DOI: 10.1111/cgf.14756
EUROGRAPHICS 2023 / K. Myszkowski and M. Nießner
(Guest Editors)

COMPUTER GRAPHICS *forum*
Volume 42 (2023), Number 2

An Optimization-based SPH Solver for Simulation of Hyperelastic Solids

Min Hyung Kee[1], Kiwon Um[2], HyunMo Kang[1] and JungHyun Han[1]

[1]Korea University, Korea
[2]LTCI, Telecom Paris, IP Paris, France

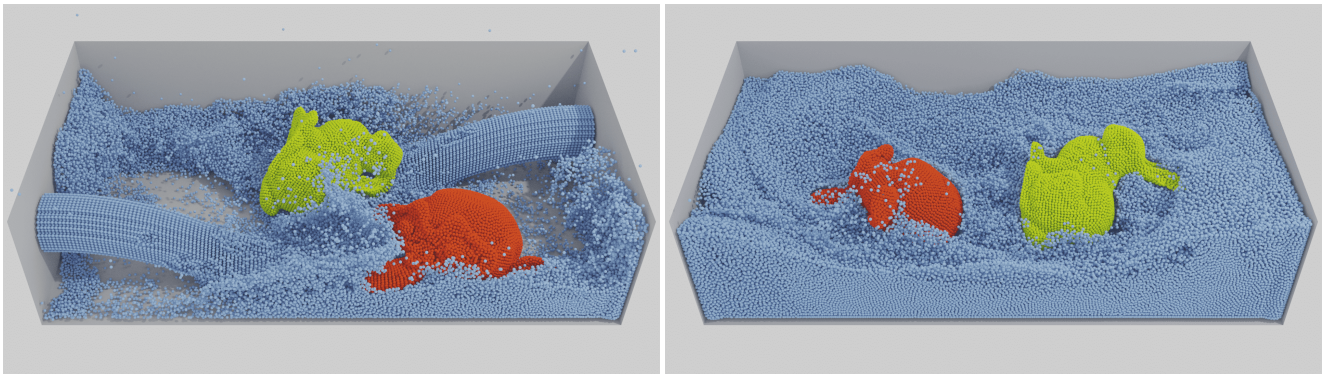

Figure 1: *Our optimization-based solver allows stable and robust simulations of coupling between hyperelastic solids and fluids in a unified SPH framework. The elastic solids of this example are simulated using the Neo-Hookean model with 39.3K particles, and the fluid is simulated using the divergence-free SPH solver with 320K particles.*

Abstract

This paper proposes a novel method for simulating hyperelastic solids with Smoothed Particle Hydrodynamics (SPH). The proposed method extends the coverage of the state-of-the-art elastic SPH solid method to include different types of hyperelastic materials, such as the Neo-Hookean and the St. Venant-Kirchoff models. To this end, we reformulate an implicit integration scheme for SPH elastic solids into an optimization problem and solve the problem using a general-purpose quasi-Newton method. Our experiments show that the Limited-memory BFGS (L-BFGS) algorithm can be employed to efficiently solve our optimization problem in the SPH framework and demonstrate its stable and efficient simulations for complex materials in the SPH framework. Thanks to the nature of our unified representation for both solids and fluids, the SPH formulation simplifies coupling between different materials and handling collisions.

CCS Concepts

• ***Computing methodologies*** → ***Physical simulation, elastic body simulation, optimization;***

1. Introduction

Simulation of elastic solids has been a popular research topic in computer graphics since the seminal work of Terzopoulos et al. [TPBF87]. Among the discretized representations of elastic solids, tetrahedral meshes have arguably been the most widely used, thanks to their efficiency and efficacy in evaluating the *deformation gradients* [SB12]. However, when the simulation setups come with different types of materials (e.g., fluids) and different discretizations for them (e.g., grids), a sophisticated coupling mechanism is often required to solve for interactions of the different materials stably.

To mitigate the challenge in coupling different simulation methods, particle-based unified frameworks have been studied; they include position-based dynamics (PBD) [MMCK14] and smoothed particle hydrodynamics (SPH) [Mon92]. A unified formulation for both solids and fluids not only simplifies the coupling mechanism but also facilitates their phase transitions (e.g., solid to fluid or vice versa). Furthermore, collisions can be easily handled by the nature

of its solving mechanism, which fundamentally aims to avoid overlapping of particles.

Unfortunately, evaluation of the deformation gradients for particles is not trivial. For example, a naive SPH evaluation for the deformation gradient leads to a misinterpretation of deformation due to the first-order inconsistency. In order to solve the problem, Bonet and Lock [BL99] introduced a corrected SPH kernel gradient formulation. Using the formulation, Peer et al. [PGBT18] proposed an implicit integration scheme for SPH elastic solids to further improve the simulation stability. However, their SPH formulation for elastic solids often suffers from instability in large deformation, due to the zero-energy modes [Gan15]. Kugelstadt et al. [KBF*21] tackled this problem by suppressing the zero-energy modes artificially. However, their elasticity solver is limited to corotated materials, which are not adequate for simulating a wide variety of real-world deformable materials such as rubber.

This paper proposes a novel SPH-based elastic solid simulation method that supports more complex hyperelastic models, such as Neo-Hookean and St. Venant-Kirchoff, with great stability. To this end, we reformulate the implicit integration as an optimization problem, as was similarly done in the mesh-based methods [MTGG11; BML*14; GSS*15]. Our experiments show that the L-BFGS solver, which is one of the most popular general-purpose quasi-Newton methods [NW99] and known to be efficient for hyperelastic solids in the mesh-based approach [LBK17], can be employed to solve our optimization problem efficiently in the context of SPH. To the best of our knowledge, our work is the first demonstration for stable simulations of hyperelastic models with the optimization-based SPH solver.

2. Related Work

The SPH framework has been one of the most popular choices for simulations of various physical phenomena [Mon92]. In computer graphics, this framework is particularly well known to be effective for fluid simulations due to the seminal work of Müller et al. [MKN*04]. (For an overview of SPH, readers are referred to the article by Koschier et al. [KBST22].) A vast amount of studies, then, has been introduced extending and improving various aspects of fluid dynamics. In addition, there have been efforts to handle other materials together with fluids. Particularly, robust interaction methods between SPH fluids and rigid bodies were proposed [AIA*12; ANZS18; GPB*19].

In order to handle deformable solids in SPH, Solenthaler et al. [SSP07] proposed a deformation gradient evaluation using SPH, particularly focusing on a linear elastic material. Since the naive evaluation of the deformation gradient with the SPH formulation is not first-order consistent [BL99], however, this formulation suffered from a misinterpretation of deformation regarding rotation. To address this problem, Becker et al. [BIT09] proposed a shape matching method and used an explicit time integration scheme with a corotated linear elasticity model, whereas Gerszewski et al. [GBB09] proposed deformation gradient based on moving least squares (MLS). On the other hand, Peer et al. [PGBT18] proposed an implicit formulation for simulating corotated SPH elastic solids in order to remedy simulation instability. However, their method often suffered from instability in large deformations due to the zero-energy modes [Gan15]. To improve the control of the zero-energy modes, Kugelstadt et al. [KBF*21] introduced a zero-energy mode penalty force that conserves linear and angular momentum. These methods, however, focused only on the corotated model whereas our method supports a wide variety of elasticity models such as Neo-Hookean and St. Venant-Kirchoff.

