



Primal Extended Position Based Dynamics for Hyperelasticity

Yizhou Chen
chenyizhou@ucla.edu
University of California, Los Angeles
Epic Games, Inc
USA

Yushan Han
yushanh1@math.ucla.edu
University of California, Los Angeles
Epic Games, Inc
USA

Jingyu Chen
chenjy@g.ucla.edu
University of California, Los Angeles
USA

Shiqian Ma
sqma@rice.edu
Rice University
USA

Ronald Fedkiw
fedkiw@cs.stanford.edu
Stanford University
USA

Joseph Teran
jteran@math.ucdavis.edu
University of California, Davis
Epic Games, Inc
USA

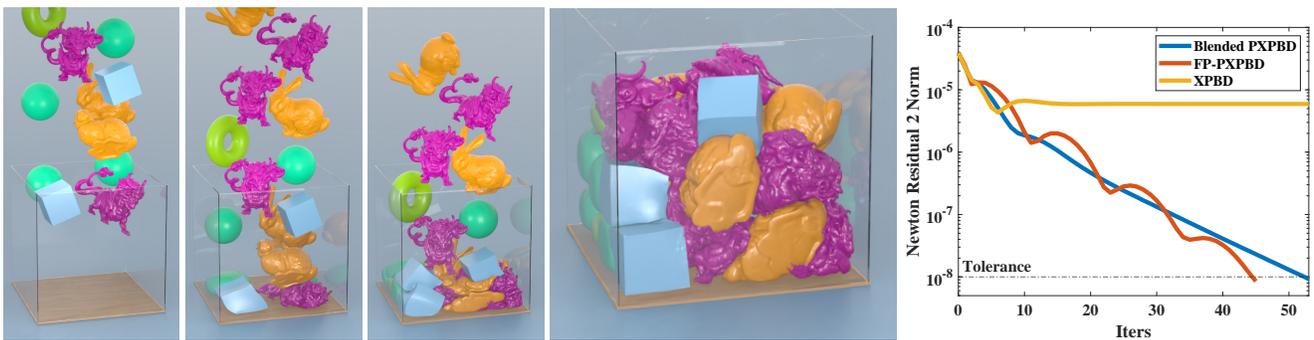


Figure 1: 30 Objects Dropping (left). Our Blended PXPBD (B-PXPBD) approach robustly handles large elastic deformations. FEM Residual Comparison (right). B-PXPBD and FP-PXPBD reduce the backward Euler residual while XPBD stagnate in a representative step of a hyperelasticity simulation.

ABSTRACT

The Extended Position Based Dynamics (XPBD) approach of Macklin et al. [2016] addresses the issues with iteration-dependent behavior in the original Position Based Dynamics [2007] (PBD) which itself is a powerful method for the real-time simulation of elastic objects. However, it is limited in its application to hyperelastic solids. It can only treat models with a strain energy density that is quadratic in some notion of constraint. Furthermore, we show that even when applicable the formulation does not always lead to convergent behaviors with hyperelasticity. We isolate the root cause in the approximate linearization of the nonlinear backward Euler systems utilized by XPBD. We provide two fixes to these terms that allow for convergent behavior. The first (B-PXPBD) is a small modification to an existing XPBD code, but can only be used with models addressable by the original XPBD. The second (FP-PXPBD) is a more general formulation that extends XPBD (and our residual correction) to arbitrary hyperelasticity. We show that our modifications allow for convergent behavior that rivals accurate

techniques like Newton’s method when the computational budget is large without sacrificing the stable and robust behavior exhibited by the original PBD and XPBD when the computational budget is limited.

CCS CONCEPTS

• **Computing methodologies** → Realtime Simulation.

KEYWORDS

Position-based dynamics, physics simulation, constrained dynamics

ACM Reference Format:

Yizhou Chen, Yushan Han, Jingyu Chen, Shiqian Ma, Ronald Fedkiw, and Joseph Teran. 2023. Primal Extended Position Based Dynamics for Hyperelasticity. In *ACM SIGGRAPH Conference on Motion, Interaction and Games (MIG ’23)*, November 15–17, 2023, Rennes, France. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3623264.3624437>

1 INTRODUCTION

In the present work, we consider large strain hyperelastic solids [Bonet and Wood 2008] whose governing equations are discretized in space with the finite element method (FEM) [Sifakis and Barbic 2012] and in time with implicit backward Euler [Baraff and Witkin 1998] or quasistatics [Bonet and Wood 2008; Rabinovich et al. 2017; Sorkine and Alexa 2007; Teran et al. 2005b; Zhu et al. 2018]. Hyperelastic solid models define continuum stresses from a

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

MIG ’23, November 15–17, 2023, Rennes, France

© 2023 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0393-5/23/11.

<https://doi.org/10.1145/3623264.3624437>

notion of elastic potential. In graphics applications, these models are commonly used for simulation-based enhancement of character flesh and musculature animation [Fan et al. 2014; McAdams et al. 2011; Modi et al. 2021; Smith et al. 2018; Teran et al. 2005a; Wang et al. 2020]. Wang et al. [2020] provide a thorough discussion of the state-of-art. In these applications, the constitutive control enabled by continuum models is essential for realism, e.g., anisotropic contraction along muscle fibers and volume preservation in soft tissues.

Various methods have been proposed for solving the FEM-discretized equations of motions for these materials [Bouaziz et al. 2014; Gast et al. 2015; Kharevych et al. 2006; Li et al. 2019; Macklin et al. 2016; Narain et al. 2016; Teran et al. 2005b; Zhao et al. 2016; Zhu et al. 2018]. Thorough summaries of the state of the art are given by Zhu et al. [2018] and Li et al. [2019]. The preferred approach in a given application generally depends on the relative importance of constitutive accuracy, robustness/stability and computational efficiency. There is no one method that is optimal in all computer graphics as different applications place different relative importance on these considerations. These equations are nonlinear, and an iterative solver must be used to improve the accuracy of an initial guess by reducing the magnitude of the system residual. While Newton's method [Nocedal and Wright 2006] generally requires the fewest iterations to reach a desired tolerance (often achieving quadratic convergence), each iteration can be costly and a line search is typically required for stability [Gast et al. 2015]. However, it is not always necessary to reduce the residual beyond a few orders of magnitude for satisfactory visual accuracy (see discussion in Liu et al. [2013], Bouaziz et al. [Bouaziz et al. 2014], Zhu et al. [2018]). In these cases, Newton's method is often outperformed by alternative techniques. Methods like the Alternating Direction Method of Multipliers (ADMM) [Boyd et al. 2011; Narain et al. 2016], the limited-memory Broyden-Fletcher-Goldfarb-Shanno algorithm (L-BFGS) [Bertsekas 1997; Liu et al. 2017; Wittemeyer et al. 2021; Zhu et al. 2018] and Sobolev preconditioned gradient descent (SGD) [Bouaziz et al. 2014; Liu et al. 2013; Neuberger 1985; Sorkine and Alexa 2007] do not require computation of the exact energy Hessian and many of these simplifications leverage direct solvers based on pre-computed matrix factorizations of simplified discrete elliptic operators to allow for reduced per-iteration cost compared to Newton's method.

Real-time applications have very modest computational budgets and this restricts which techniques can be used. The ability to reduce the nonlinear system residuals degrades considerably and many methods that behave admirably with large computational budgets can give unstable and/or visually implausible results when computation is limited. Wang [2015] provides a thorough discussion of these issues. For example, methods that make use of a linear solver at each iteration (Newton, ADMM, SGD, some L-BFGS) can only do so at either extremely low resolution or with significant constraints on the accuracy of their approximation. In particular, those that leverage a pre-factored discrete elliptic operator are limited since forward and backward substitutions do not parallelize. Furthermore, the memory requirements for storing the factors are often too burdensome. Moreover, simply replacing the direct solver with an iterative solver with minimal iteration counts can lead to visually implausible or even unstable behaviors (see Figure 2

and [Wang 2015]). At a minimum, real-time applications require methods that will remain stable and be as visually plausible as possible when the compute budget is limited.

The Position Based Dynamics (PBD) approach of Müller et al. [2007] is remarkably powerful due to its robust and stable behavior in applications with minimal computational budgets. PBD has gained wide adoption since there are often no other methods that can provide comparably reliable behavior under extreme computation constraints. For elastic materials, PBD uses a constraint view of the material resistance to deformation and is similar to strain limiting [Provot 1995] and shape matching [Müller et al. 2005] techniques. In the context of elasticity, this has been shown to be equivalent to a Gauss-Seidel minimization of an elastic potential that is quadratic in the constraints [Bouaziz et al. 2014; Liu et al. 2013; Macklin et al. 2016]. However, constitutive control over PBD behavior is challenging as effective material stiffnesses etc. vary with iteration count and time step size. The Extended Position Based Dynamics (XPBD) approach of Macklin et al. [2016] addresses these issues by reformulating the original PBD approach in terms of a Gauss-Seidel technique for discretizing a total Lagrange multiplier formulation of the backward Euler system for implicit time stepping. This formulation has similarities to PBD, but with the elastic terms handled properly where PBD can be seen as the extreme case of infinite elastic modulus (or hard strain constraints). In this case, the Lagrange multiplier terms can be interpreted as stress-like and associated with enforcing the constraints (e.g. pressure in an incompressible fluid [Batty et al. 2007]). With this view, XPBD handles these terms correctly as weak constraints (e.g. as in weakly compressible materials [Kwatra et al. 2009; Stomakhin et al. 2014]) where PBD can be interpreted as a splitting scheme where non-stress based forces are first integrated, followed by a projection step.

