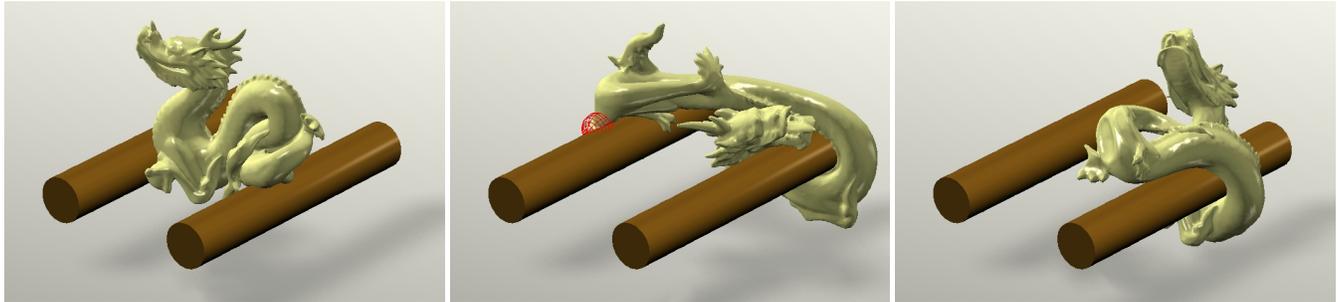


Descent Methods for Elastic Body Simulation on the GPU



(a) A resting dragon

(b) A stretched dragon

(c) A spiral dragon

Figure 1: *The dragon example. This model contains 16K vertices and 58K tetrahedra. Our elastic body simulator animates this example on the GPU at 30.5FPS, under the Mooney-Rivlin model. Thanks to a series of techniques we developed in this paper, the simulator can robustly handle very large time steps (such as $h = 1/30s$) and deformations.*

Abstract

In this paper, we show that many existing elastic body simulation approaches can be interpreted as descent methods, under a nonlinear optimization framework derived from implicit time integration. The key question is how to find an effective descent direction with a low cost. Based on this observation, we propose a novel gradient descent method using Jacobi preconditioning and Chebyshev acceleration. The convergence rate of this method is comparable to that of L-BFGS or nonlinear conjugate gradient. But unlike other methods, it requires no dot product operation, making it suitable for GPU implementation. To further ensure its convergence and performance, we develop a series of step length adjustment, initialization, and invertible model conversion techniques, all of which are compatible with GPU acceleration. Our experiment shows that the resulting simulator is simple, fast, scalable, memory-efficient, and robust against very large time steps and deformations. It can correctly simulate the deformation behaviors of many elastic materials, as long as their energy functions are second-order differentiable and the Hessian matrices can be quickly evaluated. For additional speedups, the method can serve as a complement to other real-time techniques as well, such as multi-grid.

Keywords: Nonlinear optimization, Newton’s method, implicit integration, gradient descent, Jacobi preconditioning, the Chebyshev semi-iterative method, GPU acceleration, hyperelasticity.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics—Animation.

1 Introduction

Solid materials often exhibit complex elastic behaviors in the real world. While we have seen a variety of models being developed to describe these behaviors over the past few decades, our ability to simulate them computationally is rather limited. Early simulation techniques often use explicit time integration, which is known for its numerical instability problem. A typical solution is to use implicit time integration instead. Given the nonlinear force-displacement relationship of an elastic material, we can formulate implicit time integration into a nonlinear system. Baraff and Witkin [1998] proposed to linearize this system at the current shape and solve the resulting linear system at each time step. Their method is equivalent to running one iteration of Newton’s method. Alternatively, we can linearize the system at the rest shape and solve a linear system with a constant matrix. While this method is fast thanks to matrix pre-factorization, the result becomes unrealistic under large deformation. To address this issue, Müller and Gross [2004] factored out the rotational component from the displacement and ended up with solving a new linear system at each time step again. Even if we accept formulating elastic simulation into a linear system, we still face the challenge of solving a large and sparse system. Unfortunately, many linear solvers are not fully compatible with parallel computing and they cannot be easily accelerated by the GPU.

In recent years, graphics researchers studied the use of geometric constraints and developed a number of constraint-based simulation techniques, such as strain limiting [Provot 1996; Thomaszewski et al. 2009; Wang et al. 2010], position-based dynamics [Müller et al. 2007; Müller 2008; Kim et al. 2012], and shape matching [Müller et al. 2005; Rivers and James 2007]. While these techniques are easy to implement and compatible with GPU acceleration, they offer little control on their underlying elastic models. To solve this problem, Liu and collaborators [2013] and Bouaziz and colleagues [2014] described geometric constraints as elastic energies in a quadratic form. This implies that their technique, known as *projective dynamics*, is not suitable for arbitrary elastic model. The recent work by Tournier and colleagues [2015] proposed to incorporate both elastic forces and compliant constraints into a single linear system. This technique is designed for highly stiff problems, where the condition number is more important than the problem size. It has to solve a linear system at each time step.

We think that a good elastic body simulation method should satisfy at least the following three requirements.

- **Generality.** A good method should be flexible enough to handle most elastic models, if not all. In particular, it should be able to simulate hyperelastic models, which use energy density functions to describe highly nonlinear force-displacement relationships.
- **Correctness.** Given sufficient computational resources, a good method should correctly simulate the behavior of a specified elastic model. In other words, the method is not just a temporary one for producing visually appealing animations. Instead, it can provide more accuracy for serious applications, once hardware becomes more powerful.
- **Efficiency.** A good method should be fast enough for real-time applications. It should also be compatible with parallel computing, so that it can benefit significantly from the use of graphics hardware and computer clusters.

While existing simulation methods can satisfy one or two of these requirements, none of them can satisfy all of the three, as far as we know. To develop a fast, flexible, and correct elastic body simulator, we made a series of technical contributions in this paper.

- **Insights.** We show that many recent methods, including position-based dynamics, projective dynamics and its accelerated version, can be viewed as descent methods under an energy minimization framework. The main question is how to find the descent direction, which differs in these methods.
- **Algorithm.** We propose to couple Jacobi preconditioning and Chebyshev acceleration with the gradient descent method. Our method offers a high convergence rate with a low computational cost. To further improve the performance of our method, we develop a number of techniques for step length adjustment, Chebyshev parameters, and initialization. The method is fully compatible with GPU acceleration.
- **Elastic model.** Many hyperelastic models were not designed for highly compressed or even inverted cases. To address this issue, we present a hybrid elastic model by mixing hyperelastic energy with projective dynamics energy. Our method can efficiently simulate this model, by interpolating forces and Hessian matrices on the fly.

In summary, our descent method handles any elastic model, if: 1) its energy function is second-order differentiable; and 2) the Hessian matrix of its energy function can be quickly evaluated. These two conditions can be satisfied by many elastic models, such as linear models, spring models, quadratic and cubic bending models [Bergou et al. 2006; Garg et al. 2007], and hyperelastic models. Given enough iterations, our method converges to exact implicit Euler integration under a given elastic model. It is robust against divergence, even when handling large time steps and deformations as Figure 1 shows. The whole method is fast, scalable and has a small memory demand. For more speedups, it can also be combined with multi-grid techniques, many of which were designed for hexahedral lattices [Zhu et al. 2010; McAdams et al. 2011b; Dick et al. 2011; Patterson et al. 2012] at this time.

2 Related Work

The simulation of elastic bodies is an important research topic in computer graphics, since the pioneer work by Terzopoulos and colleagues [1987]. Many early techniques use explicit time integration schemes, which are easy to implement but require sufficiently small time steps to avoid numerical instability. To simulate cloth and thin shells using a large time step, Baraff and Witkin [1998] advocated

the use of implicit time integration schemes. If we assume that elastic force is a linear function of vertex displacement, the implicit Euler scheme forms a linear system with a constant matrix, which can be pre-factorized for fast linear solve. Since linear elastic force is not rotation-invariant, it can cause unrealistic volume growth when an object is under large rotation. Müller and Gross [2004] alleviated this problem by factoring out the rotational component in their corotational method. For more accurate simulation of real-world elastic bodies, we must use nonlinear elastic force and form the implicit scheme into a nonlinear system. A typical solution to a nonlinear system is Newton’s method, which needs a large computational cost to evaluate the Hessian matrix and solve a linearized system in every iteration. Teran and colleagues [2005] developed a technique to evaluate the Hessian matrix under a hyperelastic model, so they can use the implicit scheme to handle hyperelastic bodies. Although the implicit scheme is more numerically stable, it suffers from artificial damping. To overcome this issue, Kharevych and colleagues [2006] suggested to use symplectic integrators. Hybrid implicit-explicit integration is another technique for reducing artificial damping, as Bridson and collaborators [2003] and Stern and Grinspun [2009] demonstrated. For a mass-spring system, Su and colleagues [2013] investigated how to track and preserve the total system energy over time. Daviet and collaborators [2011] studied the development of a fast iterative solver for handling Coulomb friction in hair dynamics.