Additionally, instead of relying entirely on SPH, Abu Rumman et al. [ANM*20] proposed a method that couples PBD-based deformable solids with SPH fluids. Although this method achieved an efficient simulation of coupling between solids and fluids, at the same time, it shared the downside of the PBD framework, i.e., the method has difficulties in handling physics-based elastic models. On the other hand, to resolve such coupling, Dagenais et al. [DGP12] used a predictor-corrector approach with shape matching, and Huber et al. [HEW15] presented a simulation method for cloth with SPH, which supports wetting and no-slip boundary conditions. Gissler et al. [GHB*20] proposed a compressible SPH pressure solver, which is coupled with a linear implicit elasticity solver, such that it can simulate plasticity.

Since the seminal work of Baraff and Witkin [BW98] for cloth simulation, many implicit methods have been proposed to improve the accuracy and stability. Thanks to the stability, backward Euler (as an implicit scheme) has been widely adopted in mesh-based elastic material simulations, whereas it is relatively new to SPH-based elastic material solvers.

On the other hand, reformulation of a numerical integration problem as an optimization problem has been studied for better performance and stability [SH98; MTGG11; GSS*15]. Here, solving for their objective functions leads to the constrained motions of simulated objects. For example, PBD method reformulated the backward Euler scheme for elastic object simulations as a constrained optimization problem, where the elastic energy for each finite element is considered as a hard constraint and resolved iteratively [MHHR07]. Macklin et al. [MMC16] alleviated the problem of iteration-dependent stiffness of PBD by adjusting the Lagrange multiplier for each iteration.

Liu et al. [LBOK13] presented a method for efficient simulations of mass-spring systems introducing two-step optimization approach with auxiliary variables. This method was generalized to projective dynamics (PD) [BML*14], which achieves stable simulations of many kinds of deformable objects such as rods, clothes and jellies. Wang et al. [Wan15] accelerated PD and PBD using a Chebyshev semi-iterative approach. Liu et al. [LBK17] interpreted PD as a quasi-Newton method, whereas Overby et al. [OBLN17] showed that PD can be interpreted as an alternating direction method of multipliers (ADMM). These two methods allowed supporting more general materials. Our work is inspired by these optimization-based approaches while focusing on the SPH framework for a similar goal, i.e., efficient and stable simulations of general elastic materials.

3. SPH Formulation

Our method is built upon two state-of-the-art SPH formulations: robust evaluation of deformation gradient [PGBT18] and efficient

handling of zero-energy modes via suppression [KBF*21]. They are briefly reviewed in Sec. 3.1 and Sec. 3.2, respectively.

3.1. Deformation Gradient

Elastic energy models (such as corotated, Neo-Hookean and St. Venant-Kirchoff) are widely used for elastic solids. In general, they are expressed in terms of deformation gradient, $\mathbf{F} = \partial \mathbf{x}/\partial \mathbf{X}$ ($\in \mathbb{R}^{3\times 3}$), where $\mathbf{x}$ ($\in \mathbb{R}^3$) and $\mathbf{X}$ ($\in \mathbb{R}^3$) denote deformed and reference (i.e., undeformed) positions, respectively.

A robust evaluation of the deformation gradient for SPH is proposed by Bonet and Lok [BL99]:

$$\mathbf{F}_i = \sum_{j\in\mathcal{N}_i^0} V_j \mathbf{x}_{ji} \otimes \left(\mathbf{L}_i \nabla W_{ji}\right) \tag{1}$$

The subscript i denotes each particle's index, $\mathcal{N}_i^0$ is the set of particle i's initial neighbours, V_j is the rest-pose volume, $\mathbf{x}_{ji} = \mathbf{x}_j - \mathbf{x}_i$, $\otimes$ is the Kronecker product operator (i.e., $\mathbf{a}\otimes\mathbf{b} = \mathbf{a}\mathbf{b}^T$), $\mathbf{L}_i$ is the kernel correction matrix, and $W_{ji} = W(\mathbf{X}_{ji}, r)$, where $\mathbf{X}_{ji} = \mathbf{X}_j - \mathbf{X}_i$ and $W(\mathbf{X}, r)$: $\mathbb{R}^3 \to \mathbb{R}$ is the SPH kernel with the kernel radius r. The kernel correction matrix, $\mathbf{L}_i$, is defined as follows:

$$\mathbf{L}_i = \left(\sum_{j\in\mathcal{N}_i^0} V_j \nabla W_{ji} \otimes \mathbf{X}_{ji} \right)^{-1} \tag{2}$$

This correction makes the deformation gradient, $\mathbf{F}_i$, satisfy the first-order consistency condition:

$$\sum_{j\in\mathcal{N}_i^0} V_j \mathbf{X}_{ji} \otimes \left(\mathbf{L}_i \nabla W_{ji}\right) = \mathbf{I} \tag{3}$$

which is important to capture rotational motions correctly.

The hyperelastic energy, denoted as $E_{\text{he}}(\mathbf{x})$, is defined using the deformation gradients, $\mathbf{F}_i$, given in Eq. (1):

$$E_{\text{he}}(\mathbf{x}) = \sum_i V_i \Psi(\mathbf{F}_i) \tag{4}$$

where Ψ represents the elastic energy density function, which may vary according to the type of the elastic material to model.

3.2. Zero-Energy Mode Suppression

A well-known problem of the deformation gradient approximation given in Eq. (1) is the zero-energy modes, which are similar to the hourglass modes in FEM simulations [Gan15]. To alleviate this problem, Kugelstadt et al. [KBF*21] proposed implicit penalty forces that suppress the zero-energy modes using an artificially designed energy:

$$E_{\text{ze}}(\mathbf{x}) = \frac{\alpha}{2} \sum_i w_i V_i \left(\sum_{j\in\mathcal{N}_i^0} W_{ji} V_j \frac{\|\mathbf{e}_{ji}\|^2}{\|\mathbf{X}_{ji}\|^2} \right) \tag{5}$$

where α is the user-defined parameter, w_i is the material parameter that makes the penalty forces proportional to the elastic forces, and $\mathbf{e}_{ji}$ defined as $\mathbf{F}_i\mathbf{X}_{ji} - \mathbf{x}_{ji}$ represents the error caused by the zero-energy modes.

