



Multi-Layer Thick Shells

Yunuo Chen
UCLA
USA
yunuoch@gmail.com

Tianyi Xie
UCLA
USA
tianyixie77@g.ucla.edu

Cem Yuksel
University of Utah
& Roblox
USA
cem@ceemyuksel.com

Danny Kaufman
Adobe Research
USA
dannykaufman@gmail.com

Yin Yang
University of Utah
USA
yin.yang@utah.edu

Chenfanfu Jiang
UCLA
USA
chenfanfu.jiang@gmail.com

Minchen Li
UCLA
USA
minchernl@gmail.com



Figure 1: Thick Fashion. Our method can realistically capture both the thickness dynamics and the fine wrinkling details of leather garments, down jackets, and jeans. A single layer of our dual-quadrature prism elements is used to discretize the interior volume of these thick garments without shear locking, and a high-resolution membrane is coupled on the surface to generate complementary high-frequency wrinkles.

ABSTRACT

We present a novel mesh-based method for simulating the intricate dynamics of (potentially multi-layered) continuum thick shells. In order to accurately represent the constitutive behavior of structural responses in the thickness direction, we develop a dual-quadrature prism finite element formulation that is free from shear locking and naturally incorporates three-dimensional elastoplastic and viscoelastic constitutive models. Additionally, we introduce a simple and effective technique for coupling a high-resolution membrane layer on top of the thick shell to enable complementary high-frequency deformation modes that generate realistic wrinkles. With our novelly designed sparse basis vectors for the high-frequency deformations, the constrained Lagrangian mechanics problem is expressed as an unconstrained optimization and then efficiently solved by a custom alternating minimization technique. Our method opens up a new possibility for fast, high-quality, and thickness-aware simulations of leather garments, pillows, mats, metal boards, and potentially a variety of other thick structures.

Permission to make digital or hard copies of all or part 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 components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SIGGRAPH '23 Conference Proceedings, August 06–10, 2023, Los Angeles, CA, USA
© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-0159-7/23/08...\$15.00
<https://doi.org/10.1145/3588432.3591489>

CCS CONCEPTS

• Computing methodologies → Physical simulation.

KEYWORDS

thickness modeling, solid shell, shear locking, reduced integration, wrinkle simulation, complementary dynamics

ACM Reference Format:

Yunuo Chen, Tianyi Xie, Cem Yuksel, Danny Kaufman, Yin Yang, Chenfanfu Jiang, and Minchen Li. 2023. Multi-Layer Thick Shells. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Conference Proceedings (SIGGRAPH '23 Conference Proceedings)*, August 06–10, 2023, Los Angeles, CA, USA. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3588432.3591489>

1 INTRODUCTION

Thin shell simulation has been extensively studied in computer graphics over the past decades. The combined use of codimensional geometry and physical modeling has allowed for realistic dynamic simulations of daily wearing fabrics, papers, cards, and metal sheets with detailed wrinkles and complex motion. In contrast, the simulation of *thick shells*, such as leather jackets and yoga mats, has received relatively limited attention within the field of computer graphics. The primary obstacle here is the need to accurately depict the dynamic evolution in thickness and wrinkles while also ensuring efficient computational performance.

A commonly adopted practice in augmenting thin shell simulations with thickness is to impose a contact offset between mid-surfaces [Li et al. 2018; Narain et al. 2012]. While superficially appearing to be powerful, this approach faces difficulties in consistently maintaining the distance offset. Additionally, it lacks a physically intuitive and controllable mechanism for modeling the evolution of thickness over time. Recently, the Incremental Potential Contact (IPC) [Li et al. 2020] model was proposed as a robust and accurate contact handling solution for nonlinear elastic solids with arbitrary codimensions [Li et al. 2021]. In addition to guaranteeing an inelastic thickness (contact offset), IPC also demonstrates a certain degree of thickness dynamics on codimensional materials (see e.g. Fig. 18 in Li et al. [2021]) due to its penalty force nature. However, it lacks meaningful constitutive models in the normal-direction response, fails to consider modes of shearing deformation, and does not offer a practical way for visualizing the explicit geometry of thick structures.

Computer graphics practitioners in the industry also explored animating thick shells as visual geometries fully embedded in a volumetric body. Despite its convenience, this approach tends to encounter obstacles in capturing accurate frictional contact. Additionally, complex and parameter intensive anisotropic elastic models are mandated for properly resolving bending behaviors while avoiding artificially stiffened bending penalties and stretchy artifacts. This issue naturally arises when using volumetric linear finite elements, where bending deformation could only be resolved via shearing, imposing additional elastic forces. This behavior is called *shear locking*, which has been extensively studied in mechanical engineering, where reduced spatial integration methods [Cardoso et al. 2008; Trinh et al. 2011] have been developed to ignore certain shearing modes and avoid locking without resorting to expensive higher-order elements or impractically fine tessellations.

We follow the idea of reduced integration and propose a dual-quadrature prism element for practical thick shells, balancing accuracy and efficiency. The two quadratures are chosen as the Gaussian quadratures in the thickness direction, both on the center line of the prism, perfectly ignoring trapezoidal shearing modes to avoid locking. Our model can capture accurate bending behaviors even with a coarse tessellation. Based on the finite strain theory, our model naturally supports common *three-dimensional* elastoplastic and viscoelastic constitutive models [Li et al. 2022] for simulating inelastic behaviors. In addition to single-layer shells, our prism elements can also serve as an effective spatial discretization for volumetric bodies that have a simple-topology medial surface.

To take advantage of the fact that our prism elements can accurately resolve bending even at a low resolution, we augment a membrane layer with higher resolution on the surface of the thick shell to efficiently capture fine wrinkling details. The membrane and thick shells are spectrally coupled, where the low-frequency displacement modes of the membrane perfectly matches that of the thick shell, ensuring close attachment while allowing high-frequency wrinkles to form. With novelly defined sparse basis vectors for both the low- and high- frequency modes, the coupling is implicitly achieved via a change of variable derived based on Lagrangian mechanics. As the low- and high- frequency deformations

are sufficiently independent, we further propose an alternating minimization framework equipped with custom solvers to efficiently step in time.

In summary, our contributions include

- a dual-quadrature prism finite element formulation for efficient and robust simulation of elastoplastic and viscoelastic multi-layer thick shells without shear locking (§ 3);
- a change-of-variable method to couple coarse thick shell interior and fine membrane surface for efficiently capturing complementary fine wrinkling details (§ 4);
- an alternating minimization method with custom nonlinear and linear solvers to efficiently solve our coupled system (§ 5).

2 RELATED WORK

Continuum Solid-Shell Elements. Solid-shell elements are initially investigated by Hauptmann and Schweizerhof [1998]; Miehe [1998]; Parisch [1995]; Schoop [1986] in mechanical engineering. Compared to classical shell models such as the Kirchoff-Love [Chen et al. 2018; Panc 1975] and the Mindlin-Reissner [Bathe and Dvorkin 1985; Guo et al. 2018], solid-shell elements also capture strains and stresses in thickness direction and allows general 3D material laws to be applied [Harnau and Schweizerhof 2002]. Furthermore, without rotational degrees of freedom, solid-shell elements avoid complicated updates of the rotation vector in geometrically nonlinear analyses, and they enable convenient treatment of boundary conditions and coupling with solid elements [Hauptmann and Schweizerhof 1998].