The force-displacement relationship of a real-world elastic material, such as human skin, is often highly nonlinear. This non-linearity makes the material difficult and expensive to handle in physics-based simulation. A simple way to generate nonlinear effects without using an elastic model is to apply geometric constraints on springs [Provot 1996], or triangular and tetrahedral elements [Thomaszewski et al. 2009; Wang et al. 2010]. Müller and colleagues [2007; 2008; 2012] pushed this idea even further, by using geometric constraints to replace elastic forces in a mass-spring system. Later they extended this position-based method to simulate fluids [Macklin and Müller 2013; Macklin et al. 2014] and deformable bodies [Müller et al. 2014]. Similar to position-based method, shape matching [Müller et al. 2005; Rivers and James 2007] also uses the difference between deformed shapes and rest shapes to simulate elastic behaviors. Instead of using geometric constraints, Perez and collaborators [2013] applied energy constraints to produce nonlinear elastic effects.

An interesting question is whether there is a connection between an elastic model and a geometric constraint. Liu and collaborators [2013] found that the elastic spring energy can be treated as a compliant spring constraint. Based on this fact, they developed an implicit mass-spring simulator, which iteratively solves a local constraint enforcement step and a global linear system step. Bouaziz and colleagues [2014] formulated this method into projective dynamics, by defining the elastic energy of a triangular or tetrahedral element as a constraint. The main advantage of projective dynamics is that the system matrix involved in the global step is constant, so it can be pre-factorized for fast solve. On the GPU, Wang [2015] proposed to solve projective dynamics by the Jacobi method and the Chebyshev semi-iterative method, both of which are highly suitable for parallel computing. Recently, Tournier and colleagues [2015] presented a stable way to solve elastic forces and compliant constraints together using a single linear system. Their method reduces the condition number of the system, at the cost of an increased system size.

3 Descent Methods

Let $\mathbf{q} \in \mathbb{R}^{3N}$ and $\mathbf{v} \in \mathbb{R}^{3N}$ be the vertex position and velocity vectors of a nonlinear elastic body. We can use implicit time integration to

Algorithm 1 Descent_Optimization

```

Initialize  $\mathbf{q}^{(0)}$ ;
for  $k = 0 \dots K - 1$  do
  Calculate the descent direction  $\Delta\mathbf{q}^{(k)}$ ;           Step 1
  Adjust the step length  $\alpha^{(k)}$ ;                 Step 2
   $\bar{\mathbf{q}}^{(k+1)} \leftarrow \mathbf{q}^{(k)} + \alpha^{(k)} \Delta\mathbf{q}^{(k)}$ ;   Step 3
   $\mathbf{q}^{(k+1)} \leftarrow \text{Acceleration}(\bar{\mathbf{q}}^{(k+1)}, \bar{\mathbf{q}}^{(k)}, \mathbf{q}^{(k)}, \mathbf{q}^{(k-1)})$ ; Step 4
return  $\mathbf{q}^{(K)}$ ;

```

189 simulate the deformation of the body from time t to $t + 1$ as:

$$\mathbf{q}_{t+1} = \mathbf{q}_t + h\mathbf{v}_{t+1}, \quad \mathbf{v}_{t+1} = \mathbf{v}_t + h\mathbf{M}^{-1}\mathbf{f}(\mathbf{q}_{t+1}), \quad (1)$$

190 in which $\mathbf{M} \in \mathbb{R}^{3N \times 3N}$ is the mass matrix, h is the time step, and
 191 $\mathbf{f} \in \mathbb{R}^{3N}$ is the total force as a function of \mathbf{q} . By combining the two
 192 equations, we obtain a single nonlinear system:

$$\mathbf{M}(\mathbf{q}_{t+1} - \mathbf{q}_t - h\mathbf{v}_t) = h^2\mathbf{f}(\mathbf{q}_{t+1}). \quad (2)$$

193 Since $\mathbf{f}(\mathbf{q}) = -\partial E(\mathbf{q})/\partial\mathbf{q}$, where $E(\mathbf{q})$ is the total potential energy
 194 evaluated at \mathbf{q} , we can convert the nonlinear system into an uncon-
 195 strained nonlinear optimization problem: $\mathbf{q}_{t+1} = \arg \min \epsilon(\mathbf{q})$,

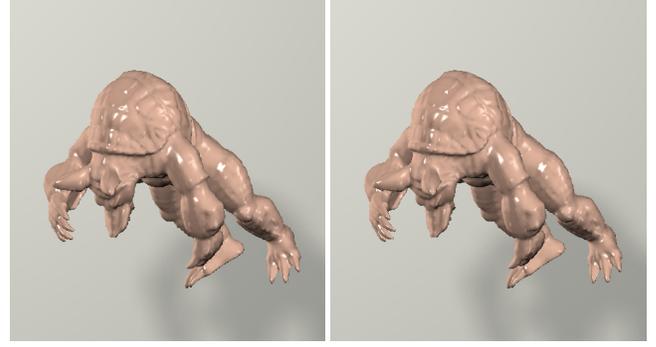
$$\epsilon(\mathbf{q}) = \frac{1}{2h^2} \|\mathbf{q} - \mathbf{q}_t - h\mathbf{v}_t\|_{\mathbf{M}}^2 + E(\mathbf{q}). \quad (3)$$

196 Nonlinear optimization is often solved by descent methods, which
 197 contain four steps in each iteration as Algorithm 1 shows. Their
 198 main difference is in how to calculate the descent direction from
 199 the gradient: $\mathbf{g}^{(k)} = \nabla\epsilon(\mathbf{q}^{(k)})$.

200 **Gradient descent.** The gradient descent method simply sets the
 201 descent direction as: $\Delta\mathbf{q}^{(k)} = -\mathbf{g}^{(k)}$, using the fact that $\epsilon(\mathbf{q})$ decreases
 202 fastest locally in the negative gradient direction. While gradient de-
 203 scent has a small computational cost per iteration, its convergence
 204 rate is only linear as shown in Figure 2c. Gradient descent can be
 205 viewed as updating \mathbf{q} by the force, since the negative gradient of
 206 the potential energy is the force. This is fundamentally similar to
 207 explicit time integration. Therefore, it is not surprising to see the
 208 step length must be small to avoid the divergence issue.

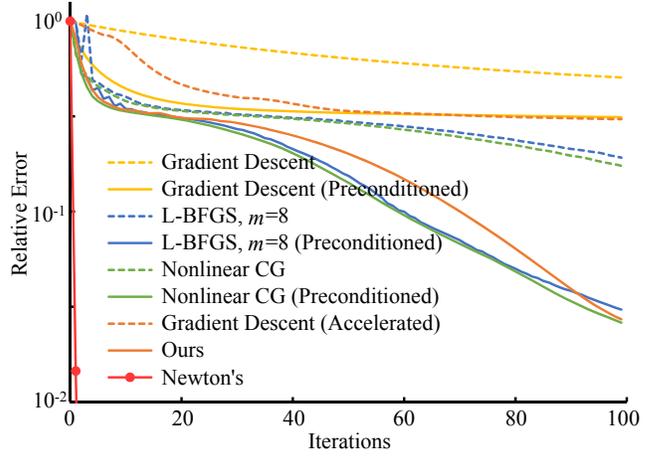
209 **Newton's method.** To achieve quadratic convergence, Newton's
 210 method approximates $\epsilon(\mathbf{q}^{(k)})$ by a quadratic function and it calcu-
 211 lates the search direction as: $\Delta\mathbf{q}^{(k)} = -(\mathbf{H}^{(k)})^{-1}\mathbf{g}^{(k)}$, where $\mathbf{H}^{(k)}$ is the
 212 Hessian matrix of $\epsilon(\mathbf{q})$ evaluated at $\mathbf{q}^{(k)}$. Figure 2c shows Newton's
 213 method converges the fastest. However, it is too computationally
 214 expensive to solve the linear system $\mathbf{H}^{(k)}\Delta\mathbf{q}^{(k)} = -\mathbf{g}^{(k)}$ involved in
 215 every iteration. To handle one linear system in the armadillo exam-
 216 ple as Figure 2 shows, the Eigen library needs 0.65 seconds by C-
 217 holesky factorization, or 2.82 seconds by preconditioned conjugate
 218 gradient with incomplete LU factorization. Unfortunately, many
 219 linear solvers cannot be easily parallelized for GPU acceleration.