Despite its many strengths, XPBD can only discretize hyperelastic models that are quadratic in some notion of strain constraint [Macklin and Muller 2021; Macklin et al. 2016]. This prevents the adoption of many models from the computational mechanics literature, e.g., for many biomechanical soft tissues. Furthermore, while XPBD is based on a Gauss-Seidel procedure for the Lagrange multiplier formulation of the backward Euler equations, it simplifies the system by omitting the Hessian of the constraints and the residual of the primary (position) equations. The omission of the primary equations is perfectly accurate in the first iteration, but as Macklin et al. [2016] point out, less so in latter iterations when constraint gradients vary significantly. We observe that this rapid variation occurs for many hyperelastic formulations and that its omission degrades residual reduction. However, the inclusion of this term introduces instabilities into XPBD.

We provide a modification to the XPBD position update that more accurately guarantees that the primary residual is zero and may be omitted. We call our approach Primal Extended Position Based Dynamics (PXPBD). It can be done in two ways. The first (B-PXPBD) uses fixed-point iteration to zero the primary residual after the Gauss-Seidel update of the Lagrange multiplier. The second (FP-PXPBD) is a reformulation of XPBD that allows for arbitrary hyperelastic models. We observe that the constraint Hessians and primary residual terms are exactly zero and can be omitted with no error if the first Piola-Kirchhoff stress [Bonet and Wood 2008] is used as the auxiliary unknown (in place of the Lagrange multipliers

in the original XPBD). We advocate for two models because each have relative strengths and weaknesses in their resolution of the primary residual omission in XPBD. B-PXPBD can be done with a simple modification to an existing XPBD code, however it requires the use of a blending parameter (see Section 3.3.1) since accurate fixed-point iteration is too costly. FP-PXPBD is a larger modification to an existing XPBD code and requires element-wise Newton solves, but it exactly resolves the the issues with both Hessian and residual omission in XPBD. Furthermore, FP-PXPBD allows for arbitrary hyperelastic models while B-PXPBD is based on constraint formulations as with XPBD.

We demonstrate our method with collision-intensive scenarios by applying it to the updated-Lagrangian formulation of hyperelasticity where the simulation mesh is embedded in a regular grid at each time step as in [Jiang et al. 2015]. We summarize our contributions as:

- B-PXPBD: A modification to the XPBD position update that improves residual reduction with hyperelasticity.
- FP-PXPBD: A first Piola-Kirchhoff formulation of the XPBD auxiliary variables that both guarantees zero primal residual for improved total residual reduction and generalizes XPBD to arbitrary hyperelastic models.
- A local affine transformation that decouples strain and translation variables in each FEM element for added efficiency with FP-PXPBD.
- A Sherman-Morrison rank-one quasi-Newton technique for each first Piola-Kirchhoff stress in FP-PXPBD.

2 PREVIOUS WORK

O'Brien et al. [2002; 1999] first demonstrated the power of FEM simulation of hyperelasticity in graphics applications. Early approaches used mostly St. Venant-Kirchhoff hyperelasticity [Barbič and James 2005; O'Brien et al. 2002; O'Brien and Hodgins 1999]. However, as noted in [Bonet and Wood 2008; Irving et al. 2004] these models weaken under compression and are ill-suited for the large deformation problems encountered in computer graphics. As rigid as possible (ARAP) [Sorkine and Alexa 2007] and corotational models behave more robustly with large deformation [Chao et al. 2010; Müller and Gross 2004; Stomakhin et al. 2012]. Extension of the elastic response to non-bijective deformation mappings was shown to be useful for graphics applications by Irving et al. [2004] with the key insight being the sign convention established in the polar singular value decomposition. This was extended to hyperelasticity by Stomakhin et al. [2012], Smith et al. [2018] and Kim et al. [2019]. Martin et al. [2011] design hyperelastic potentials from examples of desired material behaviors. Hyperelasticity and continuum modeling are useful for designing anisotropic models that capture the fiber-driven dynamics of muscle and tendon [Fan et al. 2014; Kim et al. 2019; Modi et al. 2021; Teran et al. 2003, 2005a]. Xu et al. [2015] develop isotropic and anisotropic hyperelastic models in terms of principle stretches. Tournier et al. [2015] develop a technique for bridging elasticity and constraint-based approaches that is robust to large stiffness. They use a similar primal/dual setup to XPBD, however unlike XPBD their approach solves the entire system at once, rather than iterating over individual constraints. Wang and Yang [2016] use a Chebyshev accelerated gradient descent approach for general hyperelasticity and FEM.

Baraff and Witkin first demonstrated that implicit time stepping with elasticity is essential for efficiency [Baraff and Witkin 1998]. Many approaches characterize implicit time stepping with hyperelasticity as a minimization of an incremental potential [Bouaziz et al. 2014; Gast et al. 2015; Liu et al. 2013; Martin et al. 2011; Narain et al. 2016; Stern and Desbrun 2006]. This is often referred to as variational implicit Euler [Martin et al. 2011; Stern and Desbrun 2006] or optimization implicit Euler [Liu et al. 2013]. Quasistatic time stepping is an extreme case where inertia terms are ignored and only the strain energy is minimized [Kovalsky et al. 2016; Liu et al. 2008; Rabinovich et al. 2017; Sorkine and Alexa 2007; Teran et al. 2005b]. Rotational invariance of the hyperelastic potential makes the energy minimization non-convex with potentially non-unique solutions in quasistatic problems [Bonet and Wood 2008]. Minimizers are usually found by setting the gradient of the energy to zero and solving the associated nonlinear system of equations with Newton's method. However, the non-convexity yields indefinite energy Hessians that can prevent convergence. Quasi-Newton methods can be used to approximate the Hessian with a symmetric semi-definite counterpart [Li et al. 2019; Nocedal and Wright 2006; Smith et al. 2019; Teran et al. 2005b; Zhu et al. 2018].

Many methods avoid the indefiniteness issue with the inclusion of auxiliary (or secondary) variables. Narain et al. [2016], Bouaziz et al. [Bouaziz et al. 2014], Liu et al. [2013] are recent examples of this, but similar approaches have been used in graphics since the local/global approach with ARAP by Sorkine et al. [2007]. Rabinovich et al. [2017] generalize this approach to a wider range of distortion energies. The global solve in these approaches requires the inversion of a constant discrete elliptic operator (component wise-Laplacian) which can be pre-factored for efficiency. Zhu et al. [2018] points out that approaches that make use of this operator can be classified as SGD [Neuberger 1985]. While this discrete operator does not suffer from indefiniteness issues, various authors note that SGD approaches may converge initially faster than Newton, but will often taper off [Bouaziz et al. 2014; Liu et al. 2013; Wang 2015; Zhu et al. 2018]. Zhu et al. [2018] tailor their approach to this observation and use SGD initially and then combine with L-BFGS to incorporate more second-order information. Liu et al. [2017] and Witemeyer et al. [2021] also use combinations of SGD and L-BFGS. Kovalsky et al. [Kovalsky et al. 2016] add Nesterov acceleration to SGD. Hecht et al. [2012] develop efficient updates for a pre-factored Hessian with corotated materials. Wang [2015] discusses the challenges of using direct solution/pre-factoring in the SGD style approaches of Narain et al. [2016], Bouaziz et al. [Bouaziz et al. 2014] and Liu et al. [2013] and develops a Chebyshev acceleration technique as an alternative.

3 METHODS

3.1 Equations

We consider implicit time stepping methods for integrating the FEM-discretized partial differential equations (PDEs) describing momentum balance with hyperelastic materials

$$\mathbf{M} \frac{\partial^2 \mathbf{x}}{\partial t^2} = -\frac{\partial PE}{\partial \mathbf{x}} + \mathbf{f}^{ext}, \quad PE(\mathbf{x}) = \sum_e \Psi(\mathbf{F}^e(\mathbf{x}))V^e. \quad (1)$$



Figure 2: Equal Budget Comparison. From left to right: Newton (converged), Newton, FP-PXPBD, B-PXPBD, XPBD. With a limited budget, XPBD-style methods are stable, whereas the Newton's method suffers from instability. Frame 0, 10, 60 are shown in the figure.

Here $\mathbf{M} \in \mathbb{R}^{3N_p \times 3N_p}$ is a lumped (diagonal) mass matrix, $\mathbf{x} \in \mathbb{R}^{3N_p}$ are the deformed positions of the FEM mesh and the potential energy PE in the system is related to the hyperelastic potential energy density as Ψ . We use linear interpolation over tetrahedron (3D) or triangle (2D) meshes in our FEM formulation. V^e is the volume (3D) or area (2D) of the undeformed e^{th} element arising from the piecewise constant terms in an integrands associated with linear interpolation. \mathbf{f}^{ext} are external forces (gravity etc.). The hyperelastic potential Ψ is a function of the deformation gradient in the e^{th} element (\mathbf{F}^e) which is related to deformed positions as

$$\mathbf{F}_{\alpha\beta}^e(\mathbf{x}) = \sum_i x_{i\alpha} \frac{\partial N_i}{\partial X_\beta}(\mathbf{X}^e) \quad (2)$$

where N_i are the piecewise linear interpolation functions in the FEM formulation and \mathbf{X}^e is the centroid of the undeformed element. We refer the reader to the Bonet and Wood [2008] and Barbič and Sifakis [2012] for more details.