To efficiently avoid various locking issues, the elasticity on the solid-shell elements are often integrated with reduced integration methods such as the assumed natural strain (ANS) method [Kim et al. 2005; Sze and Yao 2000; Trinh et al. 2011] and the enhanced assumed strain (EAS) method [Cardoso et al. 2008; Schwarze and Reese 2009]. The key insight for both types of methods is that parts of the strains in the element are assumed to be in their natural state or represented by a higher-order polynomial, rather than being calculated from the nodal displacements. These methods are particularly useful for problems involving large deformations, as certain deformation modes can be ignored to avoid locking. Inspired by this line of research, we apply a prismatic element to efficiently discretize multi-layer thick shells, and integrate the elasticity using only 2 quadratures on the central vertical line per element to conveniently avoid shear locking.

Wrinkle Simulation. Efficiently augmenting coarse base animation with fine wrinkling details has drawn lots of attention in computer graphics. Grinspun et al. [2002] proposed an adaptive refinement method to simulate wrinkles and folds using finer basis functions obtained from subdivisions. Bergou et al. [2007] employed constrained Lagrangian mechanics to simulate a high-resolution surface following the dynamics of an art-directed low-resolution target. Rohmer et al. [2010] added dynamic wrinkles with user-guided wrinkle size to a coarse-scale simulation by analyzing the strain tensor in real-time. Müller and Chentanez [2010] applied position-based dynamics to compute fine wrinkles by simulating high-resolution patches constrained to an input low-resolution simulation. Chen et al. [2013] simulated fine wrinkles in the cloth-body

interactions by capturing the friction and air pressure between the cloth and body. [Zuenko and Harders \[2019\]](#) solved for an amplitude and phase field to simulate the wrinkling of human skin and other materials where a stiff membrane is attached to a soft body.

Our complementary wrinkle coupling idea resembles a lot [Rémillard and Kry \[2013\]](#) and [Casafranca and Otaduy \[2022\]](#), where they both couple a fine membrane surface to a coarse mesh by matching the low-frequency displacements. However, [Rémillard and Kry \[2013\]](#) define the constraints using the averaged positions of both meshes, and they directly solve a KKT system to augment static high-frequency wrinkles. Later, [Li and Kry \[2014\]](#) extended this idea to simulate wrinkles of multi-layer skin, with different layers coupled via similar constraints. [Casafranca and Otaduy \[2022\]](#) define the constraints using low-frequency basis as we do. They focus on a special discretization where the coarse mesh is the centroidal Voronoi tessellation of the fine mesh, so that 0-th order basis can be used to make their modified preconditioned conjugate gradient solver tractable with better sparsity of the preconditioning matrices. Instead, our method supports a hierarchy of tessellations and couples different levels of dynamics in a monolithic fashion. By defining the null-space basis for subdivision meshes, our unconstrained system has no redundant degrees of freedom, and its sparsity remains nice even with multiple subdivisions.

Others such as the data-driven methods and those based on the tension field theory (TFT) have also shown effective in enhancing coarse simulation with fine wrinkles. For instance, [Chen et al. \[2021\]](#) applied the TFT to simulate efficiently parameterized wrinkles based on compression-free coarse cloth simulations. TFT was also successfully applied to simulate wrinkles on inflatable structures [[Skouras et al. 2014](#)]. For Data-driven methods, [Wang et al. \[2010\]](#) added fine details to relatively skin-tight garments by learning the cloth deformation on the pose of an underlying mannequin. [Kavan et al. \[2011\]](#) learned a dense upsampling operator to obtain more geometric details on a coarse simulated mesh. [Zurdo et al. \[2012\]](#) proposed an algorithm for synthesizing cloth wrinkles as a function of the deformation of a low-resolution cloth and a set of example poses.

3 REDUCED PRISM ELEMENTS

Inspired by [Trinh et al. \[2011\]](#), we design a reduced prism finite element to explicitly track the thickness evolution of multi-layer shells. With dual-quadrature reduced integration, shear locking can be avoided even at coarse discretization with low-order basis functions.

3.1 Kinematics

In the material space, for a prism with midsurface nodes $X_1, X_2, X_3 \in \mathbb{R}^3$ and corresponding directors $N_1, N_2, N_3 \in \mathbb{R}^3$, any point $X \in \mathbb{R}^3$ inside the prism can be parameterized by $\mathbf{q} = (\lambda_1, \lambda_2, \gamma) \in [0, 1] \times [0, 1 - \lambda_1] \times [-1, 1]$ (Fig. 2 left), which gives its material space position

$$X(\mathbf{q}) = X_1 + \lambda_1(X_2 - X_1) + \lambda_2(X_3 - X_1) + \gamma(N_1 + \lambda_1(N_2 - N_1) + \lambda_2(N_3 - N_1)). \quad (1)$$

Here the 3 directors can take any vector in \mathbb{R}^3 to construct a wide variety of thick shells with curvy rest configurations of varying

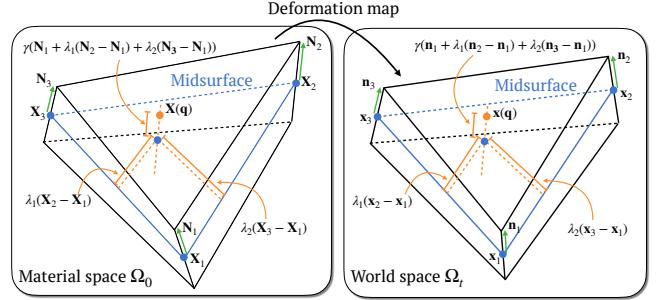


Figure 2: Parameterization and interpolation of our prism element. Spatial coordinates are parameterized with λ_1 , λ_2 and γ .

thicknesses, as long as the prism is not inverted or degenerated. With barycentric weights and bilinear interpolation, we construct the prism element in an isoparametric fashion and obtain the world space mapping of X as

$$x(\mathbf{q}) = x_1 + \lambda_1(x_2 - x_1) + \lambda_2(x_3 - x_1) + \gamma(n_1 + \lambda_1(n_2 - n_1) + \lambda_2(n_3 - n_1)), \quad (2)$$

where $x_1, x_2, x_3 \in \mathbb{R}^3$ and $n_1, n_2, n_3 \in \mathbb{R}^3$ are the world space counterparts of the midsurface nodes and directors respectively (Fig. 2 right). Now we can evaluate the deformation gradient at X as

$$F(\mathbf{q}) = \frac{\partial \mathbf{x}}{\partial \mathbf{q}}(\mathbf{q}) = \frac{\partial \mathbf{x}}{\partial \mathbf{q}} \left(\frac{\partial \mathbf{X}}{\partial \mathbf{q}} \right)^{-1}, \quad (3)$$

where $\frac{\partial \mathbf{x}}{\partial \mathbf{q}}$ and $\frac{\partial \mathbf{X}}{\partial \mathbf{q}}$ are both linear functions of $\lambda_1, \lambda_2, \gamma$; see the supplementary for details.

3.2 Spatial Discretization

We follow [Li et al. \[2022\]](#) to adopt a unified energy-based formulation of hyperelasticity, elastoplasticity and viscoplasticity. The potential energy on the entire thick shell domain Ω^0 (material space) discretized by our prism elements is integrated as

$$\int_{\Omega^0} \Psi(F) dX \approx \sum_i \int_{\Omega_i^0} \Psi(F) dX \\ = \sum_i \int_0^1 \int_0^{1-\lambda_1} \int_{-1}^1 \Psi(F_i(\mathbf{q})) \left| \left(\frac{\partial \mathbf{X}}{\partial \mathbf{q}} \right)_i \right| dy d\lambda_2 d\lambda_1, \quad (4)$$

where subscript i refers to quantities of the i -th prism element. Applying Gaussian quadrature in each dimension, the triple integral here can be properly approximated.

