

Optimization Integrator for Large Time Steps

Theodore F. Gast, Craig Schroeder, Alexey Stomakhin, Chenfanfu Jiang, and Joseph M. Teran

Abstract—Practical time steps in today’s state-of-the-art simulators typically rely on Newton’s method to solve large systems of nonlinear equations. In practice, this works well for small time steps but is unreliable at large time steps at or near the frame rate, particularly for difficult or stiff simulations. We show that recasting backward Euler as a minimization problem allows Newton’s method to be stabilized by standard optimization techniques with some novel improvements of our own. The resulting solver is capable of solving even the toughest simulations at the 24 Hz frame rate and beyond. We show how simple collisions can be incorporated directly into the solver through constrained minimization without sacrificing efficiency. We also present novel penalty collision formulations for self collisions and collisions against scripted bodies designed for the unique demands of this solver. Finally, we show that these techniques improve the behavior of Material Point Method (MPM) simulations by recasting it as an optimization problem.

Index Terms—Computer graphics, three-dimensional graphics and realism, animation

1 INTRODUCTION

THE most commonly used time integration schemes in use today for graphics applications are implicit methods. Among these, backward Euler [1], [2], [3], [4], [5] or variants on Newmark methods [6], [7], [8] are the most common, though even more sophisticated schemes like BDF-2 [9], [10], implicit-explicit schemes [11], [12], or even the more exotic exponential integrators [13] have received consideration. Integrators have been the subject of comparison before (see for example [3], [9], [14]), seeking good compromises between speed, accuracy, robustness, and dynamic behavior.

These integrators require the solution to one or more nonlinear systems of equations each time step. These systems are typically solved by some variation on Newton’s method. Even the most stable simulators are typically run several time steps per 24 Hz frame of simulation. There is growing interest in running simulations at larger time steps [15], so that the selection of Δt can be made based on other factors, such as damping or runtime, and not only on whether the simulator works at all. One of the major factors that limits time step sizes is the inability of Newton’s method to converge reliably at large time steps (See Figs. 3, 2, and 4), or if a fixed number of Newton iterations are taken, the stability of the resulting simulation. We address this by formulating our nonlinear system of equations as a minimization problem, which we demonstrate can be solved more robustly. The idea that dynamics, energy, and minimization are related has been known since antiquity and is commonly leveraged in variational

integrators [6], [12], [16], [17], [18], [19], [20]. The idea that the nonlinear system that occurs from methods like backward Euler can be formulated as a minimization problem has appeared many times in graphics in various forms [2], [4], [5], [13], [19]. [19] point out that minimization leads to a method that is both simpler and faster than the equivalent nonlinear root-finding problem, and [5] show that a minimization formulation can be used to solve mass-spring systems more efficiently. Kane et al. [17] use a minimization formulation as a means of ensuring that a solution to their nonlinear system can be found assuming one exists. Goldenthal et al. [21] shows that a minimization formulation can be used to enforce constraints robustly and efficiently. Hirota et al. [2] shows that supplementing Newton’s method with a line search greatly improves robustness. Martin et al. [4] also shows that supplementing Newton’s method with a line search and a definiteness correction leads to a robust solution procedure. Following their example, we show that recasting the solution of the nonlinear systems that result from implicit time integration schemes as a nonlinear optimization problem results in substantial robustness improvements. We also show that additional improvements can be realized by incorporating additional techniques like Wolfe condition line searches which curve around collision bodies, conjugate gradient with early termination on indefiniteness, and choosing conjugate gradient tolerances based on the current degree of convergence.

This publication is an extended version of [22] in which we have applied the optimization integrator approach to the MPM snow simulator of [23]. This allows us to take much larger time steps than the original method and results in a significant speedup.

2 TIME INTEGRATION

The equations of motion for simulating solids are

$$\dot{x} = v \quad M\dot{v} = f \quad f = f(x, v),$$

where f are forces. As is common in graphics we assume M is a diagonal lumped-mass matrix. Since we are interested

- T.F. Gast, C. Schroeder, and C. Jiang are with the University of California Los Angeles, Los Angeles, CA 90095. E-mail: {tfg, Craig}@math.ucla.edu, cffjiang@cs.ucla.edu.
- A. Stomakhin is with the Research, Walt Disney Animation Studios, Burbank, CA. E-mail: st.alexey@gmail.com.
- J.M. Teran is with the University of California Los Angeles, Los Angeles, CA 90095, and the Walt Disney Animation Studios, Burbank, CA. E-mail: jteran@math.ucla.edu.

Manuscript received 13 Nov. 2014; revised 7 July 2015; accepted 12 July 2015.

Date of publication 21 July 2015; date of current version 4 Sept. 2015.

Recommended for acceptance by E. Sifakis and V. Koltun.

For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below.

Digital Object Identifier no. 10.1109/TVCG.2015.2459687

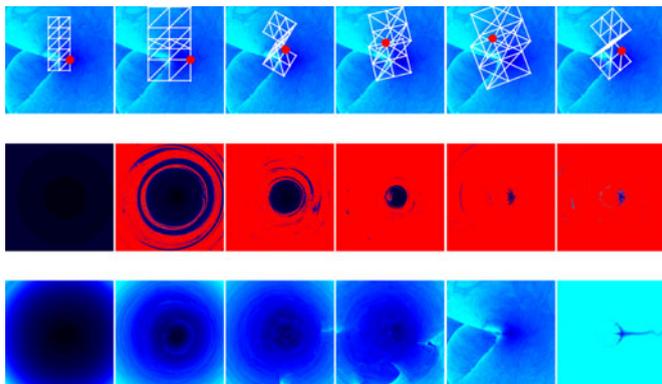


Fig. 1. Convergence of Newton's method (middle) and our stabilized optimization formulation (bottom) for a simple 36-dof simulation in 2D. The initial configuration (top) is parameterized in terms of a pixel location, with the rest configuration occurring at $(\frac{3}{5}, \frac{1}{5})$. Initial velocity is zero, and one time step is attempted. Time steps are (left to right) 170, 40, 20, 10, and 1 steps per 24 Hz frame, with the rightmost image being $\Delta t = 1$ s. Color indicates convergence in 0 iterations (black), 15 iterations (blue), 30 or more iterations (cyan), or failure to converge in 500 iterations (red). Note that Newton's method tends to converge rapidly or not at all, depending strongly on problem difficulty and initial guess.

in robustness and large time steps, we follow a backward Euler discretization. This leads to

$$\frac{x^{n+1} - x^n}{\Delta t} = v^{n+1} \quad M \frac{v^{n+1} - v^n}{\Delta t} = f^{n+1} = f(x^{n+1}, v^{n+1}).$$

Eliminating v^{n+1} yields

$$M \frac{x^{n+1} - x^n - \Delta t v^n}{\Delta t^2} = f\left(x^{n+1}, \frac{x^{n+1} - x^n}{\Delta t}\right),$$

which is a nonlinear system of equations in the unknown positions x^{n+1} . This system of nonlinear equations is normally solved with Newton's method. If we define

$$h(x^{n+1}) = M \frac{x^{n+1} - x^n - \Delta t v^n}{\Delta t^2} - f\left(x^{n+1}, \frac{x^{n+1} - x^n}{\Delta t}\right), \quad (1)$$

then our nonlinear problem is one of finding a solution to $h(x) = 0$. To do this, one would start with an initial guess $x^{(0)}$, such as the value predicted by forward Euler. This estimate is then iteratively improved using the update rule

$$x^{(i+1)} = x^{(i)} - \left(\frac{\partial h}{\partial x}(x^{(i)})\right)^{-1} h(x^{(i)}).$$

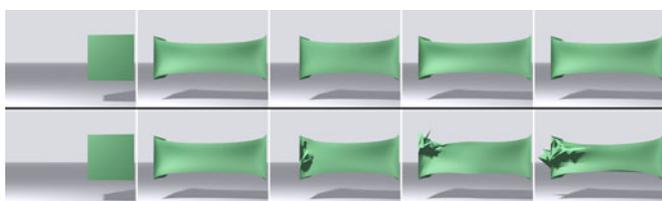


Fig. 2. Cube being stretched and then given a small compressive pulse, shown with our method (top) and standard Newton's method (bottom). Both simulations were run with one time step per 24 Hz frame. In this simulation, Newton's method is able to converge during the stretch phase, but a simple pulse of compression, as would normally occur due to a collision, causes it to fail to converge and never recover. Newton's method requires five time steps per frame to converge on this simple example.

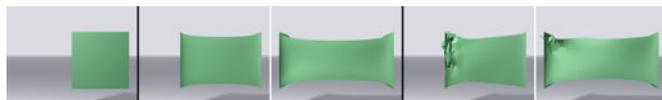


Fig. 3. Cube being stretched: initial configuration (left), our method at $t = 0.4$ s and $t = 3.0$ s (middle), and standard Newton's method at $t = 0.4$ s and $t = 3.0$ s (right). Both simulations were run with one time step per 24 Hz frame. Newton's method requires three time steps per frame to converge on this simple example.

Each step requires the solution of a linear system, which is usually symmetric and positive definite and solved with a Krylov solver such as conjugate gradient or MINRES.

If the function $h(x)$ is well-behaved and the initial guess sufficiently close to the solution, Newton's method will converge very rapidly (quadratically). If the initial guess is not close enough, Newton's method may converge slowly or not at all. For small enough time steps, the forward and backward Euler time steps will be very similar (they differ by $O(\Delta t^2)$), so a good initial guess is available. For large time steps, forward Euler will be unstable, so it will not provide a good initial guess. Further, as the time step grows larger, Newton's method may become more sensitive to the initial guess (see Fig. 1). The result is that Newton's method will often fail to converge if the time step is too large. Figs. 2, 3, and 4 show examples of simulations that ought to be routine but where Newton fails to converge at $\Delta t = 1/24$ s.

Sometimes, only one, or a small fixed number, of Newton steps are taken rather than trying to solve the nonlinear equation to a tolerance. The idea is that a small number of Newton steps is sufficient to get most of the benefit from doing an implicit method while limiting its cost. Indeed, even a single Newton step with backward Euler can allow time steps orders of magnitude higher than explicit methods. Linearizing the problem only goes so far, though, and even these solvers tend to have time step restrictions for tough problems.

