

Towards Real-time Simulation of Hyperelastic Materials

Tiantian Liu
University of Pennsylvania

Sofien Bouaziz
EPFL

Ladislav Kavan
University of Utah




Figure 1: Our method enables fast simulation of many different types of hyperelastic materials. Compared to the commonly-applied Newton’s method, our method is about 10 times faster, while achieving even higher accuracy and being simpler to implement. The Polynomial and Spline-based materials are models recently introduced by Xu et al. [2015]. Spline-based material A is a modified Neo-Hookean material with stronger resistance to compression; spline-based material B is a modified Neo-Hookean material with stronger resistance to tension.

Abstract

We present a new method for real-time physics-based simulation supporting many different types of hyperelastic materials. Previous methods such as Position Based or Projective Dynamics are fast, but support only limited selection of materials; even classical materials such as the Neo-Hookean elasticity are not supported. Recently, Xu et al. [2015] introduced new “spline-based materials” which can be easily controlled by artists to achieve desired animation effects. Simulation of these types of materials currently relies on Newton’s method, which is slow, even with only one iteration per timestep. In this paper, we show that Projective Dynamics can be interpreted as a quasi-Newton method. This insight enables very efficient simulation of a large class of hyperelastic materials, including the Neo-Hookean, spline-based materials, and others. The quasi-Newton interpretation also allows us to leverage ideas from numerical optimization. In particular, we show that our solver can be further accelerated using L-BFGS updates (Limited-memory Broyden-Fletcher-Goldfarb-Shanno algorithm). Our final method is typically more than 10 times faster than one iteration of Newton’s method without compromising quality. In fact, our result is often more accurate than the result obtained with one iteration of Newton’s method. Our method is also easier to implement, implying reduced software development costs.

Keywords: Physics-based animation, material models, numerical optimization.

1 Introduction

Physics-based animation is an important tool in computer graphics even though creating visually compelling simulations often requires a lot of patience. Waiting for results is not an option in real-time simulations, which are necessary in applications such as computer games and training simulators, e.g., surgery simulators. Previous methods for real-time physics such as Position Based Dynamics [Müller et al. 2007] or Projective Dynamics [Bouaziz et al. 2014] have been successfully used in many applications and commercial products, despite the fact that these methods support only a restricted set of material models. Even classical models from continuum mechanics, such as the Neo-Hookean, St. Venant-Kirchhoff, or Mooney-Rivlin materials, are not supported by Projective Dynamics. We tried to emulate their behavior with Projective Dynamics, but despite our best efforts, there are still obvious visual differences when compared to simulations with the original non-linear materials.

The advantages of more general material models were nicely demonstrated in the recent work of Xu et al. [2015], who proposed a new class of spline-based materials particularly suitable for physics-based animation. Their user-friendly spline interface enables artists to easily modify material properties in order to achieve desired animation effects. However, their system relies on Newton’s method, which is slow, even if the number of Newton’s iter-

ations per frame is limited to one. Our method enables fast simulation of spline-based materials, combining the benefits of artist-friendly material interfaces with the advantages of fast simulation, such as rapid iterations and/or higher resolutions.

Physics-based simulation can be formulated as an optimization problem where we minimize a multi-variate function g . Newton’s method minimizes g by performing descent along direction $-(\nabla^2 g)^{-1} \nabla g$, where $\nabla^2 g$ is the Hessian matrix, and ∇g is the gradient. One problem of Newton’s method is that the Hessian $\nabla^2 g$ can be indefinite, in which case the Newton’s direction could erroneously *increase* g . This undesired behavior can be prevented by so-called “definiteness fixes” [Teran et al. 2005; Nocedal and Wright 2006]. While definiteness fixes require some computational overheads, the slow speed of Newton’s method is mainly caused by the fact that the Hessian changes at every iteration, i.e., we need to solve a *new* linear system for every Newton step.

The point of departure for our method is the insight that Projective Dynamics can be interpreted as a special type of quasi-Newton method. In general, quasi-Newton methods [Nocedal and Wright 2006] work by replacing the Hessian $\nabla^2 g$ with a linear operator \mathbf{A} , where \mathbf{A} is positive definite and solving linear systems $\mathbf{A}\mathbf{x} = \mathbf{b}$ is fast. The descent directions are then computed as $-\mathbf{A}^{-1} \nabla g$ (where the inverse is of course not explicitly evaluated, in fact, \mathbf{A} is often not even represented with a matrix). The trade-off is that if \mathbf{A} is a poor approximation of the Hessian, the quasi-Newton method may converge slowly. Unfortunately, coming up with an effective approximation of the Hessian is not easy. We tried many previous quasi-Newton methods, but even after boosting their performance with L-BFGS [Nocedal and Wright 2006], we were unable to obtain an effective method for real-time physics. We show that Projective Dynamics can be re-formulated as a quasi-Newton method with some remarkable properties, in particular, the resulting \mathbf{A}_{our} matrix is *constant* and *positive definite*. This re-formulation enables us to generalize the method to hyperelastic materials not supported by Projective Dynamics, such as the Neo-Hookean or spline-based materials. Even though the resulting solver is slightly more complicated than Projective Dynamics (in particular, we must employ a line search to ensure stability), the computational overhead required to support more general materials is rather small.

The quasi-Newton formulation also allows us to further improve convergence of our solver. We propose using L-BFGS, which uses curvature information estimated from a certain number of previous iterates to improve the accuracy of our Hessian approximation \mathbf{A}_{our} . Adding the L-BFGS Hessian updates introduces only a small computational overhead while accelerating the convergence of our method. However, this is not a silver bullet, because the performance of L-BFGS highly depends on the quality of the initial Hessian approximation. With previous quasi-Newton methods, we observed rather disappointing convergence properties (see Figure 7). However, the combination of our Hessian approximation \mathbf{A}_{our} with L-BFGS is quite effective and can be interpreted as a generalization of the recently proposed Chebyshev Semi-Iterative method for accelerating Projective Dynamics [Wang 2015].

The L-BFGS convergence boosting is compatible with our first contribution, i.e., fast simulation of complex non-linear materials. Specifically, we can simulate any materials satisfying the Valanis-Landel assumption [Valanis and Landel 1967] which includes many classical materials, such as St. Venant-Kirchhoff, Neo-Hookean, Mooney-Rivlin, and also the recently proposed spline-based materials [Xu et al. 2015] (none of which is supported by Projective Dynamics). In summary, our final method achieves faster convergence than Projective Dynamics while being able to simulate a large variety of hyperelastic materials.

2 Related Work

The work of Terzopoulos et al. [1987] pioneered physics-based animation, nowadays an indispensable tool in feature animation and visual effects. Real-time physics became widespread only more recently, with first success stories represented by real-time rigid body simulators, commercially offered by companies such as Havok since early 2000s. Fast simulation of deformable objects is more challenging because they feature many more degrees of freedom than rigid bodies. Fast simulations of deformable objects using shape matching [Müller et al. 2005; Rivers and James 2007] paved the way towards more general Position Based Dynamics methods [Müller et al. 2007; Stam 2009]. The past decade witnessed rapid development of Position Based methods, including improvements of the convergence [Müller 2008; Kim et al. 2012], robust simulation of elastic models [Müller and Chentanez 2011], generalization to fluids [Macklin and Müller 2013] and continuum-based materials [Müller et al. 2014; Bender et al. 2014a], unified solvers including multiple phases of matter [Macklin et al. 2014], and most recently, methods to avoid element inversion [Müller et al. 2015]. We refer to a recent survey [Bender et al. 2014b] for a more detailed summary of Position Based methods.

A new interpretation of Position Based methods was offered by Liu et al. [2013], observing that Position Based Dynamics can be interpreted as an approximate solver for Implicit Euler time-stepping. The same paper introduces a fast local/global solver for mass-spring systems integrated using Implicit Euler. This method was later generalized to Projective Dynamics [Bouaziz et al. 2014] by combining the ideas of [Liu et al. 2013] with a shape editing system “Shape-Up” [Bouaziz et al. 2012]. Recently, a Chebyshev Semi-Iterative method [Wang 2015] has been proposed to accelerate convergence of Projective Dynamics, while exploring also highly parallel GPU implementations of real-time physics.

Multi-grid methods represent another approach to accelerate physics-based simulations [Georgii and Westermann 2006; Müller 2008; Wang et al. 2010; McAdams et al. 2011; Tamstorf et al. 2015]. Multi-grid methods are attractive especially for highly detailed meshes where sparse direct solvers become hindered by high memory requirements. However, constructing multi-resolution data structures and picking suitable parameters is not a trivial task. Another way to speed up FEM is by using subspace simulation where the nodal degrees of freedom are replaced with a low-dimensional linear subspace [Barbič and James 2005; An et al. 2008; Li et al. 2014]. These methods can be very efficient; however, deformations that were not accounted for during the subspace construction may not be well represented. A variety of approaches have been designed to address this limitation while trying to preserve efficiency [Harmon and Zorin 2013; Teng et al. 2014; Teng et al. 2015]. Simulating at coarser resolutions is also possible, while crafting special data-driven materials which avoid the loss of accuracy typically associated with lower resolutions [Chen et al. 2015].

The concept of constraint projection, which appears in both Position Based and Projective Dynamics, is also central to the Fast Projection method [Goldenthal et al. 2007] and strain-limiting techniques [Thomaszewski et al. 2009; Narain et al. 2012]. The Fast Projection method and Position Based Dynamics formulate physics simulation as a constrained optimization problem that is solved by linearizing the constraints in the spirit of sequential quadratic programming [Macklin et al. 2014]. The resulting Karush-Kuhn-Tucker (KKT) equation system is then solved using a direct solver [Goldenthal et al. 2007] or an iterative method such as Gauss-Seidel [Müller et al. 2007; Stam 2009; Fratarcangeli and Pellacini 2015], Jacobi [Macklin and Müller 2013], or their under/over-relaxation counterparts [Macklin et al. 2014]. By using a constrained optimization formulation the Fast Projection method and Position Based Dynamics are designed for solving infinitely

stiff systems but are not appropriate to handle soft materials. This problem can be overcome by regularizing the KKT system [Servin et al. 2006; Tournier et al. 2015], leading to approaches that can accurately handle extremely high tensile forces (e.g., string of a bow) but also support soft (compliant) constraints. However, these methods are slower than Projective Dynamics because a new linear system has to be solved at each iteration.