220 **Quasi-Newton methods.** Since it is too expensive to solve a
 221 linear system or even just evaluate the Hessian matrix, a natural
 222 idea is to approximate the Hessian matrix or its inverse. For exam-
 223 ple, quasi-Newton methods, such as BFGS, use previous gradient
 224 vectors to approximate the inverse Hessian matrix directly. To avoid
 225 storing a dense inverse matrix, L-BFGS defines the approximation
 226 by m gradient vectors, each of which provides rank-one updates to
 227 the inverse matrix sequentially. While L-BFGS converges slow-
 228 er than Newton's method, it has better performance thanks to its
 229 reduced cost per iteration. Unfortunately, the sequential nature of
 230 L-BFGS makes it difficult to run on the GPU, unless the problem is
 231 also subject to box constraints [Fei et al. 2014].



(a) Ground truth

(b) Our result



(c) The convergence plot

Figure 2: The outcomes of descent methods applied to the armadillo example. Thanks to preconditioning and momentum-based acceleration, our method converges as fast as nonlinear conjugate gradient and it needs a much smaller GPU cost. Our result in (b) is visually indistinguishable from the ground truth in (a) generated by Newton's method. In the plot, we define the relative error as $(\epsilon(\mathbf{q}^{(k)}) - \epsilon(\mathbf{q}^*)) / (\epsilon(\mathbf{q}^{(k)}) - \epsilon(\mathbf{q}^{(0)}))$, where $\mathbf{q}^{(k)}$ is the result in the k -th iteration and \mathbf{q}^* is the ground truth.

232 **Nonlinear conjugate gradient (CG).** The nonlinear conjugate
 233 gradient method generalizes the conjugate gradient method to non-
 234 linear optimization problems. Based on the Fletcher–Reeves for-
 235 mula, it calculates the descent direction as:

$$\Delta\mathbf{q}^{(k)} = -\mathbf{g}^{(k)} + \frac{z^{(k)}}{z^{(k-1)}}\Delta\mathbf{q}^{(k-1)}, \quad z^{(k)} = \mathbf{g}^{(k)} \cdot \mathbf{g}^{(k)}. \quad (4)$$

236 Nonlinear CG is highly similar to L-BFGS with $m = 1$. The reason
 237 it converges slightly faster than L-BFGS in our experiment is be-
 238 cause our implementation uses the exact Hessian matrix to estimate
 239 the step length. Intuitively, this is identical to conjugate gradient,
 240 except that the residual vector, i.e., the gradient, is recalculated in
 241 every iteration. Nonlinear CG is much more friendly with GPU ac-
 242 celeration than quasi-Newton methods. But it still requires multiple
 243 dot product operations, which restrict its performance on the GPU.

4 Our Descent Method

244 In this section, we will describe the technique used in our descent
 245 method. We will also evaluate their performance and compare them
 246 with alternatives.
 247

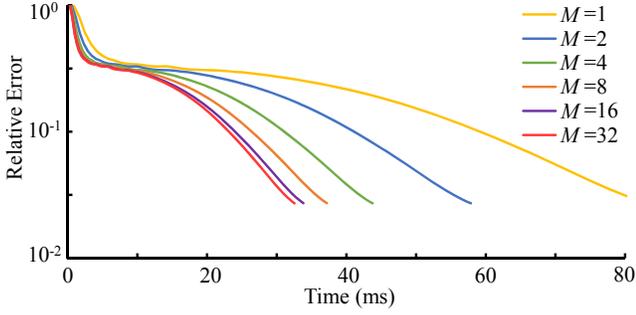


Figure 3: The convergence of our method, when using different M values to delay matrix evaluation. This plot shows that the method can reach the same residual error within 96 iterations, regardless of M . So we can use a larger M to reduce the matrix evaluation cost and improve the system performance.

4.1 Descent Direction

The idea behind our method is originated from preconditioned conjugate gradient. To achieve faster convergence, preconditioning converts the optimization problem into a well conditioned one:

$$\bar{\mathbf{q}} = \arg \min \epsilon(\mathbf{P}^{-1/2}\bar{\mathbf{q}}), \quad \text{for } \bar{\mathbf{q}} = \mathbf{P}^{1/2}\mathbf{q}, \quad (5)$$

where \mathbf{P} is the preconditioner matrix. Mathematically, doing this is equivalent¹ to replacing $\mathbf{g}^{(k)}$ by $\mathbf{P}^{-1}\mathbf{g}^{(k)}$ in Equation 4. Among all of the preconditioners, we favor the Jacobi preconditioner the most, since it is easy to implement and friendly with GPU acceleration. When an optimization problem is quadratic, conjugate gradient defines the Jacobi preconditioner as a constant matrix: $\mathbf{P} = \text{diag}(\mathbf{H})$, where \mathbf{H} is the constant Hessian matrix. To solve a general nonlinear optimization problem, if the Hessian matrix can be quickly evaluated in every iteration, we can treat $\mathbf{P}(\mathbf{q}^{(k)}) = \text{diag}(\mathbf{H}^{(k)})$ as the Jacobi preconditioner for nonlinear CG, which now varies from iteration to iteration. Such a Jacobi preconditioner significantly improves the convergence rate of nonlinear CG, as Figure 2c shows.

This Jacobi preconditioner can be effectively applied to L-BFGS and gradient descent as well. Preconditioning in L-BFGS is essentially defining $\text{diag}^{-1}(\mathbf{H}^{(k)})$ as the initial inverse Hessian estimate. Meanwhile, preconditioned gradient descent simply defines its new descent direction as: $\Delta\mathbf{q}^{(k)} = -\text{diag}^{-1}(\mathbf{H}^{(k)})\mathbf{g}^{(k)}$. While preconditioned gradient descent does not converge as fast as other preconditioned methods, it owns a unique and critical property: its convergence rate can be well improved by momentum-based techniques. So we propose to formulate our basic method as accelerated, Jacobi preconditioned gradient descent. Figure 2 demonstrates that the convergence rate of our method is comparable to that of preconditioned nonlinear CG, and our result is visually similar to the ground truth after 96 iterations.

Why is our method special? While both Jacobi preconditioning and momentum-based acceleration are popular techniques, it is uncommon to see them working with gradient descent. There are reasons for this. The use of Jacobi preconditioning destroys the advantage of gradient descent in requiring no matrix evaluation. Meanwhile, Chebyshev acceleration is effective only when the problem is mildly nonlinear [Gutknecht and Röllin 2002]. So our method is not suitable for general nonlinear optimization problems. Fortunately, it works well with elastic body simulation.

Convergence and performance. The calculation of our descent direction has two obvious advantages. First, the diagonal

¹The calculation of $\mathbf{z}^{(k)}$ should be updated as: $\mathbf{z}^{(k)} = \mathbf{g}^{(k)} \cdot \mathbf{P}^{-1}\mathbf{g}^{(k)}$.

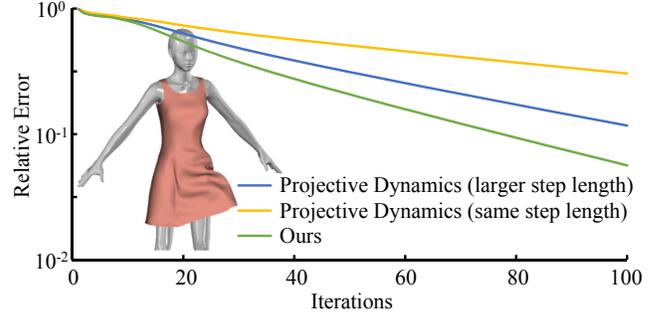


Figure 4: The convergence of our method and projective dynamics. Although projective dynamics can use a large step length, it cannot converge as fast as our method.

entries of the Hessian matrix are typically positive. As a result, $\text{diag}^{-1}(\mathbf{H}^{(k)})$ is positive definite and $\Delta\mathbf{q}^{(k)} \cdot \mathbf{g}^{(k)} < 0$. Within a bounded deformation space, the Hessian matrix of $\epsilon(\mathbf{q})$ is also bounded: $\mathbf{H} \leq \mathbf{B}\mathbf{I}$. We have:

$$\epsilon(\mathbf{q}^{(k)} + \alpha^{(k)}\Delta\mathbf{q}^{(k)}) \leq \epsilon(\mathbf{q}^{(k)}) + \alpha^{(k)}\Delta\mathbf{q}^{(k)} \cdot \mathbf{g}^{(k)} + \frac{\mathbf{B}}{2} \|\alpha^{(k)}\Delta\mathbf{q}^{(k)}\|_2^2. \quad (6)$$

So there must exist a sufficiently small step length $\alpha^{(k)}$ that ensures the energy decrease and eliminates the divergence issue. Second, both the Jacobi preconditioner and gradient descent are computationally inexpensive and suitable for parallelization. In particular, it requires zero reduction operation.