Eq. (5) can be represented in a matrix-vector product form:

$$E_{\text{ze}}(\mathbf{x}) = \frac{1}{2}\mathbf{x}^T \mathbf{H}\mathbf{x} \tag{6}$$

We refer the readers to the work of Kugelstadt et al. [KBF*21] for the details on formulation of the matrix, $\mathbf{H}$.

4. Optimization Problem

Given a physical system that consists of m particles, its state can be described as a set of the particle positions $\mathbf{x}(\in \mathbb{R}^{3m})$ and velocities $\mathbf{v}(\in \mathbb{R}^{3m})$. The system evolves at discrete time samples, $\{t_1, t_2, \cdots, t_N\}$. Given a state at t_n, the implicit Euler scheme approximates the next state at t_{n+1} as follows:

$$\begin{aligned} \mathbf{x}_{n+1} &= \mathbf{x}_n + h\mathbf{v}_{n+1} \\ \mathbf{v}_{n+1} &= \mathbf{v}_n + h\mathbf{M}^{-1}\left(\mathbf{f}_{\text{ext}} + \mathbf{f}_{\text{int}}(\mathbf{x}_{n+1})\right) \end{aligned} \tag{7}$$

where h is the time step size, $\mathbf{M}$ is the mass matrix, $\mathbf{f}_{\text{ext}}$ represents the external forces (e.g., gravity), and $\mathbf{f}_{\text{int}}(\mathbf{x})$ represents the internal forces, which include the elastic forces and the zero-energy suppression forces in the SPH framework.

Since the internal forces can be evaluated as the negative gradients of internal energy, i.e., $\mathbf{f}_{\text{int}}(\mathbf{x}) = -(\nabla E_{\text{he}}(\mathbf{x}) + \nabla E_{\text{ze}}(\mathbf{x}))$, we can rewrite Eq. (7) as follows:

$$\frac{\mathbf{M}}{h^2}(\mathbf{x}_{n+1} - \mathbf{y}_n) + \nabla E_{\text{he}}(\mathbf{x}_{n+1}) + \nabla E_{\text{ze}}(\mathbf{x}_{n+1}) = \mathbf{0} \tag{8}$$

where $\mathbf{y}_n = \mathbf{x}_n + h\mathbf{v}_n + h^2\mathbf{M}^{-1}\mathbf{f}_{\text{ext}}$. Solving for the new state, $\mathbf{x}_{n+1}$, in Eq. (8) is reformulated to an optimization problem finding $\mathbf{x}_{n+1}$ that minimizes the following objective function $g(\mathbf{x}) : \mathbb{R}^{3m} \to \mathbb{R}$:

$$g(\mathbf{x}_{n+1}) = \frac{1}{2h^2}\|\mathbf{x}_{n+1} - \mathbf{y}_n\|_{\mathbf{M}}^2 + E_{\text{he}}(\mathbf{x}_{n+1}) + E_{\text{ze}}(\mathbf{x}_{n+1}) \tag{9}$$

where $\|\cdot\|_{\mathbf{M}}^2$ denotes weighted Frobenius norm. For the sake of notation simplicity, we will henceforth use $\mathbf{x}$ instead of $\mathbf{x}_{n+1}$.

4.1. Numerical Solution

Our optimization problem with the objective given in Eq. (9) can be solved using an iterative approach such as Newton's method:

$$\mathbf{x}_{k+1} = \mathbf{x}_k - \left(\nabla^2 g(\mathbf{x}_k)\right)^{-1} \nabla g(\mathbf{x}_k) \tag{10}$$

where the subscript denotes the iteration count and $\mathbf{x}_0 = \mathbf{y}_n$. Despite the fast convergence nature of Newton's method, it requires the Hessian matrix, $\nabla^2 g(\mathbf{x}_k)$, to be computed and factorized per solver iteration, thus degrading the efficiency.

Inspired by the work of Liu et al. [LBK17], which demonstrated the efficiency of a quasi-Newton method in a similar optimization problem for mesh-based deformable solid simulations, we adopt the L-BFGS algorithm [NW99]. For the sake of efficiency, it approximates the Hessian matrix, $\nabla^2 g(\mathbf{x}_k)$, using the curvature information of l most recent iterates, i.e., $\mathbf{x}_{k-l}, \cdots, \mathbf{x}_{k-1}$. In the following, we will show an effective approximation of this matrix for our method and demonstrate its efficiency for our optimization-based elastic SPH solid solver in the experiments.

Algorithm 1: An iteration of L-BFGS

$\mathbf{q} \coloneqq -\nabla g(\mathbf{x}_k)$
for $i = k-1, \cdots, k-l$ **do**
  $\mathbf{s}_i \coloneqq \mathbf{x}_{i+1} - \mathbf{x}_i$
  $\mathbf{t}_i \coloneqq \nabla g(\mathbf{x}_{i+1}) - \nabla g(\mathbf{x}_i)$
  $\rho_i \coloneqq tr(\mathbf{t}_i^T \mathbf{s}_i)$
  $\varsigma_i = tr(\mathbf{s}_i^T \mathbf{q})/\rho_i$
  $\mathbf{q} \coloneqq \mathbf{q} - \varsigma_i \mathbf{t}_i$
end
$\mathbf{r} \coloneqq \mathbf{A}_0^{-1}\mathbf{q}$ // initial approximation of Hessian
for $i = k-m, \cdots, k-1$ **do**
  $\eta \coloneqq tr(\mathbf{t}_i^T \mathbf{r})/\rho_i$
  $\mathbf{r} \coloneqq \mathbf{r} + \mathbf{s}_i(\varsigma_i - \eta)$
end
$\mathbf{d}(\mathbf{x}_k) \coloneqq -\mathbf{r}$ // descent direction
$\mathbf{x}_{k+1} \coloneqq \mathbf{x}_k + \mathbf{d}(\mathbf{x}_k)$

4.2. Initial Approximation of Hessian

The initial approximation of the Hessian matrix, which we denote as $\mathbf{A}_0$, significantly affects the convergence rate of L-BFGS, as discussed in [LBK17]. Therefore, we search for $\mathbf{A}_0$ that is close to the Hessian, $\nabla^2 g(\mathbf{x}_k)$, and remains *constant* such that the expensive Cholesky factorization can be avoided during the simulation. $\nabla^2 g(\mathbf{x}_k) = \mathbf{M}/h^2 + \mathbf{H} + \nabla^2 E_{\text{he}}(\mathbf{x}_k)$, and $\nabla^2 E_{\text{he}}(\mathbf{x}_k)$ is the only non-constant term. Hence, we can achieve our goal if we replace the non-constant term, $\nabla^2 E_{\text{he}}(\mathbf{x})$, with a "constant approximation."