3.1.1 Hyperelastic Energy Density. The hyperelastic potential defines the constitutive response of the material. We demonstrate our method with the fixed corotated potential from Stomakhin et al. [2012]

$$\Psi^{\text{cor}}(\mathbf{F}) = \mu |\mathbf{F} - \mathbf{R}(\mathbf{F})|_F^2 + \frac{\lambda}{2} (\det(\mathbf{F}) - 1)^2. \quad (3)$$

Here $\mathbf{R}(\mathbf{F})$ is the closest rotation to \mathbf{F} which we compute from the polar singular value decomposition [Gast et al. 2016] and μ and λ are the Lamé coefficients. XPBD assumes that the potential is of the form

$$PE(\mathbf{x}) = \sum_c \frac{1}{2} C_c(\mathbf{x}) \frac{1}{a_c} C_c(\mathbf{x}) \quad (4)$$

The corotated potential Ψ^{cor} can be adapted to this from in terms of the following constraints (in element) on the deformation gradient

$$\hat{C}_1(\mathbf{F}) = |\mathbf{F} - \mathbf{R}(\mathbf{F})|_F, \hat{C}_2(\mathbf{F}) = \det(\mathbf{F}) - 1. \quad (5)$$

The gradient of \hat{C}_1 is not defined when $\mathbf{F} = \mathbf{R}(\mathbf{F})$ (a common occurrence) and we use the modification $\tilde{C}_1 = \sqrt{\hat{C}_1^2 + \epsilon}$, where ϵ is an arbitrary positive constant to ensure that the gradient is always defined. We use constraints $C_1^e(\mathbf{x}) = \tilde{C}_1(\mathbf{F}^e(\mathbf{x}))$ and $C_2^e(\mathbf{x}) = \hat{C}_2(\mathbf{F}^e(\mathbf{x}))$ with weighting μ and λ respectively (in element e). This is equivalent to using the hyperelastic potential $\Psi^{\text{cor}} + \epsilon$ so it produces the same behavior as the corotated model.

We also demonstrate our method with an anisotropic model for muscle contraction (see Figure 4). Here the potential is

$$\Psi^{\text{aniso}}(\mathbf{F}) = \Psi^{\text{cor}} + \frac{\sigma_{\text{max}}}{\lambda_{\text{ofl}}} (f_a + \alpha_{\text{act}} f_p) \quad (6)$$

where the parameter $\alpha_{\text{act}} \in [0, 1]$ controls the degree of active contractile tension and f_a and f_p are based on the anisotropic fiber terms in Blemker et al. [2005]. We refer the reader to the supplemental technical document for a detailed description of these terms [Y.Chen et al. 2023].

3.1.2 Implicit Time Stepping. We consider both backward Euler and quasistatic time stepping schemes

$$\mathbf{M} \left(\frac{\mathbf{x}^{n+1} - \mathbf{x}^n}{\Delta t} - \mathbf{v}^n \right) = - \frac{\partial PE}{\partial \mathbf{x}}(\mathbf{x}^{n+1}) + \mathbf{f}^{\text{ext}}. \quad (7)$$

Here $\mathbf{x}^n, \mathbf{v}^n$ represent the time $t^n = n\Delta t$ position and velocities. Quasistatic time stepping is the same but with the left hand side of Equation (7) replaced with $\mathbf{0}$. Note that we also may constrain some vertices $\mathbf{x}_i^n, 0 \leq i < N_p$ in practice to enforce boundary conditions and these equations are removed from Equation (7), however we omit the explicit representation of this for concise exposition.

3.2 XPBD

Macklin et al. [2016] solve Equation (7) with the introduction of a Lagrange multiplier λ_c associated with each constraint C_c . They assume the potential energy gradient is of the form

$$\Delta t^2 \frac{\partial PE}{\partial \mathbf{x}_{i\alpha}} = - \sum_c \frac{\partial C_c}{\partial \mathbf{x}_{i\alpha}}(\mathbf{x}) \lambda_c \quad (8)$$

where they introduce $\lambda_c = - \frac{\Delta t^2}{a_c} C_c$ as an additional unknown which converts Equation (7) into the system

$$\mathbf{g}(\mathbf{x}^{n+1}, \boldsymbol{\lambda}) = \mathbf{M} \left(\mathbf{x}^{n+1} - \tilde{\mathbf{x}} \right) - \sum_c \lambda_c^T \nabla C_c(\mathbf{x}^{n+1}) = \mathbf{0} \quad (9)$$

$$\mathbf{h}(\mathbf{x}^{n+1}, \boldsymbol{\lambda}) = \mathbf{C}(\mathbf{x}^{n+1}) + \frac{\mathbf{A}}{\Delta t^2} \boldsymbol{\lambda} = \mathbf{0}. \quad (10)$$

Here $\tilde{\mathbf{x}} = \mathbf{x}^n + \Delta t(\mathbf{v}^n + \mathbf{M}^{-1} \mathbf{f}^{\text{ext}})$ are the positions updated under the influence of inertia and external forces, $\boldsymbol{\lambda}$ is the vector of all Lagrange multipliers and \mathbf{A} is a diagonal matrix with entries equal to a_c . The solution is approximated iteratively with \mathbf{x}_k^{n+1} and $\boldsymbol{\lambda}_k$ denoting the k^{th} iterates. $\mathbf{g}(\mathbf{x}_k^{n+1}, \boldsymbol{\lambda}_k)$ is used to denote the residual of the position (primary) unknowns and $\mathbf{h}(\mathbf{x}_k^{n+1}, \boldsymbol{\lambda}_k)$ to denote the residual of the Lagrange multiplier (secondary) unknowns.

XPBD uses a nonlinear Gauss-Seidel procedure based on the linearization

$$\begin{aligned} & \begin{pmatrix} \mathbf{M} + \sum_c \lambda_{ck} \frac{\partial^2 C_c}{\partial \mathbf{x}_k^2}(\mathbf{x}_k^{n+1}) & -\nabla C_c^T(\mathbf{x}_k^{n+1}) \\ \nabla C_c(\mathbf{x}_k^{n+1}) & \frac{\mathbf{A}}{\Delta t^2} \end{pmatrix} \begin{pmatrix} \Delta \mathbf{x}_{k+1} \\ \Delta \boldsymbol{\lambda}_{k+1} \end{pmatrix} \\ & = - \begin{pmatrix} \mathbf{g}(\mathbf{x}_k^{n+1}, \boldsymbol{\lambda}_k) \\ \mathbf{h}(\mathbf{x}_k^{n+1}, \boldsymbol{\lambda}_k) \end{pmatrix}. \end{aligned} \quad (11)$$

In XPBD, the red terms are omitted to enable the update

$$\left(\mathbf{C}^T(\mathbf{x}_k^{n+1})\mathbf{M}^{-1}\mathbf{C}(\mathbf{x}_k^{n+1}) + \frac{\mathbf{A}}{\Delta t^2} \right) \Delta \boldsymbol{\lambda}_{k+1} = -\mathbf{h}(\mathbf{x}_k^{n+1}, \boldsymbol{\lambda}_k) \quad (12)$$

$$\Delta \mathbf{x}_{k+1} = \mathbf{M}^{-1} \nabla \mathbf{C}(\mathbf{x}_k^{n+1}) \Delta \boldsymbol{\lambda}_{k+1}. \quad (13)$$

Furthermore, Equation (12) is updated in a Gauss-Seidel fashion where the d^{th} Lagrange multiplier is updated via

$$\Delta \tilde{\lambda}_{k+1d} = \frac{-h_d(\mathbf{x}_k^{n+1}, \lambda_{kd})}{\nabla C_d^T(\mathbf{x}_k^{n+1})\mathbf{M}^{-1}\nabla C_d(\mathbf{x}_k^{n+1}) + \frac{a_d}{\Delta t^2} \lambda_{kd}}, \quad (14)$$

$$\lambda_{k+1d} = \lambda_{kd} + \Delta \tilde{\lambda}_{k+1d}.$$

Note that we distinguish $\Delta \tilde{\lambda}_{k+1d}$ in Equation (14) from $\Delta \lambda_{k+1d}$ in Equation (12) since only one step of Gauss-Seidel iteration is performed on the linear system. Then the positions associated with the constraint are updated via Equation (13) to create

$$\mathbf{x}_{k+1}^{n+1} = \mathbf{x}_k^{n+1} + \mathbf{M}^{-1} \nabla \mathbf{C}_d(\mathbf{x}_k^{n+1}) \Delta \tilde{\lambda}_{k+1d}. \quad (15)$$

The system (Equations (9)-(10)) is then re-linearized (Equation (11)) and the process (Equations (14)-(15)) is repeated iteratively.

3.3 Primary residual XPBD (PXPBD)

The motivation for the omission of the residual and constraint Hessian terms (red) in Equation (11) is natural. The constraint Hessian is non-diagonal and its retention would preclude the decoupling of primary variables from the Lagrange multipliers in Equation (12). Furthermore, the primary residual term $\mathbf{g}(\mathbf{x}_k^{n+1}, \boldsymbol{\lambda}_k)$ requires more floating point operations and generally a gather operation for efficient parallel evaluation. As Macklin et al. [2016] point out, the initial guess of $\boldsymbol{\lambda}_0 = \mathbf{0}$ and $\mathbf{x}_0^{n+1} = \bar{\mathbf{x}}$ means that $\mathbf{g}(\mathbf{x}_0^{n+1}, \boldsymbol{\lambda}_0) = \mathbf{0}$. However, its omission is harder to justify in latter iterates, though Macklin et al. [2016] argue that it is justified when the constraint gradients vary slowly and further that its omission makes the approach similar to that of Goldenthal et al. [2007]. While omission of secondary information is commonly done in quasi-Newton approaches, we observe that omission of the primary residual terms can lead to stagnation in residual reduction (see Figure 3(a)). Unfortunately, we also notice that inclusion of this term can cause XPBD to lose its favorable stability properties (see Figure 3(b)). We note though that if the global system in Equation (11) is solved with sufficient accuracy (e.g. with a Krylov method and without omission of the red terms), then stability and residual reduction can be achieved, however this is more costly than Newton's method for Equation (7) since the system size is larger with the inclusion of the $\boldsymbol{\lambda}$ unknowns.