The use of the Jacobi preconditioner demands the evaluation of the Hessian matrix. This can become a computational bottleneck if it is done in every iteration. Fortunately, we found that it is acceptable to evaluate the Hessian matrix once every M iterations and use the last matrix for the preconditioner. Figure 3 shows that this strategy has little effect on the convergence rate, but significantly reduces the computational cost per iteration.

Comparison to projective dynamics. The recent projective dynamics technique [Liu et al. 2013; Bouaziz et al. 2014] solves the optimization problem by interleaving a local constraint step and a global solve step. If we view the local step as calculating the gradient and the global step as calculating the descent direction, we can interpret projective dynamics as a preconditioned gradient descent method as well. Here the preconditioner matrix is constant, so it can be pre-factored for fast solve in every iteration. But this is not the only advantage of projective dynamics. Bouaziz and collaborators [2014] pointed out that projective dynamics is guaranteed to converge by setting $\alpha^{(k)} \equiv 1$, if the elastic energy of every element has a quadratic form $\|\mathbf{A}\mathbf{q} - \mathbf{B}\mathbf{p}(\mathbf{q})\|^2$, where \mathbf{A} and \mathbf{B} are two constant matrices and $\mathbf{p}(\mathbf{q})$ is the geometric projection of \mathbf{q} according to that element. Therefore, projective dynamics does not need to adjust the step length in every iteration.

Projective dynamics was originally not suitable for GPU acceleration. Wang [2015] addressed this problem by removing off-diagonal entries of the preconditioner matrix. In this regard, that method is highly related to our method. Since both methods can handle mass-spring systems, we compare their convergence rates as shown in Figure 4. When both methods use the same step length: $\alpha^{(k)} \equiv 0.5$, our method converges significantly faster. This is not a surprise, given the fact that our method uses the diagonal of the exact Hessian matrix and Newton’s method converges faster than original projective dynamics. The strength of projective dynamics allows it to use $\alpha^{(k)} \equiv 1$. But even so, it is still not comparable to our method. Interestingly, we do not observe substantial difference in animation results of the two methods. We guess this is because

332 the stiffness in this example is too large. As a result, small energy
 333 difference cannot cause noticeable difference in vertex positions.

334 **Comparison to a single linear solve.** Figure 2 may leave an
 335 impression that it is always acceptable to solve just one Newton’s iter-
 336 ation, as did in many existing simulators [Baraff and Witkin 1998;
 337 Dick et al. 2011]. Mathematically, it is equivalent to approximating
 338 the energy by a quadratic function and solving the resulting linear
 339 system. In that case, our method is simplified to the accelerated
 340 Jacobi method. Doing this has a clear advantage: the gradient does
 341 not need to be reevaluated in every iteration, which can be costly
 342 for tetrahedral elements. However, Newton’s method may diverge,
 343 especially if the time step is large and the initialization is bad. This
 344 problem can be lessened by using a small step length. But then
 345 it becomes pointless to waste computational resources within one
 346 Newton’s iteration. In contrast, gradient descent still converges
 347 reasonably well under the same situation. So we decide not to use
 348 quadratic approximation, i.e., one Newton’s iteration.

349 **Comparison to nonlinear CG.** The biggest competitor of our
 350 method is actually nonlinear CG. Figure 2c shows that the two
 351 methods have similar convergence rates. So the difference in their
 352 performance is mainly determined by the computational cost per
 353 iteration. While the two methods have similar performance on the
 354 CPU, our method runs three to four times faster than nonlinear CG
 355 on the GPU. This is because nonlinear CG must perform at least two
 356 dot product operations, each of which takes 0.41ms in the armadillo
 357 example using the CUDA thrust library. In contrast, the cost of our
 358 method is largely due to gradient evaluation, which takes 0.17ms
 359 per iteration and is required by nonlinear CG as well.

360 Similar to our method, nonlinear CG also needs to use a smaller
 361 step length when the energy function becomes highly nonlinear.
 362 But unlike our method, it does not need momentum-based accel-
 363 eration or parameter tuning. So if parallel architecture can allow dot
 364 products to be quickly handled in the future, it may be preferable to
 365 use nonlinear CG instead.

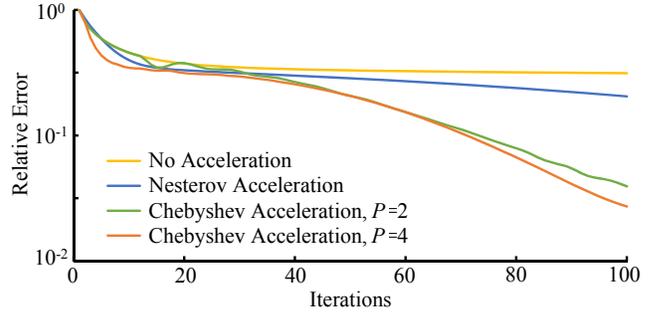
366 4.2 Step Length Adjustment

367 Given the search direction $\Delta\mathbf{q}^{(k)}$, the next question is how to calcu-
 368 late a suitable step length $\alpha^{(k)}$. A simple yet effective approach,
 369 known as *backtracking line search*, gradually reduces the step
 370 length, until the first Wolfe condition gets satisfied:

$$\epsilon(\mathbf{q}^{(k)} + \alpha^{(k)}\Delta\mathbf{q}^{(k)}) < \epsilon(\mathbf{q}^{(k)}) + c^{(k)}\alpha^{(k)}\Delta\mathbf{q}^{(k)} \cdot \mathbf{g}^{(k)}, \quad (7)$$

371 in which c is a control parameter. The Wolfe condition is straight-
 372 forward to evaluate on the CPU. However, it becomes problematic
 373 on the GPU, due to expensive energy summation and dot product
 374 operations. To reduce the computational cost, we propose to elim-
 375 inate the dot product by setting $c = 0$. Intuitively, it means we just
 376 search for the largest $\alpha^{(k)}$ that ensures monotonic energy decrease:
 377 $\epsilon(\mathbf{q}^{(k)} + \alpha^{(k)}\Delta\mathbf{q}^{(k)}) < \epsilon(\mathbf{q}^{(k)})$. We also propose to evaluate the energy
 378 every eight iterations only. Doing this can waste more iterations,
 379 before a suitable step length is found. But once it gets found, the
 380 method needs only a small energy summation cost afterwards.

381 Our simulator explores the continuity of α between two successive
 382 time steps. Specifically, it initializes the step length at time $t + 1$
 383 as $\alpha = \alpha_t/\gamma$, in which α_t is the ending step length at time t . After
 384 that, the simulator gradually reduces α by $\alpha := \gamma\alpha$, until the Wolfe
 385 condition gets satisfied. In our experiment, we use $\gamma = 0.7$. When
 386 the step length is too small, our method converges slowly and it is
 387 not worthwhile to spend more iterations. So if the Wolfe condition
 388 still cannot be satisfied once the step length reaches a minimum
 389 value, we simply end that time step and start the next one.



390 **Figure 5:** The convergence of our method with different accelera-
 391 tion techniques. By using multiple phases, the Chebyshev method
 392 can more effectively accelerate the convergence process.

4.3 Momentum-based Acceleration

391 An important strength of our method is that it can benefit from
 392 the use of momentum-based acceleration techniques, such as the
 393 Chebyshev semi-iterative method [Golub and Van Loan 1996] and
 394 the Nesterov’s method [Nesterov 2004]. Both methods use the
 395 “momentum”, the result change between the last two iterations, to
 396 improve the current search. Because the result change is calculated
 397 independently for every vertex, both methods are naturally compat-
 398 ible with parallel computing.