The idea of quasi-Newton methods in elasticity is not new and has been studied long time before real-time simulations were feasible. Several research papers have been done to accelerate Newton’s method in FEM simulations by updating the Hessian approximation only once every frame [Bathe and Cimento 1980; Fish et al. 1995]. However, even one Hessian update is usually so computationally expensive that can not fit into the computing time limit of real-time applications. Deuffhard [2011] minimizes the number of Hessian factorizations by carefully scheduled Hessian updates. But the update rate will heavily depend on the deformation. A good Hessian approximation suitable for realtime applications should be easy to refactorize or capable of prefactorization. One straightforward constant approximation which is good for prefactorization is the Hessian evaluated at the rest-pose (undeformed configuration). The rest-pose is positive semi-definite and its use at any configuration enables pre-factorization. Unfortunately, the actual Hessian of deformed configurations is often very different from the rest-pose Hessian and this approximation is therefore not satisfactory for larger deformations [Müller et al. 2002].

To improve upon this, Müller et al. [2002] introduced per-vertex “stiffness warping” of the rest-pose Hessian, which is more accurate and can still leverage pre-factorized rest-pose Hessian. Unfortunately, the per-vertex stiffness warping approach can introduce non-physical ghost forces which violate momentum conservation and can lead to instabilities [Müller and Gross 2004]. This problem was addressed by *per-element* stiffness warping [Müller and Gross 2004] which avoids the ghost forces but, unfortunately, the per-element-warped stiffness matrices need to be re-factorized, introducing computational overheads which are prohibitive in real-time simulation. For corotated elasticity, Hecht et al. [2012] proposed an improved method which can re-use previously computed Hessian factorization. Specifically, the sparse Cholesky factors are updated only *when* necessary and also only *where* necessary. This spatio-temporal staging of Cholesky updates improves run-time performance, however, the Cholesky updates are still costly and their scheduling can be problematic especially in real-time applications, which require approximately constant per-frame computing costs. Also, the frequency of Cholesky updates depends on the simulation: fast motion with large deformations will require more frequent updates and thus more computation, or risking ghost forces and potential instabilities. Neither is an option in real-time simulators.

Our re-formulation of Projective Dynamics as a quasi-Newton method reveals relationships to so called “Sobolev gradient methods”, which have been studied since the 1980s in the continuous setting [Neuberger 1983]; see also the more recent monograph [Neuberger 2009]. The idea of quasi-Newton methods appears already in [Desbrun et al. 1999; Hauth and Eitzmuss 2001] in the context of mass-spring systems and, more recently, in [Martin et al. 2013] in the context of geometry processing. Martin et al. [2013] also propose multi-scale extensions and discuss an application in physics-based simulation, but consider only the case of thin shells and their numerical method alters the physics of the simulated system. Quasi-Newton methods are also useful in situations where computation of the Hessian would be expensive or impractical [Nocedal and Wright 2006]. In character animation, Hahn et al. [2012] used BFGS to simulate physics in “Rig Space”, which is challenging because the rig is a black box function and its derivatives are approximated using finite differences.

3 Background

Projective Dynamics. We start by introducing our notation and recapitulating the key concepts of Projective Dynamics. Let $\mathbf{x} \in \mathbb{R}^{n \times 3}$ be the current (deformed) state of our system containing n nodes, each with three spatial dimensions. Projective Dynamics requires a special form of elastic potential energies, based on the concept of constraint projection. Specifically, Projective Dynamics energy for element number i is defined as:

$$E_i(\mathbf{x}) = \min_{\mathbf{p}_i \in \mathcal{M}_i} \tilde{E}_i(\mathbf{x}, \mathbf{p}_i), \quad \tilde{E}_i(\mathbf{x}, \mathbf{z}) = \|\mathbf{G}_i \mathbf{x} - \mathbf{z}\|_F^2 \quad (1)$$

where $\|\cdot\|_F$ is the Frobenius norm, \mathcal{M}_i is a constraint manifold, \mathbf{p}_i is an auxiliary “projection variable”, and \mathbf{G}_i is a discrete differential operator represented, e.g., by a sparse matrix. For example, if element number i is a tetrahedron, \mathcal{M}_i is $SO(3)$, and \mathbf{G}_i is deformation gradient operator [Sifakis and Barbič 2012], we obtain the well-known as-rigid-as-possible material model [Chao et al. 2010]. Another elementary example is a spring, where the element is an edge, \mathcal{M}_i is a sphere, and \mathbf{G}_i subtracts two endpoints of the spring. If all elements are springs, Projective Dynamics becomes equivalent to the work of Liu et al. [2013]. The key property of \mathbf{G}_i is that constant vectors are in its nullspace, which makes E_i translation invariant. The total energy of the system is:

$$E(\mathbf{x}) = \sum_i w_i E_i(\mathbf{x}) \quad (2)$$

where i indexes elements and $w_i > 0$ is a positive weight, typically defined as the product of undeformed volume and stiffness.

Time integration. As discussed by Martin et al. [2011], Backward Euler time integration can be expressed as a minimization of:

$$g(\mathbf{x}) = \frac{1}{2h^2} \underbrace{\text{tr}((\mathbf{x} - \mathbf{y})^\top \mathbf{M}(\mathbf{x} - \mathbf{y}))}_{\text{inertia}} + \underbrace{E(\mathbf{x})}_{\text{elasticity}} \quad (3)$$

where \mathbf{y} is a constant depending only on previously computed states, \mathbf{M} is a positive definite mass matrix (typically diagonal – mass lumping), and $h > 0$ is the time step (we use fixed h corresponding to the frame rate of $30fps$, i.e., $h = 1/30s$). The trace (tr) reflects the fact that there are no dependencies between the x, y, z coordinates, which enables us to work only with $n \times n$ matrices (as opposed to more general $3n \times 3n$ matrices). This is somewhat moot in the context of the mass matrix \mathbf{M} , but it will be more important in the following. The constant \mathbf{y} is defined as $\mathbf{y} := 2\mathbf{q}_t - \mathbf{q}_{t-1} + h^2 \mathbf{M}^{-1} \mathbf{f}_{\text{ext}}$, where $\mathbf{q}_t \in \mathbb{R}^{n \times 3}$ is the current state, \mathbf{q}_{t-1} the previous state, and $\mathbf{f}_{\text{ext}} \in \mathbb{R}^{n \times 3}$ are external forces such as gravity. The minimizer of $g(\mathbf{x})$ will become the next state, \mathbf{q}_{t+1} . Intuitively, the first term in Eq. 3 can be interpreted as “inertial potential,” attracting \mathbf{x} towards \mathbf{y} , where \mathbf{y} corresponds to state predicted by Newton’s first law – motion without the presence of any internal forces. The second term penalizes states \mathbf{x} with large elastic deformations. Minimization of $g(\mathbf{x})$ corresponds to finding balance between the two terms. Note that many other implicit integration schemes can also be expressed as minimization problems similar to Eq. 3. In particular, we have implemented Implicit Midpoint, which has the desirable feature of being symplectic [Hairer et al. 2002; Kharevych et al. 2006]. Unfortunately, in our experiments we found Implicit Midpoint to be markedly less stable than Backward Euler and, therefore, we continue to use Backward Euler despite its numerical damping.

Local/global solver. The key idea of Projective Dynamics is to expose the auxiliary projection variables \mathbf{p}_i , taking advantage of the special energy form according to Eq. 1. To simplify notation, we stack all projection variables into $\mathbf{p} \in \mathbb{R}^{e \times 3}$ and define binary

selector matrices \mathbf{S}_i such that $\mathbf{p}_i = \mathbf{S}_i \mathbf{p}$. Projective Dynamics uses the augmented objective:

$$\tilde{g}(\mathbf{x}, \mathbf{p}) = \frac{1}{2h^2} \text{tr}((\mathbf{x} - \mathbf{y})^\top \mathbf{M}(\mathbf{x} - \mathbf{y})) + \sum_i w_i \tilde{E}(\mathbf{x}, \mathbf{S}_i \mathbf{p}) \quad (4)$$

which is minimized over both \mathbf{x} and \mathbf{p} , subject to the constraint $\mathbf{p} \in \mathcal{M}$, where \mathcal{M} is a cartesian product of the individual constraint manifolds. The optimization is solved using an alternating (local/global) solver. In the local step, \mathbf{x} is assumed to be fixed; the optimal \mathbf{p} are given by projections on individual constraint manifolds, e.g., projecting each deformation gradient (a 3×3 matrix) on $SO(3)$. In the global step, \mathbf{p} is assumed to be fixed and we rewrite the objective $\tilde{g}(\mathbf{x}, \mathbf{p})$ in matrix form:

$$\frac{1}{2h^2} \text{tr}((\mathbf{x} - \mathbf{y})^\top \mathbf{M}(\mathbf{x} - \mathbf{y})) + \frac{1}{2} \text{tr}(\mathbf{x}^\top \mathbf{L} \mathbf{x}) - \text{tr}(\mathbf{x}^\top \mathbf{J} \mathbf{p}) + C \quad (5)$$

where $\mathbf{L} := \sum w_i \mathbf{G}_i^\top \mathbf{G}_i$, $\mathbf{J} := \sum w_i \mathbf{G}_i^\top \mathbf{S}_i$, and the constant C is irrelevant for optimization. For a fixed \mathbf{p} , the minimization of $\tilde{g}(\mathbf{x}, \mathbf{p})$ can be accomplished by finding \mathbf{x} with a vanishing gradient, i.e., $\nabla_{\mathbf{x}} \tilde{g}(\mathbf{x}, \mathbf{p}) = 0$. Computing the gradient yields some convenient simplifications (the traces disappear):

$$\nabla_{\mathbf{x}} \tilde{g}(\mathbf{x}, \mathbf{p}) = \frac{1}{h^2} \mathbf{M}(\mathbf{x} - \mathbf{y}) + \mathbf{L} \mathbf{x} - \mathbf{J} \mathbf{p} \quad (6)$$

Equating the gradient to zero leads to the solution:

$$\mathbf{x}^* = (\mathbf{M}/h^2 + \mathbf{L})^{-1} (\mathbf{J} \mathbf{p} + \mathbf{M} \mathbf{y}/h^2) \quad (7)$$

The matrix $\mathbf{M}/h^2 + \mathbf{L}$ is symmetric positive definite and therefore \mathbf{x}^* is a global minimum (for fixed \mathbf{p}). The key computational advantage of Projective Dynamics is that $\mathbf{M}/h^2 + \mathbf{L}$ does not depend on \mathbf{x} , which allows us to pre-compute and repeatedly reuse its sparse Cholesky factorization to quickly solve for \mathbf{x}^* , which is the result after one local and global step. The local and global steps are repeated for a fixed number of iterations (typically 10 or 20).