2.1 Assumptions

We have found that when trying to be very robust, assumptions matter. Before introducing our formulation in detail, we begin by summarizing some idealized

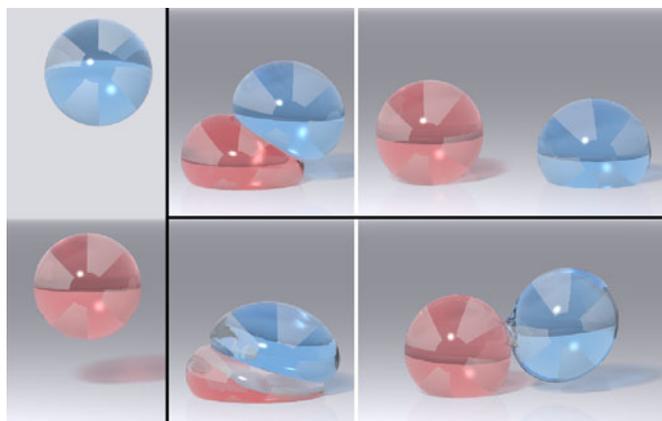


Fig. 4. Two spheres fall and collide with one another with $\Delta t = 1/24$ s: initial configuration (left), our method (top), and Newton's method (bottom). Notice the artifacts caused by Newton not converging. Newton's method requires six time steps per frame to converge on this example.

assumptions we will make. In practice, we will relax some of these as we go along.

- A1: Masses are positive
- A2: $f = -\frac{\partial\Phi}{\partial x}$ for some function Φ
- A3: Φ is bounded from below
- A4: Φ is C^1

Assumption (A1) implies that M is symmetric and positive definite and is useful for theoretical considerations; scripted objects violate this assumption, but they do not cause problems in practice.

Conservative forces always satisfy assumption (A2), and most practical elastic force models will satisfy this. We will show in Section 4.2 that even some damping models can be put into the required form. Friction can be given an approximate potential which is valid for small Δt (See [24]). Since our examples focus on taking larger time steps we address the problem by incorporating friction explicitly after the Newton solve.

Assumption (A3) is generally valid for constitutive models, with the global minimum occurring at the rest configuration. Gravity is an important example of a force that violates this assumption. In Section 2.3, we show that assumption (A3) can be safely relaxed to include forces like gravity.

Assumption (A4) is a difficult assumption. Technically, this assumption is a show-stopper, since we know of no constitutive model that is both robust and satisfies it everywhere. To be practical, this must be immediately loosened to C^0 , along with a restriction on the types of kinks that are permitted in Φ . The practical aspects of this are discussed in Section 3.3.

2.2 Minimization Problem

The solution to making Newton’s method converge reliably is to recast the equation solving problem as an optimization problem, for which robust and efficient methods exist. In principle, that can always be done, since solving $h(x) = 0$ is equivalent to minimizing $\|h(x)\|$ assuming a solution exists. This approach is not very convenient, though, since it requires a global minimum of $\|h(x)\|$. Further minimization using Newton’s method would require the Hessian of $\|h(x)\|$, which involves the second derivatives of our forces. The standard approach only requires first derivatives. What we really want is a quantity E that we can minimize whose second derivatives only require the first derivatives of our forces. That is, we need to *integrate* our system of nonlinear equations $h(x)$. Assumption (A2) allows us to do this. This way of recasting the problem also requires only a local minimum be found.

We can write (1) as

$$h(x) = M \frac{x - x^n - \Delta t v^n}{\Delta t^2} + \frac{\partial\Phi}{\partial x}.$$

We note that if we set

$$\hat{x} = x^n + \Delta t v^n \quad E(x) = \frac{1}{2\Delta t^2} (x - \hat{x})^T M (x - \hat{x}) + \Phi,$$

then we have $h = \frac{\partial E}{\partial x}$. If the required assumptions are met, a global minimum of E always exists.¹ By assumption (A4), $E(x^{n+1})$ is smooth at its minima, so $\frac{\partial E}{\partial x}(x^{n+1}) = 0$ or equivalently $h(x^{n+1}) = 0$.² Any local minimum is a solution to our original nonlinear equation (1). Although we are now doing minimization rather than root finding, we are still solving exactly the same equations. The discretization and dynamics will be the same, but the solver will be more robust. In particular, we are not making a quasistatic approximation.

2.3 Gravity

A graphics simulation would not be very useful without gravity. Gravity has the potential energy function $-Mg^T x$, where g is the gravitational acceleration vector, but this function is not bounded. An object can fall arbitrarily far and liberate a limitless supply of energy, though in practice this fall will be stopped by the ground or some other object. Adding the gravity force to our nonlinear system yields

$$h(x) = M \frac{x - x^n - \Delta t v^n}{\Delta t^2} - Mg + \frac{\partial\Phi}{\partial x},$$

which can be obtained from the bounded minimization objective

$$E(x) = \frac{1}{2\Delta t^2} (x - \hat{x} - \Delta t^2 g)^T M (x - \hat{x} - \Delta t^2 g) + \Phi.$$

A more convenient choice of E , and the one we use in practice, is obtained by simply adding the effects of gravity $\Phi_g = -Mg^T x$ into Φ . Since all choices E will differ by a constant shift, this more convenient minimization objective will also be bounded from below.

3 MINIMIZATION

The heart of our simulator is our algorithm for solving optimization problems, which we derived primarily from [25], though most of the techniques we apply are well-known. We begin by describing our method as it applies to unconstrained minimization and then show how to modify it to handle the constrained case.

3.1 Unconstrained Minimization

Our optimization routine begins with an initial guess, $x^{(0)}$. Each iteration consists of the following steps:

- 1) * Register active set
- 2) Compute gradient ∇E and Hessian H of E at $x^{(i)}$
- 3) Terminate successfully if $\|\nabla E\| < \tau$
- 4) Compute Newton step $\Delta x = -H^{-1} \nabla E$
- 5) Make sure Δx is a downhill direction

1. Assumptions (A1) and (A3) ensure that E is bounded from below. Let B be a lower bound on Φ . Then, let $L = \Phi(\hat{x}) - B + 1$ and Ω be the region where $\frac{1}{2\Delta t^2} (x - \hat{x})^T M (x - \hat{x}) \leq L$. Note that Ω is a closed and bounded ellipsoid centered at \hat{x} . E must have a global minimum when restricted to the set Ω since it is a continuous function on a closed and bounded domain. Outside Ω , we have $E(x) > L + B = E(\hat{x}) + 1$, so that the global minimum inside Ω is in fact a global minimum over all possible values of x .

2. Relaxation of assumption (A4) is discussed in Section 3.3, where Φ is allowed to have ridge-type kinks. Since these can never occur at a relative minimum, the conclusion here is unaffected.

- 6) Clamp the magnitude of Δx to ℓ if $\|\Delta x\| > \ell$
- 7) Choose step size α in direction Δx using a line search
- 8) Take the step: $x^{(i+1)} = x^{(i)} + \alpha \Delta x$
- 9) ★ Project $x^{(i+1)}$.

Here, τ is the termination criterion, which controls how accurately the system must be solved. The length clamp ℓ guards against the possibility of the Newton step being enormous (if $\|\Delta x\| = 10^{100}$, computing $\Phi(x^{(i)} + \Delta x)$ is unlikely to work well). Its value should be very large. Our line search is capable of choosing $\alpha > 1$, so the algorithm is very insensitive with respect to the choice ℓ . We normally use $\ell = 10^3$. Steps beginning with★ are only performed for constrained optimization and will be discussed later. A few of the remaining steps require further elaboration here.

Linear solver considerations. Computing the Newton step requires solving a symmetric linear system. The obvious candidate solver for this is MINRES that can handle indefinite systems, and indeed this will work. However, there are many tradeoffs to be made here. In contrast to a normal Newton solve, an accurate estimate for Δx is not necessary for convergence. Indeed, we would still converge with high probability if we chose Δx to be a random vector. The point of using the Newton direction is that convergence will typically be much more rapid, particularly when the super-convergence of Newton's method kicks in. (Choosing $\Delta x = -\nabla E$ leads to gradient descent, for example, which can display notoriously poor convergence rates.) When the current estimate is far from the solution, the exact Newton direction tends to be little better than a very approximate one. Thus, the idea is to spend little time on computing Δx when $\|\nabla E\|$ is large and more time when it is small. We do this by solving the system to a relative tolerance of $\min(\frac{1}{2}, \sigma \sqrt{\max(\|\nabla E\|, \tau)})$. The $\frac{1}{2}$ ensures that we always reduce the residual by at least a constant factor, which guarantees convergence. The scale σ adjusts for the fact that ∇E is not unitless (we usually use $\sigma = 1$). If our initial guess is naive, we must make sure we take at least one minimization iteration, even if ∇E is very small. Using τ here ensures that we do not waste time solving to a tiny tolerance in this case.

Conjugate gradient. One further optimization is to use conjugate gradient as the solver with a zero initial guess. If indefiniteness is encountered during the conjugate gradient solve, return the last iterate computed. If this occurs on the first step, return the right hand side. If this is done, Δx is guaranteed to be a downhill direction, though it might not be sufficiently downhill for our purposes. In practice, indefiniteness will only occur if far from converged, in which case little time is wasted in computing an accurate Δx that is unlikely to be very useful anyway. Indeed, if the system is detectably indefinite and Δx is computed exactly, it might not even point downhill. Since we are searching for a minimum of E (even a local one), the Hessian of E will be symmetric and positive definite near this solution. (Technically, it need only be positive semidefinite, but in practice this is of little consequence.) Thus, when we are close enough to the solution for an accurate Newton step to be useful, conjugate gradient will suffice to compute it. This is very different from the normal situation, where a solver like MINRES or an indefiniteness correction are employed to deal with the possibility of indefiniteness.

In the case of our solver, neither strategy is necessary, and both make the algorithm slower.