Nevertheless, approximating this triple integral with high accuracy can lead to shear locking [[Belytschko et al. 2014](#); [Bletzinger et al. 2000](#)]. To illustrate this in 2D, in the ideal setting, a cantilever beam has nearly no shearing anywhere, with the majority of bending force originating from the tangential stretch or compression at different levels along the thickness direction (Fig. 3a). To capture the same deformation, however, coarse linear elements must undergo significant shearing away from the central vertical line, which can exert additional elastic forces to resist bending (Fig. 3b).

Therefore, we propose a reduced integration scheme to completely ignore the trapezoidal shearing mode while evaluating the

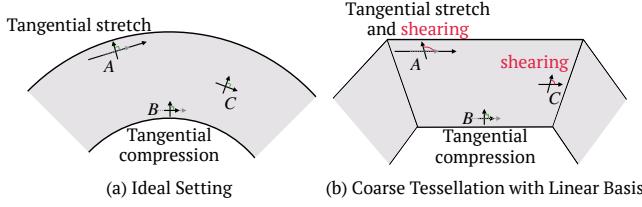


Figure 3: 2D Illustration of shear locking. Linear elements artificially incur shearing strain, limiting bending displacements.

deformation gradient. This can be conveniently achieved by only placing the Gaussian quadrature along the central vertical line, which gives our reduced potential energy

$$P_e = \sum_i \frac{1}{2} \left(\Psi(F_i(\mathbf{q}_1)) \left| \frac{\partial \mathbf{X}}{\partial \mathbf{q}}(\mathbf{q}_1) \right| + \Psi(F_i(\mathbf{q}_2)) \left| \frac{\partial \mathbf{X}}{\partial \mathbf{q}}(\mathbf{q}_2) \right| \right), \quad (5)$$

where $\mathbf{q}_1 = (1/3, 1/3, -\sqrt{3}/3)$, $\mathbf{q}_2 = (1/3, 1/3, \sqrt{3}/3)$ are the two quadratures. Note that only using a single quadrature will completely ignore the thickness modes and lead to unstable dynamics. Additionally, we do not apply the assumed strain method as in Trinh et al. [2011] to scale down the transverse shear component, since we find that ignoring trapezoidal modes alone is already sufficient to avoid shear locking under bending deformations.

Our method can also be used to discretize thicker volumes utilizing multiple layers of prism elements. However, due to our reduced integration, the discretized volume will not exhibit isotropic elastic behaviors. Thus, our elements are particularly effective for the efficient discretization of special-structured volumes, such as a pillow (Fig. 11) or a down jacket (Fig. 1).

4 COMPLEMENTARY WRINKLE COUPLING

To efficiently capture fine wrinkling details on the surface of the coarse thick shell, we attach a fine membrane to it, matching their low-frequency displacements in the coupled region and allowing high-frequency wrinkles to form.

4.1 Problem Formulation

We first transform the degrees of freedom of the thick shells from the midsurface to the outer surface nodal positions $\mathbf{y}_{2j-1} = \mathbf{x}_j^{\text{mid}} - \mathbf{n}_j$ and $\mathbf{y}_{2j} = \mathbf{x}_j^{\text{mid}} + \mathbf{n}_j$, and stack them together to form \mathbf{y} . Then, we introduce a high-resolution codimension-1 membrane mesh with all its nodal degrees of freedom stacked as \mathbf{x} . At the rest configuration, the membrane collocates with the thick shell surface (Fig. 4 left), and the coordinates of each membrane node can be expressed as a linear transformation of \mathbf{Y} , namely

$$\mathbf{X} = \mathbf{P}\mathbf{Y}, \quad (6)$$

where \mathbf{X} and \mathbf{Y} are the material space counterparts of \mathbf{x} and \mathbf{y} , and each row of the matrix \mathbf{P} contains the interpolation weights. The column vectors of the tall matrix \mathbf{P} contain low-frequency modes of the membrane's displacement.

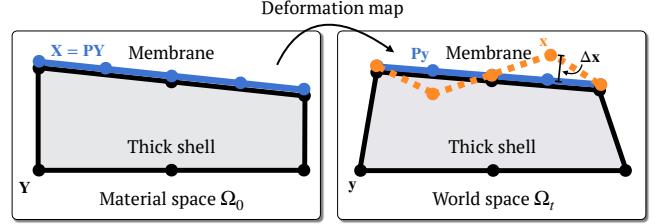


Figure 4: 2D Illustration of complementary wrinkle coupling. Note that the membrane can also be coupled with the entire thick shell surface.

In the world space, we only constrain the low-frequency displacement between \mathbf{x} and \mathbf{y} to match, which gives the constraint

$$\mathbf{x} = \mathbf{P}\mathbf{y} + \Delta\mathbf{x}, \quad (7)$$

where $\Delta\mathbf{x}$ is the high-frequency displacements (Fig. 4 right) with $\mathbf{P}^T \Delta\mathbf{x} = 0$. If we rewrite this constraint as $\mathbf{P}^T(\mathbf{x} - \mathbf{P}\mathbf{y}) = 0$ and apply Lagrangian mechanics, we obtain the spatially discretized dynamical system

$$\begin{cases} \frac{d\mathbf{x}}{dt} = \mathbf{M}_x^{-1}(-\nabla E_x(\mathbf{x}) + \mathbf{P}\lambda), & \frac{d\mathbf{x}}{dt} = \mathbf{v}_x \\ \frac{d\mathbf{y}}{dt} = \mathbf{M}_y^{-1}(-\nabla E_y(\mathbf{y}) - \mathbf{P}^T \mathbf{P}\lambda), & \frac{d\mathbf{y}}{dt} = \mathbf{v}_y, \quad \mathbf{P}^T(\mathbf{x} - \mathbf{P}\mathbf{y}) = 0 \end{cases} \quad (8)$$

where E_x and E_y are the potential energies of the membrane and thick shell including strain energy, gravity, frictional contact, etc., and λ is the Lagrange multiplier vector for the coupling constraints. This is essentially a KKT system which can be expensive to solve directly or via Schur complement. By exploring a special structure of \mathbf{P} , we propose a change-of-variable method to implicitly handle the constraints.

4.2 Change of Variable

If we assume \mathbf{X} can be expressed as barycentric coordinates of \mathbf{Y} , e.g. through subdivision, we will be able to directly define a set of sparse vector basis for the null space of \mathbf{P} . To illustrate this in 1D, consider \mathbf{Y} as a 2-segment polyline with 3 nodes $\mathbf{Y}_1, \mathbf{Y}_2, \mathbf{Y}_3$, then subdividing once gives $\mathbf{X}_1 = \mathbf{Y}_1, \mathbf{X}_2 = (\mathbf{Y}_1 + \mathbf{Y}_2)/2, \mathbf{X}_3 = \mathbf{Y}_2, \mathbf{X}_4 = (\mathbf{Y}_2 + \mathbf{Y}_3)/2, \mathbf{X}_5 = \mathbf{Y}_3$. We set

$$\mathbf{P} = \begin{bmatrix} 1 & 0 & 0 \\ 1/2 & 1/2 & 0 \\ 0 & 1 & 0 \\ 0 & 1/2 & 1/2 \\ 0 & 0 & 1 \end{bmatrix}, \quad \mathbf{Q} = \begin{bmatrix} -1/2 & 0 \\ 1 & 0 \\ -1/2 & -1/2 \\ 0 & 1 \\ 0 & -1/2 \end{bmatrix}, \quad (9)$$

where each row of $\mathbf{P} \in \mathbb{R}^{nx \times ny}$ stores the interpolation weights of both low- and high-res points, and each column of $\mathbf{Q} \in \mathbb{R}^{nx \times (nx - ny)}$ corresponds to a high-res point, with coefficients carefully chosen to ensure orthogonality to all columns of \mathbf{P} , i.e. $\mathbf{P}^T \mathbf{Q} = 0$. For 3D cases, please refer to our supplemental document for pseudo-code. With the sparse null-space matrix \mathbf{Q} , we can