To this end, we adopt the strategy taken by Bouaziz et al. [BML*14]. First, we approximate the hyperelastic energy per particle, $E_{\text{he},i}(\mathbf{x})$, as follows:

$$E_{\text{he},i}(\mathbf{x}) \approx V_i \frac{w_i}{2} \|\text{vec}(\mathbf{F}_i) - \mathbf{z}_i\|^2 \tag{11}$$

where w_i consists of material parameters such as Poisson's ratio and Young's modulus, vec$(\cdot)$ vectorizes a 3×3 matrix into a 9×1 vector, and $\mathbf{z}_i$ is the projection of vec$(\mathbf{F}_i)$ onto the manifold satisfying $E_{\text{he},i}(\mathbf{x}) = 0$. It is worth noting that vec$(\mathbf{F}_i)$ can be represented as a matrix-vector multiplication form [KBF*21]:

$$\text{vec}(\mathbf{F}_i) = \mathbf{D}_i \mathbf{x} \tag{12}$$

We refer the readers to the work of [KBF*21] for the details on formulation of the matrix $\mathbf{D}_i$. Aggregating the second derivative of Eq. (11), the desired "constant approximation," which we denote as $\mathbf{B}$, is defined as follows:

$$\mathbf{B} = \sum_i V_i w_i \mathbf{D}_i^T \mathbf{D}_i \tag{13}$$

Here, w_i is defined in the same way as done by Liu et al. [LBK17]. Finally, the initial approximation of the Hessian matrix, $\mathbf{A}_0$, is completed with $\mathbf{B}$:

$$\mathbf{A}_0 = \frac{\mathbf{M}}{h^2} + \mathbf{H} + \mathbf{B} \tag{14}$$

Algorithm 1 presents the L-BFGS algorithm based on $\mathbf{A}_0$.

Algorithm 2: Coupled solid-fluid solver

perform neighborhood search
compute external forces $\mathbf{f}_{\text{ext}}$
$\mathbf{y}_n \coloneqq \mathbf{x}_n + h_{\text{el}}\mathbf{v}_n + h_{\text{el}}^2 \mathbf{M}^{-1}\mathbf{f}_{\text{ext}}$
solve Eq. (9) for $\mathbf{x}_*$ with L-BFGS (using **Algorithm** 1)
$\mathbf{v} \coloneqq (\mathbf{x}_* - \mathbf{x}_n)/h_{\text{el}}$
$\mathbf{a}_{\text{el}} \coloneqq (\mathbf{v} - \mathbf{v}_n)/h_{\text{el}}$
$\mathbf{v}_{n+1} \coloneqq \mathbf{v}_n + h_{\text{fl}}\mathbf{a}_{\text{el}}$
pressure solve using DFSPH

5. Solid-Fluid Coupling and Collision

Our method is implemented with the divergence-free SPH (DFSPH) solver [BK16]. Coupling between solid and fluid particles and self-collisions are naturally handled in the pressure solver by which the particles are kept from interpenetration. See **Algorithm** 2.

Adaptive time steps based on the Courant-Friedrichs-Lewy (CFL) condition are commonly used in SPH. However, they drastically degrade the performance of our elasticity solver since changing the time steps requires changing the initial approximation of Hessian, $\mathbf{A}_0$. To tackle this performance issue, we use different time steps for solids and fluids, as in the works of Peer et al. [PGBT18] and Kugelstadt et al. [KBF*21]. While the time step size of our elasticity solver, h_{el}, is kept constant, that for the fluid solver, h_{fl}, is determined using the CFL condition as follows:

$$h_{\text{fl}} = 0.4d/\|\mathbf{v}\|_\infty \tag{15}$$

where d is the particle diameter. We evaluate the acceleration, $\mathbf{a}_{\text{el}}$, using the velocity change calculated in the elasticity solver and then evaluate the final velocity, i.e., $\mathbf{v}_{n+1} = \mathbf{v}_n + h_{\text{fl}}\mathbf{a}_{\text{el}}$.

6. Experimental Results

All our experiments are made on an *AMD Ryzen 7 5800X 3.8 GHz* processor with *OpenMP*-based parallelization. We run our solver with various hyperelastic materials including corotated, Neo-Hookean and St. Venant-Kirchoff models as well as for interactions of multiple hyperelastic solids and their two-way coupling with fluids. To this end, we employ the open-source SPH-based simulation framework *SPlisHSPlasH* [BKWK22] and the linear algebra library *Eigen* [GJ*10].

We will first present the results of two experiments, named **Swinging cloth** and **Stretching cube**, which demonstrate the stability and robustness of our optimization-based solver. For both experiments, corotated materials are used. The results will be compared with those of a state-of-the-art SPH-based elasticity solver, called operator-splitting solver [KBF*21], which we abbreviate to OSS. Table 1 lists the scene parameters and performance data for each experiment.

Swinging cloth: In this experiment, we set up an elastic cloth, two corners of which are attached to fixed points. The cloth is swinging under the gravitational force. As shown in Fig. 2-(a), OSS suffers from oscillation artifacts, particularly around the

(a) *OSS*, $\alpha = 0.1$

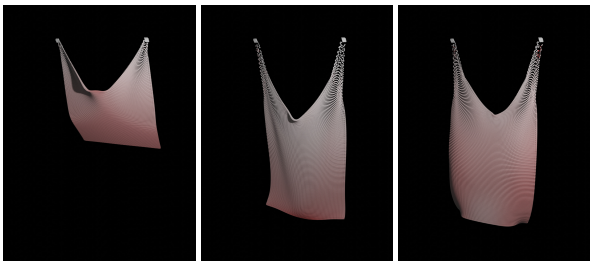

(b) *OSS*, $\alpha = 10$

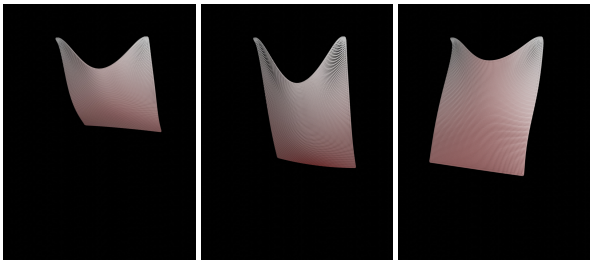

(c) *OSS*, $\alpha = 1000$

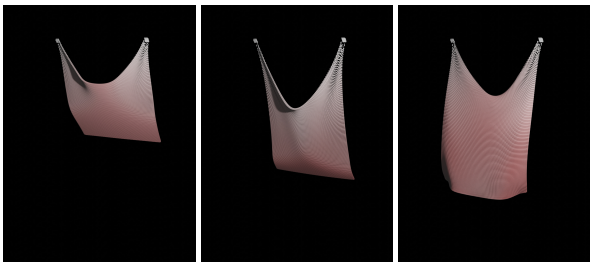

(d) *Ours*, $\alpha = 0.1$

Figure 2: ***Swinging cloth****: Increasing the zero-energy suppression in OSS alleviates the oscillating motions at the cost of introducing undesired artificial stiffness. In contrast, our solver successfully yields a stable and desired simulation for large deformations, with neither oscillation artifacts nor artificial stiffness. The particles are visualized in different colors indicating high speed in red and low speed in white.*

attached points where large deformations occur. The operator-splitting scheme jointly solves the stretching term of the corotated energy and the zero-energy term using a pre-factorized Cholesky solver, and then solves the volume-conservation term of the corotated energy using a Conjugate Gradient solver (with termination criteria, $\varepsilon < 10^{-4}$). In this alternating scheme, the second step deteriorates the first step's solution, resulting in the oscillation artifacts.