Downhill direction. Making sure Δx points downhill is fairly straightforward. If $\Delta x \cdot \nabla E < -\kappa \|\Delta x\| \|\nabla E\|$, then we consider Δx to be suitable. Otherwise, if $-\Delta x$ is suitable, use it instead. If neither Δx nor $-\Delta x$ are suitable, then we use the gradient descent direction $-\nabla E$. Note that if the conjugate gradient strategy is used for computing the Newton direction, then $-\Delta x$ will never be chosen as the search direction at this stage. We have found $\kappa = 10^{-2}$ to work well.

Line search. For our line search procedure, we use an algorithm for computing α such that the strong Wolfe Conditions are satisfied. See [25] for details. The line search procedure guarantees that E never increases from one iteration to the next and that, provided certain conditions are met, sufficient progress is always made. One important attribute of this line search algorithm is that it first checks to see if Δx itself is a suitable step. In this way, the line search is almost entirely avoided when Newton is converging properly.

Initial guess. A good initial guess is important for efficient simulation under normal circumstances. Under low- Δt or low-stress conditions, a good initial guess is obtained by replacing f^{n+1} by f^n resulting in

$$M \frac{x^{(0)} - x^n - \Delta t v^n}{\Delta t^2} = f(x^n).$$

Solving for x^{n+1} yields the initial guess

$$x^{(0)} = x^n + \Delta t v^n + \Delta t^2 M^{-1} f(x^n).$$

This initial guess is particularly effective under free fall, since here the initial guess is correct and no Newton iterations are required. On the other hand, this initial guess is the result of an explicit method, which will be unstable at large time steps or high stress. Under these conditions, this is unlikely to be a good initial guess and may in fact be very far from the solution. Under these situations, a better initial guess is obtained from $x^{(0)} = x^n + \Delta t v^n$. In practice, we compute both initial guesses and choose the one which produces the smaller value of E . This way, we get competitive performance under easy circumstances and rugged reliability under tough circumstances.

3.2 Constrained Minimization

We use constrained minimization for some of our collisions, which may result in a large active set of constraints, such as when a ball is bouncing on the ground. As the ball rises, constraints become deactivated. As the ball hits the ground, more constraints become activated. The change in the number of active constraints from iteration to iteration may be quite significant. This would render a traditional active set method impractical, since constraints are activated or deactivated one at a time. Instead, we use the gradient-projection method as our starting point, since it allows the number of active constraints to change quickly. The downside to this choice is that its reliance on the ability to efficiently project to the feasible region limits its applicability to simple collision objects.

Projections. Let $P(x)$ be the projection that applies P_{bp} to x_p for all body-particle pairs (b, p) that are labeled as active or are violated ($\phi_b(x_p) < 0$). Note that pairs such that

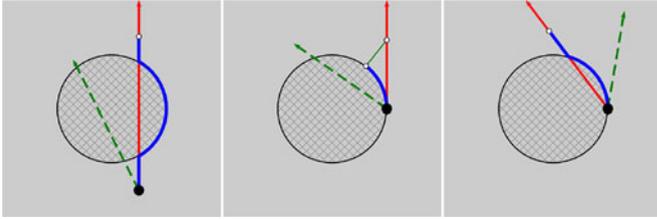


Fig. 5. Line search showing the gradient descent direction (green), Newton direction (red), and effective line search path (blue). The constraint is initially feasible (left), active (middle), and touching but inactive (right). Constraints are projected if violated or active, but only inactive constraints may separate.

$\phi_b(x_p) = 0$ (as would be the case once projected) are considered to be touching but not violated. The iterates $x^{(i)}$ obtained at the end of each Newton step, as well as the initial guess, are projected with P .

Register active set. Let E' be the objective that would be computed in the unconstrained case. The objective function for constrained optimization is $E(x) = E'(P(x))$. Compute the gradient $\nabla E'$. Constraints that are touching and for which $\nabla E' \cdot \nabla \phi_b \geq 0$ are labeled as active for the remainder of the Newton step. All others are labeled as inactive. No constraint should be violated at this stage. Note that $E'(x^{(i)}) = E(x^{(i)})$ is true before and after every Newton step, since constraints are never violated there.

Curved paths. Note that configurations are always projected to the feasible region before E is computed. One may interpret this as performing line searches along curved paths, as illustrated in Fig. 5.

When the unprojected line search curve passes through the medial axis of an object, it is possible for the search curve to be disconnected. This causes a discontinuity in the energy as seen from the line search. If the line search does not stop at the discontinuity, the discontinuity has no effect. If it does, the constraint causing the discontinuity will be active (in which case the discontinuity is projected out) or separating (in which case we move away from the discontinuity) in the next Newton step. Thus a disconnected search curve is not a problem for our method.

Discretized level sets. While discontinuities in the curved paths do not pose a problem when the level set is computed correctly, the situation can be quite different when the level set is approximated. This occurs when a grid-based level set is used to approximate a collision object. As a particle moves from cell to cell, the level set approximation (and thus projected location) changes slightly but unpredictably. The resulting kinks or discontinuities in the search path produce kinks or discontinuities in the objective function along the search line, which may cause the integrator to get stuck. For this reason, we restrict our use of optimization constraints to analytic level sets.

Derivatives. Note also that E must be differentiated twice, and that involves differentiating the projection function P twice. Since P depends on the first derivatives of ϕ_b , the Hessian H of E would seem to require third derivatives. We note, however, that the only occurrence of the third derivative of ϕ_b occurs multiplied by ϕ_b . Since H is used only at the beginning of the Newton step when the configuration is feasible, $\phi_b(x_p) = 0$ or P_{bp} is the identity function.

The third derivative term is zero either way, so only the second derivatives of ϕ_b are required.

3.3 Practical Considerations

There are a few matters of practicality relating to assumption (A4) that are worth mentioning regarding the effective use of this method. The most important of these is that the method does not tolerate discontinuities in E , not even very minute ones, except under some special circumstances that we mention below. In practice, what tends to happen is that a line search encounters a discontinuity in E , where E rises abruptly. The line search dutifully advances the configuration right up to location of this discontinuity. If in the next Newton iteration the descent direction points into the discontinuity, no progress can be made. The solver is stuck.

Discontinuities in ∇E can also cause problems and are impossible to avoid in general. These are kinks in E , which can be broken down into two types: valleys and ridges. The classification is based on whether the kink is ridge-like or valley-like. Ridge-type kinks are acceptable in practice. Valley-type kinks must be avoided, since they can also cause the solver to become stuck for the same reason. A minimum that occurs at a valley-type kink is also problematic since it does not correspond to a solution of (1). Thus, the corotated constitutive model, though not completely unusable with this solver, is ill-advised (the fixed variant has no such valleys [26] and is fine). Mass-spring systems are also fine. In practice, we have only encountered problems when evaluating self-collision models. The self-collision model we propose works well with the method.

The second practical consideration is that E can be somewhat noisy. This is particularly true with forces that involve an SVD, since its computation often involves a balance between speed and accuracy. If the Newton tolerance τ is set too low, the solver will be forced to optimize an objective E where the actual change in E is hidden by the noise. Even with our noisy SVD, we found there is typically at least a three-order-of-magnitude range between the largest value of τ below which no change in output is visually observed and the smallest value above which E is not too noisy to optimize reliably. If we make the E computation robust, E can be optimized down to roundoff level.

Another practical consideration is that occasionally very large changes in the configuration are considered by the line search. For most forces, this is of little consequence. For self-collisions, however, this poses a major performance hazard. We note that when this occurs, the other components of E become very large, too. We first compute all contributions to E except self-collisions. Since our self-collision potential has a global minimum of zero, the real E will be at least as large as the estimate. If this partial E is larger than $E(x^{(i)})$, we do not compute self-collisions at all. While this presents a discontinuity in E to the optimizer, it is safe to do so under these conditions, since the optimizer will avoid the large value in E by taking a smaller step along the search line.

4 FORCES

Our formulation is fairly insensitive to the underlying forces, provided it has a continuous potential energy function. We use five forces in our simulations. The simplest of

these is gravity, which we addressed in Section 2.2. We also employ a hyperelastic constitutive model (Section 4.1), a Rayleigh damping model (Section 4.2), and two collision penalty force models (Sections 5.2 and 5.3).

4.1 Elastic

A suitable hyperelastic constitutive model must have a few key properties to be suitable for this integrator. The most important is that it must have a potential energy function defined everywhere, and this function must be continuous. The constitutive model must be well-defined for any configuration, including configurations that are degenerate or inverted. This is true even if objects do not invert during the simulation, since the minimization procedure may still encounter such states. Examples of suitable constitutive models are those defined by the corotated hyperelasticity energy [27], [28], [29], [30], [31], [32] (but see Section 3.3), and the fixed corotated hyperelasticity variant [26]. Stress-based extrapolated models [33], [34] are unsuitable due to the lack of a potential energy function in the extrapolated regime, but energy-based extrapolation models [26] are fine. We use the fixed corotated variant [26] for all of our simulations for its combination of simplicity and robustness.

4.2 Damping

At first, one might conclude that requiring a potential energy may limit our method's applicability, since damping forces cannot be defined by a potential energy function. A very simple damping model is given by $\mathbf{f} = -kM\mathbf{v}^{n+1}$. Eliminating the velocity from the equation yields

$$\mathbf{f}(\mathbf{x}^{n+1}) = -kM \frac{\mathbf{x}^{n+1} - \mathbf{x}^n}{\Delta t} \quad (k > 0).$$

The scalar function

$$\Phi(\mathbf{x}^{n+1}) = \frac{k}{2\Delta t} (\mathbf{x}^{n+1} - \mathbf{x}^n)^T M (\mathbf{x}^{n+1} - \mathbf{x}^n)$$

has the necessary property that $\mathbf{f} = -\frac{\partial \Phi}{\partial \mathbf{x}}$. Note that this Φ looks very similar to our inertial term in E , and it is similarly bounded from below. That this Φ is not a real potential energy function is evident from its dependence on \mathbf{x}^n and Δt , but it is nevertheless suitable for use in our integrator. This simple drag force is not very realistic, though, so we do not use it in our simulations.

