

# Position-based Methods for the Simulation of Solid Objects in Computer Graphics

Jan Bender<sup>1</sup>, Matthias Müller<sup>2</sup>, Miguel A. Otaduy<sup>3</sup> and Matthias Teschner<sup>4</sup>

<sup>1</sup>Graduate School CE, TU Darmstadt

<sup>2</sup>NVIDIA PhysX Research

<sup>3</sup>URJC Madrid

<sup>4</sup>University of Freiburg

---

## Abstract

*The dynamic simulation of solids has a long history in computer graphics. The classical methods in this field are based on the use of forces or impulses to simulate joints between rigid bodies as well as the stretching, shearing and bending stiffness of deformable objects. In the last years the class of position-based methods has become popular in the graphics community. These kinds of methods are fast, unconditionally stable and controllable which make them well-suited for the use in interactive environments. Position-based methods are not as accurate as force based methods in general but they provide visual plausibility. Therefore, the main application areas of these approaches are virtual reality, computer games and special effects in movies.*

*This state of the art report covers the large variety of position-based methods that were developed in the field of deformable solids. We will introduce the concept of position-based dynamics, present dynamic simulation based on shape matching and discuss data-driven approaches. Furthermore, we will present several applications for these methods.*

## Keywords:

Categories and Subject Descriptors (according to ACM CCS): Computer Graphics [I.3.7]: Three-Dimensional Graphics and Realism—Animation

---

## 1. Introduction

The simulation of solid objects such as rigid bodies, soft bodies or cloth has been an important and active research topic in computer graphics for more than 30 years. The field was introduced to graphics by Terzopoulos and his colleagues in the late eighties [TPBF87a]. Since then, a large body of work has been published and the list is growing rapidly. There exists a variety of survey papers [GM97, MTV05, NMK\*06, MSJT08, BETC12] which document this development. However, due to the sheer number of papers, it has become difficult to cover the entire field in one survey.

In this state of the art report we focus on a special class of simulation methods for solids, namely position-based approaches. Classical dynamics simulation methods formulate the change of momentum of a system as a function of applied forces, and evolve positions through numerical integration of accelerations and velocities. Position-based approaches, in-

stead, compute positions directly, based on the solution to a quasi-static problem.

Physical simulation is a well studied problem in the computational sciences and therefore, many of the well established methods have been adopted in graphics such as the Finite Element Method (FEM) [OH99], the Finite Differences Method [TPBF87b], the Finite Volume Method [TBHF03], the boundary element method [JP99] or particle-based approaches [DSB99, THMG04]. The main goal of computer simulations in computational physics and chemistry is to replace real-world experiments and thus, to be as accurate as possible. In contrast, the main applications of physically based simulation methods in computer graphics are special effects in movies and commercials and more recently, computer games and other interactive systems. Here, speed and controllability are the most important factors and all that is required in terms of accuracy is visual plausibility. This is especially true for real-time applications.

Position-based methods are tailored particularly for the use in interactive environments. They provide a high level of control and are unconditionally stable even when simple and fast explicit time integration schemes are used. Due to their simplicity, robustness and speed these approaches have become very popular in computer graphics and in the game industry recently. Since there exists no survey paper on position-based simulation methods so far, we decided to provide an overview of this field of research in a state of the art report and hope it will be useful for the growing number of people interested in real-time physically-based simulations.

In this state of the art report, we focus the discussion on several geometrically motivated methods, in particular those that directly compute the positions of the solids being simulated. We start with a description of position-based dynamics methods, which compute equilibrium positions by iteratively resolving geometric constraints. Then, we cover shape matching methods, and we conclude with data-driven methods, which compute positions as a function of data from pre-captured deformation examples.

Collision detection is an important part of any simulation system. However, an adequate discussion of this topic is beyond the scope of this report. Therefore, we refer the reader to the surveys of Lin and Gottschalk [LG98] and the one of Teschner et al. [TKH\*05].