399 The two methods differ in how they define and weight the result
 400 change. The weight used by the Chebyshev method is calculated
 401 from the gradient decrease rate, which can be tuned for differ-
 402 ent problems as shown in [Wang 2015]. On the other hand, the
 403 control parameter used by the Nesterov’s method is related to the
 404 strong convexity of the Hessian matrix. Since this parameter is
 405 not easy to find, it is often set to zero for simplicity. Because of
 406 such a difference, the Chebyshev method typically outperforms the
 407 Nesterov’s method, as shown in Figure 5. Our experiment shows
 408 that the Chebyshev method is also more reliable, as long as the
 409 gradient decrease rate is underestimated. In contrast, the Nes-
 410 terov’s method may need multiple restarts to avoid the divergence
 411 issue [O’donoghue and Candès 2015].

412 We note that neither of the techniques was designed for general de-
 413 scendent methods. The Chebyshev method was initially developed for
 414 linear solvers, while the Nesterov’s method was proposed for speed-
 415 ing up the gradient descent method. Since our method is highly
 416 related to linear solvers² and gradient descent, it can be effectively
 417 accelerated by momentum-based acceleration techniques. Neither
 418 L-BFGS nor nonlinear CG can be accelerated by these techniques,
 419 according to our experiment.

420 **Adaptive parameters.** When Wang [2015] adopted the Cheby-
 421 shev method for accelerating projective dynamics, he defined the
 422 gradient decrease rate ρ as a constant:

$$\rho \approx \frac{\|\nabla\epsilon(\mathbf{q}^{k+1})\|}{\|\nabla\epsilon(\mathbf{q}^k)\|}. \quad (8)$$

423 This is a reasonable practice, since the rate is related to the spectral
 424 radius of the global matrix, which stays the same through the whole
 425 simulation process. The simulation of generic elastic materials,
 426 however, can exhibit more complex convergence behaviors. So if
 427 a constant ρ is still used, it must be kept at the minimum level to
 428 avoid oscillation or even divergence issues, especially in the first
 429 few iterations. To make Chebyshev acceleration more effective,
 430 we propose to use a varying ρ instead. Specifically, we divide the

²Our method can also be viewed as solving each Newton’s iteration by
 only one iteration of the Jacobi method.

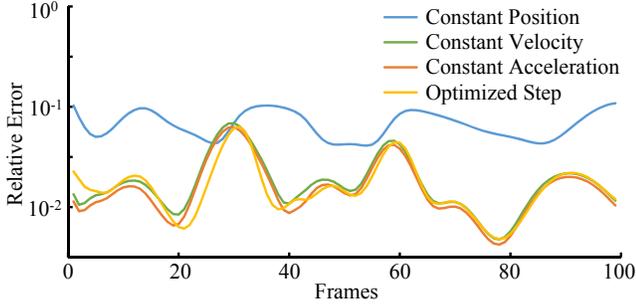


Figure 6: The convergence of our method using different initialization approaches. This plot shows that the constant acceleration approach works the best in most cases.

iterations into P phases and assign each phase with its own ρ . We can then perform the transition from one phase to another by simply restarting the Chebyshev method. The question is: *how can we tune the phases and their parameters?* Our idea is to change the length and the parameter of a phase each time, and then test whether that helps the algorithm reduce the residual error in pre-simulation. We slightly increase or decrease ρ each time by:

$$\rho^{\text{new}} = 1 - (1 \pm \epsilon)(1 - \rho), \quad (9)$$

in which ϵ is typically set to 0.05. We accept the change that causes the most significant error decrease, and then start another tuning cycle. The tuning process terminates once the error cannot be reduced any further. Figure 5 compares the convergence rates of our method, by using two and four Chebyshev phases respectively.

4.4 Initialization

The initialization of $\mathbf{q}^{(0)}$ is also an important component in our algorithm. It helps the descent method reduce the total energy to a low level, after a fixed number of iterations. Intuitively, the initialization works as a prediction on the solution \mathbf{q}_{t+1} . Here we test four different prediction approaches. The first three assume that vertex positions, velocities, and accelerations are constant, respectively: $\mathbf{q}_{t+1} \approx \mathbf{q}^{(0)} = \mathbf{q}_t$; $\mathbf{q}_{t+1} \approx \mathbf{q}^{(0)} = \mathbf{q}_t + h\mathbf{v}_t$; $\mathbf{q}_{t+1} \approx \mathbf{q}^{(0)} = \mathbf{q}_t + h\mathbf{v}_t + \eta h(\mathbf{v}_t - \mathbf{v}_{t-1})$. We use the parameter η to damp the acceleration effect, which is typically set to 0.2. The fourth approach assumes that vertices move in the \mathbf{v}_t direction with an unknown step distance d : $\mathbf{q}^{(0)} = \mathbf{q}_t + d\mathbf{v}_t$. We then optimize d by minimizing a quadratic approximation of $\epsilon(\mathbf{q}_t + d\mathbf{v}_t)$:

$$d = \arg \min_d \left\{ \epsilon(\mathbf{q}_t) + (d\mathbf{v}_t) \cdot \nabla \epsilon(\mathbf{q}_t) + \frac{1}{2} (d\mathbf{v}_t) \cdot \mathbf{H}(\mathbf{q}_t) (d\mathbf{v}_t) \right\}, \quad (10)$$

which can be solved as a simple linear equation. This is similar to how the conjugate gradient method determines the optimal step in a search direction.

Figure 6 compares the effects of the four approaches on the convergence of our method, over a precomputed sequence with 100 frames. It shows that the optimized step approach does not outperform the constant acceleration approach in most cases, even though it is the most complex one. Because of this, our system chooses the constant acceleration approach to initialize $\mathbf{q}^{(0)}$ by default. We note that Figure 6 illustrates the errors during a single frame only. These errors can be accumulated over time, causing slightly larger differences in simulation results. These differences are often manifested as small artificial damping artifacts, as shown in our experiment.

5 Nonlinear Elastic Models

Our new descent method can handle any elastic model, if: 1) its energy function is second-order differentiable; and 2) the Hessian matrix of its energy function can be quickly evaluated. These two conditions are satisfied by many elastic models, such as spring model under Hooke's law, quadratic or cubic bending models [Bergou et al. 2006; Garg et al. 2007], and hyperelastic models. In this section, we would like to specifically discuss hyperelastic models, some of which are not suitable for immediate use in simulation.

5.1 Hyperelasticity

Hyperelastic models are developed by researchers in mechanical engineering and computational physics to model complex force-displacement relationships of real-world materials. The energy density function of an isotropic hyperelastic material is typically defined by the three invariants³ of the right Cauchy-Green deformation tensor $\mathbf{C} = \mathbf{F}^T \mathbf{F}$:

$$I = \text{tr}(\mathbf{C}), \quad II = \text{tr}(\mathbf{C}^2), \quad III = \det(\mathbf{C}). \quad (11)$$

Here \mathbf{F} is the deformation gradient. For example, the St. Venant-Kirchhoff model has the following strain energy density function:

$$W = \frac{s_0}{2}(I - 3)^2 + \frac{s_1}{4}(II - 2I + 3), \quad (12)$$

where s_0 and s_1 are the two elastic moduli controlling the resistance to deformation, also known as the Lamé parameters. The compressible neo-Hookean model [Ogden 1997] defines its strain energy density function as:

$$W = s_0(III^{-1/3} \cdot I - 3) + s_1(III^{-1/2} - 1), \quad (13)$$

in which s_0 is the shear modulus and s_1 is the bulk modulus. Many hyperelastic models can be considered as extensions of the neo-Hookean model. For example, the compressible Mooney-Rivlin model for rubber-like materials uses the following strain energy density function [Macosko 1994]:

$$W = s_0(III^{-1/3} \cdot I - 3) + s_1(III^{-1/2} - 1) + s_2 \left(\frac{1}{2} III^{-2/3} (I^2 - II) - 3 \right). \quad (14)$$

To model the growing stiffness of soft tissues, the isotropic Fung model [Fung 1993] uses an exponential term:

$$W = s_0(III^{-1/3} \cdot I - 3) + s_1(III^{-1/2} - 1) + s_2 \left(e^{s_3(III^{-1/3} \cdot I - 3)} - 1 \right), \quad (15)$$

in which s_3 controls the speed of the exponential growth.

Invertible model conversion. A practical problem associated with the use of hyperelastic models is that they are not designed for highly compressed or inverted cases. As a result, a simulated hyperelastic body can become unnecessarily stiff, or even stuck in an inverted shape. A common solution to this problem is to set a limit on the compression rate or the stress, as described by Irving and colleagues [2004]. Since such a limit will cause C^2 discontinuity in the deformation energy, we choose not to do so in our system.