4 Method

As described in the previous section, Projective Dynamics relies on the special type of elastic energies according to Eq. 1. Let us now describe how Projective Dynamics can be interpreted as a quasi-Newton method. The first step is to compute the gradient of the objective $g(\mathbf{x})$ from Eq. 3. The energy $E(\mathbf{x})$ used in this objective contains constrained minimization over the projection variables $\mathbf{p}_i \in \mathcal{M}_i$ (see Eq. 1 and Eq. 2). Equivalently, we can interpret the \mathbf{p}_i as functions of \mathbf{x} realizing the projections, according to Eq. 1. Nevertheless, the gradient $\nabla g(\mathbf{x})$ can still be computed easily – in fact, it is exactly equivalent to $\nabla_{\mathbf{x}} \tilde{g}(\mathbf{x}, \mathbf{p})$ from Eq. 6 where we assumed that \mathbf{p} is constant. This at first surprising fact has been observed in previous work [Chao et al. 2010; Bouaziz et al. 2012]. Intuitively, the reason is that if we infinitesimally perturb \mathbf{x} , its projection $\mathbf{p}_i(\mathbf{x})$ can move only in the tangent space of \mathcal{M}_i and therefore, the differential $\delta \mathbf{p}_i(\mathbf{x})$ has no effect on $\delta \|\mathbf{x} - \mathbf{p}_i(\mathbf{x})\|^2$. As an intuitive explanation, imagine that \mathbf{x} is a space shuttle projected to its closest point on Earth $\mathbf{p}_i(\mathbf{x})$; to first order, the distance of the space shuttle from Earth does not depend on the tangent motion $\delta \mathbf{p}_i(\mathbf{x})$. Please see the Appendix for a more formal discussion. In summary, the gradient of Eq. 3 is:

$$\nabla g(\mathbf{x}) = \frac{1}{h^2} \mathbf{M}(\mathbf{x} - \mathbf{y}) + \mathbf{L} \mathbf{x} - \mathbf{J} \mathbf{p}(\mathbf{x}) \quad (8)$$

where $\mathbf{p}(\mathbf{x})$ is a function stacking all of the individual projections $\mathbf{p}_i(\mathbf{x})$. Newton’s method would proceed by computing second derivatives, i.e., the Hessian matrix $\nabla^2 g(\mathbf{x})$, and use it to compute

a descent direction $-(\nabla^2 g(\mathbf{x}))^{-1} \nabla g(\mathbf{x})$. Note that definiteness fixes may be necessary to guarantee this will really be a *descent* direction [Gast et al. 2015].

What happens if we modify Newton’s method by using $\mathbf{M}/h^2 + \mathbf{L}$ instead of the Hessian $\nabla^2 g(\mathbf{x})$? Simple algebra reveals:

$$(\mathbf{M}/h^2 + \mathbf{L})^{-1} \nabla g(\mathbf{x}) = \mathbf{x} - (\mathbf{M}/h^2 + \mathbf{L})^{-1} (\mathbf{J} \mathbf{p}(\mathbf{x}) + \mathbf{M} \mathbf{y}/h^2)$$

However, the latter term is equivalent to the result of one iteration of the local/global steps of Projective Dynamics, see Eq. 7. Therefore, $(\mathbf{M}/h^2 + \mathbf{L})^{-1} \nabla g(\mathbf{x}) = \mathbf{x} - \mathbf{x}^*$ and we can interpret $\mathbf{d}_{\text{PD}} := -(\mathbf{M}/h^2 + \mathbf{L})^{-1} \nabla g(\mathbf{x})$ as a descent direction (this time there is no need for any definiteness fixes). Projective Dynamics can be therefore understood as a quasi-Newton method which computes the next iterate as $\mathbf{x} + \mathbf{d}_{\text{PD}}$. Typically, quasi-Newton methods use line search techniques [Nocedal and Wright 2006] to find parameter $\alpha > 0$ such that $\mathbf{x} + \alpha \mathbf{d}_{\text{PD}}$ reduces the objective as much as possible. However, with Projective Dynamics energies according to Eq. 1, the optimal value is always $\alpha = 1$.

4.1 More general materials

The interpretation of Projective Dynamics as a quasi-Newton method suggests that a similar optimization strategy might be effective for more general elastic potential energies. First, let us focus on isotropic materials, deferring the discussion of anisotropy to Section 4.4. The assumption of isotropy (material-space rotation invariance) together with world-space rotation invariance means that we can express elastic energy density function Ψ as a function of singular values of the deformation gradient [Irving et al. 2004; Sifakis and Barbič 2012]. In the volumetric case, we have three singular values $\sigma_1, \sigma_2, \sigma_3 \in \mathbb{R}$, also known as “principal stretches”. The function $\Psi(\sigma_1, \sigma_2, \sigma_3)$ must be invariant to any permutation of the principal stretches, e.g., $\Psi(\sigma_1, \sigma_2, \sigma_3) = \Psi(\sigma_2, \sigma_1, \sigma_3)$ etc. Because directly working with such functions Ψ could be cumbersome, we instead use the Valanis-Landel hypothesis [Valanis and Landel 1967], which assumes that:

$$\Psi(\sigma_1, \sigma_2, \sigma_3) = a(\sigma_1) + a(\sigma_2) + a(\sigma_3) + b(\sigma_1 \sigma_2) + b(\sigma_2 \sigma_3) + b(\sigma_1 \sigma_3) + c(\sigma_1 \sigma_2 \sigma_3) \quad (9)$$

where $a, b, c : \mathbb{R} \rightarrow \mathbb{R}$. Many material models can be written in the Valanis-Landel form, including linear corotated material [Sifakis and Barbič 2012], St. Venant-Kirchhoff, Neo-Hookean, and Mooney-Rivlin. The recently proposed spline-based materials [Xu et al. 2015] are also based on the Valanis-Landel assumption. How can we generalize Projective Dynamics to these types of materials? Invoking the quasi-Newton interpretation discussed above, our method will minimize the objective g by performing descent along direction $\mathbf{d}(\mathbf{x}) := -(\mathbf{M}/h^2 + \mathbf{L})^{-1} \nabla g(\mathbf{x})$. The mass matrix \mathbf{M} and step size h are defined as before, and computing $\nabla g(\mathbf{x})$ is straightforward. But how to define a matrix \mathbf{L} for a given material model? This matrix can still have the form $\mathbf{L} := \sum w_i \mathbf{G}_i^\top \mathbf{G}_i$, but the question is how to choose the weights w_i . In Projective Dynamics, we assumed the weights are given as $w_i = V_i k_i$, where $V_i > 0$ is rest-pose volume of i -th element, and $k_i > 0$ is a stiffness parameter provided by the user. In our case, the user instead specifies a material model according to Eq. 9 from which we have to infer the appropriate k_i value. In the following we drop the subscript i for ease of notation.

For linear materials (Hooke’s law), stiffness is given as the second derivative of elastic energy. Therefore, it would be tempting to set k equal to the second derivative of Ψ at the rest pose (corresponding to $\sigma_1 = \sigma_2 = \sigma_3 = 1$), which evaluates to $a''(1) + 2b''(1) + c''(1)$, regardless of whether we differentiate with respect to σ_1, σ_2 , or σ_3 .

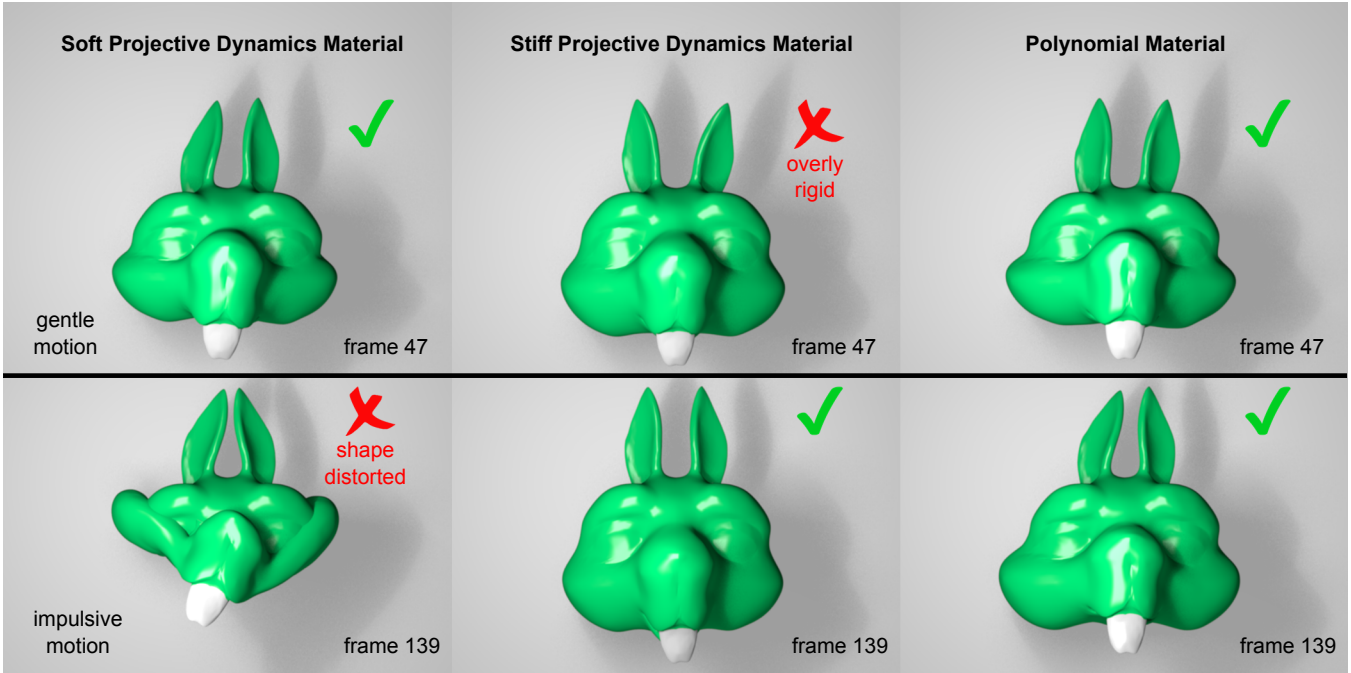


Figure 2: Animating jiggly squirrel head. The squirrel head is driven by a gentle keyframed motion in the top row, and by a faster, impulsive motion in the bottom row. Soft Projective Dynamics material (left column) creates nice secondary motion, but does not prevent large distortions of the shape. If we stiffen the Projective Dynamics material (middle column), we prevent the distortions, but also kill the secondary motion. Our polynomial material $a(x) = \mu(x - 1)^4$, $b(x) = 0$, $c(x) = 0$ (right column) achieves the desired effect of jiggling without large shape distortions.