## 2. Background

The most popular approaches for the simulation of dynamic systems in computer graphics are force based. Internal and external forces are accumulated from which accelerations are computed based on Newton's second law of motion. A time integration method is then used to update the velocities and finally the positions of the object. A few simulation methods (most rigid body simulators) use impulse based dynamics and directly manipulate velocities. In contrast, geometry-based methods omit the velocity layer as well and immediately work on the positions. The main advantage of a position-based approach is its controllability. Overshooting problems of explicit integration schemes in force based systems can be avoided. In addition, collision constraints can be handled easily and penetrations can be resolved completely by projecting points to valid locations.

Among the force based approaches, one of the simplest methods is to represent and simulate solids with mass-spring networks. A mass spring system consists of a set of point masses that are connected by springs. The physics of such a system is straightforward and a simulator is easy to implement. However, there are some significant drawbacks of the simple method.

- The behavior of the object depends on the way the spring network is set up.
- It can be difficult to tune the spring constants to get the desired behavior.

- Mass spring networks cannot capture volumetric effects directly such as volume conservation or prevention of volume inversions.

The Finite Element Method solves all of the above problems because it considers the entire volume of a solid instead of replacing it with a finite number of point masses. Here, the object is discretized by splitting the volume into a number of elements with finite size. This discretization yields a mesh as in the mass spring approach in which the vertices play the role of the mass points and the elements, typically tetrahedra, can be viewed as generalized springs acting on multiple points at the same time. In both cases, forces at the mass points or mesh vertices are computed due to their velocities and the actual deformation of the mesh.

## 3. Position-Based Dynamics

In this section we present an approach which omits the velocity and acceleration layer and immediately works on the positions [MHHR07].

### 3.1. Overview

The objects to be simulated are represented by a set of  $N$  particles and a set of  $M$  constraints. Each particle  $i$  has three attributes, namely

$m_i$	mass
$\mathbf{x}_i$	position
$\mathbf{v}_i$	velocity

A constraint  $j$  is defined by the five attributes

$n_j$	cardinality
$C_j : \mathbb{R}^{3n_j} \rightarrow \mathbb{R}$	scalar constraint function
$\{i_1, \dots, i_{n_j}\}, i_k \in [1, \dots, N]$	set of indices
$k_j \in [0 \dots 1]$	stiffness parameter
<i>unilateral</i> or <i>bilateral</i>	type

Constraint  $j$  with type *bilateral* is satisfied if  $C_j(\mathbf{x}_{i_1}, \dots, \mathbf{x}_{i_{n_j}}) = 0$ . If its type is *unilateral* then it is satisfied if  $C_j(\mathbf{x}_{i_1}, \dots, \mathbf{x}_{i_{n_j}}) \geq 0$ . The stiffness parameter  $k_j$  defines the strength of the constraint in a range from zero to one.

Given this data and a time step  $\Delta t$ , the simulation proceeds as described by Algorithm 1. Since the algorithm simulates a system which is second order in time, both the positions and the velocities of the particles need to be specified in (1)-(3) before the simulation loop starts. Lines (5)-(6) perform a simple explicit forward Euler integration step on the velocities and the positions. The new locations  $\mathbf{p}_i$  are not assigned to the positions directly but are only used as predictions. Non-permanent external constraints such as collision

**Algorithm 1** Position-based dynamics

---

```

1: for all vertices  $i$  do
2:   initialize  $\mathbf{x}_i = \mathbf{x}_i^0$ ,  $\mathbf{v}_i = \mathbf{v}_i^0$ ,  $w_i = 1/m_i$ ,
3: end for
4: loop
5:   for all vertices  $i$  do  $\mathbf{v}_i \leftarrow \mathbf{v}_i + \Delta t w_i \mathbf{f}_{\text{ext}}(\mathbf{x}_i)$ 
6:   for all vertices  $i$  do  $\mathbf{p}_i \leftarrow \mathbf{x}_i + \Delta t \mathbf{v}_i$ 
7:   for all vertices  $i$  do genCollConstraints( $\mathbf{x}_i \rightarrow \mathbf{p}_i$ )
8:   loop solverIteration times
9:     projectConstraints( $C_1, \dots, C_{M+M_{\text{coll}}}, \mathbf{p}_1, \dots, \mathbf{p}_N$ )
10:  end loop
11:  for all vertices  $i$  do
12:     $\mathbf{v}_i \leftarrow (\mathbf{p}_i - \mathbf{x}_i) / \Delta t$ 
13:     $\mathbf{x}_i \leftarrow \mathbf{p}_i$ 
14:  end for
15:  velocityUpdate( $\mathbf{v}_1, \dots, \mathbf{v}_N$ )
16: end loop