A more realistic damping force is Rayleigh damping. Let ψ be an elastic potential energy function. The stiffness matrix corresponding to this force is $-\frac{\partial^2 \psi}{\partial \mathbf{x} \partial \mathbf{x}}$ and the Rayleigh damping force and associated objective are

$$\mathbf{f} = -k \left(\frac{\partial^2 \psi}{\partial \mathbf{x} \partial \mathbf{x}} (\mathbf{x}^{n+1}) \right) \mathbf{v}^{n+1} \quad \Phi_c = \frac{k}{\Delta t} \left((\mathbf{x}^{n+1} - \mathbf{x}^n)^T \frac{\partial \psi}{\partial \mathbf{x}} - \psi \right).$$

This candidate Φ_c has at least two serious problems. The first is that second derivatives of Φ_c involve third derivatives of ψ . The second is that $\frac{\partial^2 \psi}{\partial \mathbf{x} \partial \mathbf{x}}$ may be indefinite, in which case the damping force may not be entirely dissipative. Instead, we approximate Rayleigh damping with a lagged version. Let $\mathbf{D} = \frac{\partial^2 \psi}{\partial \mathbf{x} \partial \mathbf{x}} (\mathbf{x}^n)$. Since \mathbf{D} does not depend on \mathbf{x}^{n+1} , the lagged Rayleigh damping force

and associated objective are

$$\mathbf{f} = -k\mathbf{D}\mathbf{v}^{n+1} \quad \Phi_d = \frac{k}{2\Delta t} (\mathbf{x}^{n+1} - \mathbf{x}^n)^T \mathbf{D} (\mathbf{x}^{n+1} - \mathbf{x}^n).$$

This solves the first problem, since the second derivative of Φ_d is just $\frac{k}{\Delta t} \mathbf{D}$. Since \mathbf{D} is not being differentiated, it is safe to modify it to eliminate indefiniteness as described in [26], [34]. This addresses the second problem. We did not use the damping model found in [35], which uses $\psi(\mathbf{x}^{n+1})$ with \mathbf{x}^n used as the rest configuration, because it is not defined when \mathbf{x}^n is degenerate.

5 COLLISIONS

Collisions are a necessary part of any practical computer graphics simulator. The simplest approach to handling collisions is to process them as a separate step in the time integration scheme. This works well for small time steps, but it causes problems when used with large time steps as seen in Fig. 10. Such arrangement often leads to the collision step flattening objects to remove penetration and the elastic solver restoring the flattened geometry by pushing it into the colliding object. To get around this problem, the backward Euler solver needs to be aware of collisions. A well-tested strategy for doing this is to use penalty collisions, and we do this for two of our three collision processing techniques. Our other approach is to add position constraints to the nonlinear solve.

Fig. 11 uses constraints for all collision body collisions and demonstrates that our constraint collisions are effective with concave and convex constraint manifolds. Fig. 13 demonstrates our constraint collisions are effective for objects with sharp corners. Fig. 7 is a classical torus drop demonstrating that our self collisions are effective at stopping collisions at the torus's hole. Fig. 12 demonstrates our self collisions method with stiffer deformable bodies with sharp corners. Finally, Fig. 9 shows a more practical example which uses all three types of collisions: self collisions, constraint collisions (with ground) and penalty collisions (against a bowl defined by a grid-based level set).

5.1 Object Collisions As Constraints

Our first collision processing technique takes advantage of our minimization framework to treat collisions with non-simulated objects as inequality constraints. Treating collisions or contacts as constraints is not new and in fact forms the basis for LCP formulations such as [36], [37]. Unlike LCP formulations, however, our formulation does not attempt to be as complete and as a result can be solved about as efficiently as a simple penalty formulation.

Our constraint collision formulation works reliably when the level set is known analytically. This limits its applicability to analytic collision objects. While this approach is feasible only under limited circumstances, these circumstances occur frequently in practice. When this approach is applicable, it is our method of choice, since it produces better results (e.g., no interpenetration) for similar cost. When this formulation is not applicable, we use a penalty collision formulation instead.

We begin by representing our collision objects (indexed with b) by a level set, which we denote ϕ_b to avoid confusion with potential energy. By convention, $\phi_b(\mathbf{x}) < 0$ for points \mathbf{x} in the interior of the collision object b . Our collision constraint is simply that $\phi_b(\mathbf{x}_p^{n+1}) \geq 0$ for each simulation particle p and every constraint collision object b . With such a formulation, we can project a particle at \mathbf{x}_p to the closest point \mathbf{x}'_p on the constraint manifold using

$$\mathbf{x}'_p = P_{b_p}(\mathbf{x}_p) = \mathbf{x}_p - \phi_b(\mathbf{x}_p) \nabla \phi_b(\mathbf{x}_p).$$

We show how to solve the resulting minimization problem in Section 3.2.

We apply friction after the Newton solve. The total collision force felt by particles is

$$\Delta t \mathbf{f}_{col} = \nabla E'(\mathbf{x}^{n+1}) - \nabla E(\mathbf{x}^{n+1}) = \nabla E'(\mathbf{x}^{n+1}) - \nabla E'(P(\mathbf{x}^{n+1})),$$

where E' is the objective in the absence of constraints (See Section 3.2). Only collision pairs that are active at the end of the minimization will be applying such forces. We use the level set's normal and the collision force to apply Coulomb friction to colliding particles. In particular, we use the rule ($\mathbf{v}_p^{n+1} \rightarrow \hat{\mathbf{v}}_p^{n+1}$)

$$\begin{aligned} \mathbf{n} &= \nabla \phi & \mathbf{v}_{p,n}^{n+1} &= (\mathbf{n} \cdot \mathbf{v}_p^{n+1}) \mathbf{n} & \mathbf{v}_{p,t}^{n+1} &= \mathbf{v}_p^{n+1} - \mathbf{v}_{p,n}^{n+1} \\ \hat{\mathbf{v}}_p^{n+1} &= \mathbf{v}_{p,n}^{n+1} + \max\left(1 - \frac{\mu \Delta t (\mathbf{n} \cdot \mathbf{f}_{p,col})}{m \|\mathbf{v}_{p,t}\|}, 0\right) \mathbf{v}_{p,t}^{n+1}. \end{aligned}$$

Our constraint collision formulation is not directly applicable to grid-based level sets, since we assume that $P_{b_p}(P_{b_p}(\mathbf{x}_p)) = P_{b_p}(\mathbf{x}_p)$ and $P_{b_p}(x)$ is continuous. Continuity of $P_{b_p}(x)$ can be achieved, for example, with C^1 cubic spline level set interpolation. However, it will not generally be true that $P_{b_p}(P_{b_p}(\mathbf{x}_p)) = P_{b_p}(\mathbf{x}_p)$. Alternatively, the projection routine can be modified to iterate the projection to convergence, but then continuity is lost.

5.2 Object Penalty Collisions

When a collision object is not analytic, as will normally be the case for characters for instance, we use a penalty formulation instead. As in the constraint formulation, we assume our collision object is represented by a level set ϕ_b . The elastic potential energy $\Phi_{b_p}(\mathbf{x}_p)$ of our penalty force is $\Phi_{b_p}(\mathbf{x}) = 0$ if $\phi_b(\mathbf{x}_p) > 0$ and $\Phi_{b_p}(\mathbf{x}_p) = k\phi_b(\mathbf{x}_p)^3$ otherwise. Since Φ_{b_p} is a potential energy, we must differentiate it twice for our solver. It is important to compute the derivatives of ϕ_b exactly by differentiating the interpolation routine rather than approximating them using central differences. While a C^1 cubic spline interpolation is probably a wiser interpolation strategy since it would avoid the energy kinks that may be caused by a piecewise linear encoding of the level set, we found linear interpolation to work well, too, and we use linear interpolation in our examples.

As in the constraint case, we apply friction after the Newton solve. The total collision force felt by a particle due to object penalty collisions is obtained by evaluating the penalty force at \mathbf{x}^{n+1} and using this force as the normal direction. That is,

$$\begin{aligned} \mathbf{f}_{col} &= -\frac{\partial \Phi_{b_p}}{\partial \mathbf{x}}(\mathbf{x}^{n+1}) & \mathbf{f}_{p,n} &= \|\mathbf{f}_{p,col}\| & \mathbf{n} &= \frac{\mathbf{f}_{p,col}}{f_{p,n}} \\ \mathbf{v}_{p,n}^{n+1} &= (\mathbf{n} \cdot \mathbf{v}_p^{n+1}) \mathbf{n} & \mathbf{v}_{p,t}^{n+1} &= \mathbf{v}_p^{n+1} - \mathbf{v}_{p,n}^{n+1} \\ \hat{\mathbf{v}}_p^{n+1} &= \mathbf{v}_{p,n}^{n+1} + \max\left(1 - \frac{\mu \Delta t f_{p,n}}{m \|\mathbf{v}_{p,t}\|}, 0\right) \mathbf{v}_{p,t}^{n+1}. \end{aligned}$$

5.3 Penalty Self-Collisions

We detect self-collisions by performing point-tetrahedron inclusion tests, which we accelerate with a bounding box hierarchy. If a point is found to be inside a tetrahedron but not one of the vertices of that tetrahedron, then we flag the particle as colliding.

Once we know a particle is involved in a self collision, we need an estimate for how close the particle is to the boundary. If this particle has collided before, we use the primitive it last collided with as our estimate. Otherwise, we compute the approximate closest primitive in the rest configuration using a level set and use the current distance to this surface element as an estimate.

Given this upper bound estimate of the distance to the boundary, we perform a bounding box search to conservatively return all surface primitives within that distance. We check these candidates to find the closest one. Now we have a point-primitive pair, where the primitive is the surface triangle, edge, or vertex that is closest to the point being processed. Let d be the square of the point-primitive distance. The penalty collision energy for this point is $\Phi = kd\sqrt{d} + \epsilon$, where ϵ is a small number (10^{-15} in our case) to prevent the singularities when differentiating. Note that this penalty function is approximately cubic in the penetration depth. This final step is the only part that must be differentiated.