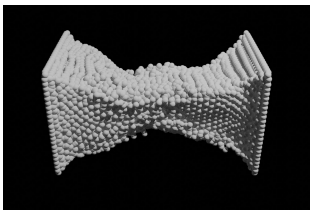

(a) *Ours*, $h = 0.03$, *1 iteration*

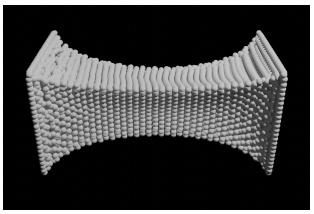

(b) *Ours*, $h = 0.03$, *2 iterations*

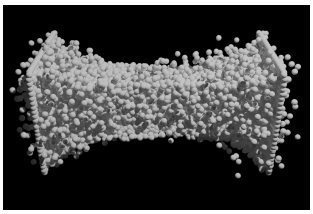

(c) *OSS*, $h = 0.03$

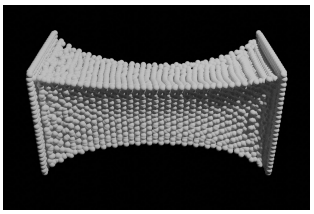

(d) *OSS*, $h = 0.003$

Figure 3: ***Stretching cube****: Our solver presents a stable simulation with only two iterations and a large time step, whereas OSS requires ten times smaller ti me step to get stabilized.*

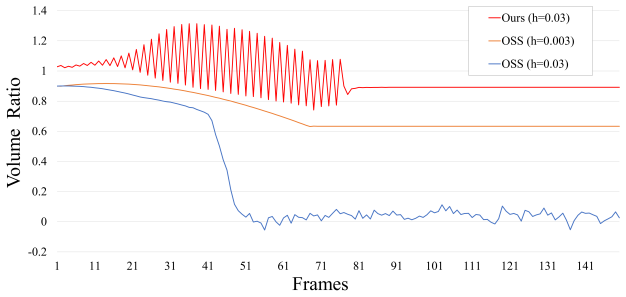

Figure 4: *The ratio of the deformed volume to the rest volume over frames of the* ***Stretching cube*** *example shown in Fig. 3. Ours took a single L-BFGS iteration up to the 80$^{\text{th}}$ frame and then two iterations afterward.*

To remedy this problem, one would consider increasing the zero-energy control parameter, α in Eq. (5), such that the zero-energy suppression becomes dominant. Unfortunately, as shown in Fig. 2-(b) and (c), strong zero-energy suppression results in undesired stiff motions.

In contrast, thanks to the optimization-based scheme, our solver successively simulates large deformations without increasing α, as shown in Fig. 2-(d). It does not suffer either from oscillation artifacts or from artificial stiffness.

Stretching cube: This example presents simulations of a nearly-incompressible elastic cuboid solid, whose Poisson's ratio is set to 0.45. We pull the two opposing sides of the solid and test its deformation. In order to see how our solver's iteration count affects the simulation, we first run our solver with a single iteration and then increase it to two on the fly. Even with a single iteration, our solver presents a rough deformation close to the expected (yet with oscillations), as shown in Fig. 3-(a). When the iteration count is

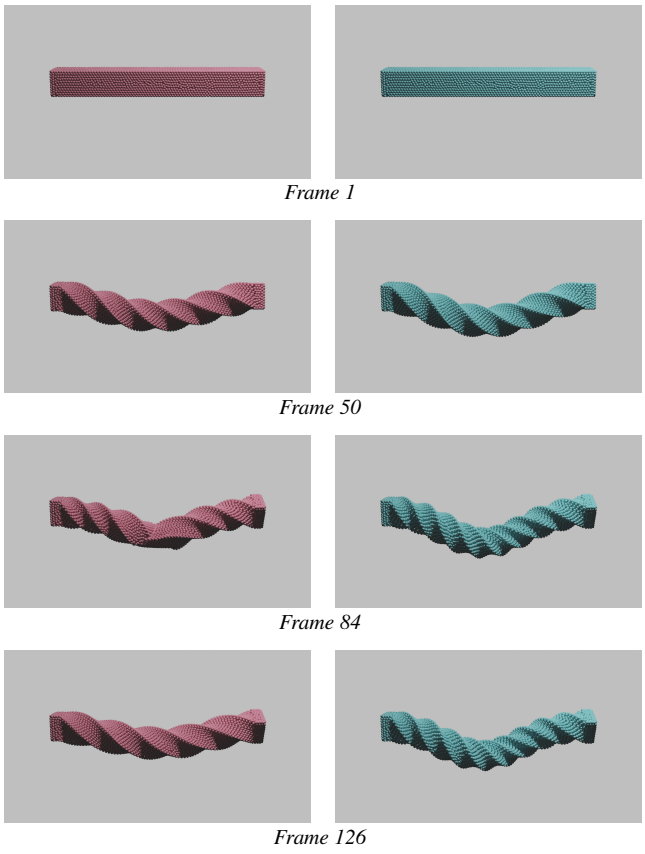

Figure 5: ***Twisting beam****: Our solver tested with different elasticity models: (left) corotated model, (right) Neo-Hookean model.*

increased just to two on the fly, our solver quickly produces a stable and desired deformation, as shown in Fig. 3-(b).

We test OSS with the same setup. With the same time step, i.e., $h = 0.03$s, OSS fails to generate a desired deformation, as shown in Fig. 3-(c). The simulation is stabilized with the ten times smaller time step, i.e., $h = 0.003$s, as shown in Fig. 3-(d); however, it takes approximately five times longer than ours.

The graphs in Fig. 4 show the ratio of the deformed volume to the rest volume, which is evaluated from the simulation results presented in Fig. 3. In our method, the oscillations are made when using a single iteration of our solver, but they immediately disappear when the iteration count is increased to two at the 80^{th} frame, and the ratio value converges to 0.89. It is significantly better than the others. Indicating that the volume ratio of one means incompressible deformation, we note that OSS fails to preserve the desired incompressbility even when using a small time step size (i.e., $h = 0.003$s). The ratio converges to 0.63.

Twisting beam: In this experiment, we set up an elastic beam that is twisted at its two ends, thus gradually causing a large deformation. See Fig. 5. As discussed in Sec. 1, corotated materials are not adequate for large compression. In Fig. 5, the left column shows the results of our simulation made with the corotated model.