```

---

constraints are generated at the beginning of each time step from scratch in line (7). Here the original and the predicted positions are used in order to perform continuous collision detection. The solver (8)-(10) then iteratively corrects the predicted positions such that they satisfy the  $M_{\text{coll}}$  external as well as the  $M$  internal constraints. Finally, the corrected positions  $\mathbf{p}_i$  are used to update the positions *and* the velocities. It is essential here to update the velocities along with the positions. If this is not done, the simulation does not produce the correct behavior of a second order system. As you can see, the integration scheme used here is very similar to the Verlet method.

### 3.2. The System to be Solved

The goal of the solver step (8)-(10) is to correct the predicted positions  $\mathbf{p}_i$  of the particles such that they satisfy all constraints. The problem that needs to be solved comprises of a set of  $M$  equations for the  $3N$  unknown position components, where  $M$  is now the total number of constraints. This system does not need to be symmetric. If  $M > 3N$  ( $M < 3N$ ) the system is over-determined (under-determined). In addition to the asymmetry, the equations are in general non-linear. The function of a simple distance constraint  $C(\mathbf{p}_1, \mathbf{p}_2) = (\mathbf{p}_1 - \mathbf{p}_2)^2 - d^2$  yields a non-linear equation. What complicates things even further is the fact that collisions produce inequalities rather than equalities. Solving a non-symmetric, non-linear system with equalities and inequalities is a tough problem.

Let  $\mathbf{p}$  be the concatenation  $[\mathbf{p}_1^T, \dots, \mathbf{p}_N^T]^T$  and let all the constraint functions  $C_j$  take the concatenated vector  $\mathbf{p}$  as input while only using the subset of coordinates they are de-

finied for. We can now write the system to be solved as

$$\begin{aligned} C_1(\mathbf{p}) &\succ 0 \\ &\dots \\ C_M(\mathbf{p}) &\succ 0, \end{aligned}$$

where the symbol  $\succ$  denotes either  $=$  or  $\geq$ . Newton-Raphson iteration is a method to solve non-linear symmetric systems with equalities only. The process starts with a first guess of a solution. Each constraint function is then linearized in the neighborhood of the current solution using

$$C(\mathbf{p} + \Delta\mathbf{p}) \approx C(\mathbf{p}) + \nabla_{\mathbf{p}} C(\mathbf{p}) \cdot \Delta\mathbf{p} + O(|\Delta\mathbf{p}|^2) = 0. \quad (1)$$

This yields a *linear* system for the global correction vector  $\Delta\mathbf{p}$

$$\begin{aligned} \nabla_{\mathbf{p}} C_1(\mathbf{p}) \cdot \Delta\mathbf{p} &= -C_1(\mathbf{p}) \\ &\dots \\ \nabla_{\mathbf{p}} C_M(\mathbf{p}) \cdot \Delta\mathbf{p} &= -C_M(\mathbf{p}), \end{aligned}$$

where  $\nabla_{\mathbf{p}} C_j(\mathbf{p})$  is the  $1 \times N$  dimensional vector containing the derivatives of the function  $C_j$  w.r.t. all its parameters, i.e. the  $N$  components of  $\mathbf{p}$ . It is also the  $j$ -th row of the linear system. Both, the rows  $\nabla_{\mathbf{p}} C_j(\mathbf{p})$  and the right hand side scalars  $-C_j(\mathbf{p})$  are constant because they are *evaluated* at the location  $\mathbf{p}$  before the system is solved. When  $M = 3N$  and only equalities are present, the system can be solved by any linear solver, e.g. a preconditioned conjugate gradient method. Once it is solved for  $\Delta\mathbf{p}$  the current solution is updated as  $\mathbf{p} \leftarrow \mathbf{p} + \Delta\mathbf{p}$ . A new linear system is generated by evaluating  $\nabla_{\mathbf{p}} C_j(\mathbf{p})$  and  $-C_j(\mathbf{p})$  at the new location after which the process repeats.