As with the other two collision models, we apply friction after the Newton solve. In the most general case, a point n_0 collides with a surface triangle with vertices n_1 , n_2 , and n_3 . As with the object penalty collision model, collision forces are computed by evaluating $\Phi(\mathbf{x}^{n+1})$ and its derivative. The force applied to n_0 is denoted \mathbf{f} ; its direction is taken to be the normal direction \mathbf{n} . The closest point on the triangle to n_0 has barycentric weights w_1 , w_2 , and w_3 . Let $w_0 = -1$ for convenience. Let $\mathbf{Q} = \mathbf{I} - \mathbf{n}\mathbf{n}^T$, noting that $\mathbf{Q}^2 = \mathbf{Q}$. If we apply a tangential impulse $\mathbf{Q}\mathbf{j}$ to these particles, their new velocities and kinetic energy will be

$$\hat{\mathbf{v}}_{n_i}^{n+1} = \mathbf{v}_{n_i}^{n+1} + w_i m_{n_i}^{-1} \mathbf{Q}\mathbf{j} \quad KE = \sum_{n=0}^3 \frac{1}{2} m_{n_i} (\hat{\mathbf{v}}_{n_i}^{n+1})^T \hat{\mathbf{v}}_{n_i}^{n+1}.$$

We want to minimize this kinetic energy to prevent friction from causing instability. Since M is positive definite, we see that KE is minimized when

$$\nabla KE = \mathbf{Q}\bar{\mathbf{v}} + \bar{\mathbf{m}}^{-1} \mathbf{Q}\mathbf{j} = 0 \bar{\mathbf{v}} = \sum_{n=0}^3 w_i \mathbf{v}_{n_i}^{n+1} \bar{\mathbf{m}}^{-1} = \sum_{n=0}^3 w_i m_{n_i}^{-1} w_i.$$

If we let $\mathbf{j} = -\bar{\mathbf{m}}\mathbf{Q}\bar{\mathbf{v}}$ then $\nabla KE = 0$ and $\mathbf{Q}\mathbf{j} = \mathbf{j}$. This leads to the friction application rule

$$\hat{\mathbf{v}}_{n_i}^{n+1} = \mathbf{v}_{n_i}^{n+1} + w_i m_{n_i}^{-1} \min\left(\frac{\mu \|\mathbf{f}\|}{\|\mathbf{j}\|}, 1\right) \mathbf{j}.$$

Note that all three friction algorithms decrease kinetic energy but do not modify positions, so none of them can add energy to the system, and thus stability ramifications are unlikely even though friction is applied explicitly. This approach to friction can have artifacts, however, since friction will be limited to removing kinetic energy from colliding particles. This limits the amount of friction that can be applied at large time steps. An approach similar to the one in [36] that uses successive Quadratic Programming solves could possibly be applied to eliminate these artifacts. However [38] found existing large-scale sparse QP solvers to be insufficiently robust, and thus we did not use this method.

6 ACCELERATING MATERIAL POINT METHOD (MPM)

In this section we describe the application of this optimization approach to the snow simulation from [23]. Their approach to simulating snow uses the material point method, a hybrid Eulerian-Lagrangian formulation that uses unstructured particles as the primary representation and a background grid for applying forces. They used an energy-based formulation to facilitate a semi-implicit treatment of MPM. While this leads to a significant time step improvement over more standard explicit treatments, it still requires a small time step in practice to remain stable. We show how to modify their original formulation so that we are able to take time steps on the order of the CFL condition. We also provide an improved treatment of collisions with solid bodies that naturally handles them as constraints in the optimization. Although the optimization solve is for grid velocities, we show that a backward Euler (rather than forward Euler) update of particle positions in the grid based velocity field automatically guarantees no particles penetrate solid bodies. In addition to the significantly improved stability, we demonstrate in Section 7.1 that in many cases a worthwhile speedup can be obtained with our new formulation.

6.1 Revised MPM Time Integration

In [23, Section 4.1], the original method is broken down into 10 steps. From the original method, Steps 3-6 and 9-10 are modified. We begin by summarizing these steps as they apply to our optimization-based MPM integrator.

- 1) *Rasterize particle data to the grid.* First, mass and momentum are transferred from particles to the grid using $m_i^n = \sum_p m_p w_{ip}^n$ and $m_i^n \mathbf{v}_i^n = \sum_p \mathbf{v}_p^n m_p w_{ip}^n$. Velocity is then obtained by division using $\mathbf{v}_i^n = m_i^n \mathbf{v}_i^n / m_i^n$. Transferring velocity in this way conserves momentum.
- 2) *Compute particle volumes.* First time step only. Our force discretization requires a notion of a particle's volume in the initial configuration. Since cells have a well-defined notion of volume and mass, we can estimate a cell's density as $\rho_i^0 = m_i^0 / h^3$ and

interpolate it back to the particle as $\rho_p^0 = \sum_i \rho_i^0 w_{ip}^0$. Finally, we can define a particle's volume to be $V_p^0 = m_p / \rho_p^0$. Though rather indirect, this approach automatically provides an estimate of the amount of volume that can be attributed to individual particles.

- 3) *Solve the optimization problem.* Minimize the objective (2) using the methods of Section 3. This produces a new velocity estimate \mathbf{v}_i^{n+1} on the grid. This step replaces Steps 3-6 of the original method.
- 4) *Update deformation gradient.* The deformation gradient for each particle is updated as $\mathbf{F}_p^{n+1} = (\mathbf{I} + \Delta t \nabla \mathbf{v}_p^{n+1}) \mathbf{F}_p^n$, where we have computed $\nabla \mathbf{v}_p^{n+1} = \sum_i \mathbf{v}_i^{n+1} (\nabla w_{ip}^n)^T$. Note that this involves updates for the elastic and plastic parts of \mathbf{F} . See [23] for details, as they are unchanged.
- 5) *Update particle velocities.* Our new particle velocities are $\mathbf{v}_p^{n+1} = (1 - \alpha) \mathbf{v}_{\text{PIC}_p}^{n+1} + \alpha \mathbf{v}_{\text{FLIP}_p}^{n+1}$, where the PIC part is $\mathbf{v}_{\text{PIC}_p}^{n+1} = \sum_i \mathbf{v}_i^{n+1} w_{ip}^n$ and the FLIP part is $\mathbf{v}_{\text{FLIP}_p}^{n+1} = \mathbf{v}_p^n + \sum_i (\mathbf{v}_i^{n+1} - \mathbf{v}_i^n) w_{ip}^n$. We typically used $\alpha = 0.95$.
- 6) *Update particle positions.* Particle positions are updated using $\mathbf{x}_p^{n+1} = \mathbf{x}_p^n + \Delta t \mathbf{v}(\mathbf{x}_p^{n+1})$ as described in Section 6.3. This step replaces steps 9-10 of the original method.

6.2 Optimization Formulation

The primary modification that we propose is to use the optimization framework in place of the original solver. For this, we must formulate their update in terms of an optimization objective E . The original formulation defined the potential energy $\Phi(\mathbf{x}_i)$ conceptually in terms of the grid node locations \mathbf{x}_i . Here we use the index i to refer to grid node indices. Their grid is a fixed Cartesian grid and never moves, and they solve for \mathbf{v}_i^{n+1} . We will follow the same conceptual formulation here. This leads to the objective

$$E(\mathbf{v}_i) = \sum_i \frac{1}{2} m_i \|\mathbf{v}_i - \mathbf{v}_i^n\|^2 + \Phi(\mathbf{x}_i^n + \Delta t \mathbf{v}_i), \quad (2)$$

where m_i is the mass assigned to grid index i . Our final \mathbf{v}_i^{n+1} is computed so that $E(\mathbf{v}_i^{n+1})$ is minimized. We solve this minimization problem as in Section 3. Note that we apply plasticity explicitly as in the original formulation.

Using larger time steps causes our linear systems to become slower to solve. In the case of MPM, we found it beneficial to use the diagonal preconditioner

$$\mathbf{L}_{ii} = \sum_p \text{diag}\left(m_p w_{ip} \mathbf{I} + \Delta t^2 V_p^0 \mathbf{H}\right),$$

where

$$\mathbf{H} = (\lambda_p + \mu_p) \nabla w_{ip} \nabla w_{ip}^T + \mu_p \nabla w_{ip}^T \nabla w_{ip} \mathbf{I}.$$

This preconditioner approximates the diagonal of the stiffness matrix at the rest configuration. This works well since snow is unable to deform much without hardening or fracturing. We use an approximation to the diagonal, rather

than the exact diagonal, because we never explicitly form the matrix. This approximation suffices for preconditioning and is more efficient.

The original method performed solid body collisions while computing new grid velocities. We treat body collisions using constraints in our optimization problem. We assume sticking collisions and let $P(v_i) = 0$ for all grid nodes i that lie inside a collision object. Note that we do not permit separation during optimization, though separation may occur during other steps in the algorithm.

6.3 Particle Position Update

One of the difficulties with running the method of [23] with larger time steps is the particle-based solid body collisions. They were needed under the old formulation to prevent settling into the ground, but at the same time they cause bunching of particles at collision objects. These problems are exacerbated at larger time steps, and another approach is required. Instead, we show that altering the way we update particle positions can avoid the need for a separate particle collision step.

For each particle position x_p we solve the backward Euler update equation

$$x_p^{n+1} = x_p^n + \Delta t v(x_p^{n+1}) \quad v(x_p^{n+1}) = \sum_i v_i^{n+1} N_i^h(x_p),$$

where $v(x_p^{n+1})$ is the interpolated grid velocity at the particle location x_p^{n+1} . These updates are independent per particle and so are relatively inexpensive. A solution to this backward Euler equation always exists nearby provided a suitable CFL condition is respected (no particle moves more than Δx in a time step). Note that pure PIC velocities are used in the particle position updates. While a combination of FLIP/PIC is still stored on particles (to avoid excessive dissipation in subsequent transfer to grid), PIC velocities for position updates lead to more stable behavior.

The motivation for our modification can be best understood in the case of sticking collisions. Inside a collision object, we will have $v_i^{n+1} = 0$ due to the collision constraints imposed during optimization. If we then assume that we will interpolate $v(x_p^{n+1}) = 0$ here, then we can see from $x_p^{n+1} = x_p^n + \Delta t v(x_p^{n+1})$ that $x_p^{n+1} = x_p^n$. Note that if a particle ends up inside the collision object, then it must have already been there. Thus, it is not possible for particles to penetrate collision objects. In our implementation, $v(x_p^{n+1}) = 0$ will only be true if we are slightly inside collision objects, but in practice this procedure actually stops particles slightly outside collision objects.