Even though this method would produce suitable k for some materials (such as corotated elasticity), it does not work e.g. for a polynomial material defined as $a(x) = \mu(x - 1)^4$, $b(x) = 0$, $c(x) = 0$. Already this relatively simple material can facilitate certain animation tasks, such as creating a cartoon squirrel head which jiggles, but does not overly distort its shape, see Figure 2. However, with this material, the second derivatives at $x = 1$ evaluates to zero regardless of the value of μ , which would lead to zero stiffness which is obviously not a good approximation. The problem is the second derivative takes into account only infinitesimally small neighborhood of $x = 1$, i.e., the rest pose. However, we need a *single* value of k which will work well in the entire range of deformations expected in our simulations. To capture this requirement, we define an interval $[x_{\text{start}}, x_{\text{end}}]$ where we expect our principal stretches to be. We consider the following stress function:

$$\left. \frac{\partial \Psi}{\partial \sigma_1} \right|_{\sigma_2=1, \sigma_3=1} = a'(\sigma_1) + 2b'(\sigma_1) + c'(\sigma_1) \quad (10)$$

and define our k as the slope of the best linear approximation of Eq. 10 for $\sigma_1 \in [x_{\text{start}}, x_{\text{end}}]$. Note that due to the symmetry of the Valanis-Landel assumption, we would obtain exactly the same result if we differentiated with respect to σ_2 or σ_3 (instead of σ_1 as above). We study different choices of $[x_{\text{start}}, x_{\text{end}}]$ intervals in Section 5. In summary, the results are not very sensitive to the particular choice of x_{start} and x_{end} . The key fact is that regardless of the specific setting of x_{start} and x_{end} , spatial variations of μ are correctly taken into account, i.e., softer and stiffer parts of the simulated object will have different μ coefficients (e.g., in our squirrel head we made the teeth more stiff). Even though all elements have the same $[x_{\text{start}}, x_{\text{end}}]$ interval, the resulting matrices \mathbf{L} and \mathbf{J} properly reflect the spatially varying stiffness.

Line search. With Projective Dynamics materials (Eq. 1), the line search parameter $\alpha = 1$ is always guaranteed to decrease the objective g (Eq. 3). Unfortunately, this is no longer true in our gen-

Algorithm 1: Quasi-Newton Solver

```

1  $\mathbf{x}_1 := \mathbf{y}$ ;  $g(\mathbf{x}_1) := \text{evalObjective}(\mathbf{x}_1)$ 
2 for  $k = 1, \dots, \text{numIterations}$  do
3    $\nabla g(\mathbf{x}_k) := \text{evalGradient}(\mathbf{x}_k)$ 
4    $\mathbf{d}(\mathbf{x}_k) := -(\mathbf{M}/h^2 + \mathbf{L})^{-1} \nabla g(\mathbf{x}_k)$ 
5    $\alpha := 2$ 
6   repeat
7      $\alpha := \alpha/2$ 
8      $\mathbf{x}_{k+1} := \mathbf{x}_k + \alpha \mathbf{d}(\mathbf{x}_k)$ 
9      $g(\mathbf{x}_{k+1}) := \text{evalObjective}(\mathbf{x}_{k+1})$ 
10  until  $g(\mathbf{x}_{k+1}) \leq g(\mathbf{x}_k) + \gamma \alpha \text{tr}((\nabla g(\mathbf{x}_k))^T \mathbf{d}(\mathbf{x}_k))$ ;
11 end

```

eralized quasi-Newton setting, where it is easy to find examples where $g(\mathbf{x} + \mathbf{d}(\mathbf{x})) > g(\mathbf{x})$, i.e., taking a step of size one actually *increases* the objective. This can lead to erroneous energy accumulation, potentially resulting in catastrophic failure of the simulation (“explosions”), as shown in Figure 3. Fortunately, thanks to the fact that $\mathbf{M}/h^2 + \mathbf{L}$ is positive definite, $\mathbf{d}(\mathbf{x})$ is guaranteed to be a descent direction. Therefore, there exists $\alpha > 0$ such that $g(\mathbf{x} + \alpha \mathbf{d}(\mathbf{x})) \leq g(\mathbf{x})$ (unless we are already at a critical point $\nabla g(\mathbf{x}) = \mathbf{0}$, at which point the optimization is finished). In fact, we can ask for even more, i.e., we can always find $\alpha > 0$ such that $g(\mathbf{x} + \alpha \mathbf{d}(\mathbf{x})) \leq g(\mathbf{x}) + \gamma \alpha \text{tr}((\nabla g(\mathbf{x}))^T \mathbf{d}(\mathbf{x}))$ for some constant $\gamma \in (0, 1)$ (we use $\gamma = 0.3$). This is known as the Armijo condition which expresses the requirement of “sufficient decrease” [Nocedal and Wright 2006], preventing the line search algorithm from reducing the objective only by a negligible amount. Another requirement for robust line search is to avoid too small steps α (even though they might satisfy the Armijo condition). We tested two possibilities: Wolfe conditions, which impose an additional “curvature condition”, and backtracking line search, which starts from large α and progressively decreases it until the Armijo condition is satisfied. We found that in our setting both approaches lead to




Figure 3: Without line search, the squirrel head animation using our polynomial material (as in Figure 2) quickly becomes unstable.

comparable error reduction, but the backtracking line search is less computationally expensive. Also, $\alpha = 1$ is an excellent initial guess for the backtracking strategy. Therefore, in our final algorithm we implement the backtracking line search; after a failed attempt, we multiply alpha by 0.5. This value worked well in our experiments, even though, in theory, any constant $\in (0, 1)$ could be used instead.

Alg. 1 summarizes the process of computing one frame of our simulation. The outer loop (lines 2-11) performs quasi-Newton iterations and the inner loop (lines 6-10) implements the line search. What is the extra computational cost required to support more general materials? With Projective Dynamics energies (Eq. 1), we do not need the line search, because $\alpha = 1$ always works. Indeed, if we drop the line search from Alg. 1, the algorithm becomes equivalent to a generalized local/global process, as discussed in Section 3 (which is unstable for non-Projective-Dynamics energies). Rejected line search attempts, i.e., additional iterations of the line search, represent the main computational overhead of our method. Fortunately, we found that in practical simulations the number of extra line search iterations is relatively small. For example, in the squirrel head example in Figure 2 using the polynomial material, we need only 4280 line search iterations for the entire sequence with 400 frames, 10 quasi-Newton iterations per frame, i.e., the average number of line search iterations per quasi-Newton iteration is only 1.07. Even though in most cases the full step ($\alpha = 1$) succeeds, the Armijo safeguard is essential for stability; if we drop it, the simulation can quickly explode, as shown in Figure 3.

4.2 Accelerating convergence

The connection between Projective Dynamics and quasi-Newton methods allows us to take advantage of further mathematical optimization techniques. In this section, we discuss how to accelerate convergence of our method using L-BFGS (Limited-memory BFGS). The BFGS algorithm (Broyden-Fletcher-Goldfarb-Shanno) is one of the most popular general purpose quasi-Newton methods; its key idea is to approximate the Hessian using curvature information calculated from previous iterates, i.e., $\mathbf{x}_1, \dots, \mathbf{x}_{k-1}$. The L-BFGS modification means that we will use only the most recent m iterates, i.e., $\mathbf{x}_{k-m}, \dots, \mathbf{x}_{k-1}$; the rationale being that too distant iterates are less relevant in estimating the Hessian at \mathbf{x}_k .

In Alg. 1, the matrix $\mathbf{M}/h^2 + \mathbf{L}$ in line 4 can be interpreted as our initial approximation of the Hessian. This matrix is constant which on one hand enables its pre-factorization, but on the other hand, $\mathbf{M}/h^2 + \mathbf{L}$ may be far from the Hessian $\nabla^2 g(\mathbf{x}_k)$, which is the reason for slower convergence compared to Newton’s method [Bouaziz et al. 2014]. L-BFGS allows us to develop a more accurate, state-dependent Hessian approximation, leading to faster convergence without too much computational overhead (in our experiments the overhead is typically less than 1% of the simulation time, see Table 1). The key to fast iterations of L-BFGS is the fact that the progressively updated approximate Hessian \mathbf{A}_k is not stored explicitly, which would require us to solve a new linear system $\mathbf{A}_k \mathbf{d}(\mathbf{x}_k) = -\nabla g(\mathbf{x}_k)$ each iteration, implying high compu-

Algorithm 2: Descent direction computation with L-BFGS

```

1  $\mathbf{q} := -\nabla g(\mathbf{x}_k)$ 
2 for  $i = k - 1, \dots, k - m$  do
3    $\mathbf{s}_i := \mathbf{x}_{i+1} - \mathbf{x}_i; \mathbf{t}_i := \nabla g(\mathbf{x}_{i+1}) - \nabla g(\mathbf{x}_i); \rho_i := \text{tr}(\mathbf{t}_i^\top \mathbf{s}_i)$ 
4    $\zeta_i := \text{tr}(\mathbf{s}_i^\top \mathbf{q}) / \rho_i$ 
5    $\mathbf{q} := \mathbf{q} - \zeta_i \mathbf{t}_i$ 
6 end
7  $\mathbf{r} := \mathbf{A}_0^{-1} \mathbf{q}$  //  $\mathbf{A}_0$  is initial Hessian approximation
8 for  $i = k - m, \dots, k - 1$  do
9    $\eta := \text{tr}(\mathbf{t}_i^\top \mathbf{r}) / \rho_i$ 
10   $\mathbf{r} := \mathbf{r} + \mathbf{s}_i (\zeta_i - \eta)$ 
11 end
12  $\mathbf{d}(\mathbf{x}_k) := \mathbf{r}$  // resulting descent direction

```

tational costs. Instead, L-BFGS implicitly represents the *inverse* of \mathbf{A}_k , i.e., linear operator \mathbf{B}_k such that the desired descent direction can be computed simply as $\mathbf{d}(\mathbf{x}_k) = -\mathbf{B}_k \nabla g(\mathbf{x}_k)$. The linear operator \mathbf{B}_k is not represented using a matrix (which would have been dense), but instead as a sequence of dot products, known as the L-BFGS two-loop recursion, see Alg. 2. For a more detailed discussion of BFGS and its variants we refer to Chapters 6 and 7 of [Nocedal and Wright 2006].