If  $M \neq 3N$  the resulting matrix of the linear system is non-symmetric and not invertible. [GHF\*07a] solve this problem by using the pseudo-inverse of the system matrix which yields the best solution in the least-squares sense. Still, handling inequalities is not possible directly.

### 3.3. The Non-Linear Gauss-Seidel Solver

In the position-based dynamics approach, non-linear Gauss-Seidel is used. It solves each constraint equation separately. Each constraint yields a single scalar equation  $C(\mathbf{p}) \succ 0$  for all the particle positions associated with it. The subsystem is therefore highly under-determined. PBD solves this problem as follows. Again, given  $\mathbf{p}$  we want to find a correction  $\Delta\mathbf{p}$  such that  $C(\mathbf{p} + \Delta\mathbf{p}) = 0$ . It is important to notice that PBD also linearizes the constraint function but individually for each constraint. The constraint equation is approximated by

$$C(\mathbf{p} + \Delta\mathbf{p}) \approx C(\mathbf{p}) + \nabla_{\mathbf{p}} C(\mathbf{p}) \cdot \Delta\mathbf{p} = 0. \quad (2)$$

The problem of the system being under-determined is solved by restricting  $\Delta\mathbf{p}$  to be in the direction of  $\nabla_{\mathbf{p}} C$  which con-

serves the linear and angular momenta. This means that only one scalar  $\lambda$  - a Lagrange multiplier - has to be found such that the correction

$$\Delta \mathbf{p} = \lambda \nabla_{\mathbf{p}} C(\mathbf{p}). \quad (3)$$

solves Equation (2). This yields the following formula for the correction vector of a single particle  $i$

$$\Delta \mathbf{p}_i = -s w_i \nabla_{\mathbf{p}_i} C(\mathbf{p}), \quad (4)$$

where

$$s = \frac{C(\mathbf{p})}{\sum_j w_j |\nabla_{\mathbf{p}_j} C(\mathbf{p})|^2} \quad (5)$$

and  $w_i = 1/m_i$ .

As mentioned above, this solver linearizes the constraint functions. However, in contrast to the Newton-Raphson method, the linearization happens individually *per constraint*. Solving the linearized constraint function of a single distance constraint for instance yields the correct result in a single step. Because the positions are immediately updated after a constraint is processed, these updates will influence the linearization of the next constraint because the linearization depends on the actual positions. Asymmetry does not pose a problem because each constraint produces one scalar equation for one unknown Lagrange multiplier  $\lambda$ . Inequalities are handled trivially by first checking whether  $C(\mathbf{p}) \geq 0$ . If this is the case, the constraint is simply skipped.


We have not considered the stiffness  $k$  of the constraint so far. There are several ways to incorporate it. The simplest variant is to multiply the corrections  $\Delta \mathbf{p}$  by  $k \in [0 \dots 1]$ . However, for multiple iteration loops of the solver, the effect of  $k$  is non-linear. The remaining error for a single distance constraint after  $n_s$  solver iterations is  $\Delta \mathbf{p}(1-k)^{n_s}$ . To get a linear relationship we multiply the corrections not by  $k$  directly but by  $k' = 1 - (1-k)^{1/n_s}$ . With this transformation the error becomes  $\Delta \mathbf{p}(1-k')^{n_s} = \Delta \mathbf{p}(1-k)$  and, thus, becomes linearly dependent on  $k$  and independent of  $n_s$  as desired. However, the resulting material stiffness is still dependent on the time step of the simulation. Real time environments typically use fixed time steps in which case this dependency is not problematic.