3.3.1 Blended Primal XPBD (B-PXPBD). We believe that the stability of XPBD is due to the omission of this primary residual term $\mathbf{g}(\mathbf{x}_k^{n+1}, \boldsymbol{\lambda}_k)$. We observe that this omission can be done without any error if the position update is chosen to guarantee that the primary residual is zero. This can be done by solving Equation (9) for \mathbf{x}_{k+1}^{n+1} with $\boldsymbol{\lambda}_{k+1}$ fixed after the update (of a single Lagrange multiplier λ_{k+1d}) in Equation (14). We again note that in this context, the Lagrange multipliers $\boldsymbol{\lambda}_{k+1}$ are similar to stresses. Indeed as the a_c are taken to infinity we can see similarities between Equations (9)-(10) and the discretized equations for incompressible fluids and for

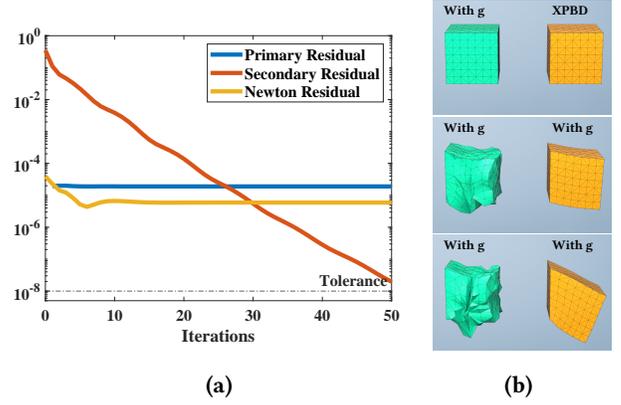


Figure 3: (a) Primal Residual Comparison: Stagnation. While XPBD reliably reduces the secondary residual, its omission of the primary residual in the linearization causes its primary residual to stagnate, making its true (Newton) residual stagnate as well. **(b) Primal Residual Inclusion: Instability.** XPBD is unstable when the primal residual term is not omitted.

finite values of a_c the formulation is similar to the compressible formulations in Stomakhin et al. [Stomakhin et al. 2014] and Kwatra et al. [2010]. Therefore, the process of solving Equation (9) for \mathbf{x}_{k+1}^{n+1} with $\boldsymbol{\lambda}_{k+1}$ fixed is akin to solving for the change in positions given a fixed stress state (that does not depend on positions).

Unfortunately, solving Equation (9) for \mathbf{x}_{k+1}^{n+1} is complicated by the dependence of the constraint gradient $\nabla \mathbf{C}(\mathbf{x}_{k+1}^{n+1})$ on positions and solving it accurately would be nearly as difficult as solving the original system in Equation (7). Furthermore, this dependence of the constraint gradient on positions means changing the stress in one constraint propagates to changes in positions in adjacent constraints and therefore throughout the mesh. For example if fixed point iteration were used to solve for \mathbf{x}_{k+1}^{n+1} given $\boldsymbol{\lambda}_{k+1}$ where the only change to $\boldsymbol{\lambda}_k$ was in a single constraint d (as in Equation (14)), then first only the positions of the vertices in the constraint would be changed, but then in the second iteration, any other constraint gradients with dependence on these positions would change, and all positions associated with those constraints would change, and so on. This would quickly become computationally inefficient, however performing one iteration results in an update that only changes the positions involved in the constraint associated with the Lagrange multiplier update in Equation (14)

$$\mathbf{x}_{k+1}^{n+1} = \bar{\mathbf{x}} + \sum_c \lambda_{kc} \mathbf{M}^{-1} \nabla \mathbf{C}_c(\mathbf{x}_k^{n+1}) + \mathbf{M}^{-1} \nabla \mathbf{C}_d(\mathbf{x}_k^{n+1}) \Delta \tilde{\lambda}_{k+1d}. \quad (16)$$

Note that when the residual $\mathbf{g}(\mathbf{x}_k^{n+1}, \boldsymbol{\lambda}_k) = \mathbf{0}$ is zero this update coincides with that of Equation (13). We found that even using this first fixed point iterate was enough to improve residual reduction, however we also found that it reduced the stability compared to Equation (13). We remedy this by taking a linear combination of

the updates in Equations (13) and (16)

$$\begin{aligned} \mathbf{x}_{k+1}^{n+1} &= \zeta \left(\mathbf{M}^{-1} \nabla C_d(\mathbf{x}_k^{n+1}) \Delta \tilde{\lambda}_{k+1d} \right) + (1 - \zeta) \Delta \mathbf{x}_{k+1}^{fp} + \mathbf{x}_k^{n+1} \quad (17) \\ \Delta \mathbf{x}_{k+1}^{fp} &= \tilde{\mathbf{x}} + \sum_c \lambda_{kc} \mathbf{M}^{-1} \nabla C_c(\mathbf{x}_k^{n+1}) + \Delta \tilde{\lambda}_{k+1d} \mathbf{M}^{-1} \nabla C_d(\mathbf{x}_k^{n+1}) - \mathbf{x}_k^{n+1}. \quad (18) \end{aligned}$$

The parameter ζ can usually chosen to be 0.5. We increase it if we observe instability and raise it if we see residual stagnation.



Figure 4: Muscle Box Activation. A rectangular bar with both ends clamped falls under gravity. Two seconds later, the muscle box is activated and contracts along the horizontal direction. The level of activation is shown on the right side of the images. $t = 0.0333, 1.2, 2.9$ seconds are shown in the footage.

ALGORITHM 1: FP-PXPBD Simulation Loop

```

while not reached maximal iterations do
  for element e do
    while not converged or reached maximal iterations do
      begin Solve Newton system
        1. Compute Newton residual via Equation (25);
        2. Compute  $\mathbf{b}_{k+1l}^e$  via Equation 28;
        3. Compute  $\delta \mathbf{F}_{k+1l}^e$  via Equation 30 with the
           approximation in Equation 33;
        4. Compute  $\delta \mathbf{x}_{k+1l}^e$  as in Equation 26 ;
        5. Update the nodes on the element with
            $\mathbf{x}_{i^e k+1l+1}^{n+1} = \mathbf{x}_{i^e k+1l}^{n+1} + \delta \mathbf{x}_{i^e k+1l}^e$ ;
        6. Update  $\mathbf{P}_{k+1l+1}^e = \frac{\partial \Psi}{\partial \mathbf{F}}(\mathbf{F}^e(\mathbf{x}_{i^e k+1l+1}^{n+1}))$ ;
      end
    end
  end
end

```

3.3.2 First Piola-Kirchhoff Primal XPBD (FP-PXPBD). Noting that the auxiliary Lagrange multiplier variables are similar to stresses, we observe some convenient properties that arise from choosing an alternative stress measure in an analogous primary/secondary formulation of Equation 7. In a general FEM-discretized hyperelastic formulation (see Barbič and Sifakis [2012]), the potential energy gradient has the expression

$$\frac{\partial PE}{\partial \mathbf{x}_{i\alpha}}(\mathbf{x}) = \sum_{e, \beta, \gamma} P_{\beta\gamma}(\mathbf{F}^e(\mathbf{x})) \delta_{\alpha\beta} \frac{\partial N_i}{\partial X_\gamma}(\mathbf{X}^e) V^e \quad (19)$$

where $\delta_{\alpha\beta}$ is the Kronecker delta tensor and $\mathbf{P} = \frac{\partial \Psi}{\partial \mathbf{F}}$ is the gradient of the hyperelastic potential energy density with respect to the deformation gradient. This is the first Piola-Kirchhoff stress [Bonet and Wood 2008]. If we introduce it as an unknown (analogous to λ_c), then tensor $B_{i\alpha\beta}^e = \delta_{\alpha\beta} \frac{\partial N_i}{\partial X_\gamma}(\mathbf{X}^e) V^e$ is analogous to the ∇C_c terms in XPBD since they convert the auxiliary (stress) terms to

force in the expression in Equation (8). With this formulation, an analogous method consists of

$$\mathbf{g}(\mathbf{x}^{n+1}, \mathbf{P}^{n+1}) = \mathbf{M} \left(\mathbf{x}^{n+1} - \tilde{\mathbf{x}} \right) + \Delta t^2 \mathbf{B} \mathbf{P} = \mathbf{0} \quad (20)$$

$$\mathbf{h}^e(\mathbf{x}^{n+1}, \mathbf{P}^{n+1}) = \frac{\partial \Psi}{\partial \mathbf{F}}(\mathbf{F}^e(\mathbf{x}^{n+1})) - \mathbf{P}^e = \mathbf{0}. \quad (21)$$

Note that with this expression, the tensor \mathbf{B} does not depend on the positional unknowns \mathbf{x}^{n+1} . In contrast, the analogous expression $\nabla C(\mathbf{x}^{n+1})$ in Equations (9) does have this dependence, and it is precisely this issue that leads to the red terms in the linearization in Equation (11). Therefore, a formulation based on Equations (20) and (21) rather than Equations (9) and (10) will automatically satisfy the constraint that $\mathbf{g} = \mathbf{0}$ at each Gauss-Newton iteration and will not require the omission of the constraint Hessian since it is exactly zero. We adopt this strategy and iteratively solve Equations (20) and (21) for primary position unknowns \mathbf{x}_k^{n+1} and secondary element stresses \mathbf{P}_k^e in a Gauss-Seidel manner analogous to that of the original XPBD. We observe that this retains the favorable stability properties of XPBD, while allowing for accurate residual reduction and application to arbitrary hyperelastic constitutive models.