Our solution is to use projective dynamics instead. Bouaziz and colleagues [2014] proved that projective dynamics is numerically robust, even against inverted cases. Its basic form uses the following energy density function:

$$W^{\text{proj}} = \sum_{i=1}^3 (\lambda_i - 1)^2, \quad (16)$$

³Tensor invariants can be formulated in other ways. For example, it is also common to define the second variant as: $II = \frac{1}{2}(\text{tr}^2(\mathbf{C}) - \text{tr}(\mathbf{C}^2))$. In this paper, we follow the definition used in [Teran et al. 2005], since we will use their formula to derive the Hessian matrix later.

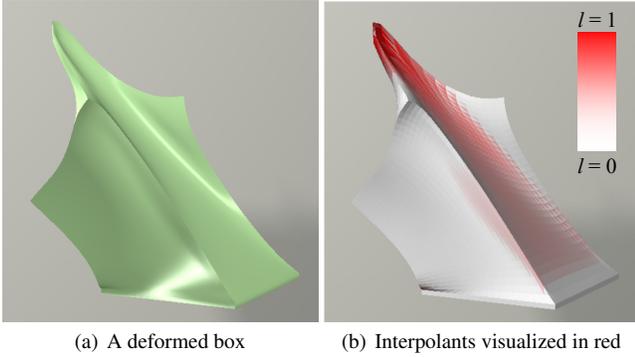


Figure 7: A deformed box and its interpolants. For the *St. Venant-Kirchhoff* model, we set $\lambda^+ = 0.5$ to address its low resistance against compression. Even so, only a small number of elements need to use invertible model conversion.

Name	#vert	#ele	CPU Cost	GPU Cost	GPU FPS
Dragon (Fig. 1)	16K	58K	6.75s	32.8ms	30.5
Armadillo (Fig. 2)	15K	55K	6.18s	31.4ms	31.8
Box (Fig. 10)	14K	72K	7.12s	37.6ms	26.6
Dress (Fig. 4)	15K	44K	1.35s	26.6ms	37.6
Double helix (Fig. 9)	13K	41K	4.72s	27.5ms	36.4
Double helix (Fig. 9)	24K	82K	9.67s	38.5ms	26.0
Double helix (Fig. 9)	48K	158K	19.8s	65.4ms	15.3
Double helix (Fig. 9)	96K	316K	38.8s	12.2ms	8.2

Table 1: Statistics and timings of our examples. The computational time depends on the number of tetrahedra and iterations.

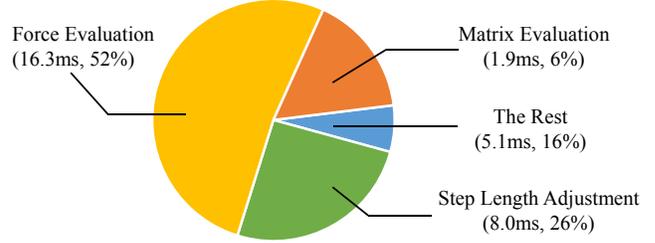


Figure 8: The Breakdown of the computational time. This pie chart reveals that force evaluation is the most expensive step.

GPU implementation. In our GPU implementation, we handle each iteration in two steps. In the first step, we evaluate the forces and the matrices of every element. We use the fast method proposed by McAdams and colleagues [2011a] for singular value decomposition. To evaluate the Hessian matrix of a hyperelastic model, we adopt the co-rotational scheme developed by Teran and collaborators [2005]. Once we obtain the results, we distribute them to the four vertices using atomic CUDA operations. In the second step, we calculate the descent direction, adjust the step length, and finally update vertex positions by Chebyshev acceleration. Our step length adjustment scheme needs the total system energy, which is computed by one CUDA thrust reduction operation.

Both air damping and viscous damping can be easily integrated into our system. Let the air damping force be:

$$\mathbf{f}^{\text{air}}(\mathbf{q}) = -\frac{c}{h}(\mathbf{q} - \mathbf{q}_t), \quad (20)$$

in which c is the air damping coefficient. The corresponding damping energy is $-\frac{c}{2h} \|\mathbf{q} - \mathbf{q}_t\|^2$ and its Hessian matrix is $-\frac{c}{h} \mathbf{I}$. Viscous damping can be implemented in a similar way, by taking the adjacency into consideration. Both air damping and viscous damping can make the Hessian matrix more diagonally dominant and reduce the condition number of the optimization problem. So to fully demonstrate the stability of our system, we typically turn damping off in our experiment. The observed energy loss effect is mainly caused by implicit time integration.

Our system can handle collisions in two ways. It can model collisions by repulsive potential energies and add them into the total energy. Alternatively, it can treat collisions as position constraints and enforce them at the end of each time step. Although the second approach requires smaller time steps, it can simulate static frictions more appropriately. So we use it for handling cloth-body collisions in the dress example.

Performance evaluation. Our algorithm is not attractive on the CPU as shown in Table 1, since forces and matrices must be evaluated multiple times. But thanks to the parallelization of tetrahedron

in which $\lambda_1, \lambda_2,$ and λ_3 are the three principal stretches, i.e., the singular values of the deformation gradient. Our basic idea is to gradually convert a hyperelastic model into projective dynamics, when an element gets highly compressed. Let $[\lambda^- = 0.05, \lambda^+ = 0.15]$ be the typical stretch interval for model conversion to happen in our experiment. For every element t in the k -th iteration, we define an interpolant $l_t^{(k)}$ as:

$$l_t^{(k)} = \min\left(1, \max\left(0, l_t^{(k-1)} - L, \max_i(\lambda^+ - \lambda_i)/(\lambda^+ - \lambda^-)\right)\right), \quad (17)$$

where $l_t^{(0)}$ is set to 0 and L is typically set to 0.05. The reason we use the $l_t^{(k-1)} - L$ term in Equation 17 is to prevent the interpolant from being rapidly changed between two time steps, which can cause oscillation artifacts in animation. We then formulate the hybrid elastic energy density of the element in the k -th iteration as:

$$W_t^{\text{hybrid}} = (1 - l_t^{(k)}) W_t + l_t^{(k)} W_t^{\text{proj}}, \quad (18)$$

where W_t is the hyperelastic energy density of element t . According to Equation 18, we calculate the total contribution of element t to the Jacobi preconditioner as:

$$\mathbf{P}_t(\mathbf{q}^{(k)}) = \text{diag}\left((1 - l_t^{(k)}) \mathbf{H}_t^{(k)} + l_t^{(k)} \mathbf{A}_t^\top \mathbf{A}_t\right), \quad (19)$$

where $\mathbf{H}_t^{(k)}$ is the Hessian matrix of W_t and $\mathbf{A}_t^\top \mathbf{A}_t$ is the constant matrix of element t used by projective dynamics. It is straightforward to implement model conversion described in Equation 19, thanks to the structural similarity between our algorithm and GPU-based projective dynamics developed by Wang [2015]. We note that the interpolant is defined for every element. This allows most elements to maintain the original hyperelastic model, even when we use a larger λ^+ as Figure 7 shows.

6 Implementation and Results

(Please watch the video for more examples. We will release our code and demos to facilitate the dissemination of this work.) We implemented and tested our system on both the CPU and the GPU. Our CPU implementation used the Eigen library (eigen.tuxfamily.org). The CPU tests ran on a single core of an Intel i7-4790K 4.0GHz processor. The GPU tests ran on an NVIDIA GeForce GTX TITAN X graphics card with 3,072 cores. The statistics and the timings of our examples are provided in Table 1. Our examples typically use $h = 1/30$ s as the time step and run 96 iterations per time step. The only exception is the dress example, which divides each time step into 8 substeps and executes 40 iterations per substep.