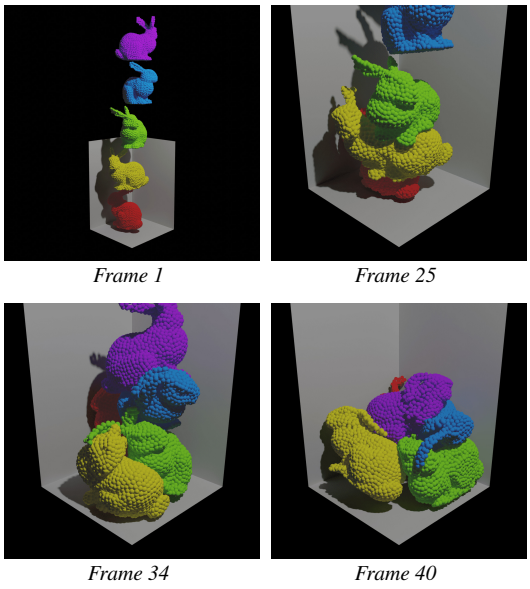

Figure 6: ***Colliding bunnies****: Collision between St. Venant-Kirchoff materials.*

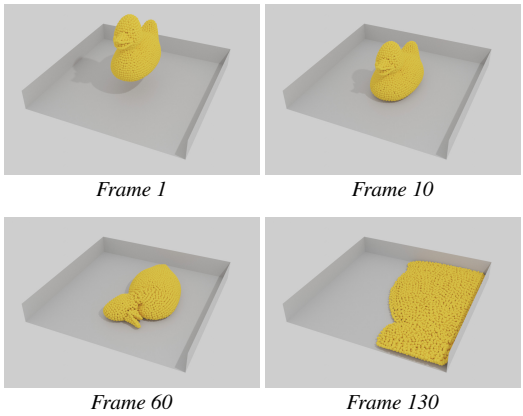

Figure 7: ***Melting duck****: A deformable duck is melted down.*

Observe the undesired inversion (at Frame 84), which indicates a failure of resolving a large compression induced by the series of twists. With the Neo-Hookean model, however, our solver produces the desired twisting motions stably, as shown on the right column in Fig. 5.

Colliding bunnies: In order to test the efficacy of collision handling in the SPH framework, this experiment is made with five elastic bunnies falling into a rigid box, as shown in Fig. 6. We use the St. Venant-Kirchoff model for the purpose of proving our solver's robustness in handling collision between complex hyperelastic materials.

Melting duck: Fig. 7 shows a deformable duck dropped onto a hot plate. This experiment presents a simple yet efficient handling of melting effects in our framework. As the heat induced from the

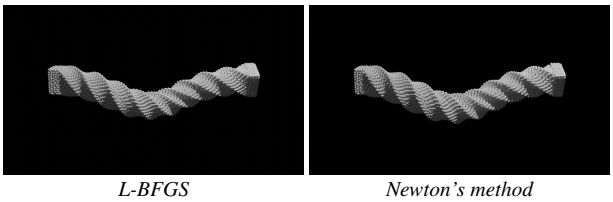

L-BFGS *Newton's method*

(a)

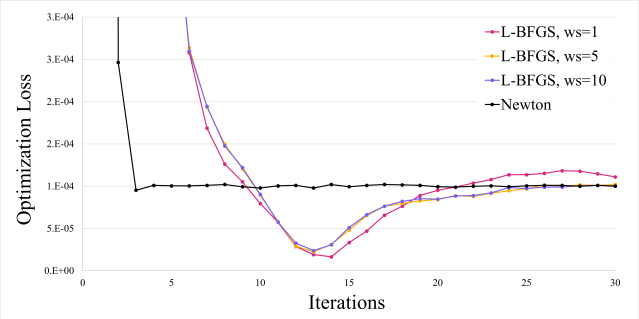

(b)

Figure 8: *Comparison between L-BFGS and Newton's method. (a) Simulation results of* ***Twisting beam****. (b) Optimization losses of L-BFGS with different window sizes (ws) and Newton's method over solver iterations.*

hot plate is diffused through the particles, we decrease the magnitude of the elastic force according to the temperatures of particles. To this end, a temperature-dependent weight function applies to each particle's energy gradient. This avoids update of the constant Hessian approximation, $\mathbf{A}_0$, while effectively taking into account the change of material property, i.e., elasticity.

Solid-fluid coupling: In Fig. 1, two Neo-Hookean elastic solids (bunny and Suzzane) are dropped into a box, and then fluids are emitted into the box from two nozzles. Our solver, jointly working with the DFSPH fluid solver, successfully presents the interactions of multiple hyperelastic solids and their two-way coupling with the fluids.

In order to see the reliability of L-BFGS, we compare the simulation results of L-BFGS with those of Newton's method. As shown in Fig. 8-(a), both methods present visually similar results, while L-BFGS is two orders of magnitude faster than Newton's method. Fig. 8-(b) shows optimization losses, $g(\mathbf{x})$, of L-BFGS with different window sizes and of Newton's method over solver iterations. These losses are evaluated at a selected frame of **Twisting beam** example. As shown in the graphs, Newton's method converges very fast (i.e., within three iterations) whereas convergence of L-BFGS depends on the window size. As it increases from one to five, the optimization loss converges faster. When it is increased to a larger value (i.e., ten), however, little difference is observed since the larger window size makes L-BFGS consider the more distant iterates' information. In our experiments, we empirically chose five for the window size.

7. Discussions

As an alternative approach for efficient simulation of elastic solids, one can consider solving Eq. (8) using the operator splitting scheme proposed by Kugelstadt et al. [**?**], which takes advantage of the corotated model thus making the system matrix for the stretching and zero-energy terms constant. Unfortunately, this scheme has to sacrifice its efficiency when considering more general hyperelastic materials, e.g., Neo-Hookean and St. Venant-Kirchoff models.

In the optimization perspective, a solution of the operator splitting scheme can be seen as an approximate solution yielded after a single iteration of Newton's method. As a consequence, the solution may end up with undesired oscillation motions thus suffering from instability. In contrast, our optimization-based solver tackles the aforementioned issues while aiming at a non-splitting scheme and converged solutions, thus enabling a wider variety of materials to be simulated more stably, as demonstrated in our experiments.

8. Conclusion and Future work

In this work, we presented an extension of the implicit integration scheme for SPH elastic solid simulations into an optimization problem. This extension supports a wide variety of hyperelastic materials such as the Neo-Hookean and St. Venant-Kirchoff models, which have not been well supported in the state-of-the-art SPH-based frameworks. Thanks to the efficiency of L-BFGS used in our optimization solver, our method presents fast and stable simulations of elastic solids from which we foresee potential applications of our solver in many real-time simulations such as surgical simulators and video games. Demonstrating our successful adaptation of the optimization formulation, which was originally presented for the mesh-based framework, into the SPH framework, we believe that our work sheds light on the great potential of adapting mesh-based methods for elastic solids into the SPH framework.