This approach shifts the difficulty from the primary Equation (20) to the secondary Equation (21). It is trivial to maintain a zero primary residual, which simply requires plugging the current guess for the element stresses \mathbf{P}_k into Equation (20) to define the current guess for \mathbf{x}_k^{n+1} . We update this guess iteratively by solving for the positions in element e that satisfy Equation (21). This is equivalent to solving the nonlinear system equations for one element with the stresses in all adjacent elements held fixed, with their dependence on the element positions ignored. We use Ω^e to denote set of the mesh vertices i^e in element e and solve

$$\mathbf{M}^e(\mathbf{x}_{k+1}^{n+1,e} - \tilde{\mathbf{x}}^e) + \Delta t^2 \mathbf{B}^e \mathbf{P}_{k+1}^e = \mathbf{f}^e \quad (22)$$

$$\frac{\partial \Psi}{\partial \mathbf{F}}(\mathbf{F}^e(\mathbf{x}_{k+1}^{n+1})) - \mathbf{P}_{k+1}^e = \mathbf{0} \quad (23)$$

where $f_{i^e \alpha k}^e = \Delta t^2 (f_{i^e \alpha}^{ext} - \sum_{\tilde{e} \neq e, \gamma, \delta} B_{i^e \alpha \gamma \delta}^{\tilde{e}} P_{k \gamma \delta}^{\tilde{e}})$, \mathbf{f}^{ext} is the external force. In index notations Equation 22 can be written as:

$$\sum_{j^e, \beta} m_{i^e j^e} \delta_{\alpha\beta} \left(x_{k+1j^e\beta}^{n+1} - \tilde{x}_{j^e\beta} \right) + \Delta t^2 \sum_{\gamma, \delta} B_{i^e \alpha \gamma \delta}^e P_{k+1\gamma\delta}^e = f_{i^e \alpha k}^e \quad (24)$$

Here Equation (23) can be satisfied trivially by setting $\mathbf{P}_{k+1}^e = \frac{\partial \Psi}{\partial \mathbf{F}}(\mathbf{F}^e(\mathbf{x}_{k+1}^{n+1}))$. With this simplification, Equations (23)-(24) can be rewritten as

$$\mathbf{M}^e(\mathbf{x}_{k+1}^{n+1,e} - \tilde{\mathbf{x}}^e) + \Delta t^2 \mathbf{B}^e \frac{\partial \Psi}{\partial \mathbf{F}}(\mathbf{F}^e(\mathbf{x}_{k+1}^{n+1,e})) - \mathbf{f}^e = \mathbf{0} \quad (25)$$

where $M_{i^e \alpha j^e \beta}^e = m_{i^e j^e} \delta_{\alpha\beta}$ is the element-wise mass matrix and $\mathbf{x}_k^{n+1,e}$ and $\tilde{\mathbf{x}}^e$ are extractions of element-wise positions from \mathbf{x}_k^{n+1} and $\tilde{\mathbf{x}}$ respectively. Note that $\frac{\partial \Psi}{\partial \mathbf{F}}(\mathbf{F}^e(\mathbf{x}_{k+1}^{n+1})) = \frac{\partial \Psi}{\partial \mathbf{F}}(\mathbf{F}^e(\mathbf{x}_{k+1}^{n+1,e}))$ since the element deformation gradient only depends on the nodes of the element. Lastly, \mathbf{B}^e has entries $B_{i^e \alpha \gamma \delta}^e = \delta_{\alpha\gamma} \frac{\partial N_i}{\partial X_\delta}(\mathbf{X}^e) V^e$ from Equation (19).

We use Newton's method to solve Equation (25). $\mathbf{x}_{i^e k+1l}^{n+1,e}$ denotes the l^{th} iteration of the local Newton procedure for computing the $k+1^{\text{th}}$ global iteration, which modifies the nodes i^e of element e . These nodes are updated in Newton's method as $\mathbf{x}_{i^e k+1l+1}^{n+1,e} =$

$\mathbf{x}_{i^e k+1}^{n+1,e} + \delta \mathbf{x}_{i^e k+1}^e$. To solve for $\delta \mathbf{x}_{i^e k+1}^e$, we need to solve a linear system of size 12×12 (6×6 in 2D). To reduce the size of the system, we use an affine basis for the change in positions determined by a Newton step:

$$\delta \mathbf{x}_{i^e k+1}^e = \delta \mathbf{F}_{k+1}^e (\mathbf{X}_{i^e}^e - \mathbf{X}_{\text{com}}^e) + \mathbf{b}_{k+1}^e \quad (26)$$

where $\delta \mathbf{F}_{k+1}^e$ are the distortional degrees of freedom in the element, \mathbf{b}_{k+1}^e are the translational degrees of freedom and $\mathbf{X}_{\text{com}}^e$ is the center of mass of the element in the undeformed configuration. We refer the readers to the technical document [2023] for details. Similarly, $\mathbf{X}_{i^e}^e$ refer to the undeformed positions of the vertices in the element. The tensor $D_{i^e \alpha \epsilon \tau}^e = \delta_{\alpha \epsilon} (X_{i^e \tau}^e - X_{\text{com}}^e)$ relates these variables with $\delta \mathbf{x}_{i^e \alpha k+1}^e = \sum_{\epsilon, \tau} D_{i^e \alpha \epsilon \tau}^e \delta F_{\epsilon \tau k+1}^e + b_{\alpha k+1}^e$. It can be shown that with this choice of basis, the translational degrees of freedom decouple, reducing the dimensionality of the system to be solved. That is, by introducing the linearization

$$\frac{\partial \Psi}{\partial \mathbf{F}}(\mathbf{F}^e(\mathbf{x}_{k+1}^{n+1,e})) \approx \frac{\partial^2 \Psi}{\partial \mathbf{F}^2}(\mathbf{F}^e(\mathbf{x}_{k+1}^{n+1,e})) : \delta \mathbf{F}_{k+1}^e + \mathbf{P}_{k+1}^e \quad (27)$$

we can solve for the translational component as

$$\mathbf{b}_{k+1}^e = \frac{-1}{m^e} \sum_{i^e \in \Omega^e} \mathbf{g}_{i^e k+1} \quad (28)$$

$$g_{i^e \alpha k+1} = m_{i^e} (\mathbf{x}_{i^e \alpha k+1}^{n+1} - \bar{\mathbf{x}}_{i^e \alpha}) - f_{i^e \alpha k}^e + \Delta t^2 \sum_{\gamma, \delta} B_{i^e \alpha \gamma \delta}^e p_{k+1 \gamma \delta}^e \quad (29)$$

where $\mathbf{P}_{k+1}^e = \frac{\partial \Psi}{\partial \mathbf{F}}(\mathbf{F}^e(\mathbf{x}_{k+1}^{n+1,e}))$. Furthermore, once the translational component \mathbf{b}_{k+1}^e is determined from Equation (28), the distortional degrees of freedom $\delta \mathbf{F}_{k+1}^e$ can be determined from the linear system

$$\left(\tilde{\mathbf{M}}^e + \Delta t^2 V^e \frac{\partial^2 \Psi}{\partial \mathbf{F}^2}(\mathbf{F}^e(\mathbf{x}_{k+1}^{n+1,e})) \right) \delta \mathbf{F}_{k+1}^e = -\mathbf{G}_{k+1}^e \quad (30)$$

where

$$\mathbf{G}_{\eta v k+1}^e = \sum_{i^e, \alpha} D_{i^e \alpha \eta v}^e (g_{i^e \alpha k+1} + m_{i^e} b_{\alpha k+1}^e) \quad (31)$$

$$\tilde{\mathbf{M}}_{\eta v \epsilon \tau}^e = \sum_{i^e, j^e \in \Omega^3} D_{i^e \alpha \eta v}^e m_{i^e} j^e D_{j^e \alpha \epsilon \tau}^e. \quad (32)$$

Note that the choice of the center of the center of the translation as well as the partition of unity/reproduction of linear functions properties of the FEM interpolation functions are the keys to this decoupling. Also note that the matrix $\tilde{\mathbf{M}}^e$ is block diagonal with the structure $\tilde{M}_{\alpha \beta \gamma \delta}^e = \delta_{\alpha \gamma} \tilde{M}_{\beta \delta}^e$, where interestingly $\tilde{\mathbf{M}}^e = \sum_{i^e} m_{i^e} (\mathbf{X}_{i^e}^e - \mathbf{X}_{\text{com}}^e) (\mathbf{X}_{i^e}^e - \mathbf{X}_{\text{com}}^e)^T$ is the affine inertia tensor used in [Jiang et al. 2015]. Derivation details are provided in the supplemental technical document [Y.Chen et al. 2023].