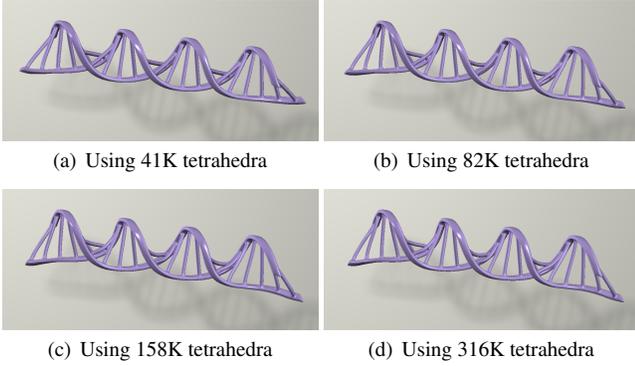


Figure 9: *The double helix example. This example indicates that our method can handle high-resolution meshes without overly stretching artifacts. All of the results use 96 iterations per time step.*

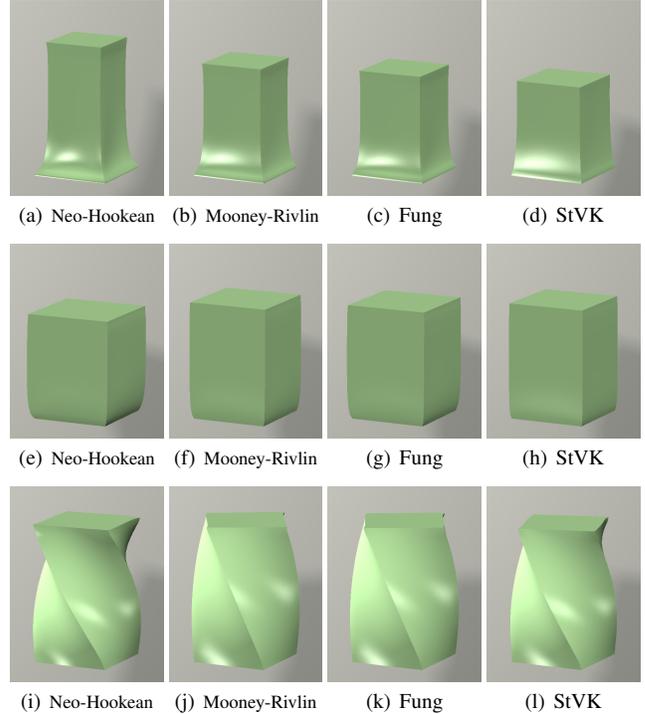


Figure 10: *The box example. Our simulator can robustly and efficiently simulate the stretching, compression, and twisting behaviors of boxes, under different hyperelastic models.*

580 threads, it can run in real time on the GPU. Figure 8 provides a
 581 typical breakdown of the computational time spent on solving a
 582 single time step. It shows that the total cost depends heavily on the
 583 force evaluation step and the step length adjustment step. Although
 584 matrix evaluation is also expensive, it contributes only 6 percent
 585 of the cost, after avoiding matrix evaluation in every iteration as
 586 discussed in Subsection 4.1.

587 To reveal the scalability of our algorithm, we simulate a double helix
 588 example at four resolutions. Table 1 shows that the computational
 589 cost is almost linearly proportional to the number of tetrahedra as
 590 expected. The high-resolution result in Figure 9d does not exhibit
 591 any overly stretching artifact, which is a common issue in position-
 592 based dynamics. Nevertheless, if computational resource permits,
 593 we still recommend the use of more iterations for high-resolution
 594 meshes, to reduce residual errors and artificial damping artifacts.

595 **Model analysis.** To evaluate the simulated behaviors of different
 596 hyperelastic models, we design a box example where the bottom
 597 face is fixed and the top face is loaded by stretching, compression,
 598 or twisting forces, as Figure 10 shows. Here we use the same s_0
 599 and s_1 for the neo-Hookean model, the Mooney-Rivlin model, and
 600 the Fung model. So the Mooney-Rivlin model and the Fung model
 601 behave stiffer than the neo-Hookean model, due to additional terms
 602 in their strain energy density functions. From our experiment, we
 603 found that the St. Venant-Kirchhoff model is more difficult to han-
 604 dle, because of its low resistance against compression. Although
 605 we can address this problem by using a larger λ^+ to make invertible
 606 model conversion earlier, it is still difficult to tune the stiffness of
 607 projective dynamics, since low stiffness cannot fix inverted ele-
 608 ments while high stiffness can cause oscillation between the two
 609 models. An alternative solution is to use isotropic strain limit-
 610 ing [Thomaszewski et al. 2009; Wang et al. 2010]. But that requires
 611 more iterations or smaller time steps, as shown in our experiment.

612 Figure 11 plots out the relationship between the stretch ratio of
 613 the box and the uplifting force applied on the top face. The na-
 614 ture of our simulator guarantees that its result is consistent with
 615 the stress-strain relationship specified by each hyperelastic model,
 616 under elastostatic situations. In particular, the stiffness of the Fung
 617 model grows more rapidly than that of the neo-Hookean model or
 618 the Mooney-Rivlin model. Meanwhile, the force is almost a cubic
 619 function of the stretch ratio under the St. Venant-Kirchhoff model.

620 **Limitations.** Our method can robustly handle high stiffness and
 621 high nonlinearity, at the expense of a lower convergence rate. So
 622 if the method does not use enough iterations, it can cause various
 623 artifacts. For example, if bending elasticity is significantly stiffer

624 than planar elasticity, it can cause cloth to be overly stretched.
 625 Meanwhile, if stiff elastic energy dominates gravitational energy, it
 626 can cause deformable bodies to fall slowly. Certain elastic models,
 627 such as the St. Venant-Kirchhoff model, do not offer sufficient
 628 stiffness against compression. In that case, the method will have
 629 difficulty in avoiding inverted elements and oscillation artifacts at
 630 the same time. The initialization approach under the constant ac-
 631 celeration assumption can also cause small oscillation artifacts, if
 632 the parameter η is not sufficiently small. The whole idea behind
 633 our method is based on the implicit time integration scheme, so
 634 it suffers from the artificial damping issue. Finally, we still need
 635 additional mechanisms for self collision detection.

7 Conclusions and Future Work

637 In this paper, we show how to improve the gradient descent method
 638 by Jacobi preconditioning and Chebyshev acceleration, for solving
 639 the nonlinear optimization problem involved in elastic body simula-
 640 tion. While the convergence rate of our method is similar to that of
 641 nonlinear conjugate gradient, it requires zero dot product operation.
 642 This characteristics allows it to run efficiently and robustly on the
 643 GPU, after applying step length adjustment, initialization, model
 644 conversion techniques.

645 Since force evaluation is the bottleneck of our simulator, we will
 646 investigate possible ways to reduce its cost, especially the cost spent
 647 on singular value decomposition. We are also interested in finding
 648 better ways for handling step lengths and inverted elements. Poten-
 649 tial solutions should have minimal impact on the simulation perfor-
 650 mance. Another interesting research direction we plan to explore
 651 is to couple our method with multi-grid techniques. The design
 652 of our method does not prevent it from using other parallelizable
 653 preconditioners. So we would like to know whether the method

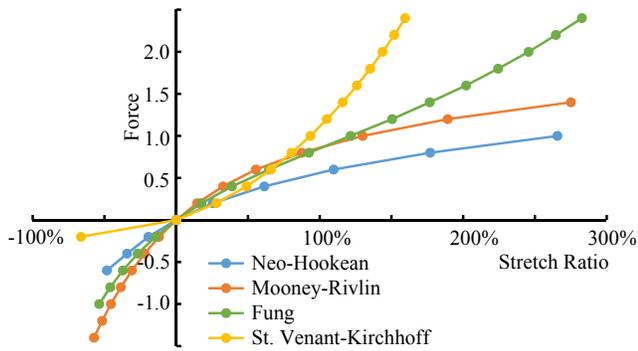


Figure 11: The force-displacement curves generated by the box example. These curves are consistent with the stress-strain relationships of the underlying hyperelastic models.

654 can work with multi-color Gauss-Seidel preconditioners as well.
 655 Finally, we will study the use of our idea in solving other simulation
 656 problems, such as material and shape design.

657 References

658 BARAFF, D., AND WITKIN, A. 1998. Large steps in cloth simulation.
 659 In *Proceedings of the 25th annual conference on Computer
 660 graphics and interactive techniques*, ACM, New York, NY, US-
 661 A, SIGGRAPH '98, 43–54.

662 BERGOU, M., WARDETZKY, M., HARMON, D., ZORIN, D., AND GRINSPUN,
 663 E. 2006. A quadratic bending model for inextensible surfaces.
 664 In *Proc. of SGP*, 227–230.

665 BOUAZIZ, S., MARTIN, S., LIU, T., KAVAN, L., AND PAULY, M. 2014.
 666 Projective dynamics: Fusing constraint projections for fast sim-
 667 ulation. *ACM Trans. Graph. (SIGGRAPH)* 33, 4 (July), 154:1–
 668 154:11.

669 BRIDSON, R., MARINO, S., AND FEDKIW, R. 2003. Simulation of
 670 clothing with folds and wrinkles. In *Proceedings of SCA*, 28–
 671 36.