### 3.4. Constraint Examples


In the following we will introduce different constraint examples. For a better readability we define  $\mathbf{p}_{i,j} = \mathbf{p}_i - \mathbf{p}_j$ .

#### 3.4.1. Stretching

To give an example, let us consider the distance constraint function  $C(\mathbf{p}_1, \mathbf{p}_2) = |\mathbf{p}_{1,2}| - d$ . The derivatives with respect to the points are  $\nabla_{\mathbf{p}_1} C(\mathbf{p}_1, \mathbf{p}_2) = \mathbf{n}$  and  $\nabla_{\mathbf{p}_2} C(\mathbf{p}_1, \mathbf{p}_2) = -\mathbf{n}$  with  $\mathbf{n} = \frac{\mathbf{p}_{1,2}}{|\mathbf{p}_{1,2}|}$ . The scaling factor  $s$  is, thus,  $s = \frac{|\mathbf{p}_{1,2}| - d}{1+1}$  and



**Figure 1:** Projection of the constraint  $C(\mathbf{p}_1, \mathbf{p}_2) = |\mathbf{p}_{1,2}| - d$ . The corrections  $\Delta \mathbf{p}_i$  are weighted according to the inverse masses  $w_i = 1/m_i$ .



**Figure 2:** For bending resistance, the constraint function  $C(\mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_3, \mathbf{p}_4) = \arccos(\mathbf{n}_1 \cdot \mathbf{n}_2) - \varphi_0$  is used. The actual dihedral angle  $\varphi$  is measured as the angle between the normals of the two triangles.

the final corrections

$$\Delta \mathbf{p}_1 = -\frac{w_1}{w_1 + w_2} (|\mathbf{p}_{1,2}| - d) \frac{\mathbf{p}_{1,2}}{|\mathbf{p}_{1,2}|} \quad (6)$$

$$\Delta \mathbf{p}_2 = +\frac{w_2}{w_1 + w_2} (|\mathbf{p}_{1,2}| - d) \frac{\mathbf{p}_{1,2}}{|\mathbf{p}_{1,2}|} \quad (7)$$

which are the formulas proposed in [Jak01] for the projection of distance constraints (see Figure 1). They pop up as a special case of the general constraint projection method.

#### 3.4.2. Bending

In cloth simulation it is important to simulate bending in addition to stretching resistance. To this end, for each pair of adjacent triangles  $(\mathbf{p}_1, \mathbf{p}_3, \mathbf{p}_2)$  and  $(\mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_4)$  a bilateral bending constraint is added with constraint function


$$C_{bend}(\mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_3, \mathbf{p}_4) = \arccos \left( \frac{\mathbf{p}_{2,1} \times \mathbf{p}_{3,1}}{|\mathbf{p}_{2,1} \times \mathbf{p}_{3,1}|} \cdot \frac{\mathbf{p}_{2,1} \times \mathbf{p}_{4,1}}{|\mathbf{p}_{2,1} \times \mathbf{p}_{4,1}|} \right) - \varphi_0$$

and stiffness  $k_{bend}$ . The scalar  $\varphi_0$  is the initial dihedral angle between the two triangles and  $k_{bend}$  is a global user parameter defining the bending stiffness of the cloth (see Figure 2). The advantage of this bending term over adding a distance constraint between points  $\mathbf{p}_3$  and  $\mathbf{p}_4$  or over the bending term proposed by [GHDS03] is that it is *independent of stretching*. This is because the term is independent of edge lengths.