Alg. 2 requires us to provide an initial Hessian approximation \mathbf{A}_0 , ideally such that the linear system $\mathbf{A}_0 \mathbf{r} = \mathbf{q}$ can be solved efficiently (line 7). In our method, we use our old friend: $\mathbf{M}/h^2 + \mathbf{L}$. At first, it may seem the initialization of the Hessian is perhaps not too important and the L-BFGS iterations quickly approach the exact Hessian. However, this intuition is not true. In Section 5 we experiment with different possible initializations of the Hessian and show that our particular choice of $\mathbf{M}/h^2 + \mathbf{L}$ outperforms alternatives such Hessian of the rest-pose and many others. In short, the reason is that the L-BFGS updates use only a very few gradient samples, which provide only a limited amount of information about the exact Hessian. The appeal of the L-BFGS strategy is that it is very fast – the compute cost of the two for-loops in Alg. 2 is negligible compared to the cost of solving the linear system in line 7 with our choice of $\mathbf{A}_0 = \mathbf{M}/h^2 + \mathbf{L}$. This is true even for high values of m . In other words, the linear solve using $\mathbf{M}/h^2 + \mathbf{L}$ (line 7) is still doing the “heavy lifting”, while the L-BFGS updates provide additional convergence boost at the cost of minimal computational overheads.

Upgrading our method with L-BFGS is simple: we only need to replace line 4 in Alg. 1 with a call of Alg. 2. Note that for $m = 0$, Alg. 2 returns exactly the same descent direction as before, i.e., $\mathbf{d}(\mathbf{x}_k) := -(\mathbf{M}/h^2 + \mathbf{L})^{-1} \nabla g(\mathbf{x}_k)$. What is the optimal m , i.e., the size of the history window? Too small m will not allow us to unlock the full potential of L-BFGS. The main problem with too high m is not the higher computational cost of the two loops in Alg. 2, but the fact that too distant iterates (such as \mathbf{x}_{k-100}) may contain information irrelevant for the Hessian at \mathbf{x}_k and the result can be even worse than with a shorter window. We found that $m = 5$ is typically a good value in our experiments.

The recently proposed Chebyshev Semi-Iterative methods for Projective Dynamics [Wang 2015] can also be interpreted as a special type of a quasi-Newton method which utilizes two previous iterates, i.e., corresponding to $m = 2$. Indeed, in our experiments L-BFGS with $m = 2$ exhibits similar convergence rate as the Chebyshev method, see Figure 7 and further discussion in Section 5. Finally, we note that even though the Wolfe conditions are the recommended line search strategy for L-BFGS, we did not observe any significant convergence benefit compared to our backtracking




Figure 4: Our method is capable of simulating complex collision scenarios, such as squeezing the Big Bunny through a torus. The Big Bunny uses corotated elasticity with $\mu = 5$ and $\lambda = 200$.

strategy. However, evaluating the Wolfe conditions increases the computational cost per iteration and therefore, we continue to rely on the backtracking strategy as described in Alg. 1.

4.3 Collisions

A classical approach to enforcing non-penetration constraints between deformable solids is to 1) detect collisions and 2) resolve them using temporarily instanced repulsion springs, which bring the volume of undesired overlaps to zero [McAdams et al. 2011; Harmon et al. 2011]. However, in Projective Dynamics the primary emphasis is on computational efficiency and therefore only simplified collision resolution strategies have been proposed by Bouaziz et al. [2014]. Specifically, Projective Dynamics offers two possible strategies. The first strategy is a two-phase method, where collisions are resolved in a separate post-processing step using projections, similar to Position Based Dynamics. The same strategy was employed also by Liu et al. [2013]. The drawback of this approach is the fact the collision projections are oblivious to elasticity and inertia of the simulated objects. The second approach used in Projective Dynamics is more physically realistic, but introduces additional computational overheads. Specifically, temporarily-instanced repulsion springs are added to the total energy. This leads to physically realistic results, but the drawback is that the matrix $\mathbf{M}/h^2 + \mathbf{L}$ needs to be re-factorized whenever the set of repulsion springs is updated – typically, at the beginning of each frame.

Our quasi-Newton interpretation invites a new approach to collision response which is physically realistic, but avoids expensive re-factorizations. Specifically, for each inter-penetration found by collision detection, we introduce an energy term $E_{\text{collision}}(\mathbf{x}) = ((\mathbf{S}\mathbf{x} - \mathbf{t})^T \mathbf{n})^2$, where \mathbf{S} is a selector matrix of the collided vertex, \mathbf{t} is its projection on the surface and \mathbf{n} is the surface normal. This constraint pushes the collided vertex to the tangent plane. It is important to add this term to our total energy $E(\mathbf{x})$ only if the vertex is in collision or contact. Whenever the relative velocity between the vertex and the collider indicates separation, the $E_{\text{collision}}(\mathbf{x})$ term is discarded (otherwise it would correspond to unrealistic “glue” forces). This is done once at the beginning of each iteration (just before line 3 in Alg. 1). The line search (lines 6-10 of Alg. 1) is unaffected by these updates, i.e., the unilateral nature of the collision constraints is handled correctly without any further processing.

The key idea of our approach is to leverage the quasi-Newton approximation for collision processing. In particular, we account for the $E_{\text{collision}}(\mathbf{x})$ terms when evaluating the energy and its gradients, but we ignore their contributions to the $\mathbf{M}/h^2 + \mathbf{L}$ matrix. This means that we form a somewhat more aggressive approximation of the Hessian, with the benefit that the system matrix will never need to be re-factorized. The line search process (lines 6-10 in Alg. 1) guarantees that energy will decrease in spite of this more aggressive approximation. The only trade-off we observed in our experiments is that the number of line search iterations may increase, which is a small cost to pay for avoiding re-factorizations. We observed that even in challenging collision scenarios, such as when squeezing a

Big Bunny through a torus, the approach behaves robustly and successfully resolves all collisions, see Figure 4.

4.4 Anisotropy

Our numerical methods, including the L-BFGS acceleration, can be directly applied also to anisotropic material models. We verified this on an elastic cube model with corotated base material ($\mu = 10$, $\lambda = 100$, referring to the notation of Sifakis and Barbič [2012]) enhanced with additional anisotropic stiffness term $\frac{\kappa}{2} (\|\mathbf{F}\mathbf{d}\| - 1)^2$, where \mathbf{F} is the deformation gradient and \mathbf{d} is the (rest-pose) direction of anisotropy. This corresponds to the directional reinforcement of the material which is common, e.g., in biological soft tissues containing collagenous fibers. The result of our method with $\kappa = 50$ can be seen in Figure 5.




Figure 5: Dropping an elastic cube on the ground. Left: deformation using isotropic elasticity (linear corotated model). Right: the result after adding anisotropic stiffness.

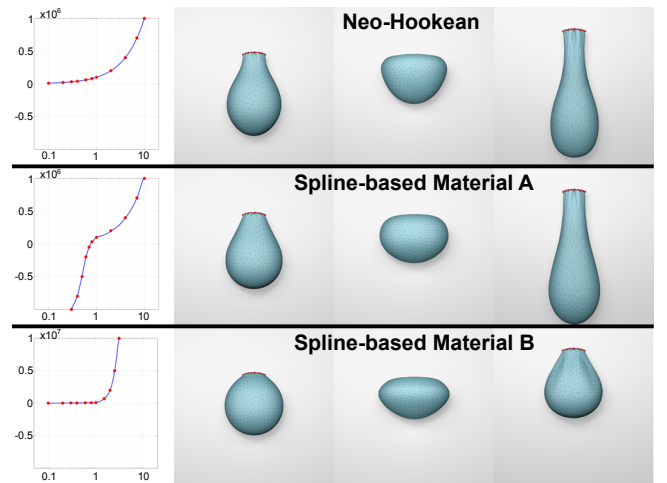


Figure 6: Elastic sphere with spline-based materials [Xu et al. 2015], simulated using our method. Spline-based material A is a modified Neo-Hookean material that resists compression more; material B is a modified Neo-Hookean material that resists tension more. The strain-stress curves are shown on the left.

5 Results

Our method supports standard elastic materials, such as corotated linear elasticity, St. Venant-Kirchhoff and the Neo-Hookean model,

model	#ver.	#ele.	material model	our method (10 iterations)				Newton (1 iteration)	
				linesearch iterations	L-BFGS overhead	per-frame time	relative error	per-frame time	relative error
Thin sheet	660	1932	Polynomial	10.8	0.026 ms	4.4 ms	2.7×10^{-8}	184 ms	8.8×10^{-4}
Sphere	889	1821	Spline-based A	24.5 [†]	0.155 ms	21.2 ms	2.7×10^{-7}	188 ms	6.9×10^{-4}
Sphere	889	1821	Spline-based B	21.8	0.156 ms	19.7 ms	6.9×10^{-6}	187 ms	2.5×10^{-4}
Shaking bar	574	1647	Corotated	10.1	0.193 ms	7.2 ms	1.6×10^{-4}	171 ms	4.4×10^{-3}
Ditto	1454	4140	Neo-Hookean	11.7	0.203 ms	17.8 ms	3.0×10^{-5}	305 ms	1.6×10^{-3}
Hippo	2387	8406	Corotated	11.9	0.555 ms	40.6 ms	2.2×10^{-3}	640 ms	3.7×10^{-2}
Twisting bar	3472	10441	Neo-Hookean	10.6	0.945 ms	45.6 ms	9.4×10^{-5}	681 ms	7.9×10^{-3}
Cloth	6561	32160	Mass-Springs	10.0	1.20 ms	42.3 ms	9.3×10^{-4}	798 ms	1.2×10^{-2}
Big Bunny	6308	26096	Corotated	49.2 [‡]	2.19 ms	623 ms	9.8×10^{-2}	2700 ms	2.8×10^{-1}
Squirrel	8395	23782	Polynomial	10.7	1.41 ms	153 ms	8.3×10^{-8}	2400 ms	9.1×10^{-6}
Squirrel	33666	125677	Polynomial	10.5	6.38 ms	706 ms	1.5×10^{-5}	15800 ms	5.4×10^{-5}

Table 1: In all examples, we execute 10 iterations of our method per frame, accelerated with L-BFGS with history window $m = 5$. Newton’s method uses 1 iteration per frame. The “linesearch iterations” reports the average number of line search iterations per frame. The “L-BFGS overhead” is the runtime overhead of L-BFGS, i.e., timing of Alg. 2 without line 7 ($m = 5$). The reported per-frame time for our method accounts for all 10 iterations. One iteration of our method is approximately 100 times faster than one iteration of Newton’s method. We use 10 iterations of our method which reduce the error more than one iteration of Newton’s method, while being about 10 times faster. [†]The higher number of line search iterations is due to the high nonlinearity of the spline-based materials and large deformations of the sphere. [‡]In this case, the higher number of line search iterations is caused by nonlinearities due to collisions (Section 4.3).