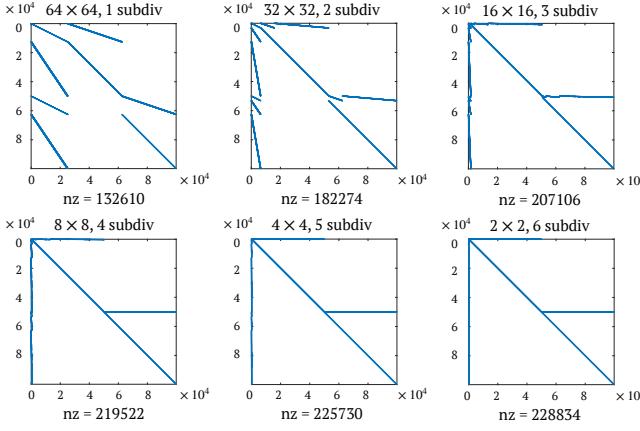


Figure 5: Sparsity pattern of $[P, Q]$. We subdivide coarse grids with different initial resolutions to achieve a final resolution of 128×128 . For example, a 64×64 grid only needs one subdivision. Here we visualize the sparsity pattern and record the number of nonzero entries nz . Regardless of the number of subdivisions, the sparsity pattern of $[P, Q]$ remains consistently nice. Note that here we did not duplicate the entries for the 3 dimensions.

rewrite the coupling constraints as

$$\mathbf{x} = \mathbf{P}\mathbf{y} + \mathbf{Q}\mathbf{z} \quad (10)$$

by introducing the high-frequency variable \mathbf{z} measuring wrinkle magnitudes. In this way, after time discretization (e.g. with implicit Euler here), our dynamical system can be reformulated as an unconstrained optimization

$$\mathbf{w}^{n+1} = \arg \min_{\mathbf{w}} \frac{1}{2} \|\mathbf{w} - \tilde{\mathbf{w}}^n\|_{\mathbf{M}_w} + h^2 \left(E_x([P, Q]\mathbf{w}) + E_y([I, 0]\mathbf{w}) \right) \quad (11)$$

followed by velocity update $\mathbf{v}_w^{n+1} = (\mathbf{w}^{n+1} - \mathbf{w}^n)/h$, where $\mathbf{w} = [\mathbf{y}^T, \mathbf{z}^T]^T$, $\tilde{\mathbf{w}}^n = \mathbf{w}^n + h\mathbf{v}_w^n$, h is the time step size, and

$$\mathbf{M}_w = \begin{bmatrix} \mathbf{M}_y + \mathbf{P}^T \mathbf{M}_x \mathbf{P} & \mathbf{P}^T \mathbf{M}_x \mathbf{Q} \\ \mathbf{Q}^T \mathbf{M}_x \mathbf{P} & \mathbf{Q}^T \mathbf{M}_x \mathbf{Q} \end{bmatrix}. \quad (12)$$

Please see our supplemental document for detailed derivations.

Our method also supports multiple subdivisions to capture a wide range of wrinkles with various wavelengths. For any node on the finer levels, its material space coordinates are still the barycentric interpolation inside one triangle element on the prism mesh surface, which at most depends on 3 prism mesh nodes. This means that \mathbf{P} will always have at most 3 nonzeros per row. Then for constructing \mathbf{P} 's null-space basis, similar to our 1D illustration, each column of \mathbf{Q} corresponding to a node on the finer levels needs at most 4 nonzeros, which are the fine node itself and the 3 fine nodes collocated with the top or bottom nodes of a prism element. Therefore, our basis matrices \mathbf{P} and \mathbf{Q} always maintain their nice sparsity patterns regardless of the number of subdivisions (Fig. 5). In practice, 2 or 3 subdivisions can already enable efficient simulation of high-quality wrinkles (Fig. 12).

Algorithm 1 Alternating Minimization

```

1:  $\mathbf{y} \leftarrow \mathbf{y}^n, \mathbf{z} \leftarrow \mathbf{z}^n, \tilde{\mathbf{y}}^n \leftarrow \mathbf{y}^n + h\mathbf{v}_y^n, \tilde{\mathbf{z}}^n \leftarrow \mathbf{z}^n + h\mathbf{v}_z^n$ 
2:  $\tilde{\mathbf{w}}^n \leftarrow [(\tilde{\mathbf{y}}^n)^T, (\tilde{\mathbf{z}}^n)^T]^T$ 
3: do
4:   // Low-frequency Step
5:    $\mathbf{x} \leftarrow \mathbf{P}\mathbf{y} + \mathbf{Q}\mathbf{z}, \mathbf{w} \leftarrow [\mathbf{y}^T, \mathbf{z}^T]^T$ 
6:    $\mathbf{H} \leftarrow \mathbf{M}_y + \mathbf{P}^T \mathbf{M}_x \mathbf{P} + h^2 \left( \nabla^2 E_y(\mathbf{y}) + \mathbf{P}^T \nabla^2 E_x(\mathbf{x}) \mathbf{P} \right)$ 
7:    $\mathbf{g} \leftarrow [\mathbf{I}, \mathbf{0}] \mathbf{M}_w (\mathbf{w} - \tilde{\mathbf{w}}^n) + h^2 \left( \nabla E_y(\mathbf{y}) + \mathbf{P}^T \nabla E_x(\mathbf{x}) \right)$ 
8:    $\mathbf{p}_y \leftarrow -\mathbf{H}^{-1} \mathbf{g}$  // via Cholesky factorization
9:    $\alpha \leftarrow$  CCD-aware Backtracking Line Search
10:   $\mathbf{y} \leftarrow \mathbf{y} + \alpha \mathbf{p}_y$ 
11:  // High-frequency Step
12:   $\mathbf{x} \leftarrow \mathbf{P}\mathbf{y} + \mathbf{Q}\mathbf{z}, \mathbf{w} \leftarrow [\mathbf{y}^T, \mathbf{z}^T]^T$ 
13:   $\mathbf{H} \leftarrow \mathbf{Q}^T \mathbf{M}_x \mathbf{Q} + h^2 \mathbf{Q}^T \nabla^2 E_x(\mathbf{x}) \mathbf{Q}$ 
14:   $\mathbf{g} \leftarrow [\mathbf{0}, \mathbf{I}] \mathbf{M}_w (\mathbf{w} - \tilde{\mathbf{w}}^n) + h^2 \mathbf{Q}^T \nabla E_x(\mathbf{x})$ 
15:   $\mathbf{p}_z \leftarrow -\mathbf{H}^{-1} \mathbf{g}$  // via matrix-free Conjugate Gradient
16:   $\alpha \leftarrow$  Backtracking Line Search
17:   $\mathbf{z} \leftarrow \mathbf{z} + \alpha \mathbf{p}_z$ 
18:  while  $\frac{1}{h} \|\mathbf{p}_y\| > \epsilon$  or  $\frac{1}{h} \|\mathbf{p}_z\| > \epsilon$ 
19:   $\mathbf{y}^{n+1} \leftarrow \mathbf{y}, \mathbf{z}^{n+1} \leftarrow \mathbf{z}$ 
20:   $\mathbf{v}_y^{n+1} \leftarrow (\mathbf{y}^{n+1} - \mathbf{y}^n)/h, \mathbf{v}_z^{n+1} \leftarrow (\mathbf{z}^{n+1} - \mathbf{z}^n)/h$ 
21: return  $\mathbf{y}^{n+1}, \mathbf{z}^{n+1}, \mathbf{v}_y^{n+1}, \mathbf{v}_z^{n+1}$ 

```

5 ALTERNATING MINIMIZATION

As the low and high frequency dynamics are sufficiently independent, we apply an alternating minimization strategy to efficiently solve the time-stepping system in Eq. 11, applying customized solvers for both subproblems.