Although our solver enhances the simulation stability as demonstrated in our experiments, the proposed method also shares drawbacks of the backward Euler time integration scheme. One well-known issue is the artificial damping of the elastic energy. As discussed in the previous studies [DLK18; KUJH21], investigating energy-momentum conserving schemes such as constrained optimization, Implicit-Midpoint, and Newmark-beta remains as a future work.

In our optimization-based solver, one central factor of efficiency comes from the constant Hessian approximation, $\mathbf{A}_0$, given in Eq. (14), which avoids repeated evaluations of the time consuming Cholesky factorization. Nonetheless, as discussed in previous studies [BML*14; LBK17; KBF*21], such a direct solver on which our method also relies is turning into a problem when the number of particles per solid is growing very large. Moreover, if the reference configuration of a solid object has to be updated in order to simulate challenging effects such as plasticity, our method sacrifices the efficiency of constant matrix in the solver. In order to facilitate more efficient simulations, extending our method with the Galerkin multigrid method [XTL19], dimension reduction technique [BEH18], or progressive Cholesky update method [LLKC21] will be an interesting future work.

Additionally, we plan to investigate a strong coupling of elastic solids and fluids in SPH. One potential approach is to fuse an optimization-based pressure solver such as Projective Fluid [WKB16] with our method such that the pressure and the elasticity can be solved simultaneously. Like the work of Gissler et al. [**?**], which solves for strong fluid-rigid coupling via the interlinked SPH pressure solver, combining the pressure solver with our elasticity solver will also be an interesting future research direction.

Acknowledgments

This research was supported by the Ministry of Science and ICT, Korea, under the ICT Creative Consilience Program (IITP-2023-2020-0-01819), ITRC (Information Technology Research Center) Support Program (IITP-2023-2020-0-01460) and the grant No.2020-0-00861.

References

[AIA*12] AKINCI, NADIR, IHMSEN, MARKUS, AKINCI, GIZEM, et al. "Versatile rigid-fluid coupling for incompressible SPH". *ACM Transactions on Graphics (TOG)* 31.4 (2012), 1–8 2.

[ANM*20] ABU RUMMAN, NADINE, NAIR, PRAPANCH, MÜLLER, PATRIC, et al. "ISPH–PBD: coupled simulation of incompressible fluids and deformable bodies". *The Visual Computer* 36.5 (2020), 893–910 2.

[ANZS18] AKBAY, MUZAFFER, NOBLES, NICHOLAS, ZORDAN, VICTOR, and SHINAR, TAMAR. "An extended partitioned method for conservative solid-fluid coupling". *ACM Transactions on Graphics (TOG)* 37.4 (2018), 1–12 2.

[BEH18] BRANDT, CHRISTOPHER, EISEMANN, ELMAR, and HILDEBRANDT, KLAUS. "Hyper-reduced projective dynamics". *ACM Transactions on Graphics (TOG)* 37.4 (2018), 1–13 7.

[BIT09] BECKER, MARKUS, IHMSEN, MARKUS, and TESCHNER, MATTHIAS. "Corotated SPH for Deformable Solids." *NPH*. 2009, 27–34 2.

[BK16] BENDER, JAN and KOSCHIER, DAN. "Divergence-free SPH for incompressible and viscous fluids". *IEEE Transactions on Visualization and Computer Graphics* 23.3 (2016), 1193–1206 4.

[BKWK22] BENDER, JAN, KUGELSTADT, TASSILO, WEILER, MARCEL, and KOSCHIER, DAN. *SPlisHSPlasH Library*. https://github.com/InteractiveComputerGraphics/SPlisHSPlasH. (Accessed December 2021). 2022 4.

[BL99] BONET, JAVIER and LOK, T-SL. "Variational and momentum preservation aspects of smooth particle hydrodynamic formulations". *Computer Methods in applied mechanics and engineering* 180.1-2 (1999), 97–115 2, 3.

[BML*14] BOUAZIZ, SOFIEN, MARTIN, SEBASTIAN, LIU, TIANTIAN, et al. "Projective dynamics: Fusing constraint projections for fast simulation". *ACM transactions on graphics (TOG)* 33.4 (2014), 1–11 2, 4, 7.

[BW98] BARAFF, DAVID and WITKIN, ANDREW. "Large steps in cloth simulation". *Proceedings of the 25th annual conference on Computer graphics and interactive techniques*. 1998, 43–54 2.

[DGP12] DAGENAIS, FRANCOIS, GAGNON, JONATHAN, and PAQUETTE, ERIC. "A prediction-correction approach for stable sph fluid simulation from liquid to rigid". (2012) 2.

[DLK18] DINEV, DIMITAR, LIU, TIANTIAN, and KAVAN, LADISLAV. "Stabilizing integrators for real-time physics". *ACM Transactions on Graphics (TOG)* 37.1 (2018), 1–19 7.

[Gan15] GANZENMÜLLER, GEORG C. "An hourglass control algorithm for Lagrangian smooth particle hydrodynamics". *Computer Methods in Applied Mechanics and Engineering* 286 (2015), 87–106 2, 3.

[GBB09] GERSZEWSKI, DAN, BHATTACHARYA, HAIMASREE, and BARGTEIL, ADAM W. "A point-based method for animating elastoplastic solids". *Proceedings of the 2009 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*. 2009, 133–138 2.

[GHB*20] GISSLER, CHRISTOPH, HENNE, ANDREAS, BAND, STEFAN, et al. "An implicit compressible SPH solver for snow simulation". *ACM Transactions on Graphics (TOG)* 39.4 (2020), 36–1 2.

[GJ*10] GUENNEBAUD, GAËL, JACOB, BENOÎT, et al. *Eigen v3*. http://eigen.tuxfamily.org. 2010 4.

[GPB*19] GISSLER, CHRISTOPH, PEER, ANDREAS, BAND, STEFAN, et al. "Interlinked SPH pressure solvers for strong fluid-rigid coupling". *ACM Transactions on Graphics (TOG)* 38.1 (2019), 1–13 2.

[GSS*15] GAST, THEODORE F, SCHROEDER, CRAIG, STOMAKHIN, ALEXEY, et al. "Optimization integrator for large time steps". *IEEE transactions on visualization and computer graphics* 21.10 (2015), 1103–1115 2.

[HEW15] HUBER, MARKUS, EBERHARDT, BERNHARD, and WEISKOPF, DANIEL. "Boundary handling at cloth–fluid contact". *Computer Graphics Forum*. Vol. 34. 1. Wiley Online Library. 2015, 14–25 2.

[KBF*21] KUGELSTADT, TASSILO, BENDER, JAN, FERNÁNDEZ-FERNÁNDEZ, JOSÉ ANTONIO, et al. "Fast Corotated Elastic SPH Solids with Implicit Zero-Energy Mode Control". *Proc. ACM Comput. Graph. Interact. Tech.* 4.3 (Sept. 2021). DOI: 10.1145/3480142. URL: https://doi.org/10.1145/3480142 2–4, 7.

