodies. In *Symposium on Interactive 3D Graphics*.
- MOORE, M., AND WILHELMS, J. 1988. Collision detection and response for computer animation. In *Computer Graphics (SIGGRAPH 88)*, 289–298.
- MOREAU, J. J. 1966. Quadratic programming in mechanics: One-sided constraints. *Journal SIAM Control* 4, 1, 153–158.
- MOREAU, J. J. 1973. On unilateral constraints, friction and plasticity. *New Variational Techniques in Mathematical Physics*, 172–322.
- MURTY, K., AND KABADI, S. 1987. Some NP-complete problems in quadratic and nonlinear programming. *Mathematical Programming* 39, 117–129.
- MURTY, K. G. 1988. *Linear Complementarity, Linear and Non-linear Programming*. Helderman Verlag.
- NEALEN, A., MLLER, M., KEISER, R., BOXERMAN, E., AND CARLSON, M. 2005. Physically based deformable models in computer graphics. In *Eurographics 2005*.
- OTADUY, M. A., GERMANN, D., REDON, S., AND GROSS, M. 2007. Adaptive deformations with fast tight bounds. In *Proc. of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation*.
- PAINLEVÉ, P. 1895. Sur le lois du frottement de glissement. *C.R. Acad. Sci. Paris* 121, 112–115.
- PAULY, M., PAI, D. K., AND GUIBAS, L. 2004. Quasi-rigid objects in contact. In *ACM SIGGRAPH Symposium on Computer Animation*, ACM/Eurographics.
- RAGHUPATHI, L., AND FAURE, F. 2006. QP-Collide: A new approach to collision treatment. In *journées du groupe de travail Animation et Simulation (GTAS)*, 91–101.
- REDON, S., KHEDDAR, A., AND COQUILLART, S. 2002. Gauss' least constraints principle and rigid body simulations. In *IEEE International Conference on Robotics and Automation*.
- SCHITTKOWSKI, K. 2005. QL: A Fortran code for convex quadratic programming - user's guide, Version 2.11. *Report, Department of Mathematics, University of Bayreuth*.
- SHABANA, A. 2005. *Dynamics of Multibody Systems*, 3rd ed. Cambridge.
- SIGNORINI, S. 1933. Sopra alcune questioni di elastostatica. *Atti della Societa Italiana per il Progresso della Scienze*.
- SMITH, R. 2006. Open Dynamics Engine, V0.5, user guide.
- SONG, P., AND KUMAR, V. 2003. Distributed compliant model for efficient dynamic simulation of systems with frictional contacts. In *The 2003 ASME Design Engineering Technical Conferences*.
- SPILLMANN, J., BECKER, M., AND TESCHNER, M. 2007. Non-iterative computation of contact forces for deformable objects. *Journal of WSCG* 15, 13.
- STEWART, D., AND TRINKLE, J. C. 1996. An implicit time-stepping scheme for rigid body dynamics with inelastic collisions and Coulomb friction. *International Journal Numerical Methods Engineering* 39, 2673–2691.
- STEWART, D. E. 2000. Rigid-body dynamics with friction and impact. *SIAM Rev.* 42, 1, 3–39.
- STEWART, D. E. 2001. Finite-dimensional contact mechanics. *Phil. Trans. R. Soc. Lond. A* 359, 2467–2482.
- TERZOPOULOS, D., AND WITKIN, A. 1988. Physically based models with rigid and deformable components. *IEEE Computer Graphics and Applications* 8, 6, 41–51.
- TRINKLE, J. C., PANG, J. S., SUDARSKY, S., AND LO, G. 1995. On dynamic multi-rigid-body contact problems with Coulomb friction. Tech. rep., Texas A&M University, Department of Computer Science.
- WASFY, T. M., AND NOOR, A. K. 2003. Computational strategies for flexible multibody systems. *Applied Mechanics Reviews* 56, 6 (November), 553–613.
- WRIGGERS, P. 2002. *Computational Contact Mechanics*. John Wiley and Sons, Ltd.