Specifically, for each time step n , we alternate between inexactly solving a low-frequency step

$$\mathbf{y}^{i+1} = \arg \min_{\mathbf{y}} \frac{1}{2} \left\| \begin{bmatrix} \mathbf{y} \\ \mathbf{z}^i \end{bmatrix} - \tilde{\mathbf{w}}^n \right\|_{\mathbf{M}_w} + h^2 \left(E_x(\mathbf{P}\mathbf{y} + \mathbf{Q}\mathbf{z}^i) + E_y(\mathbf{y}) \right) \quad (13)$$

and a high-frequency step

$$\mathbf{z}^{i+1} = \arg \min_{\mathbf{z}} \frac{1}{2} \left\| \begin{bmatrix} \mathbf{y}^{i+1} \\ \mathbf{z} \end{bmatrix} - \tilde{\mathbf{w}}^n \right\|_{\mathbf{M}_w} + h^2 E_x(\mathbf{P}\mathbf{y}^{i+1} + \mathbf{Q}\mathbf{z}) \quad (14)$$

until convergence, where i is the number of alternating passes. We apply the line search method with projected Newton (PN) [Li et al. 2020] to robustly solve for both steps.

For the low-frequency steps, since the system for \mathbf{y} is usually in a small-to-medium scale, we apply direct Cholesky factorization to solve for the search direction in each PN iteration.

For the medium-to-large scale system in the high-frequency steps, we instead apply a matrix-free Jacobi-preconditioned conjugate gradient method to solve for the search direction. For the cases with multiple subdivisions, we have tested further splitting the solve per level of subdivision but observed similar performance. This is because we did not enforce orthogonality between modes from different levels of subdivision when constructing \mathbf{Q} , since this preserves the nice sparsity pattern. Please see our pseudo-code in Alg. 1 for more details.

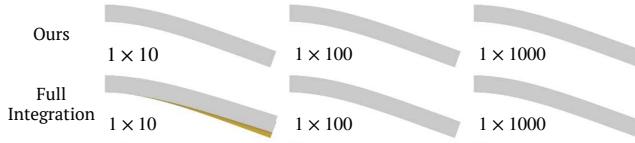


Figure 6: 2D cantilever beam. Static equilibrium of a 2D cantilever beam (left end fixed) under gravity simulated using our element with different integration schemes, where higher full integration is adopted as the reference. Our results converge early even with a single layer of 10 elements, while full integration converges slower.

6 EXPERIMENTS

We implement our method in C++ with Eigen for basic linear algebra operations, TBB for CPU multi-threading, and Thrust for GPU parallelization. We use CHOLMOD [Chen et al. 2008] to solve the thick shell systems on CPU, while developing our own matrix-free Conjugate Gradient solver on GPU. We use neo-Hookean and ECI [Li et al. 2022] for the elasticity and inelasticity of thick shells, ARAP and discrete hinge [Grinspun et al. 2003] for the coupled membrane, and IPC [Li et al. 2020] for frictional contact. Note that some of our examples only contain thick shells. For the examples with complementary wrinkle coupling, since the membrane is closely attaching the thick shell, it is sufficient to only handle contact on the thick shells, which provides better efficiency. This is an option we expose to users. All our experiments are performed on a 24-core 3.50GHz Intel i9-10920X machine with an Nvidia RTX 3090 GPU. The parameters and timing of our simulations are provided in Table 2.

As follows, we start by validating our method on 2D cantilever beam tests (Fig. 6 and 7). We then evaluate different design choices of our method on discretizing volumetric objects (Fig. 11), and capturing fine wrinkling details (Fig. 12). We also compare our method with linear tetrahedral FEM and discrete shell on effectively capturing thickness deformations (§ 6.1). Then we demonstrate our method’s capability of realistically simulating thick garments (Fig. 1) and its potential to be used as wrinkle enhancement tools (§ 6.2).

6.1 Ablation and Comparative Studies

2D cantilever beam. We begin with validating the effectiveness of our reduced integration by comparing its results to those obtained through full integration (4 quadratures in 2D). We use the classic cantilever beam experiment, fixing the left end of a beam in 2D and letting it bend under gravity (Fig. 6). As a reference, we use high-resolution full integration and overlay it with our results. With our dual-quadrature scheme, our results quickly converge under refinement even with only a single layer of 10 elements (Fig. 6 top). However, we observe that full integration converges at a much higher resolution than ours (Fig. 6 bottom).

To further verify our method, we compare our results with theoretical solutions by superimposing the data from our simulation onto the master curve presented in Romero et al. [2021]. The master curve uniquely determines the aspect ratio (H/W , see Fig. 7 bottom right) of the cantilever beam under static equilibrium as a function of the the dimensionless gravito-bending parameter

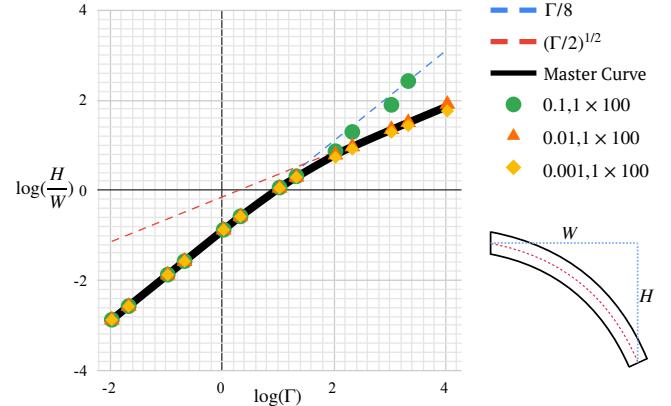


Figure 7: Comparison with theoretical solutions. We superimpose our results of the 2D cantilever beam simulation onto the master curve presented in Romero et al. [2021]. The blue and red dashed lines correspond to the asymptotic expressions for small and large Γ . Our results quickly converge to the master curve as the thickness of the beam decreases.

$\Gamma = \frac{12(1-v^2)\rho g L^3}{Eh^2}$ computed using the beam’s length (L), thickness (h), density (ρ), Young’s modulus (E), Poisson’s ratio (v), and gravitational acceleration (g). We conduct the experiment using three 1m-long beams, with thicknesses of 0.1m, 0.01m, and 0.001m, respectively. All beams are discretized with a 1 × 100 element grid. Using different Young’s modulus, we simulate with 13 Γ values per thickness. Our simulation results quickly converges to the master curve as thickness of the beam decreases (Fig. 7).

Yoga mat. In this example, we capture thickness deformation while rolling a yoga mat from one side on the floor and then tie it up with a buckle (Fig. 8). After becoming static, we can observe that the radius of the circular cross section grows from the middle to both sides, and the buckle in the middle creates sharp deformations on the mat (Fig. 8 bottom right). These are all thickness deformations that codimensional models are not able to capture. Here, accurate frictional contact provided by IPC is also crucial in simulating the rolling dynamics.