[KBST22] KOSCHIER, DAN, BENDER, JAN, SOLENTHALER, BARBARA, and TESCHNER, MATTHIAS. "A Survey on SPH Methods in Computer Graphics". *Computer Graphics Forum*. Vol. 41. 2. Wiley Online Library. 2022, 737–760 2.

[KUJH21] KEE, MIN HYUNG, UM, KIWON, JEONG, WOOSEOK, and HAN, JUNGHYUN. "Constrained projective dynamics: real-time simulation of deformable objects with energy-momentum conservation". *ACM Transactions on Graphics (TOG)* 40.4 (2021), 1–12 7.

[LBK17] LIU, TIANTIAN, BOUAZIZ, SOFIEN, and KAVAN, LADISLAV. "Quasi-newton methods for real-time simulation of hyperelastic materials". *Acm Transactions on Graphics (TOG)* 36.3 (2017), 1–16 2–4, 7.

[LBOK13] LIU, TIANTIAN, BARGTEIL, ADAM W, O'BRIEN, JAMES F, and KAVAN, LADISLAV. "Fast simulation of mass-spring systems". *ACM Transactions on Graphics (TOG)* 32.6 (2013), 1–7 2.

[LLKC21] LI, JING, LIU, TIANTIAN, KAVAN, LADISLAV, and CHEN, BAOQUAN. "Interactive Cutting and Tearing in Projective Dynamics with Progressive Cholesky Updates". *ACM Trans. Graph.* 40.6 (Dec. 2021). ISSN: 0730-0301. DOI: 10.1145/3478513.3480505. URL: https://doi.org/10.1145/3478513.3480505 7.

[MHHR07] MÜLLER, MATTHIAS, HEIDELBERGER, BRUNO, HENNIX, MARCUS, and RATCLIFF, JOHN. "Position based dynamics". *Journal of Visual Communication and Image Representation* 18.2 (2007), 109–118 2.

[MKN*04] MÜLLER, MATTHIAS, KEISER, RICHARD, NEALEN, ANDREW, et al. "Point based animation of elastic, plastic and melting objects". *Proceedings of the 2004 ACM SIGGRAPH/Eurographics symposium on Computer animation*. 2004, 141–151 2.

[MMC16] MACKLIN, MILES, MÜLLER, MATTHIAS, and CHENTANEZ, NUTTAPONG. "XPBD: position-based simulation of compliant constrained dynamics". *Proceedings of the 9th International Conference on Motion in Games*. 2016, 49–54 2.

[MMCK14] MACKLIN, MILES, MÜLLER, MATTHIAS, CHENTANEZ, NUTTAPONG, and KIM, TAE-YONG. "Unified particle physics for real-time applications". *ACM Transactions on Graphics (TOG)* 33.4 (2014), 1–12 1.

[Mon92] MONAGHAN, JOE J. "Smoothed particle hydrodynamics". *Annual review of astronomy and astrophysics* 30.1 (1992), 543–574 1, 2.

Table 1: *Scene parameters and performance data of the experiments made with our method. Here, the timings for pressure and elasticity are measured per simulation step.*

experiments	# of particles	material type	# of iterations	pressure time [ms]	elasticity time [ms]	parameters			
						h [ms]	E	ν	α
Solid-fluid coupling (Fig. 1)	359,256	Neo-Hookean	20	1566.7	1195.0	1	5×10^4	0.33	10.0
Swinging cloth (Fig. 2)	22,260	Corotated	1	69.2	42.3	10	8×10^4	0.2	0.1
Stretching cube (Fig. 3)	10,852	Corotated	2	78.6	72.1	30	5×10^5	0.45	1.0
Twisting beam (Fig. 5)	8,982	Neo-Hookean	10	188.4	179.4	30	5×10^5	0.33	10.0
Colliding bunnies (Fig. 6)	14,070	St. Venant-Kirchoff	10	132.9	114.5	1	5×10^4	0.3	10.0
Melting duck (Fig. 7)	5,292	Corotated	5	110.3	26.5	10	2.5×10^6	0.33	0.1

[MTGG11] MARTIN, SEBASTIAN, THOMASZEWSKI, BERNHARD, GRINSPUN, EITAN, and GROSS, MARKUS. "Example-based elastic materials". *ACM SIGGRAPH 2011 papers*. 2011, 1–8 2.

[NW99] NOCEDAL, JORGE and WRIGHT, STEPHEN J. *Numerical optimization*. Springer, 1999 2, 3.

[OBLN17] OVERBY, MATTHEW, BROWN, GEORGE E, LI, JIE, and NARAIN, RAHUL. "ADMM ⊇ projective dynamics: Fast simulation of hyperelastic models with dynamic constraints". *IEEE Transactions on Visualization and Computer Graphics* 23.10 (2017), 2222–2234 2.

[PGBT18] PEER, ANDREAS, GISSLER, CHRISTOPH, BAND, STEFAN, and TESCHNER, MATTHIAS. "An implicit SPH formulation for incompressible linearly elastic solids". *Computer Graphics Forum*. Vol. 37. 6. Wiley Online Library. 2018, 135–148 2, 4.

[SB12] SIFAKIS, EFTYCHIOS and BARBIC, JERNEJ. "FEM simulation of 3D deformable solids: a practitioner's guide to theory, discretization and model reduction". *Acm siggraph 2012 courses*. 2012, 1–50 1.

[SH98] STUART, A. and HUMPHRIES, A.R. *Dynamical Systems and Numerical Analysis*. Cambridge Monographs on Applie vol. 8. Cambridge University Press, 1998. ISBN: 9780521645638. URL: https://books.google.fr/books?id=ymoQA8s5pNIC 2.

[SSP07] SOLENTHALER, BARBARA, SCHLÄFLI, JÜRG, and PAJAROLA, RENATO. "A unified particle model for fluid–solid interactions". *Computer Animation and Virtual Worlds* 18.1 (2007), 69–82 2.

[TPBF87] TERZOPOULOS, DEMETRI, PLATT, JOHN, BARR, ALAN, and FLEISCHER, KURT. "Elastically deformable models". *Proceedings of the 14th annual conference on Computer graphics and interactive techniques*. 1987, 205–214 1.

[Wan15] WANG, HUAMIN. "A chebyshev semi-iterative approach for accelerating projective and position-based dynamics". *ACM Transactions on Graphics (TOG)* 34.6 (2015), 1–9 2.

[WKB16] WEILER, MARCEL, KOSCHIER, DAN, and BENDER, JAN. "Projective fluids". *Proceedings of the 9th International Conference on Motion in Games*. 2016, 79–84 8.

[XTL19] XIAN, ZANGYUEYANG, TONG, XIN, and LIU, TIANTIAN. "A Scalable Galerkin Multigrid Method for Real-time Simulation of Deformable Objects". *ACM Transactions on Graphics (TOG)* 38.6 (2019) 7.