3.3.3 Quasi-Newton. In general, solving Equation (30) for the distortional $\delta \mathbf{F}_{k+1}^e$ requires the solution of a 9×9 linear system (4×4 in 2D). However, we generally know (or can compute with minimal effort) the eigen decomposition of $\frac{\partial^2 \Psi}{\partial \mathbf{F}^2}(\mathbf{F}^e(\mathbf{x}_{k+1}^{n+1,e}))$ [Smith et al. 2019; Teran et al. 2005b]. Since $\tilde{\mathbf{M}}^e$ is constant and block diagonal, its inverse can be precomputed with minimal storage and the inverse of $\tilde{\mathbf{M}}^e + \Delta t^2 \frac{\partial^2 \Psi}{\partial \mathbf{F}^2}(\mathbf{F}^e(\mathbf{x}_{k+1}^{n+1,e}))$ can be approximated using the Sherman-Morrison rank-1 update formula [Hager 1989]. However, if all eigen modes are used, this computation can be costly. We

therefore use just a few modes in a quasi-Newton strategy, where the cost of the slow down in Newton convergence must be balanced against higher computational time per iteration, brought by using more modes in the Sherman-Morrison formula. In the case of the corotated model, we can use

$$\frac{\partial^2 \Psi}{\partial \mathbf{F}^2}(\mathbf{F}_e) \approx 2\mu \mathbf{I} + \lambda \frac{\partial \det(\mathbf{F}^e)}{\partial \mathbf{F}^e} \otimes \frac{\partial \det(\mathbf{F}^e)}{\partial \mathbf{F}^e} \quad (33)$$

where \mathbf{I} is the 9×9 (4×4 in 2D) identity matrix (see the supplementary technical document for more detail [Y.Chen et al. 2023]). With this approximation, we can use the Sherman-Morrison formula for

$$\left(\tilde{\mathbf{M}}^e + \Delta t^2 V^e \frac{\partial^2 \Psi}{\partial \mathbf{F}^2}(\mathbf{F}_e(\mathbf{x}_{k+1}^{n+1,e})) \right)^{-1} \approx \mathbf{Z}^e - \frac{\mathbf{Z}^e (\mathbf{W}^e \otimes \mathbf{W}^e) \mathbf{Z}^e}{1 + \mathbf{W}^e : (\mathbf{Z}^e \mathbf{W}^e)} \quad (34)$$

where $\mathbf{W}^e = \frac{\partial \det(\mathbf{F}^e)}{\partial \mathbf{F}^e}$ and $\mathbf{Z}^e = (\tilde{\mathbf{M}}^e + 2V^e \Delta t^2 \mu \mathbf{I})^{-1}$. Note that \mathbf{Z}^e is constant and has the same symmetric block diagonal structure as $\tilde{\mathbf{M}}^e$ so its inverse can be precomputed and stored with only 6 floats (3 in 2D).

While the procedure outlined in Section 3.3.3 requires some elaborate notation, we note that it is effectively a standard Newton's method for FEM-discretized hyperelasticity on a single element. The only difference is that the stresses from the neighboring elements do not change when the element nodal positions change. This is inherent in the introduction of the stresses \mathbf{P}^e as additional variables in Equations (20)-(21). The stresses from the neighboring elements just contribute forces that effect the right hand side of the Newton procedure, but not the matrix in the linearization. We summarize the process in Algorithm 1. In general, we run with 1-5 Newton iterations. As discussed in Section 3.3.3, with our novel approximation of the Hessian, the cost of solving the linear system becomes trivial. The major cost of the computation time for both XPBD and FP-PXPBD is computing the singular value decomposition of \mathbf{F}^e . As shown in Section 5.1, 5.2, 5.3 and 5.4 the speed of FP-PXPBD is comparable to XPBD.

4 PARALLELISM

Computation of the element-wise updates must be done in parallel for optimal efficiency. Even though we use a Gauss-Seidel (as opposed to Jacobi) approach, we can achieve this with careful ordering of element-wise updates. This was discussed by Macklin and Müller [2021], however their approach is limited to tetrahedral meshes created from hexahedral meshes. We provide a simple coloring scheme that works for all tetrahedron meshes. The coloring is done so that elements in the same color do not share vertices and can be updated in parallel without race conditions. For each vertex \mathbf{x}_i in the mesh we maintain a set $S_{\mathbf{x}_i}$ that stores the colors used by its incident elements. For each mesh element e , we find the minimal color that is not contained in the set $\cup_{\mathbf{x}_i, e \in e} S_{\mathbf{x}_i}$. Then, we register the color by adding it into $S_{\mathbf{x}_i}$ for each $\mathbf{x}_i \in e$. This coloring strategy does not depend on the topology of the mesh and requires only a one-time cost at the beginning of the simulation. For the grid-based variation mentioned in Section 5.5, we do a similar process as the coloring scheme for the mesh, except the incident points of an element are now a subset of the grid nodes. The grid-based variation requires coloring at the beginning of every time step. For more detailed descriptions and illustrations we refer the readers to [Y.Chen et al. 2023]. We note that Fratarcangeli et al.

[2016] develop a randomized and effective ordering technique that could be used here as well.

5 EXAMPLES

We demonstrate our methods in a variety of representative scenarios with elastic deformation. Our approach has comparable computational complexity to XPBD, so we only provide limited run-time statistics in Table 1. Examples run with the corotated model (Equation 3) use the algorithm from [Gast et al. 2016] for its accuracy and efficiency. All the examples were run on an AMD Ryzen Threadripper PRO 3995WX CPU with 64 cores and 128 threads.

In each of the examples, we compute the mass m_i of node \mathbf{x}_i from a user-specified density ρ . We denote \mathcal{I} to be the set of elements that contain node i . We define $m_i = \sum_{e \in \mathcal{I}} \frac{V_e \rho}{n_e}$, where $n_e = 4$ in 3D and 3 in 2D. Then the mass matrix is set with $M_{i\alpha j\beta} = \delta_{ij} \delta_{\alpha\beta} m_i$. We compute Lamé parameters μ and λ with Poisson ratio ν and Young’s modulus E . They are computed as $\mu = \frac{E}{2(1+\nu)}$, $\lambda = \frac{E\nu}{(1+\nu)(1-2\nu)}$. For all the examples in this paper we set Poisson ratio $\nu = 0.3$ and density $\rho = 10$.

5.1 Residual Comparison

Residual reduction between XPBD, B-PXPBD and FP-PXPBD is compared. Figure 1 (right) shows the residual reduction in a representative time step of a simple elasticity simulation. While B-PXBD and FP-PXBD continually reduce the nonlinear backward Euler residual, XPBD stagnates. XPBD effectively reduces the auxiliary residual, but not the primary residual. The example setup is the same as the one shown in Figure 3(b). B-PXPBD has blending parameter $\zeta = 0.5$.

5.2 Equal Budget Comparison

In Figure 2 we compare methods when simulated with a restricted computational budget. At the left we show Newton’s method run to full-convergence (residual of Equation (7) less than $1e^{-8}$), which is computationally expensive. Then, we compare (from left to right) Newton’s method, FP-PXPBD, B-PXPBD and XPBD when only allowed 200ms of computation time per frame. Newton’s method is remarkably unstable, but the XPBD-style methods are stable and visually plausible. Here we fix the left side of the tetrahedron mesh created from a $32 \times 32 \times 32$ grid and apply gravity. The Young’s modulus is $E = 1000$ and the time step is $\Delta t = 0.01$. B-PXPBD has blending parameter $\zeta = 0.1$.

5.3 XPBD Hyperelastic

We demonstrate that XPBD is incapable of dealing with certain hyperelastic models. The top bar is simulated with XPBD and the corotated model, where the constraint is reformulated as $C_e(\mathbf{x}) = \sqrt{\Psi^{cor}(\mathbf{F}^e)}$. The middle bar is simulated with FP-PXPBD and the bottom bar XPBD as formulated in Equation (5). As demonstrated in Figure 5, the top bar becomes unstable after a couple of time steps. The reformulation at the top is a simple means of addressing general hyperelasticity with a XPBD formulation, however it does not behave stably. For this example we take a rectangular mesh and clamp both ends which are then stretched and then squeezed. We set Young’s modulus as $E = 1e^4$ and time step $\Delta t = 0.01$.

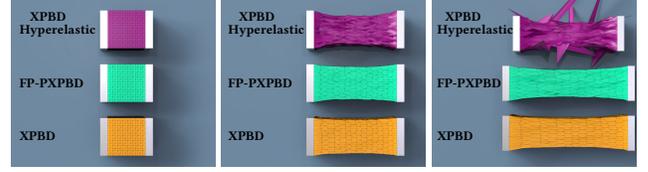


Figure 5: XPBD Hyperelastic. Defining the XPBD constraint as the square root of the hyperelastic potential is not stable (top). Results at frame 0, 10, 30 are shown.

5.4 XPBD NeoHookean

In this example we compare XPBD and FP-PXPBD when used with the Neo-Hookean model proposed in Macklin et al.[2021]. We generalize the low-rank approximation to the Hessian from Equation 33 to this model as $\frac{\partial^2 \Psi}{\partial \mathbf{F}^e}(\mathbf{F}^e) \approx \mu \mathbf{I} + \lambda \frac{\partial \det(\mathbf{F}^e)}{\partial \mathbf{F}^e} \otimes \frac{\partial \det(\mathbf{F}^e)}{\partial \mathbf{F}^e}$. Similarly, we can approximate the Hessian inverse as in Equation 34 with $\mathbf{Z}^e = (\tilde{\mathbf{M}}^e + V^e \Delta t^2 \mu \mathbf{I})^{-1}$. The test scenario is similar to that in Section 5.3. We use Young’s modulus $E = 1000$ and time step $\Delta t = 0.01$. Results are shown in Figure 6. The top bar is simulated with XPBD and is run with 100 iterations per time step. However, it does not converge to the ground truth run with Newton’s method, which is shown in the bottom row. It is also visibly less volume conserving. On the other hand, FP-PXPBD converges to the ground truth with 10 iterations per time step.