Metal board. Our method can also effectively capture thickness evolution under plasticity. In Fig. 9, we simulate a stretching metal board with a single layer of prism elements. The visualization of the



Figure 8: Yoga mat. A yoga mat rolled from one side on the floor and then tied up by a buckle.

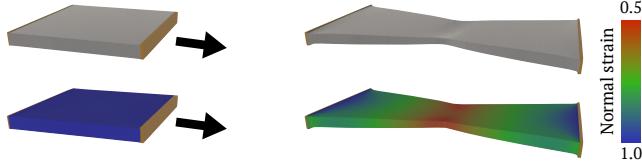


Figure 9: Metal board. Elongating a metal board, the smooth thickness decreases and plastic deformations are effectively captured by our method.

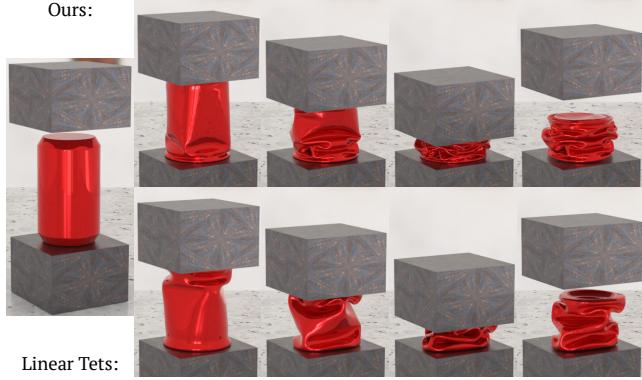


Figure 10: Coke can. Using real-world aluminium parameters, our method (top) can generate sharp wrinkles and buckling on the compressed coke can, while Li et al. [2022] (bottom) suffers from severe shear locking issues.

normal strain at each node demonstrates the smooth propagation of thickness changes resulted from plastic deformation.

Coke can. We use the coke can example from Li et al. [2022] to further demonstrate the shear locking free property of our method under plasticity. We setup the same scene in our code but using real-world aluminum parameters ($E = 70\text{GPa}$, yield stress= 0.3GPa) for the can. Our method generates compressed coke cans with sharp wrinkles and buckling while Li et al. [2022] results in much thicker behaviors due to shear locking (Fig. 10). (In their paper, $0.1\times$ real parameters are used to avoid shear locking.)

Hand over pillow. Here we show that our prism element is also capable of discretizing volumetric objects. We discretize the viscoelastic pillow from Li et al. [2022] using 1, 2, and 3 layers of prism elements, where we extrude a rounded square mesh with spatially varying thickness to accurately represent the curvy profile (Fig. 11). Pressing the pillow using a hand-shaped collision object, all our results are left with a hand print, which gradually disappears as the hand moves away (see our video). Here we setup the prism mesh to match the surface resolution of the input in Li et al. [2022]. With 2 or 3 layers, the magnitude of our print can already well-match their result, achieving a $4.7\times$ or $2.5\times$ faster run time as we have much less interior elements but no shear locking.

Wrinkle Mat. Next, we study the performance of our method on capturing fine wrinkling details with different number of subdivisions on the membrane. We drop 2 small rigid cubes onto a soft mat

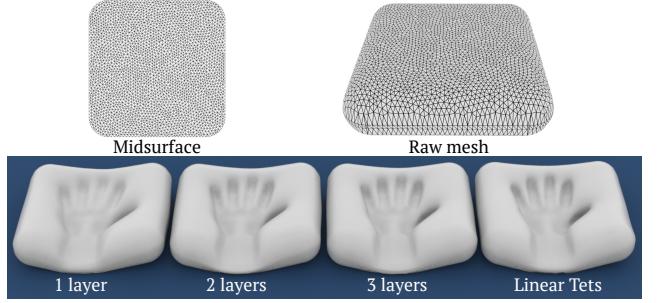


Figure 11: Hand over pillow. Discretizing the pillow using 1, 2, and 3 layers of prism elements, our method can more efficiently simulate the hand-printing viscoelastic effects demonstrated in Li et al. [2022]. The 3-layer shell can reproduce high-resolution linear tets results with much fewer DOF, while a single layer would produce a shallower handprint.

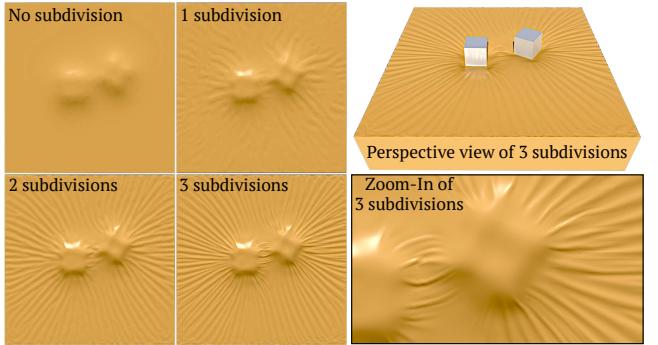


Figure 12: Wrinkle mat. 2 rigid cubes dropped onto a soft mat simulated as a single-layer thick shell coupled with a membrane on its top. With more subdivisions on the membrane in the setup, more significant and diverse wrinkles are captured.

simulated as a single-layer thick shell coupled with a membrane on its top (Fig. 12 top right). The thick shell only contains 2K prism elements, and we compare results by setting up the membrane with 0, 1, 2, and 3 subdivisions. For 0 subdivision, the membrane shares the same coarse mesh with the thick shell, and we can barely observe any wrinkles (Fig. 12 top left). For 1 subdivision, only disconnected wrinkles with small magnitude are generated (Fig. 12 top middle). For 2 or 3 subdivisions, we can observe smooth and significant wrinkles with different wavelengths in the same scene (Fig. 12 bottom).

We then compare directly solving Eq. 11 using projected Newton and Cholesky factorization (Joint) to our methods with CPU (') or GPU (*) version of the Conjugate Gradient (CG) method on the wrinkle mat example (see Table 1). On the CPU, our alternating minimization method with matrix-free CG can already achieve an up to $4\times$ speedup and faster convergence compared to the Joint method. Implementing our CG solver on the GPU, we gain another up to $5\times$ speedup especially when using high-resolution membrane mesh.

Table 1: Solver comparison. Statistics of Fig. 12 with different solvers. Ours'/Ours* refers to the CPU/GPU versions of our method respectively.

Example	Solver	Min./Frame	#Iter./Frame
1 subdiv	Joint/Ours'/Ours*	0.09/0.05/0.04	23.1/14.3/13.8
2 subdiv	Joint/Ours'/Ours*	1.17/0.26/0.13	32.6/13.6/14.1
3 subdiv	Joint/Ours'/Ours*	13.99/3.77/0.74	30.0/21.3/21.2

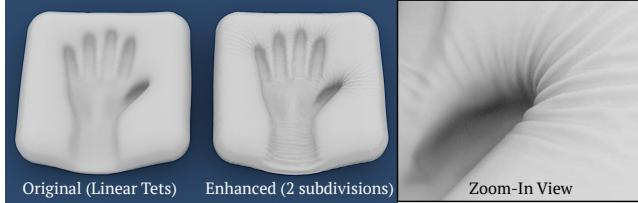


Figure 13: Animation processing. Given an input surface mesh sequence (left), our method can treat it as a low-frequency boundary condition and augment it with realistic fine wrinkles (middle and right).

6.2 Wrinkle Enhancement

Our method can also serve as a wrinkle enhancement tool for an existing animation sequence or 3D model.