Figure 7: Convergence of our method with different L-BFGS history settings, compared to Chebyshev Semi-Iterative method and Newton’s method (baseline). The model is “Twisting bar” with Neo-Hookean elasticity.

see Figure 1. None of these materials is supported by Projective Dynamics (note that Projective Dynamics supports a special sub-class of corotated linear materials, specifically, ones with $\lambda = 0$). Our method also supports the recently introduced spline-based materials proposed by Xu et al. [2015], as shown in Figure 1 and Figure 6.

Table 1 reports our testing scenarios and compares the run time of our method with Newton’s method, both executed on an Intel i7-4910MQ CPU at 2.90GHz. All scenarios are produced with a fixed timestep of 1/30 seconds. Because Newton’s method is not guaranteed to work with indefinite Hessians, we employ the standard definiteness fix [Teran et al. 2005], i.e., we project the Hessian of each element to its closest positive definite component. We found this method works better than other definiteness fixes, such as adding a multiple of the identity matrix [Martin et al. 2011], which affects the *entire* simulation even if there are just a few problematic elements. The approximately 100 times faster run-time of one iteration of our method compared to one iteration of Newton’s method is due to the following facts: 1) we use pre-computed sparse Cholesky factorization, because our matrix $\mathbf{M}/h^2 + \mathbf{L}$ is constant, 2) the size of our matrix is $n \times n$, whereas the Hessian used in Newton’s method is a $3n \times 3n$ matrix, i.e., the x, y, z coordinates

are no longer decoupled, 3) the computation of SVD derivatives, necessary to evaluate the Hessians of materials based on principal stretches [Xu et al. 2015], is expensive. Note that our method is also simpler to implement, as no SVD derivatives or definiteness fixes are necessary.

Comparison to Chebyshev Semi-Iterative method. We compared the convergence of our method with various lengths of the L-BFGS window to the recently introduced Chebyshev Semi-Iterative method [Wang 2015]. We also plot results obtained with Newton’s method as a baseline, see Figure 7.

Even though the Chebyshev method was originally proposed only for Projective Dynamics energies, our generalization to arbitrary materials is compatible with the Chebyshev Semi-Iterative acceleration, see Alg. 3. The Alg. 3 computes a descent direction which can be used in line 4 of Alg. 1. As discussed by Wang [2015], the Chebyshev acceleration should be disabled during the first S iterations, where the recommended value is $S = 10$. Another parameter which is essential for the Chebyshev method is an estimate of spectral radius ρ , which is calculated from training simulations [Wang 2015]. This parameter must be estimated carefully, because underestimated ρ can lead to the Chebyshev method producing *ascent* di-




Figure 8: Convergence comparison of L-BFGS methods (all using $m = 5$) initialized with different Hessian approximations, along with Newton’s method (baseline). The model is “Twisting bar” with Neo-Hookean elasticity.

rections (as opposed to descent directions). Without line search, the ascent directions manifest themselves as oscillations [Wang 2015]. For the purpose of comparisons, we implemented the Chebyshev method with direct solver which is the fastest method on the CPU [Wang 2015].

Algorithm 3: Descent direction computation using Chebyshev Semi-Iterative Method [Wang 2015]

```

1 // S ... Chebyshev disabled for the first S iterates, default S = 10
2 // ρ ... approximated spectral radius
3  $\mathbf{q} := -(\mathbf{M}/h^2 + \mathbf{L})^{-1} \nabla g(\mathbf{x}_k)$ 
4  $\hat{\mathbf{x}}_{k+1} := \mathbf{x}_k + \mathbf{q}$ 
5 if  $k < S$  then  $\omega_{k+1} := 1$  if  $k = S$  then  $\omega_{k+1} := 2/(2 - \rho^2)$  if
    $k > S$  then  $\omega_{k+1} := 4/(4 - \rho^2 \omega_k)$ 
6  $\mathbf{d}(\mathbf{x}_k) := \omega_{k+1}(\hat{\mathbf{x}}_{k+1} - \mathbf{x}_{k-1}) + \mathbf{x}_{k-1} - \mathbf{x}_k$ 

```

We compare the convergence of all methods using relative error, defined as:

$$\frac{g(\mathbf{x}_k) - g(\mathbf{x}^*)}{g(\mathbf{x}_0) - g(\mathbf{x}^*)} \quad (11)$$

where \mathbf{x}_0 is the initial guess (we use $\mathbf{x}_0 := \mathbf{y}$ for all methods), \mathbf{x}_k is the k -th iterate, and \mathbf{x}^* is the exact solution computed using Newton’s method (iterated until convergence). The decrease of relative error for one example frame is shown in Figure 7, where all methods are using the backtracking line search outlined in Alg. 1. As expected, descent directions computed using Newton’s method are the most effective ones, as can be seen in Figure 7 (right). However, each iteration of Newton’s method is computationally expensive, and therefore other methods can realize faster error reduction with respect to computational time, as shown in Figure 7 (left). All of the remaining methods are based on the constant Hessian approximation $\mathbf{M}/h^2 + \mathbf{L}$ which leads to much faster convergence. Out of these methods, classical Projective Dynamics converges slowest. The Chebyshev Semi-Iterative method improves the convergence; we also confirmed that disabling the Chebyshev method during the first 10 iterations indeed helps, as recommended by Wang [2015]. Our method aided with L-BFGS improves convergence even further. Already with $m = 2$ (where m is the size of the history window), we obtain slightly faster convergence than with the Chebyshev method. One reason is that it is not necessary to disable L-BFGS in the first several iterates, because L-BFGS is effective as soon as the previous iterates become available. Also, we do not have to estimate the spectral radius which is required by the Chebyshev

shev method. With L-BFGS, we can also increase the history window, e.g., to $m = 5$, obtaining even more rapid convergence.

L-BFGS with different initial Hessian estimates. Our method can be interpreted as providing a particularly good initial estimate of the Hessian for L-BFGS. Therefore, it is important to compare to other possible Hessian initializations. In a general setting, Nocedal and Wright [2006] recommend to bootstrap L-BFGS using a scaled identity matrix:

$$\mathbf{A}_0 := \frac{\text{tr}(\mathbf{s}_{k-1}^T \mathbf{y}_{k-1})}{\text{tr}(\mathbf{y}_{k-1}^T \mathbf{y}_{k-1})} \mathbf{I} \quad (12)$$

We experimented with this approach, but we found that our choice $\mathbf{A}_0 := \mathbf{M}/h^2 + \mathbf{L}$ leads to much faster convergence, trumping the computational overhead associated with solving the pre-factorized system $\mathbf{A}_0 \mathbf{r} = \mathbf{q}$ (see Figure 8, red graph).

Another possibility would be to set \mathbf{A}_0 equal to the rest pose Hessian, which can of course also be pre-factorized. As shown in Figure 8 (yellow graph), this is a slightly better approximation than scaled identity, but still not very effective. This is because the actual Hessian depends on world-space rotations of the model, deviating significantly from the rest-pose Hessian. This issue was observed by Müller et al. [2002], who proposed per-vertex *stiffness warping* as a possible remedy. Per-vertex stiffness warping still allows us to leverage pre-factorization of the rest-pose Hessian and results in better convergence than pure rest-pose Hessian, see Figure 8 (purple graph). However, per-vertex stiffness warping may introduce ghost forces, because stiffness warping uses different rotation matrices for each vertex, which means that internal forces in one element no longer have to sum to zero. The ghost forces disappear in a fully converged solution, however, this would require a prohibitively high number of iterations.

Yet another possibility is to completely re-evaluate the Hessian at the beginning of each frame. This requires re-factorization, however, the remaining 10 (or so) iterations can reuse the factorization, relying only on L-BFGS updates. When measuring convergence with respect to number of iterations, this approach is very effective, as shown in Figure 8 (right, green graph). However, the cost of the initial Hessian factorization is significant, as obvious from Figure 8 (left, green graph). Our method uses the same Hessian factorization for all frames, avoiding the per-frame factorization costs, while featuring excellent convergence properties, see Figure 8 (blue graph).

The overheads of per-frame Hessian factorizations can be mitigated by carefully scheduled Hessian updates. In particular, the Hessian can be reused for multiple subsequent frames if the state is




Figure 9: Simulation of a bar with corotated elasticity, constrained in the middle and rapidly shaken. The method of Hecht et al. [2012] with full Hessian updates every other frame explodes due to large ghost forces (top). Our method does not introduce any ghost forces and remains stable (bottom).

not changing too much [Deuffhard 2011]. Assuming the corotated elastic model, Hecht et al [2012] push this idea even further by proposing a warp-cancelling form of the Hessian which allows not only for *temporal* schedule, but also for *spatially* localized updates. Specifically, a nested dissection tree allows for recomputing only parts of the mesh, which is particularly advantageous in situations where only small part of the object is undergoing large deformations. However, the updates are still costly, and the frequency of the updates depends on the simulation. Similarly to per-vertex stiffness warping, insufficiently frequent update may produce ghost forces and consequent instabilities. This can be a problem when simulating quickly moving elastic objects. To illustrate this, in Figure 9 we show a simulation of shaking an elastic bar. Even if we schedule the Hessian updates every other frame and recompute the entire domain, this method still generates too large ghost forces and becomes unstable. In contrast, our method remains stable and does not require *any* run-time Hessian updates.