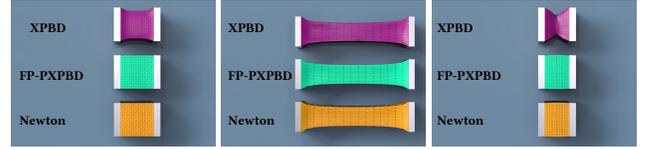


Figure 6: XPBD NeoHookean. XPBD is less volume-conserving than FP-PXPBD when the cube is squeezed. Results at frame 1, 25, 52 are shown.

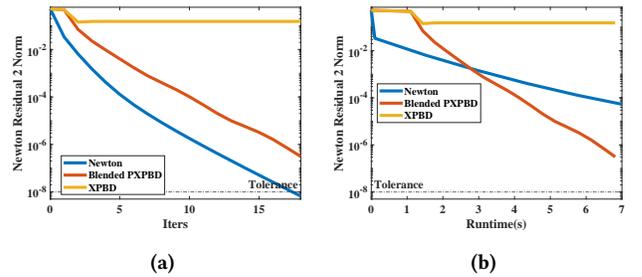


Figure 7: (a) Grid-Based Residual vs. Iterations. Newton’s method and B-PXPBD reliably reduce the residual, but XPBD stagnates. (b) Grid-based Residual vs. Runtime. Grid-based B-PXPBD and grid-based XPBD take an extra 1 second at the beginning of each timestep to compute preprocessing data. Note that B-PXPBD achieves faster convergence than Newton’s method.

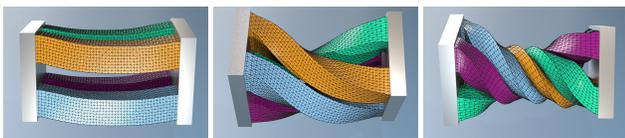
Table 1: Timing Comparisons: runtime is measured for each frame (averaged over the course of the simulation). Each frame is run after advancing time .033.

Example	# Vertices	# Elements	XPBD Runtime	B-PXPBD Runtime	FP-PXPBD Runtime	XPBD # iter	B-PXPBD # iter	FP-PXPBD # iter
Residual Reduction (Figure 3(b))	4K	17K	200ms	200ms	216ms	40	40	40
Equal Budget Comparison (Figure 2)	33K	149K	210ms	210ms	200ms	7	7	5
XPBD Hyperelastic (Figure 5)	4K	17K	22ms	-	44ms	4	-	4
XPBD NeoHookean (Figure 6)	4K	17K	795ms	-	345ms	400	-	40
Simple Muscle (Figure 4)	5k	20k	-	-	160ms	-	-	4

5.5 Grid-Based B-PXPBD Examples

We showcase the versatility and the robustness of B-PXPBD through a variety of collision intensive examples. We use the grid-based approach of Jiang et al. [2015] since this approach does not require modification of the potential energy to address collision/contact and therefore clearly demonstrates our solver performance. The backward Euler degrees of freedom are over a regular grid where the tetrahedron mesh is embedded/interpolated from its motion. We use B-PXPBD to solve the system of equations for implicit time stepping outlined in Jiang et al. [2015], but the energy is written in the XPBD way using the constraints Equation (5). This requires a modification to the coloring strategy used for parallel implementation (see supplementary technical document for specifics [Y.Chen et al. 2023]) but is otherwise a straightforward application of our techniques so we omit explicit detail.

In Figure 1, we drop 30 objects stacked on top of each other into a glass box. The objects include bunnies, dragons, balls, boxes and tori. The bunny mesh used is obtained from [Turk and Levoy 1994] The total vertex count of the mesh is around 800,000. We visualize the convergence behaviors of grid-based XPBD, grid-based B-PXPBD and Newton’s method in Figure 7. While the residual of grid-based XPBD stagnates, grid-based B-PXPBD continually reduces the non-linear residual. Though grid-based B-PXPBD has a convergence rate that is slower than Newton’s method, it has a much lower computational budget than Newton’s method. As the right plot in Figure 7 indicates, given the same computational budget, grid-based B-PXPBD would reduce residual more than Newton’s method. On average, the grid-based B-PXPBD takes 17.6s/frame, whereas Newton’s method takes 58.9s/frame. We demonstrate more collision-intensive scenarios in Figures 8, 10 and Figure 9. For these examples, the Young’s modulus is $E = 1000$ and CFL number is .4. B-PXPBD has blending parameter $\zeta = 0.5$.

**Figure 8: Four Bars Twisting. Grid-based B-PXPBD is capable of handling large deformation and complex collisions.**

6 DISCUSSION AND LIMITATIONS

Our framework effectively addresses XPBD convergence issues with hyperelasticity and allows for generalization to arbitrary constitutive models. However, our approach does have limitations.

**Figure 9: Muscle. Large-scale muscle simulation using grid-based B-PXPBD. Frames 30, 60, 140 are shown.**

With B-PXPBD the blending parameter ζ can require numerous simulations to establish a useful value. FP-PXPBD is more general, but the local step may be more costly and the Sherman-Morrison formula must be applied on a case-by-case basis for different constitutive models. It also requires a larger deviation from an existing XPBD code. In practice, these considerations must be weighed when deciding which approach to use.

**Figure 10: Dropping Dragons. Grid-based simulation with B-PXPBD exhibits many collision-driven large deformations.**

REFERENCES

- D. Baraff and A. Witkin. 1998. Large Steps in Cloth Simulation. In *Proc ACM SIGGRAPH (SIGGRAPH '98)*. 43–54.
- J. Barbić and D. James. 2005. Real-Time Subspace Integration for St. Venant-Kirchhoff Deformable Models. *ACM Trans. Graph.* 24, 3 (2005), 982?990.
- C. Batty, F. Bertails, and R. Bridson. 2007. A fast variational framework for accurate solid-fluid coupling. *ACM Trans Graph* 26, 3 (2007).
- D. Bertsekas. 1997. Nonlinear programming. *J Op Res Soc* 48, 3 (1997), 334–334.
- S. Blemler, P. Pinsky, and S. Delp. 2005. A 3D model of muscle reveals the causes of nonuniform strains in the biceps brachii. *J. Biomech* 38, 4 (2005), 657–665.
- J. Bonet and R. Wood. 2008. *Nonlinear continuum mechanics for finite element analysis*. Cambridge University Press.
- S. Bouaziz, S. Martin, T. Liu, L. Kavan, and M. Pauly. 2014. Projective Dynamics: Fusing Constraint Projections for Fast Simulation. *ACM Trans Graph* 33, 4 (2014), 154:1–154:11.
- S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. 2011. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning* 3, 1 (2011), 1–122.
- I. Chao, U. Pinkall, P. Sanan, and P. Schröder. 2010. A Simple Geometric Model for Elastic Deformations. *ACM Trans Graph* 29, 4, Article 38 (2010), 6 pages.
- Y. Fan, J. Litven, and D. Pai. 2014. Active Volumetric Musculoskeletal Systems. *ACM Trans Graph* 33, 4 (2014), 152:1–152:9.
- M. Fratarcangeli, T. Valentina, and F. Pellacini. 2016. Vivace: a practical gauss-seidel method for stable soft body dynamics. *ACM Trans Graph* 35, 6 (Nov 2016), 1–9. <https://doi.org/10.1145/2980179.2982437>
- T. Gast, C. Fu, C. Jiang, and J. Teran. 2016. *Implicit-shifted Symmetric QR Singular Value Decomposition of 3x3 Matrices*. Technical Report. University of California Los Angeles.