Animation processing. Given an input animation sequence (surface meshes alone is sufficient), we can take the input meshes as the low-frequency degrees of freedom \mathbf{y} , and subdivide its rest configuration to construct the high-frequency membrane mesh \mathbf{X} and the basis $[\mathbf{P}, \mathbf{Q}]$. Then we run our method with only the high-frequency steps, setting the position and velocity of \mathbf{y} in each time step as boundary conditions. Taking the hand over pillow sequence generated using Li et al. [2022] as an example, as the stretch and compression on \mathbf{y} guide the low-frequency deformations of \mathbf{x} , we efficiently enrich the animation with fine wrinkling details (Fig. 13). Under this setting, our method is in a similar spirit to complementary dynamics [Benchekroun et al. 2023; Zhang et al. 2020].

Geometric modeling. Similarly, if we take a single 3D surface mesh, subdivide it to construct our fine membrane mesh and the basis, our method can also augment wrinkles on it, setting the input mesh as boundary conditions. Here, the key is to treat the initial state of the fine membrane elements as pre-compressed. In Fig. 14, the rest shape of each fine element is scaled by $1.5\times$ along the principal curvature direction computed using libigl [Jacobson et al. 2018]. After simulating for several time steps, interesting wrinkles tracing the surface profile are generated.

7 CONCLUSION AND FUTURE WORKS

We presented a novel method for simulating the intricate dynamics of multi-layer continuum thick shells. Discretized with our dual-quadrature prism finite elements, the thick shell model is free from

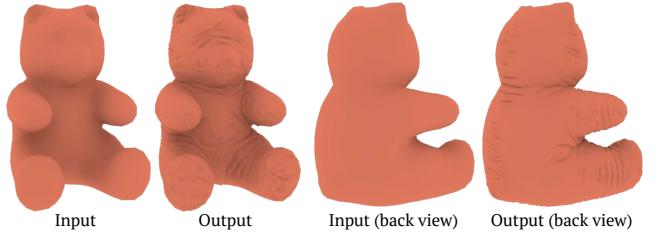


Figure 14: Geometric modeling. Given an input surface mesh, our method can augment it with realistic fine wrinkles by simulating pre-compressed fine membranes attached to it.

shear locking and naturally incorporates three-dimensional elasto-plastic and viscoelastic constitutive models. By coupling a high-resolution membrane layer on top of the thick shell, we enabled the efficient capture of complementary high-frequency wrinkles. In addition to opening up a new possibility for fast, high-quality, and thickness-aware simulations of a wide variety of thick structures, our method also has the potential to serve as a wrinkle enhancement tool for animation processing and geometric modeling.

Our method inspires numerous future works. For instance, one major limitation of our complementary wrinkle coupling is that the wrinkle profiles that can be captured may be mesh dependent as the mesh connectivity at the coarsest level directly determines the low-frequency basis under linear subdivision. This is also because we used all modes of the base mesh to constrain the low-frequency displacement of the high-resolution membrane. Thus, enabling user-oriented low-frequency mode design or selection to mitigate the mesh dependency issue will be very meaningful.

Additionally, we are interested in exploring an automatic way of generating multi-layer prism meshes from arbitrary input surface geometries. This will improve the accessibility of our method for simulating specially structured volumetric objects. Last but not least, our change-of-variable strategy based on defining sparse null-space basis vectors could lead to solver innovations for a variety of equality-constrained problems in physically-based animation and geometry processing.

ACKNOWLEDGMENTS

We sincerely thank the anonymous reviewers for their valuable feedback. This work is partially supported by NSF 2153851, 2153863, 2023780, 2301040, 2008915, 2244651, and 2008564.

REFERENCES

- Klaus-Jürgen Bathe and Eduardo N Dvorkin. 1985. A four-node plate bending element based on Mindlin/Reissner plate theory and a mixed interpolation. *Internat. J. Numer. Methods Engng.* 21, 2 (1985), 367–383.
- Ted Belytschko, Wing Kam Liu, Brian Moran, and Khalil Elkhodary. 2014. *Nonlinear finite elements for continua and structures*. John wiley & sons.
- Otman Benchekroun, Jiayi Eris Zhang, Siddhartha Chaudhuri, Eitan Grinspun, Yi Zhou, and Alec Jacobson. 2023. Fast Complementary Dynamics via Skinning Eigenmodes. *arXiv preprint arXiv:2303.11886* (2023).
- Miklós Bergou, Saurabh Mathur, Max Wardetzky, and Eitan Grinspun. 2007. Tracks: toward directable thin shells. *ACM Transactions on Graphics (TOG)* 26, 3 (2007), 50–es.
- Kai-Uwe Bletzinger, Manfred Bischoff, and Ekkehard Ramm. 2000. A unified approach for shear-locking-free triangular and rectangular shell finite elements. *Computers & Structures* 75, 3 (2000), 321–334.

Table 2: Simulation Statistics. All examples are simulated with $\Delta t = 0.01$. Here, thick shell parameters include E , ν , and ρ , while wrinkle surface parameters include E_{membrane} , ν_{membrane} , E_{bending} , ν_{bending} , and ρ_{membrane} .

Scene	N_Y	Prism Shell Param.	N_X	Wrinkle Surface Param.	T (Min./Frame)
Leather jacket & pants	18.8K	$(5 \times 10^5, 0.3, 860.0)$	43.0K	$(5 \times 10^5, 0.4, 10^5, 0.4, 472.6)$	5.842
Down jacket & jeans	18.8K	Down $(10^5, 0.3, 500)$	43.0K	$(10^6, 0.4, 10^4, 0.4, 472.6)$	26.666
/	/	Jeans $(2.4 \times 10^6, 0.3, 400.0)$	/	/	/
Yoga mat	29.0K	$(10^6, 0.4, 1000.0)$	/	/	0.052
Hand over pillow	9.3K/12.0K/14.7K	$(10^3, 0.3, 1.0)$	/	/	0.058/0.069/0.124
Coke can	42.2K	$(7 \times 10^{10}, 0.3, 2700.0)$	/	/	0.489
Wrinkle mat	2.0K	$(5 \times 10^5, 0.4, 472.6)$	3.9K/15.5K/61.4K	$(4 \times 10^8, 0.4, 0.0, 0.0, 1000.0)$	0.043/0.129/0.742
Animation processing	9.3K	/	85.9K	$(4 \times 10^6, 0.3, 5 \times 10^4, 0.4, 472.6)$	0.352
Geometric modeling	1.2K	/	19.1K	$(10^6, 0.4, 10^4, 0.4, 472.6)$	0.095 (2 frames)
Metal Board	25.1K	$(7 \times 10^8, 0.3, 2700.0)$	/	/	0.091

Rui PR Cardoso, Jeong Whan Yoon, Made Mahardika, S Choudhry, RJ Alves de Sousa, and RA Fontes Valente. 2008. Enhanced assumed strain (EAS) and assumed natural strain (ANS) methods for one-point quadrature solid-shell elements. *Internat. J. Numer. Methods Engrg.* 75, 2 (2008), 156–187.

Juan J. Casafranca and Miguel A. Otaduy. 2022. Voronoi Filters for Simulation Enrichment. *Computer Graphics Forum (Proc. SCA)* 41, 8 (2022).

Hsiao-Yu Chen, Arnav Sastry, Wim M van Rees, and Etienne Vouga. 2018. Physical simulation of environmentally induced thin shell deformation. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 1–13.