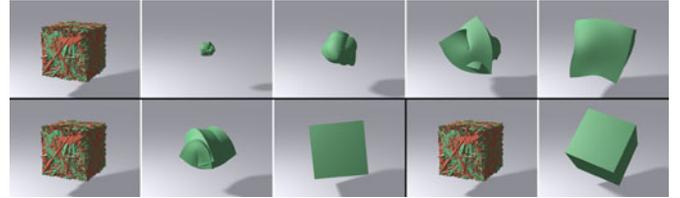


Fig. 6. Random test with $65 \times 65 \times 65$ particles simulated with $\Delta t = 1/24$ s for three stiffnesses: low stiffness recovering over 100 time steps (top), medium stiffness recovering over 40 time steps (bottom left), and high stiffness recovering in a single time step (bottom right). The red tetrahedra are inverted, while the green are uninverted.

We solve this equation with Newton’s method. Since Newton’s method need not converge, some care is required, though in practice nothing as sophisticated as Section 3 is needed. We always use the Newton direction but repeatedly halve the length of the Newton step until the objective $E = \|x_p^{n+1} - x_p^n - \Delta t v(x_p^{n+1})\|$ no longer increases. (If halving the step size 14 times does not suffice, we take the reduced step anyway.) Typically, only one Newton step is required for convergence. We have never observed this to fail.

We use a quadratic spline rather than the cubic of the original formulation to reduce stencil width and improve the effectiveness of the modified position update. That is, we let

$$N(x) = \begin{cases} -x^2 & |x| < \frac{1}{2}, \\ \frac{1}{2}x^2 - \frac{3}{2}|x| + \frac{9}{8} & \frac{1}{2} \leq |x| < \frac{3}{2}, \\ 0 & |x| \geq \frac{3}{2}. \end{cases}$$

Using a quadratic stencil also has the advantage of being more efficient. We do not use a linear spline since it is not smooth enough for Newton’s method to be effective in the particle position update.

Since MPM involves a grid, we limit our time step so that particles do not travel more than one grid spacing per time step. That is, we choose Δt so that $v \frac{\Delta t}{\Delta x} \geq \max_p \|v_p^n\|$ for some $v < 1$. We chose $v = 0.6$ for our examples. Although the time step restriction is computed based on v_p^n rather than v_p^{n+1} , this suffices in practice.

7 RESULTS

We begin by demonstrating how robust our solver is by considering the two most difficult constitutive model tests we are aware of: total randomness and total degeneracy. The attributes that make them tough constitutive model tests also make them tough solver tests: high stress, terrible initial guess, tangled configurations, and the need to dissipate massive amounts of unwanted energy. Fig. 6 shows the recovery of a $65 \times 65 \times 65$ cube (824k dofs) from a randomized initial configuration for



Fig. 7. A torus falls on the ground (constraint collisions) and collides with itself (penalty collisions).

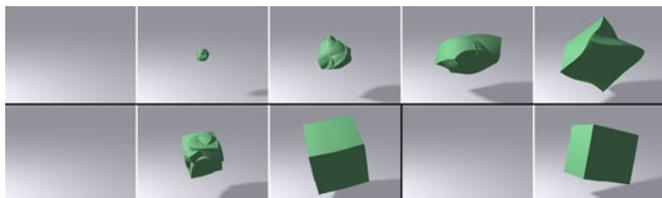


Fig. 8. Point test with $65 \times 65 \times 65$ particles simulated with $\Delta t = 1/24$ s for three stiffnesses: low stiffness recovering over 120 time steps (top), medium stiffness recovering in five time steps (bottom left), and high stiffness recovering in a single time step (bottom right).

three different stiffnesses with $\Delta t = 1/24$ s. Fig. 8 repeats the tests with all points starting at the origin. The recovery times vary from about 3 s for the softest to a single time step for the stiffest. We were surprised to find that a single step of backward Euler could untangle a randomized cube, even at high resolution.

7.1 MPM Results

We demonstrate the advantages of using our optimization integrator by applying it to the MPM snow formulation from [23]. We run three examples using both the original formulation and our modified formulation. We compare with the snowball examples from the original paper. In each case, for our formulation we use the CFL $\nu = 0.6$. Fig. 17 shows a snowball hitting a wall using sticky collisions, which causes the snow to stick to the wall. Fig. 14 shows a dropped snowball hitting the ground with sticky collisions. Fig. 16 shows two snowballs colliding in mid air with sticky collisions against the ground. On average, we get a speed up of 3.5 times over the original method. These results are tabulated in

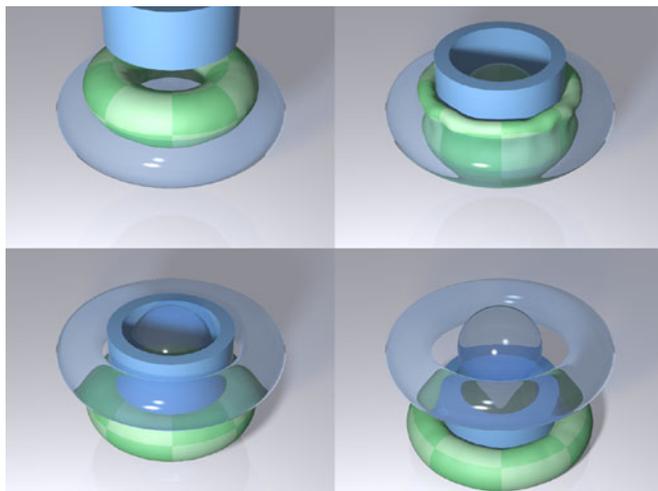


Fig. 11. A torus is pushed through a hole (constraint collisions).

Fig. 18. Notably, we are able to take significantly larger time steps, however some of the potential gains from this are lost to an increased complexity per time step. Nonetheless, we provide a significant computational savings with minimal modification to the original approach.

8 CONCLUSIONS

We have demonstrated that backward Euler solved with Newton's method can be made more robust by recasting the resulting system of nonlinear equations as a nonlinear optimization problem so that robust optimization techniques can be employed. The resulting method is extremely robust to large time step sizes, high stress, and tangled configurations.

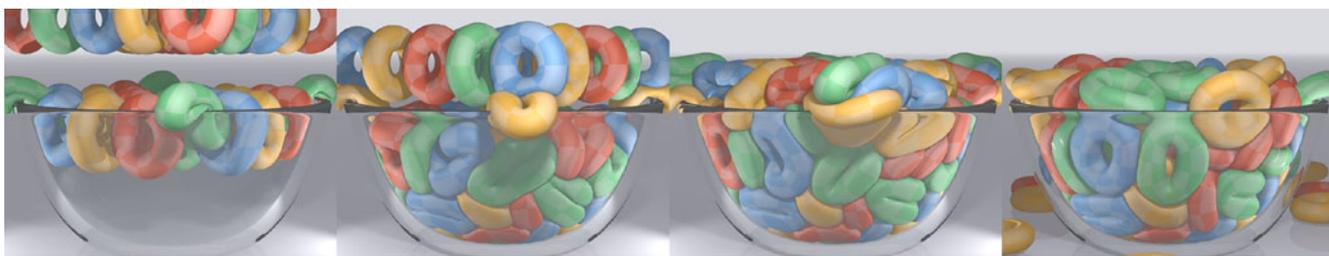


Fig. 9. 125 tori are dropped into a bowl at five time steps per frame, resulting in significant deformation and tough collisions.

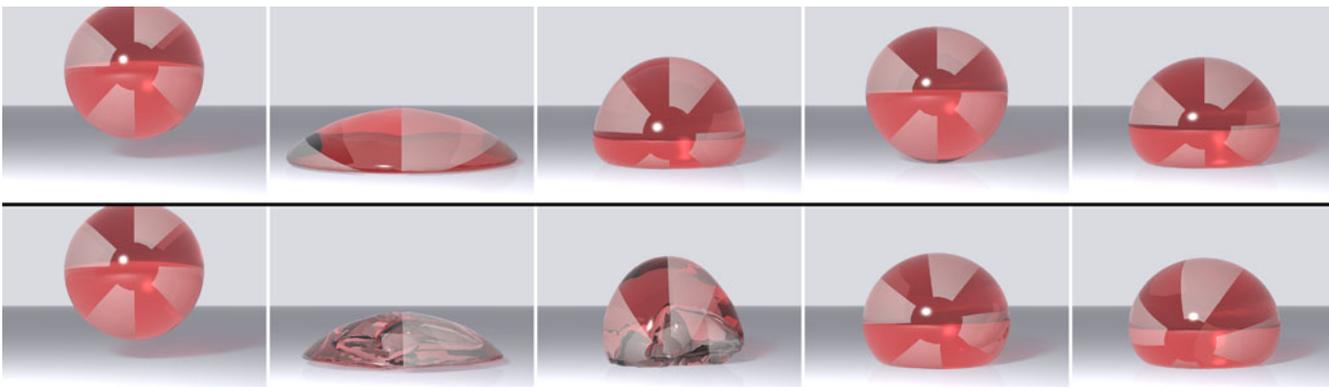


Fig. 10. Sphere dropping hard on the ground with $\Delta t = 1/24$ s with constraint collisions (top) and collisions as a post-process (bottom). Penalty collisions produce a result very similar to constraint collisions, though some penetration with the ground occurs. Note that the post-processing approach leads to inversion during recovery from the collision.

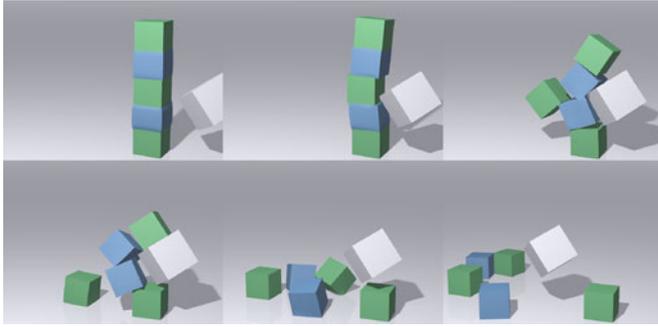


Fig. 12. A stack of deformable boxes of varying stiffness is struck with a rigid kinematic cube (constraint collisions) with $\Delta t = 1/24$ s. The green boxes are 10 times as stiff as the blue boxes.