672 DAVIET, G., BERTAILS-DESCOUBES, F., AND BOISSIEUX, L. 2011. A
 673 hybrid iterative solver for robustly capturing Coulomb friction
 674 in hair dynamics. *ACM Trans. Graph. (SIGGRAPH Asia)* 30, 6
 675 (Dec.), 139:1–139:12.

676 DICK, C., GEORGIL, J., AND WESTERMANN, R. 2011. A real-time
 677 multigrid finite hexahedra method for elasticity simulation using
 678 CUDA. *Simulation Modelling Practice and Theory* 19, 2, 801–
 679 816.

680 FEI, Y., RONG, G., WANG, B., AND WANG, W. 2014. Parallel L-BFGS-
 681 B algorithm on GPU. *Comput. Graph.* 40 (May), 1–9.

682 FUNG, Y.-C. 1993. *Biomechanics: Mechanical properties of living
 683 tissues*. Springer-Verlag.

684 GARG, A., GRINSPUN, E., WARDETZKY, M., AND ZORIN, D. 2007. Cubic
 685 shells. In *Proc. of SCA*, 91–98.

686 GOLUB, G. H., AND VAN LOAN, C. F. 1996. *Matrix computations (3rd
 687 Ed.)*. Johns Hopkins University Press, Baltimore, MD, USA.

688 GUTKNECHT, M. H., AND RÖLLIN, S. 2002. The Chebyshev iteration
 689 revisited. *Parallel Computing*, 28, 263–283.

690 IRVING, G., TERAN, J., AND FEDKIW, R. 2004. Invertible finite ele-
 691 ments for robust simulation of large deformation. In *Proceedings
 692 of SCA*, 131–140.

693 KHAREVYCH, L., YANG, W., TONG, Y., KANSO, E., MARSDEN, J. E.,
 694 SCHRÖDER, P., AND DESBRUN, M. 2006. Geometric, variational
 695 integrators for computer animation. In *Proceedings of SCA*, 43–
 696 51.

697 KIM, T.-Y., CHENTANEZ, N., AND MÜLLER-FISCHER, M. 2012. Long
 698 range attachments - A method to simulate inextensible clothing
 699 in computer games. In *Proceedings of SCA*, 305–310.

700 LIU, T., BARGTEIL, A. W., O'BRIEN, J. F., AND KAVAN, L. 2013.
 701 Fast simulation of mass-spring systems. *ACM Trans. Graph. (SIGGRAPH Asia)* 32, 6 (Nov.), 214:1–214:7.

702

703 MACKLIN, M., AND MÜLLER, M. 2013. Position based fluids. *ACM
 704 Trans. Graph. (SIGGRAPH)* 32, 4 (July), 104:1–104:12.

705 MACKLIN, M., MÜLLER, M., CHENTANEZ, N., AND KIM, T.-Y. 2014.
 706 Unified particle physics for real-time applications. *ACM Trans.
 707 Graph. (SIGGRAPH)* 33, 4 (July), 153:1–153:12.

708 MACOSKO, C. W. 1994. *Rheology: Principles, measurement and
 709 applications*. VCH Publishers.

710 McADAMS, A., SELLE, A., TAMSTORF, R., TERAN, J., AND SIFAKIS,
 711 E. 2011. Computing the singular value decomposition of 3x3
 712 matrices with minimal branching and elementary floating point
 713 operations. Technical report, University of Wisconsin - Madison.

714 McADAMS, A., ZHU, Y., SELLE, A., EMPEY, M., TAMSTORF, R., TERAN,
 715 J., AND SIFAKIS, E. 2011. Efficient elasticity for character skinning
 716 with contact and collisions. *ACM Trans. Graph. (SIGGRAPH)*
 717 30, 4 (July), 37:1–37:12.

718 MÜLLER, M., AND GROSS, M. 2004. Interactive virtual materials. In
 719 *Proceedings of Graphics Interface*, 239–246.

720 MÜLLER, M., HEIDELBERGER, B., TESCHNER, M., AND GROSS, M. 2005.
 721 Meshless deformations based on shape matching. *ACM Trans.
 722 Graph. (SIGGRAPH)* 24, 3 (July), 471–478.

723 MÜLLER, M., HEIDELBERGER, B., HENNIX, M., AND RATCLIFF, J. 2007.
 724 Position based dynamics. *J. Vis. Comun. Image Represent.* 18, 2
 725 (Apr.), 109–118.

726 MÜLLER, M., CHENTANEZ, N., KIM, T., AND MACKLIN, M. 2014. Strain
 727 based dynamics. In *Proceedings of SCA*, 21–23.

728 MÜLLER, M. 2008. Hierarchical position based dynamics. In
 729 *Proceedings of VRIPHYS*, 1–10.

730 NESTEROV, Y. 2004. *Introductory lectures on convex optimization:
 731 A basic course*. Applied optimization. Kluwer Academic Publ.,
 732 Boston, Dordrecht, London.

733 O'DONOGHUE, B., AND CANDÈS, E. 2015. Adaptive restart for accel-
 734 erated gradient schemes. *Found. Comput. Math.* 15, 3 (June),
 735 715–732.

736 OGDEN, R. W. 1997. *Non-linear elastic deformations*. Dover Civil
 737 and Mechanical Engineering. Dover Publications, Inc.

738 PATTERSON, T., MITCHELL, N., AND SIFAKIS, E. 2012. Simulation of
 739 complex nonlinear elastic bodies using lattice deformer. *ACM
 740 Trans. Graph. (SIGGRAPH Asia)* 31, 6 (Nov.), 197:1–197:10.

741 PEREZ, J., PEREZ, A. G., AND OTADUY, M. A. 2013. Simulation of
 742 hyperelastic materials using energy constraints. In *Proceedings
 743 of the XXIII CEIG (Spanish Conference on Computer Graphics)*.

744 PROVOT, X. 1996. Deformation constraints in a mass-spring model
 745 to describe rigid cloth behavior. In *Proceedings of Graphics
 746 Interface*, 147–154.

- 747 RIVERS, A. R., AND JAMES, D. L. 2007. FastLSM: Fast lattice shape
748 matching for robust real-time deformation. *ACM Trans. Graph.*
749 (*SIGGRAPH*) 26, 3 (July).
- 750 STERN, A., AND GRINSPUN, E. 2009. Implicit-explicit variational
751 integration of highly oscillatory problems. *Multiscale Model.*
752 *Simul.* 7, 4, 1779–1794.
- 753 SU, J., SHETH, R., AND FEDKIW, R. 2013. Energy conservation
754 for the simulation of deformable bodies. *IEEE Transactions on*
755 *Visualization and Computer Graphics* 19, 2 (Feb.), 189–200.
- 756 TERAN, J., SIFAKIS, E., IRVING, G., AND FEDKIW, R. 2005. Robust
757 quasistatic finite elements and flesh simulation. In *Proceedings*
758 *of SCA*, 181–190.
- 759 TERZOPOULOS, D., PLATT, J., BARR, A., AND FLEISCHER, K. 1987.
760 Elastically deformable models. *SIGGRAPH Comput. Graph.* 21,
761 4 (Aug.), 205–214.
- 762 THOMASZEWSKI, B., PABST, S., AND STRASSER, W. 2009. Continuum-
763 based strain limiting. *Computer Graphics Forum (Eurographics)*
764 28, 2, 569–576.
- 765 TOURNIER, M., NESME, M., GILLES, B., AND FAURE, F. 2015. Stable
766 constrained dynamics. *ACM Trans. Graph. (SIGGRAPH)* 34, 4
767 (July), 132:1–132:10.
- 768 WANG, H., O'BRIEN, J., AND RAMAMOORTHI, R. 2010. Multi-
769 resolution isotropic strain limiting. *ACM Trans. Graph. (SIG-*
770 *GRAPH Asia)* 29, 6 (Dec.), 156:1–156:10.
- 771 WANG, H. 2015. A Chebyshev semi-iterative approach for accelerat-
772 ing projective and position-based dynamics. *ACM Trans. Graph.*
773 (*SIGGRAPH Asia*) 34, 6 (Oct.), 246:1–246:9.
- 774 ZHU, Y., SIFAKIS, E., TERAN, J., AND BRANDT, A. 2010. An efficient
775 multigrid method for the simulation of high-resolution elastic
776 solids. *ACM Trans. Graph.* 29, 2 (Apr.), 16:1–16:18.