Yanqing Chen, Timothy A Davis, William W Hager, and Sivasankaran Rajamanickam. 2008. Algorithm 887: CHOLMOD, supernodal sparse Cholesky factorization and update/downdate. *ACM Transactions on Mathematical Software (TOMS)* 35, 3 (2008), 1–14.

Zhen Chen, Hsiao-Yu Chen, Danny M Kaufman, Mélina Skouras, and Etienne Vouga. 2021. Fine Wrinkling on Coarsely Meshed Thin Shells. *ACM Transactions on Graphics (TOG)* 40, 5 (2021), 1–32.

Zhili Chen, Renguo Feng, and Huamin Wang. 2013. Modeling friction and air effects between cloth and deformable bodies. *ACM Transactions on Graphics (TOG)* 32, 4 (2013), 1–8.

Eitan Grinspun, Anil N Hirani, Mathieu Desbrun, and Peter Schröder. 2003. Discrete shells. In *Proceedings of the 2003 ACM SIGGRAPH/Eurographics symposium on Computer animation*. Citeseer, 62–67.

Eitan Grinspun, Petr Krysl, and Peter Schröder. 2002. CHARMS: A simple framework for adaptive simulation. *ACM transactions on graphics (TOG)* 21, 3 (2002), 281–290.

Qi Guo, Xuchen Han, Chuyuan Fu, Theodore Gast, Rasmus Tamstorf, and Joseph Teran. 2018. A material point method for thin shells with frictional contact. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 1–15.

Matthias Harnau and Karl Schweizerhof. 2002. About linear and quadratic “solid-shell” elements at large deformations. *Computers & Structures* 80, 9–10 (2002), 805–817.

Ralf Hauptmann and Karl Schweizerhof. 1998. A systematic development of ‘solid-shell’ element formulations for linear and non-linear analyses employing only displacement degrees of freedom. *Internat. J. Numer. Methods Engrg.* 42, 1 (1998), 49–69.

Alec Jacobson, Daniele Panozzo, et al. 2018. libigl: A simple C++ geometry processing library. <https://libigl.github.io/>.

Ladislav Kavan, Dan Gerszewski, Adam W Bargteil, and Peter-Pike Sloan. 2011. Physics-inspired upsampling for cloth simulation in games. In *ACM SIGGRAPH 2011 papers*, 1–10.

Ki-Du Kim, GZ Liu, and Sung-Cheon Han. 2005. A resultant 8-node solid-shell element for geometrically nonlinear analysis. *Computational Mechanics* 35, 5 (2005), 315–331.

Jie Li, Gilles Daviet, Rahul Narain, Florence Bertails-Descoubes, Matthew Overby, George E Brown, and Laurence Boissieux. 2018. An implicit frictional contact solver for adaptive cloth simulation. *ACM Transactions on Graphics (TOG)* 37, 4 (2018), 1–15.

Minchen Li, Zachary Ferguson, Teseo Schneider, Timothy R Langlois, Denis Zorin, Daniele Panozzo, Chenfanfu Jiang, and Danny M Kaufman. 2020. Incremental potential contact: intersection-and inversion-free, large-deformation dynamics. *ACM Trans. Graph.* 39, 4 (2020), 49.

Minchen Li, Danny M. Kaufman, and Chenfanfu Jiang. 2021. Codimensional Incremental Potential Contact. *ACM Trans. Graph. (SIGGRAPH)* 40, 4, Article 170 (2021).

Pengbo Li and Paul G Kry. 2014. Multi-layer skin simulation with adaptive constraints. In *Proceedings of the Seventh International Conference on Motion in Games*. 171–176.

Xuan Li, Minchen Li, and Chenfanfu Jiang. 2022. Energetically consistent inelasticity for optimization time integration. *ACM Transactions on Graphics (TOG)* 41, 4 (2022), 1–16.

Christian Miehe. 1998. A theoretical and computational model for isotropic elastoplastic stress analysis in shells at large strains. *Computer Methods in Applied Mechanics and Engineering* 155, 3-4 (1998), 193–233.

Matthias Müller and Nuttапong Chentanez. 2010. Wrinkle Meshes. In *Symposium on Computer Animation*. Madrid, Spain, 85–91.

Rahul Narain, Armin Samii, and James F O’Brien. 2012. Adaptive anisotropic remeshing for cloth simulation. *ACM transactions on graphics (TOG)* 31, 6 (2012), 1–10.

Vladimir Panc. 1975. *Theories of elastic plates*. Vol. 2. Springer Science & Business Media.

Horst Parisch. 1995. A continuum-based shell theory for non-linear applications. *Internat. J. Numer. Methods Engrg.* 38, 11 (1995), 1855–1883.

Oliver Rémyillard and Paul G Kry. 2013. Embedded thin shells for wrinkle simulation. *ACM Transactions on Graphics (TOG)* 32, 4 (2013), 1–8.

Damien Rohmer, Tiberiu Popa, Marie-Paule Cani, Stefanie Hahmann, and Alla Sheffer. 2010. Animation wrinkling: augmenting coarse cloth simulations with realistic-looking wrinkles. *ACM Transactions on Graphics (TOG)* 29, 6 (2010), 1–8.

Victor Romero, Mickaël Ly, Abdullah Haroon Rasheed, Raphaël Charrodière, Arnaud Lazarus, Sébastien Neukirch, and Florence Bertails-Descoubes. 2021. Physical validation of simulators in Computer Graphics: A new framework dedicated to slender elastic structures and frictional contact. *ACM Transactions on Graphics (TOG)* 40, 4 (2021), 1–19.

Heinrich Schoop. 1986. Oberflächenorientierte Schalentheorien endlicher Verschiebungen. *Ingénieur-Archiv* 56, 6 (1986), 427–437.

Marco Schwarze and Stefanie Reese. 2009. A reduced integration solid-shell finite element based on the EAS and the ANS concept—geometrically linear problems. *Internat. J. Numer. Methods Engrg.* 80, 10 (2009), 1322–1355.

Mélina Skouras, Bernhard Thomaszewski, Peter Kauffmann, Akash Garg, Bernd Bickel, Eitan Grinspun, and Markus Gross. 2014. Designing inflatable structures. *ACM Transactions on Graphics (TOG)* 33, 4 (2014), 1–10.

K. Y. Sze and L. Q. Yao. 2000. A hybrid stress ANS solid-shell element and its generalization for smart structure modelling. Part I—solid-shell element formulation. *Internat. J. Numer. Methods Engrg.* 48, 4 (2000), 545–564.

Vuong-Dieu Trinh, Farid Abed-Meraim, and Alain Combescure. 2011. A new assumed strain solid-shell formulation “SHB6” for the six-node prismatic finite element. *Journal of mechanical science and technology* 25, 9 (2011), 2345–2364.

Huamin Wang, Florian Hecht, Ravi Ramamoorthi, and James F O’Brien. 2010. Example-based wrinkle synthesis for clothing animation. In *ACM SIGGRAPH 2010 papers*, 1–8.

Jiayi Eris Zhang, Seungbae Bang, David I.W. Levin, and Alec Jacobson. 2020. Complementary Dynamics. *ACM Transactions on Graphics* (2020).

Evgeny Zuenko and Matthias Harders. 2019. Wrinkles, folds, creases, buckles: Small-scale surface deformations as periodic functions on 3D meshes. *IEEE Transactions on Visualization and Computer Graphics* 26, 10 (2019), 3077–3088.

Javier S Zurdo, Juan P Brito, and Miguel A Otaduy. 2012. Animating wrinkles by example on non-skinned cloth. *IEEE Transactions on Visualization and Computer Graphics* 19, 1 (2012), 149–158.