- T. Gast, C. Schroeder, A. Stomakhin, C. Jiang, and J. Teran. 2015. Optimization Integrator for Large Time Steps. *IEEE Trans Vis Comp Graph* 21, 10 (2015), 1103–1115.
- R. Goldenthal, D. Harmon, R. Fattal, M. Bercovier, and E. Grinspun. 2007. Efficient Simulation of Inextensible Cloth. *ACM Trans Graph* 26, 3, Article 49 (July 2007).
- W. Hager. 1989. Updating the inverse of a matrix. *SIAM review* 31, 2 (1989), 221–239.
- F. Hecht, Y. Lee, J. Shewchuk, and J. O'Brien. 2012. Updated Sparse Cholesky Factors for Corotational Elastodynamics. *ACM Trans Graph* 31, 5, Article 123 (2012), 13 pages.
- G. Irving, J. Teran, and R. Fedkiw. 2004. Invertible Finite Elements for Robust Simulation of Large Deformation. In *Proc ACM SIGGRAPH/Eurograph Symp Comp Anim*. 131–140.
- C. Jiang, C. Schroeder, A. Selle, J. Teran, and A. Stomakhin. 2015. The Affine Particle-In-Cell Method. *ACM Trans Graph* 34, 4 (2015), 51:1–51:10.
- L. Kharevych, W. Yang, Y. Tong, E. Kanso, J.E. Marsden, P. Schröder, and M. Desbrun. 2006. Geometric, Variational Integrators for Computer Animation. In *Proc 2006 ACM SIGGRAPH/Eurograph Symp Comp Anim* (Vienna, Austria) (SCA '06). Eurographics Association, 437:51.
- T. Kim, F. De Goes, and H. Iben. 2019. Anisotropic Elasticity for Inversion-Safety and Element Rehabilitation. *ACM Trans. Graph.* 38, 4 (2019), 15 pages.
- S. Kovalsky, M. Galun, and Y. Lipman. 2016. Accelerated Quadratic Proxy for Geometric Optimization. *ACM Trans Graph* 35, 4, Article 134 (2016), 11 pages.
- N. Kwatra, J. Grétarsson, and R. Fedkiw. 2010. Practical Animation of Compressible Flow for ShockWaves and Related Phenomena.. In *Symp Comp Anim*. 207–215.
- N. Kwatra, J. Su, J. Grétarsson, and R. Fedkiw. 2009. A method for avoiding the acoustic time step restriction in compressible flow. *J Comp Phys* 228, 11 (2009), 4146–4161.
- M. Li, M. Gao, T. Langlois, C. Jiang, and D. Kaufman. 2019. Decomposed Optimization Time Integrator for Large-Step Elastodynamics. *ACM Trans Graph* 38, 4 (jul 2019), 10 pages. <https://doi.org/10.1145/3306346.3322951>
- L. Liu, L. Zhang, Y. Xu, C. Gotsman, and S. Gortler. 2008. A Local/Global Approach to Mesh Parameterization. In *Proc Symp Geom Proc (SGP '08)*. Eurograph Assoc, 14957:1504.
- T. Liu, A. Bargteil, J. O'Brien, and L. Kavan. 2013. Fast Simulation of Mass-Spring Systems. *ACM Trans Graph* 32, 6 (2013), 209:1–7.
- T. Liu, S. Bouaziz, and L. Kavan. 2017. Quasi-Newton Methods for Real-Time Simulation of Hyperelastic Materials. *ACM Trans Graph* 36, 4, Article 116a (2017), 16 pages.
- M. Macklin and M. Müller. 2021. A Constraint-based Formulation of Stable Neo-Hookean Materials. In *Motion, Interaction and Games*. ACM, 1–7. <https://doi.org/10.1145/3487983.3488289>
- M. Macklin, M. Müller, and N. Chentanez. 2016. XPBD: Position-Based Simulation of Compliant Constrained Dynamics. In *Proc 9th Int Conf Motion Games* (Burlingame, California) (MIG '16). ACM, 497:54.
- S. Martin, B. Thomaszewski, E. Grinspun, and M.Gross. 2011. Example-Based Elastic Materials. In *ACM SIGGRAPH 2011* (Vancouver, British Columbia, Canada) (SIGGRAPH '11). ACM, Article 72, 8 pages.
- A. McAdams, Y. Zhu, A. Selle, M. Empey, R. Tamstorf, J. Teran, and E. Sifakis. 2011. Efficient Elasticity for Character Skinning with Contact and Collisions. *ACM Trans Graph* 30, 4 (2011), 37:1–37:12.
- V. Modi, L. Fulton, A. Jacobson, S. Sueda, and D. Levin. 2021. Emu: Efficient muscle simulation in deformation space. In *Comp Graph Forum*, Vol. 40. Wiley Online Library, 234–248.
- M.Tournier, M. Nesme, B. Gilles, and F. Faure. 2015. Stable Constrained Dynamics. *ACM Trans Graph* (2015), 1–10.
- M. Müller and M. Gross. 2004. Interactive virtual materials. In *Proc Graph Int*. Canadian Human-Computer Communications Society, 239–246.
- M. Müller, B. Heidelberger, M. Hennix, and J. Ratcliff. 2007. Position based dynamics. *J Vis Comm Im Rep* 18, 2 (2007), 109–118.
- M. Müller, B. Heidelberger, M. Teschner, and M. Gross. 2005. Meshless deformations based on shape matching. In *ACM transactions on graphics (TOG)*, Vol. 24. ACM, 471–478.
- R. Narain, M. Overby, and G. Brown. 2016. ADMM Projective Dynamics: Fast Simulation of General Constitutive Models. In *Proc ACM SIGGRAPH/Eurograph Symp Comp Anim* (Zurich, Switzerland) (SCA '16). Eurograph Assoc, 217:28.
- J. Neuberger. 1985. Steepest descent and differential equations. *J Math Soc Japan* 37, 2 (1985), 187–195.
- J. Nocedal and S. Wright. 2006. Conjugate gradient methods. *Num Opt* (2006), 101–134.
- J. O'Brien, A. Bargteil, and J. Hodgins. 2002. Graphical modeling and animation of ductile fracture. In *Proc ACM SIGGRAPH 2002*. 291–294.
- J. O'Brien and J. Hodgins. 1999. Graphical modeling and animation of brittle fracture. In *Proc 26th Conf Comp Graph Int Tech (SIGGRAPH '99)*. ACM Press/Addison-Wesley Publishing Co., 137–146. <https://doi.org/10.1145/311535.311550>
- X. Provat. 1995. Deformation constraints in a mass-spring model to describe rigid cloth behaviour. In *Graph Int*. Canadian Information Processing Society, 147–155.
- M. Rabinovich, R. Poranne, D. Panozzo, and O. Sorkine-Hornung. 2017. Scalable Locally Injective Mappings. *ACM Trans Graph* 36, 2, Article 16 (2017), 16 pages.
- E. Sifakis and J. Barbic. 2012. FEM simulation of 3D deformable solids: a practitioner's guide to theory, discretization and model reduction. In *ACM SIGGRAPH 2012 Courses (SIGGRAPH '12)*. ACM, 20:1–20:50.
- B. Smith, F. Goes, and T. Kim. 2019. Analytic eigensystems for isotropic distortion energies. *ACM Trans Graph (TOG)* 38, 1 (2019), 1–15.
- B. Smith, F. De Goes, and T. Kim. 2018. Stable neo-hookean flesh simulation. *ACM Trans Graph (TOG)* 37, 2 (2018), 1–15.
- O. Sorkine and M. Alexa. 2007. As-Rigid-As-Possible Surface Modeling. In *EUROGRAPHICS SYMPOSIUM ON GEOMETRY PROCESSING*.
- A. Stern and M. Desbrun. 2006. Discrete Geometric Mechanics for Variational Time Integrators. In *ACM SIGGRAPH 2006 Courses* (Boston, Massachusetts) (SIGGRAPH '06). ACM, 757:80.
- A. Stomakhin, R. Howes, C. Schroeder, and J. Teran. 2012. Energetically consistent invertible elasticity. In *Proc Symp Comp Anim*. 25–32.
- A. Stomakhin, C. Schroeder, C. Jiang, L. Chai, J. Teran, and A. Selle. 2014. Augmented MPM for phase-change and varied materials. *ACM Trans Graph* 33, 4 (2014), 138:1–138:11.
- J. Teran, S. Blemker, V. Hing, and R. Fedkiw. 2003. Finite volume methods for the simulation of skeletal muscle. In *Proc 2003 ACM SIGGRAPH/Eurograph Symp Comp Anim*. Eurographics Association, 68–74.
- J. Teran, E. Sifakis, S. Blemker, V. Ng-Thow-Hing, C. Lau, and R. Fedkiw. 2005a. Creating and simulating skeletal muscle from the visible human data set. *IEEE Trans Vis Comp Graph* 11, 3 (2005), 317–328.
- J. Teran, E. Sifakis, G. Irving, and R. Fedkiw. 2005b. Robust quasistatic finite elements and flesh simulation. In *Proc 2005 ACM SIGGRAPH/Eurograph Symp Comp Anim*. 181–190.
- G. Turk and M. Levoy. 1994. Stanford Bunny. <http://graphics.stanford.edu/data/3Dscanrep/> Stanford University Computer Graphics Laboratory.
- B. Wang, M. Zheng, and J. Barbič. 2020. Adjustable Constrained Soft-Tissue Dynamics. *Pac Graph 2020 and Comp Graph Forum* 39, 7 (2020).
- H. Wang. 2015. A Chebyshev Semi-Iterative Approach for Accelerating Projective and Position-Based Dynamics. *ACM Trans Graph* 34, 6, Article 246 (nov 2015), 9 pages.
- H. Wang and Y. Yang. 2016. Descent methods for elastic body simulation on the GPU. *ACM Trans Graph* 35, 6 (Nov 2016), 1–10. <https://doi.org/10.1145/2980179.2980236>
- B. Wittemeyer, N. Weidner, T. Davis, T. Kim, and S. Sueda. 2021. QLB: Collision-Aware Quasi-Newton Solver with Cholesky and L-BFGS for Nonlinear Time Integration. In *Proc 14th ACM SIGGRAPH Conf Mot Int Games (MIG '21)*. ACM, Article 14, 7 pages.
- H. Xu, F. Sin, Y. Zhu, and J. Barbič. 2015. Nonlinear Material Design Using Principal Stretches. *ACM Trans Graph* 34, 4 (2015).
- Y.Chen, Y.Han, J.Chen, S.Ma, R.Fedkiw, and J.Teran. 2023. *Supplementary Technical Document* (2023).
- D. Zhao, Y. Li, and J. Barbič. 2016. Asynchronous Implicit Backward Euler Integration. (2016).
- Y. Zhu, R. Bridson, and D. Kaufman. 2018. Blended Cured Quasi-Newton for Distortion Optimization. *ACM Trans Graph* 37, 4, Article 40 (jul 2018), 14 pages.