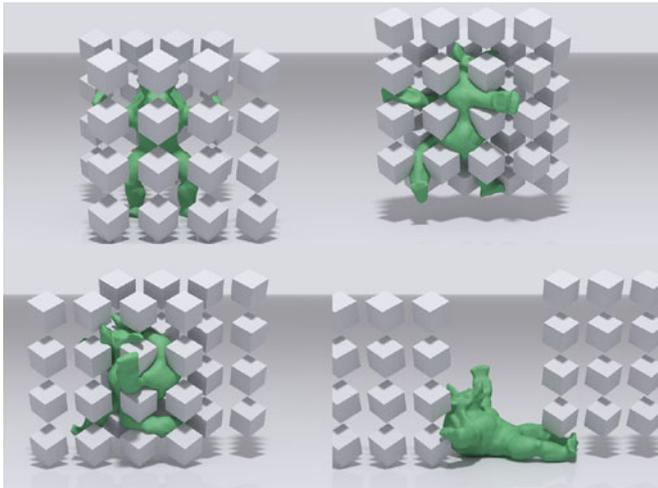


Fig. 13. An armadillo is squeezed between 32 rigid cubes (constraint collisions) with $\Delta t = 1/24$ s. When this torture test is run at one, two, four and eight steps per frame the average runtime per frame is 46, 58, 88, and 117 seconds respectively.

Runtimes and other performance-related information for all of our sims are provided in Fig. 15. All Lagrangian simulations were run single-threaded on a 3.1-3.5 GHz Xeon core, the MPM simulations were run with 10 threads for Fig. 16 and 12 threads for Fig. 14 and 17. Our solver’s performance is competitive with a standard Newton solver for those examples where both were run. In general, we take more Newton steps but spend less time on each, and the resulting runtime for typical examples is about the same for the two solvers, though our solver is faster for all of the difficult examples in this paper. Taking a large time step size

| Figure | Ours? | Steps frame | Time frame (s) | # dofs | Solves step |
|--------|-------|----------------|-------------------|--------|----------------|
| 9 | Y | 5 | 200 | 984k | 2.2 |
| 3 mid | Y | 1 | 0.51 | 18.5k | 2.8 |
| 3 rt | N | 1/3 | 8.7/1.1 | 18.5k | 15/0.7 |
| 2 top | Y | 1 | 0.52 | 18.5k | 2.9 |
| 2 bot | N | 1/5 | 3.8/1.3 | 18.5k | 6.6/0.6 |
| 4 top | Y | 1 | 4.25 | 28.0k | 8.1 |
| 4 bot | N | 1/6 | 33/7.3 | 28.0k | 26/0.8 |
| 7 | Y | 5 | 1.13 | 7.9k | 2.1 |
| 6 top | Y | 1 | 68.0 ⁺ | 824k | 12.3 |
| 6 lt | Y | 1 | 1470 ⁺ | 824k | 236.8 |
| 6 rt | Y | 1 | 667 ⁺ | 824k | 109.6 |
| 8 top | Y | 1 | 43.1 ⁺ | 824k | 10.7 |
| 8 lt | Y | 1 | 831 ⁺ | 824k | 155.9 |
| 8 rt | Y | 1 | 444 ⁺ | 824k | 88.8 |
| 10 top | Y | 1 | 0.42 | 14.0k | 3.8 |
| 10 bot | N | 1* | 1.13 | 14.0k | 9.8 |
| 11 | Y | 1 | 0.45 | 7.9k | 8.6 |
| 13 | Y | 1 | 46.1 | 73.8k | 34.7 |
| 12 | Y | 1 | 17.1 | 138k | 6.9 |

Fig. 15. Time step sizes and average running times for the examples in the paper. The last column shows the average number of linear solves per time step. Each of the Newton’s method examples fails to converge at the frame rate. For fairer comparison, timing information for all but the one marked * is shown at the frame rate and the stable time step size. The stress tests marked ⁺ spend the majority of their time on the first frame or two due to the difficult initial state.

can actually be slower than taking a smaller one, even with the same solver. For time integrators (like backward Euler) that have a significant amount of damping at large time steps, constitutive models are often tuned to take into account the numerical damping. If the integrator is forced to simulate a portion of a simulation at a smaller time step, the dynamic behavior can change noticeably. Solving with constraints is about the same speed as using penalty collisions.

Note that Figs. 9 and 7 were run with smaller time steps sizes to avoid collision artifacts. This indicates that a self-collision scheme that is more tolerant of large time steps is required. The scheme does not have problems with collisions between different objects at the frame rate as long as they are not too thin. Continuous collision detection could perhaps be used. We leave both of these problems for future work.

The current method has a couple disadvantages compared with current techniques. It requires a potential



Fig. 14. Our approach works naturally with the material point method simulations from [23]. Here we demonstrate with a snowball that drops to the ground and fractures. Notably, we provide a new treatment of particle position updates that naturally prevents penetration in solid objects like the ground.

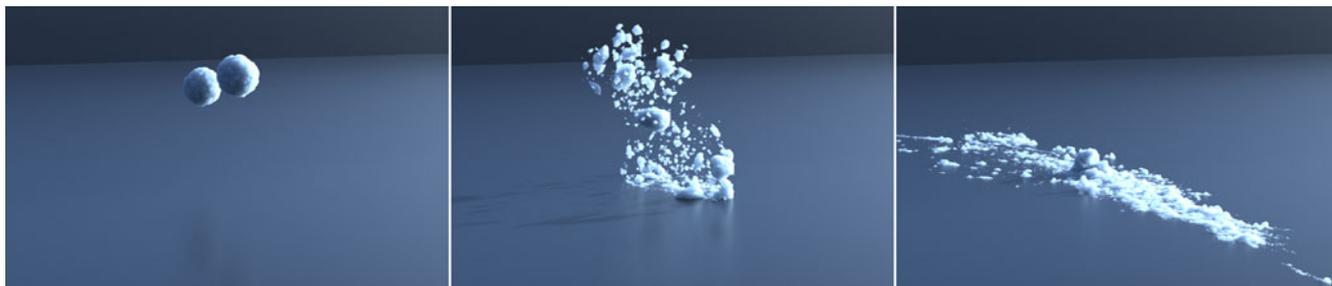


Fig. 16. The extension of our method to [23] is robust to large deformation and collisions scenarios. Here we demonstrate this for with two snowballs that smash into each other and fall to the ground.

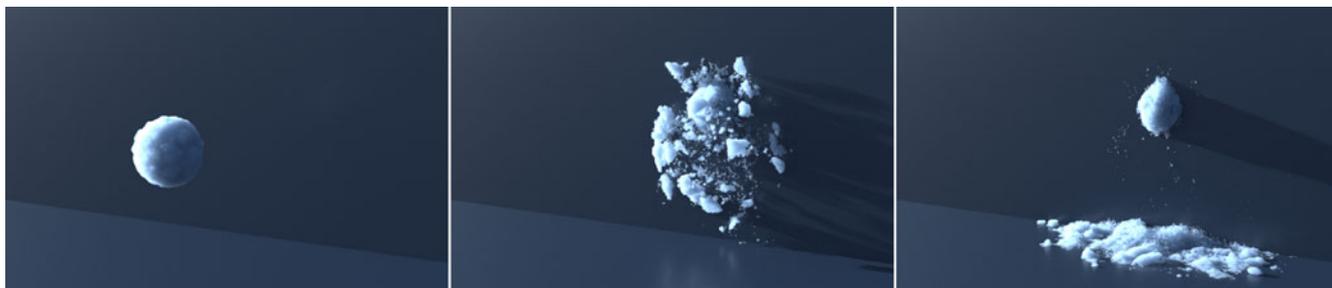


Fig. 17. A snowball smashes into a wall and sticks to it.

| Figure | Min Δt (s) | | Average Δt (s) | | Time/frame (s) | | Speedup factor | Grid size | # of particles |
|--------|----------------------|----------------------|------------------------|----------------------|----------------|------|----------------|-----------------------------|-------------------|
| | Ours | Orig | Ours | Orig | Ours | Orig | | | |
| 14 | 6.0×10^{-4} | 5.2×10^{-5} | 3.4×10^{-3} | 4.4×10^{-4} | 59 | 184 | 3.1 | $600 \times 300 \times 600$ | 2.8×10^5 |
| 17 | 5.4×10^{-4} | 1.7×10^{-6} | 1.5×10^{-3} | 1.4×10^{-4} | 85 | 431 | 5.1 | $200 \times 240 \times 600$ | 2.8×10^5 |
| 16 | 3.5×10^{-4} | 3.7×10^{-5} | 1.6×10^{-3} | 1.7×10^{-4} | 288 | 780 | 2.7 | $800 \times 300 \times 800$ | 5.6×10^5 |

Fig. 18. Performance comparison of our modified MPM snow formulation (“Ours”) with the original formulation (“Orig”).

energy to exist (which is how most constitutive models are defined anyway) and is sensitive to discontinuities in this energy. The method also occasionally fails to make progress due to valley shaped kinks in our collision processing. In practice, this only occurs when the system is already fairly close to a solution, since otherwise any energy kinks are overwhelmed by the strong gradients in the objective. From a practical perspective, this means this sort of breakdown can be dealt with by simply ignoring it. This does, however, prevent the method from being absolutely robust. We leave this weakness to be addressed in future work.

Our method was derived and implemented on top of a backward Euler integrator, which is known for being very stable but quite damped. The nonlinear system of equations for other A-stable integrators such as trapezoid rule and BDF-2 can also be readily converted into minimization form and solved similarly. Being second order schemes, their use would reduce damping at large time steps, though trapezoid rule’s oscillatory properties should be taken into account.

ACKNOWLEDGMENTS

The authors would like to acknowledge Shunsuke Saito and Yuwei Jiang for their suggestions regarding optimization. They were partially supported by the US National Science Foundation (NSF) (CCF-1422795), the US

Department of Energy (DOE) (09-LR-04-116741-BERA), ONR (N000140310071, N000141010730, N000141210834) and Intel STCVISUAL Computing Grant (20112360).