Comparison to Projective Dynamics. One possible alternative to our method would be to apply regular Projective Dynamics with additional strain-limiting constraints [Bouaziz et al. 2014], enabling us to construct piece-wise linear approximations of the strain-stress curves of more general materials. We tried to use this approach to approximate the polynomial material ($a(x) = \mu(x - 1)^4, b(x) = 0, c(x) = 0$) discussed in Section 4.1, see Figure 10. Even though we obtain similar overall behavior, there are two types of artifacts associated with this approximation. First, the strain-limiting constraints introduce damping when they are *not* activated. This is because the projection terms still exist in our constant matrix $\mathbf{M}/h^2 + \mathbf{L}$; if the strain-limiting is not activated, the deformation gradients project to their current values, which produces the undesired damping. The second problem is due to the non-smooth nature of the piece-wise linear approximation, i.e., the stiffness of the simulated object is abruptly changed when the strain-limiting constraints become activated. As shown in the accompanying video, our method avoids both of these issues.

The L-BFGS acceleration benefits also simulations which use only Projective Dynamics materials (Eq. 1). The most elementary example of these materials are mass-spring systems. In Figure 11, we can see that the L-BFGS acceleration applied to a mass-spring system simulation results in more realistic wrinkles.

Choice of L-BFGS history window size. Which history window (m) is the best? We experimented with different values of m , see




Figure 10: The strain-stress curve of a polynomial material can be approximated piece-wise linearly with two Projective Dynamics constraints.




Figure 11: Mass-spring system simulation using our method with L-BFGS (left) and without, i.e., using pure Projective Dynamics (right). The L-BFGS acceleration results in more realistic wrinkles.

Figure 13. Too large m takes into account too distant iterates which can lead to worse approximation of the Hessian. In Figure 13, we see the optimal value is $m = 5$, which is also our recommended default setting. However, it is comforting that the algorithm is not particularly sensitive to the setting of m – even large values such as $m = 100$ produce only slightly worse convergence. In Figure 7 we can notice that the convergence rate of the Chebyshev method is similar to our method with L-BFGS using $m = 2$. We believe this is not a coincidence, because the Chebyshev method uses *two* previous iterates, just like L-BFGS with $m = 2$.

Choice of stiffness parameters. As discussed in Section 4.1, we Eq. 10 and define our stiffness parameter k as the slope of the best linear approximation of Eq. 10 for $\sigma_1 \in [x_{\text{start}}, x_{\text{end}}]$. What is the best $[x_{\text{start}}, x_{\text{end}}]$ interval to use? In the limit, with $[x_{\text{start}}, x_{\text{end}}] \rightarrow [1, 1]$, our k would converge to the second derivative. However, a finite interval $[x_{\text{start}}, x_{\text{end}}]$ guarantees that our k is meaningful even for materials such as the polynomial material $a(x) = \mu(x - 1)^4, b(x) = 0, c(x) = 0$; in this case, we obtain a k which depends linearly on μ . We argue the convergence of our algorithm is not very sensitive to a particular choice of the $[x_{\text{start}}, x_{\text{end}}]$ interval. In Figure 14, we show convergence graphs of a twisting bar with Neo-Hookean material using different intervals to compute the stiffness parameter k . Although Neo-Hookean material is highly non-linear, the convergence rates for different interval choices are quite similar. Therefore, we decided not to investigate more sophisticated strategies and we set $x_{\text{start}} = 0.5, x_{\text{end}} = 1.5$ in all of our simulations.

Comparison with iterative solvers. Sparse iterative solvers do not require expensive factorizations and are therefore attractive in interactive applications. A particularly popular iterative method are Conjugate gradients (CG) [Shewchuk 1994]. An additional ad-




Figure 12: Convergence comparison of various methods using sparse direct solvers and conjugate gradients. The model is “Twisting bar” with Neo-Hookean elasticity.

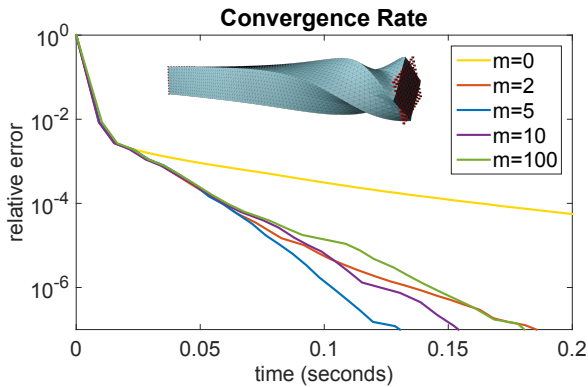


Figure 13: Comparison of L-BFGS convergence rate with different history window sizes (m).




Figure 14: The convergence rate for different stiffness parameters chosen from different regions.

vantage is that CG can be implemented in a matrix-free fashion, i.e., without explicitly forming the sparse system matrix. Gast et al. [2015] further accelerate the CG solver used in Newton’s method by proposing a CG-friendly definiteness fix. Specifically, the CG iterations are terminated whenever the maximum number of iteration is reached or indefiniteness of the Hessian matrix is detected.

While iterative methods can be the only possible choice in high-resolution simulations (e.g., in scientific computing), in real-time simulation scales, sparse direct solvers with pre-computed factorization are hard to beat, as we show in Figure 12. Specifically, we test Newton’s method with linear systems solved using CG with 5 and 15 iterations, using Jacobi preconditioner. Even with 15 CG iterations, the accuracy is still not the same as with the direct solver the computational cost becomes high. If we use only five CG iterations the running time improves, but the convergence rate suffers because the descent directions are not sufficiently effective. The method of Gast et al. [2015] initially outperforms Newton with CG, however, the convergence slows down in subsequent iterations. We also tried to apply CG to our method, in lieu of the direct solver. With 15 CG iterations the convergence is competitive, however, the CG solver is slower.

Robustness. We demonstrate that our proposed extensions to more general materials and the L-BFGS solver upgrade do not compromise simulation robustness. In Figure 15, we show an elastic hippo




Figure 15: Our method is robust despite extreme initial conditions: a randomly initialized hippo returns back to its rest pose.

which recovers from an extreme (randomized) deformation with many inverted elements. Specifically, the hippo model uses L-BFGS with $m = 5$ and corotated linear elasticity with $\mu = 20$ and $\lambda = 100$ (note that Projective Dynamics supports only corotated materials with $\lambda = 0$).

6 Limitations and Future Work

Our method is currently limited only to hyperelastic materials satisfying the Valanis-Landel assumption. Even though this assumption covers many practical models, including the recently proposed spline-based materials [Xu et al. 2015], it would be interesting to study the further generalization of our method. Perhaps even more interesting would be to remove the assumption of hyperelasticity. Can we develop fast algorithms for simulating non-hyperelastic materials, including the effects such as relaxation, creep, and hys-

teresis [Bargteil et al. 2007]? Inspired by the recent work of Wang [2015], we would like to explore GPU implementations of physics-based simulations. Our current method is derived from the Implicit Euler time integration method and therefore inherits its artificial damping drawbacks. We experimented with Implicit Midpoint – a symplectic integrator which does not suffer from this problem. However, we found that Implicit Midpoint is much less stable. In the future we would like to explore fast numerical solvers for symplectic yet stable integration methods. Finally, we plan to investigate specific physics-based applications which require *both* high accuracy and speed, such as interactive surgery simulation.

7 Conclusions

We have presented a method for fast physics-based simulation of a large class of hyperelastic materials. The key to our approach is the insight that Projective Dynamics [Bouaziz et al. 2014] can be re-formulated as a quasi-Newton method. Aided with line search, we obtain a robust simulator supporting many practical material models. Our quasi-Newton formulation also allows us to further accelerate convergence by combining our method with L-BFGS. Even though L-BFGS is sensitive to initial Hessian approximation, our method suggests a particularly effective Hessian initialization which yields fast convergence. Most of our experiments use ten iterations of our method which is typically more accurate than one iteration of Newton’s method, while being about ten times faster and easier to implement. Traditionally, real-time physics is considered to be approximate but fast, while off-line physics is accurate but slow. We hope that our method will help to blur the boundaries between real-time and off-line physics-based animation.

References

- AN, S. S., KIM, T., AND JAMES, D. L. 2008. Optimizing cubature for efficient integration of subspace deformations. *ACM Trans. Graph.*
- BARBIČ, J., AND JAMES, D. L. 2005. Real-time subspace integration for st. venant-kirchhoff deformable models. *ACM Trans. Graph.*
- BARGTEIL, A. W., WOJTAN, C., HODGINS, J. K., AND TURK, G. 2007. A finite element method for animating large viscoplastic flow. *ACM Trans. Graph.*
- BATHE, K. J., AND CIMENTO, A. P. 1980. Some practical procedures for the solution of nonlinear finite element equations. *Computer Methods in Applied Mechanics and Engineering.*
- BENDER, J., KOSCHIER, D., CHARRIER, P., AND WEBER, D. 2014. Position-based simulation of continuous materials. *Computers & Graphics.*
- BENDER, J., MÜLLER, M., OTADUY, M. A., TESCHNER, M., AND MACKLIN, M. 2014. A survey on position-based simulation methods in computer graphics. In *Comput. Graph. Forum.*
- BOUAZIZ, S., DEUSS, M., SCHWARTZBURG, Y., WEISE, T., AND PAULY, M. 2012. Shape-up: Shaping discrete geometry with projections. In *Comput. Graph. Forum.*
- BOUAZIZ, S., MARTIN, S., LIU, T., KAVAN, L., AND PAULY, M. 2014. Projective dynamics: Fusing constraint projections for fast simulation. *ACM Trans. Graph.*
- CHAO, I., PINKALL, U., SANAN, P., AND SCHRÖDER, P. 2010. A simple geometric model for elastic deformations. *ACM Trans. Graph.*
- CHEN, D., LEVIN, D., SUEDA, S., AND MATUSIK, W. 2015. Data-driven finite elements for geometry and material design. *ACM Trans. Graph.*
- DESBRUN, M., SCHRÖDER, P., AND BARR, A. 1999. Interactive animation of structured deformable objects. *Graphics Interface.*
- DEUFLHARD, P. 2011. *Newton methods for nonlinear problems: affine invariance and adaptive algorithms.* Springer Science & Business Media.
- FISH, J., PANDHEERADI, M., AND BELSKY, V. 1995. An efficient multilevel solution scheme for large scale non-linear systems. *International Journal for Numerical Methods in Engineering.*
- FRATARCANGELI, M., AND PELLACINI, F. 2015. Scalable partitioning for parallel position based dynamics. *Comput. Graph. Forum.*
- GAST, T. F., SCHROEDER, C., STOMAKHIN, A., JIANG, C., AND TERAN, J. M. 2015. Optimization integrator for large time steps. *Visualization and Comp. Graph., IEEE Transactions on.*
- GEORGII, J., AND WESTERMANN, R. 2006. A multigrid framework for real-time simulation of deformable bodies. *Computers & Graphics 30, 3, 408–415.*
- GOLDENTHAL, R., HARMON, D., FATTAL, R., BERCOVIER, M., AND GRINSPUN, E. 2007. Efficient simulation of inextensible cloth. *ACM Trans. Graph.*
- HAHN, F., MARTIN, S., THOMASZEWSKI, B., SUMNER, R., COROS, S., AND GROSS, M. 2012. Rig-space physics. *ACM Trans. Graph.*
- HAIRER, E., LUBICH, C., AND WANNER, G. 2002. *Geometric Numerical Integration: Structure-Preserving Algorithms for Ordinary Differential Equations.* Springer.
- HARMON, D., AND ZORIN, D. 2013. Subspace integration with local deformations. *ACM Trans. Graph.*
- HARMON, D., PANOZZO, D., SORKINE, O., AND ZORIN, D. 2011. Interference-aware geometric modeling. In *ACM Trans. Graph.*
- HAUTH, M., AND ETZMUSS, O. 2001. A high performance solver for the animation of deformable objects using advanced numerical methods. In *Comput. Graph. Forum.*
- HECHT, F., LEE, Y. J., SHEWCHUK, J. R., AND O’BRIEN, J. F. 2012. Updated sparse cholesky factors for corotational elastodynamics. *ACM Trans. Graph.*
- IRVING, G., TERAN, J., AND FEDKIW, R. 2004. Invertible finite elements for robust simulation of large deformation. In *Proc. EG Symp. Computer Animation.*
- KHAREVYCH, L., YANG, W., TONG, Y., KANSO, E., MARSDEN, J. E., SCHRÖDER, P., AND DESBRUN, M. 2006. Geometric, variational integrators for computer animation. In *Proc. EG Symp. Computer Animation.*
- KIM, T.-Y., CHENTANEZ, N., AND MÜLLER-FISCHER, M. 2012. Long range attachments-a method to simulate inextensible clothing in computer games. In *Proc. EG Symp. Computer Animation.*
- LI, S., HUANG, J., DE GOES, F., JIN, X., BAO, H., AND DESBRUN, M. 2014. Space-time editing of elastic motion through material optimization and reduction. *ACM Trans. Graph.*
- LIU, T., BARGTEIL, A. W., O’BRIEN, J. F., AND KAVAN, L. 2013. Fast simulation of mass-spring systems. *ACM Trans. Graph.*

- MACKLIN, M., AND MÜLLER, M. 2013. Position based fluids. *ACM Trans. Graph.*
- MACKLIN, M., MÜLLER, M., CHENTANEZ, N., AND KIM, T.-Y. 2014. Unified particle physics for real-time applications. *ACM Trans. Graph.*
- MARTIN, S., THOMASZEWSKI, B., GRINSPUN, E., AND GROSS, M. 2011. Example-based elastic materials. In *ACM Trans. Graph.*
- MARTIN, T., JOSHI, P., BERGOU, M., AND CARR, N. 2013. Efficient non-linear optimization via multi-scale gradient filtering. In *Comput. Graph. Forum.*
- MCADAMS, A., ZHU, Y., SELLE, A., EMPEY, M., TAMSTORF, R., TERAN, J., AND SIFAKIS, E. 2011. Efficient elasticity for character skinning with contact and collisions. *ACM Trans. Graph.*
- MÜLLER, M., AND CHENTANEZ, N. 2011. Solid simulation with oriented particles. In *ACM Trans. Graph.*
- MÜLLER, M., AND GROSS, M. 2004. Interactive virtual materials. In *Proceedings of Graphics Interface.*
- MÜLLER, M., DORSEY, J., MCMILLAN, L., JAGNOW, R., AND CUTLER, B. 2002. Stable real-time deformations. In *Proc. EG Symp. Computer Animation.*
- MÜLLER, M., HEIDELBERGER, B., TESCHNER, M., AND GROSS, M. 2005. Meshless deformations based on shape matching. In *ACM Trans. Graph.*
- MÜLLER, M., HEIDELBERGER, B., HENNIX, M., AND RATCLIFF, J. 2007. Position based dynamics. *J. Vis. Comun. Image Represent.*
- MÜLLER, M., CHENTANEZ, N., KIM, T.-Y., AND MACKLIN, M. 2014. Strain based dynamics. In *Proc. EG Symp. Computer Animation*, vol. 2.
- MÜLLER, M., CHENTANEZ, N., KIM, T.-Y., AND MACKLIN, M. 2015. Air meshes for robust collision handling. *ACM Trans. Graph.*
- MÜLLER, M. 2008. Hierarchical Position Based Dynamics. In *Workshop in Virtual Reality Interactions and Physical Simulation "VRIPHYS" (2008)*, The Eurographics Association.
- NARAIN, R., SAMII, A., AND O'BRIEN, J. F. 2012. Adaptive anisotropic remeshing for cloth simulation. *ACM Trans. Graph.*
- NEUBERGER, J. 1983. Steepest descent for general systems of linear differential equations in hilbert space. In *Ordinary differential equations and operators*. Springer.
- NEUBERGER, J. 2009. *Sobolev gradients and differential equations*. Springer Science & Business Media.
- NOCEDAL, J., AND WRIGHT, S. J. 2006. *Numerical optimization*. Springer Verlag.
- RIVERS, A., AND JAMES, D. 2007. FastLSM: fast lattice shape matching for robust real-time deformation. *ACM Trans. Graph.*
- SERVIN, M., LACOURSIÈRE, C., AND MELIN, N. 2006. Interactive simulation of elastic deformable materials. In *SIGRAD*.
- SHEWCHUK, J. R., 1994. An introduction to the conjugate gradient method without the agonizing pain.
- SIFAKIS, E., AND BARBIČ, J. 2012. Fem simulation of 3d deformable solids: A practitioner's guide to theory, discretization and model reduction. In *ACM SIGGRAPH Courses*.
- STAM, J. 2009. Nucleus: towards a unified dynamics solver for computer graphics. In *IEEE Int. Conf. on CAD and Comput. Graph.*
- TAMSTORF, R., JONES, T., AND MCCORMICK, S. F. 2015. Smoothed aggregation multigrid for cloth simulation. *ACM Trans. Graph.*
- TENG, Y., OTADUY, M. A., AND KIM, T. 2014. Simulating articulated subspace self-contact. *ACM Trans. Graph.*
- TENG, Y., MEYER, M., DEROSE, T., AND KIM, T. 2015. Subspace condensation: Full space adaptivity for subspace deformations. *ACM Trans. Graph.*
- TERAN, J., SIFAKIS, E., IRVING, G., AND FEDKIW, R. 2005. Robust quasistatic finite elements and flesh simulation. In *Proc. EG Symp. Computer Animation*.
- TERZOPOULOS, D., PLATT, J., BARR, A., AND FLEISCHER, K. 1987. Elastically deformable models. In *Computer Graphics (Proceedings of SIGGRAPH)*.
- THOMASZEWSKI, B., PABST, S., AND STRASSER, W. 2009. Continuum-based strain limiting. In *Comput. Graph. Forum*.
- TOURNIER, M., NESME, M., GILLES, B., AND FAURE, F. 2015. Stable constrained dynamics. *ACM Trans. Graph.*
- VALANIS, K., AND LANDEL, R. 1967. The strain-energy function of a hyperelastic material in terms of the extension ratios. *Journal of Applied Physics*.
- WANG, H., O'BRIEN, J., AND RAMAMOORTHI, R. 2010. Multi-resolution isotropic strain limiting. *ACM Trans. Graph.*
- WANG, H. 2015. A chebyshev semi-iterative approach for accelerating projective and position-based dynamics. *ACM Trans. Graph.*
- XU, H., SIN, F., ZHU, Y., AND BARBIČ, J. 2015. Nonlinear material design using principal stretches. *ACM Trans. Graph.*

Appendix

In this Appendix we compute the gradient $\nabla_{\mathbf{x}} E_i(\mathbf{x})$, where $E_i(\mathbf{x})$ is defined according to Eq. 1. Dropping the subscript i for clarity, we define the projection:

$$\mathbf{p}(\mathbf{x}) = \operatorname{argmin}_{\mathbf{z} \in \mathcal{M}} \|\mathbf{G}\mathbf{x} - \mathbf{z}\|^2 \quad (13)$$

where \mathbf{G} represents a discrete differential operator and \mathcal{M} is a constraint manifold. We need to compute the differential:

$$\begin{aligned} \frac{1}{2} \delta \|\mathbf{G}\mathbf{x} - \mathbf{p}(\mathbf{x})\|^2 &= (\mathbf{G}\mathbf{x} - \mathbf{p}(\mathbf{x}))^\top (\mathbf{G}\delta\mathbf{x} - \delta\mathbf{p}(\mathbf{x})) \quad (14) \\ &= (\mathbf{G}\mathbf{x} - \mathbf{p}(\mathbf{x}))^\top \mathbf{G}\delta\mathbf{x} \quad (15) \end{aligned}$$

because the second term $(\mathbf{G}\mathbf{x} - \mathbf{p}(\mathbf{x}))^\top \delta\mathbf{p}(\mathbf{x})$ vanishes. This is due to the fact that $\delta\mathbf{p}(\mathbf{x}) \in T_{\mathbf{p}(\mathbf{x})}\mathcal{M}$, where $T_{\mathbf{p}(\mathbf{x})}\mathcal{M}$ denotes the tangent space at point $\mathbf{p}(\mathbf{x}) \in \mathcal{M}$. The vector $\mathbf{G}\mathbf{x} - \mathbf{p}(\mathbf{x})$ must be orthogonal to $T_{\mathbf{p}(\mathbf{x})}\mathcal{M}$, otherwise $\mathbf{p}(\mathbf{x})$ could not be the minimizer according to Eq. 13.