REFERENCES

- [1] D. Baraff and A. Witkin, “Large steps in cloth simulation,” in *Proc. SIGGRAPH*, 1998, pp. 43–54.
- [2] G. Hirota, S. Fisher, C. Lee, H. Fuchs et al., “An implicit finite element method for elastic solids in contact,” in *Proc. 14th Conf. Comput. Animation*, 2001, pp. 136–254.
- [3] P. Volino and N. Magnenat-Thalmann, “Comparing efficiency of integration methods for cloth simulation,” in *Proc. Comput. Graph. Int.*, 2001, pp. 265–272.
- [4] S. Martin, B. Thomaszewski, E. Grinspun, and M. Gross, “Example-based elastic materials,” *ACM Trans. Graph.*, vol. 30, no. 4, p. 72, 2011.
- [5] T. Liu, A. W. Bargteil, J. F. O’Brien, and L. Kavan, “Fast simulation of mass-spring systems,” *ACM Trans. Graph.*, vol. 32, no. 6, p. 214, 2013.
- [6] C. Kane, “Variational integrators and the newmark algorithm for conservative and dissipative mechanical systems,” Ph.D. dissertation, Caltech, Pasadena, CA, USA, 1999.
- [7] R. Bridson, R. Fedkiw, and J. Anderson, “Robust treatment of collisions, contact and friction for cloth animation,” *ACM Trans. Graph.*, vol. 21, no. 3, pp. 594–603, 2002.
- [8] R. Bridson, S. Marino, and R. Fedkiw, “Simulation of clothing with folds and wrinkles,” in *Proc. Symp. Comput. Animation*, 2003, pp. 28–36.
- [9] M. Hauth and O. Eitz, “A high performance solver for the animation of deformable objects using advanced numerical methods,” *Comput. Graph. Forum*, vol. 20, no. 3, pp. 319–328, 2001.
- [10] K.-J. Choi and H.-S. Ko, “Stable but responsive cloth,” in *Proc. ACM SIGGRAPH Courses*, 2005, p. 1.

- [11] B. Eberhardt, O. Eitzmuß, and M. Hauth, *Implicit-Explicit Schemes for Fast Animation with Particle Systems*. New York, NY, USA: Springer, 2000.
- [12] A. Stern and E. Grinspun, "Implicit-explicit variational integration of highly oscillatory problems," *Multiscale Model. Simul.*, vol. 7, no. 4, pp. 1779–1794, 2009.
- [13] D. Michels, G. Sobottka, and A. Weber, "Exponential integrators for stiff elastodynamic problems," *ACM Trans. Graph.*, vol. 33, pp. 7:1–7:20, 2013.
- [14] D. Parks and D. Forsyth, "Improved integration for cloth simulation," in *Proc. Eurographics*, 2002, <http://diglib2.eg.org/EG/DL/Conf/EG2002/short>
- [15] J. Su, R. Sheth, and R. Fedkiw, "Energy conservation for the simulation of deformable bodies," *IEEE Trans. Vis. Comput. Graph.*, vol. 19, no. 2, pp. 189–200, Feb. 2013.
- [16] J. C. Simo, N. Tarnow, and K. Wong, "Exact energy-momentum conserving algorithms and symplectic schemes for nonlinear dynamics," *Comput. Methods Appl. Mech. Eng.*, vol. 100, no. 1, pp. 63–116, 1992.
- [17] C. Kane, J. E. Marsden, and M. Ortiz, "Symplectic-energy-momentum preserving variational integrators," *J. Math. Phys.*, vol. 40, p. 3353, 1999.
- [18] A. Lew, J. Marsden, M. Ortiz, and M. West, "Variational time integrators," *Int. J. Num. Meth. Eng.*, vol. 60, no. 1, pp. 153–212, 2004.
- [19] L. Kharevych, W. Yang, Y. Tong, E. Kanso, J. E. Marsden, P. Schröder, and M. Desbrun, "Geometric, variational integrators for computer animation," in *Proc. Symp. Comput. Animation*, 2006, pp. 43–51.
- [20] M. Gonzalez, B. Schmidt, and M. Ortiz, "Force-stepping integrators in lagrangian mechanics," *Int. J. Num. Meth. Eng.*, vol. 84, no. 12, pp. 1407–1450, 2010.
- [21] R. Goldenthal, D. Harmon, R. Fattal, M. Bercovier, and E. Grinspun, "Efficient simulation of inextensible cloth," *ACM Trans. Graph.*, vol. 26, no. 3, p. 49, 2007.
- [22] T. F. Gast and C. Schroeder, "Optimization integrator for large time steps," in *Proc. Eurograph./ACM SIGGRAPH Symp. Comput. Animation*, pp. 31–40, 2014.
- [23] A. Stomakhin, C. Schroeder, L. Chai, J. Teran, and A. Selle, "A material point method for snow simulation," *ACM Trans. Graph.*, vol. 32, no. 4, pp. 102:1–102:10, Jul. 2013.
- [24] A. Pandolfi, C. Kane, J. Marsden, and M. Ortiz, "Time-discretized variational formulation of non-smooth frictional contact," *Int. J. Num. Meth. Engng.*, vol. 53, pp. 1801–1829, 2002.
- [25] J. Nocedal and S. Wright, *Numerical Optimization*, series Springer series in operations research and financial engineering. New York, NY, USA: Springer, 2006.
- [26] A. Stomakhin, R. Howes, C. Schroeder, and J. M. Teran, "Energetically consistent invertible elasticity," in *Proc. Symp. Comput. Animation*, 2012, pp. 25–32.
- [27] R. Schmedding and M. Teschner, "Inversion handling for stable deformable modeling," *Vis. Comput.*, vol. 24, pp. 625–633, 2008.
- [28] Y. Zhu, E. Sifakis, J. Teran, and A. Brandt, "An efficient and parallelizable multigrid framework for the simulation of elastic solids," *ACM Trans. Graph.*, vol. 29, pp. 16:1–16:18, 2010.
- [29] M. Müller and M. Gross, "Interactive virtual materials," in *Proc. Graph. Interface*, 2004, pp. 239–246.
- [30] O. Eitzmuss, M. Keckeisen, and W. Strasser, "A fast finite element solution for cloth modeling," in *Proc. 11th Pac. Conf. Comput. Graph. Appl.*, 2003, pp. 244–251.
- [31] I. Chao, U. Pinkall, P. Sanan, and P. Schröder, "A simple geometric model for elastic deformations," *ACM Trans. Graph.*, vol. 29, pp. 38:1–38:6, 2010.
- [32] A. McAdams, Y. Zhu, A. Selle, M. Empey, R. Tamstorf, J. Teran, and E. Sifakis, "Efficient elasticity for character skinning with contact and collisions," *ACM Trans. Graph.*, vol. 30, pp. 37:1–37:12, 2011.
- [33] G. Irving, J. Teran, and R. Fedkiw, "Invertible finite elements for robust simulation of large deformation," in *Proc. Symp. Comput. Animation*, 2004, pp. 131–140.
- [34] J. Teran, E. Sifakis, G. Irving, and R. Fedkiw, "Robust quasistatic finite elements and flesh simulation," in *Proc. Symp. Comput. Animation*, 2005, pp. 181–190.
- [35] L. Kharevych, W. Yang, Y. Tong, E. Kanso, J. Marsden, and P. Schröder, "Geometric, variational integrators for computer animation," in *Proc. Symp. Comput. Animation*, 2006, pp. 43–51.
- [36] D. M. Kaufman, S. Sueda, D. L. James, and D. K. Pai, "Staggered projections for frictional contact in multibody systems," in *ACM Trans. Graph.*, vol. 27, no. 5, p. 164, 2008.
- [37] J. Gascón, J. S. Zurdo, and M. A. Otaduy, "Constraint-based simulation of adhesive contact," in *Proc. Symp. Comput. Animation*, 2010, pp. 39–44.
- [38] C. Zheng and D. L. James, "Toward high-quality modal contact sound," *ACM Trans. Graph.*, vol. 30, no. 4, p. 38, 2011.



Theodore F. Gast received the BS degree in mathematics from Carnegie Mellon University in 2010. He is currently working toward the PhD degree at the University of California, Los Angeles (UCLA). He is also at Walt Disney Animation Studios, where he is putting the techniques used in this paper into production.



Craig Schroeder received the PhD degree in computer science from Stanford University in 2011 and is currently a postdoctoral scholar at the University of California, Los Angeles (UCLA). He received the Chancellor's Award for Postdoctoral Research in 2013, recognizing research impact and value to the UCLA community. He actively publishes in both computer graphics and computational physics. His primary areas of interest are solid mechanics and computational fluid dynamics and their applications to physically based animation for computer graphics. He began collaborating with Pixar Animation Studios during the PhD degree and later collaborated with Walt Disney Animation Studios during his postdoctoral studies. For his research contributions he received screen credits in Pixar's "Up" and Disney's "Frozen."



Alexey Stomakhin received the PhD degree in mathematics from the University of California, Los Angeles (UCLA) in 2013. His interest is primarily physics-based simulation for special effects, including simulation of fluids, solids and multimaterial interactions. He is currently employed at Walt Disney Animation Studios in Burbank, CA. He has received screen credits for his work on *Frozen* (2013) and *Big Hero 6* (2014).



Chenfanfu Jiang received the PhD degree in computer science from the University of California, Los Angeles (UCLA) in 2015. He is currently a postdoctoral researcher at UCLA, jointly appointed in the Departments of Mathematics and Computer Science. His primary research interests include solid/fluid mechanics and physics based animation.



Joseph M. Teran is a professor of applied mathematics at the University of California, Los Angeles (UCLA). His research focuses on numerical methods for partial differential equations in classical physics, including computational solids and fluids, multi-material interactions, fracture dynamics and computational biomechanics. He is also with Walt Disney Animation applying scientific computing techniques to simulate the dynamics of virtual materials like skin/soft tissue, water, smoke and recently, snow for the movie *Frozen*. He received a 2011 Presidential Early Career Award for Scientists and Engineers (PECASE) and a 2010 Young Investigator award from the Office of Naval Research.

▷ For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/publications/dlib.