#### 3.4.3. Triangle Collisions

The handling of self collisions within cloth can be handled by an additional unilateral constraint. For vertex  $\mathbf{q}$  moving through a triangle  $\mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_3$ , the constraint function reads

$$C(\mathbf{q}, \mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_3) = (\mathbf{q} - \mathbf{p}_1) \cdot \frac{\mathbf{p}_{2,1} \times \mathbf{p}_{3,1}}{|\mathbf{p}_{2,1} \times \mathbf{p}_{3,1}|} - h, \quad (8)$$



**Figure 3:** Constraint function  $C(\mathbf{q}, \mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_3) = (\mathbf{q} - \mathbf{p}_1) \cdot \mathbf{n} - h$  makes sure that  $\mathbf{q}$  stays above the triangle  $\mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_3$  by the cloth thickness  $h$ .

where  $h$  is the cloth thickness. If the vertex enters from below with respect to the triangle normal, the constraint function has to be

$$C(\mathbf{q}, \mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_3) = (\mathbf{q} - \mathbf{p}_1) \cdot \frac{\mathbf{p}_{3,1} \times \mathbf{p}_{2,1}}{|\mathbf{p}_{3,1} \times \mathbf{p}_{2,1}|} - h. \quad (9)$$

#### 3.4.4. Volume Conservation

For tetrahedral meshes it is useful to have a constraint that conserves the volume of single tetrahedron. Such a constraint has the form

$$C(\mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_3, \mathbf{p}_4) = \frac{1}{6} (\mathbf{p}_{2,1} \times \mathbf{p}_{3,1}) \cdot \mathbf{p}_{4,1} - V_0, \quad (10)$$

where  $\mathbf{p}_1, \mathbf{p}_2, \mathbf{p}_3$  and  $\mathbf{p}_4$  are the four corners of the tetrahedron and  $V_0$  is its rest volume.

#### 3.4.5. Cloth Balloons

For closed triangle meshes, overpressure inside the mesh as shown in Figure 5 can easily be modeled with an *equality* constraint concerning all  $N$  vertices of the mesh:

$$C(\mathbf{p}_1, \dots, \mathbf{p}_N) = \left( \sum_{i=1}^{n_{\text{triangles}}} (\mathbf{p}_{t_1^i} \times \mathbf{p}_{t_2^i}) \cdot \mathbf{p}_{t_3^i} \right) - k_{\text{pressure}} V_0. \quad (11)$$

Here  $t_1^i, t_2^i$  and  $t_3^i$  are the three indices of the vertices belonging to triangle  $i$ . The sum computes the actual volume of the closed mesh. It is compared against the original volume  $V_0$  times the overpressure factor  $k_{\text{pressure}}$ . This constraint function yields the gradients

$$\nabla_{\mathbf{p}_i} C = \sum_{j:t_1^j=i} (\mathbf{p}_{t_2^j} \times \mathbf{p}_{t_3^j}) + \sum_{j:t_2^j=i} (\mathbf{p}_{t_3^j} \times \mathbf{p}_{t_1^j}) + \sum_{j:t_3^j=i} (\mathbf{p}_{t_1^j} \times \mathbf{p}_{t_2^j}) \quad (12)$$

These gradients have to be scaled by the scaling factor given in Equation (5) and weighted by the masses according to Equation (4) to get the final projection offsets  $\Delta \mathbf{p}_i$ .

### 3.5. Strain Limiting

Strain limiting is an important topic in the field of cloth simulation. The reason is that the low solver iteration counts used in real-time applications yield stretchy cloth. Since



**Figure 5:** Simulation of overpressure inside a character.


most cloth types are perceived by the human eye as completely inextensible, it is important to make simulated cloth inextensible in order to avoid disturbing visual artifacts.

A strain limiting method makes sure that the overall stretch of the cloth stays below a certain threshold. In force based simulations, strain limiting is a separate pass which is executed before or after the regular cloth solver. In most cases, this pass moves the positions of vertices directly, even in force based simulations. Therefore, most strain limiting methods fall under the category of position-based methods.

A straightforward way of limiting strain is to iterate through all edges of a cloth mesh and project the adjacent particles of overstretched edges as shown in Figure 1 so that the stretch of the edge does not exceed the stretch limit. Provot [Pro95] was among the first to use this method in the context of cloth simulation. He performs a single iteration through all cloth edges after a force based solver. Desbrun et al. [DSB99] and Bridson et al. [BMR03] later used the same post solver strain limiter but with multiple iterations through all edges. Due to its simplicity, this method is still one of the most popular strain limiting methods used in cloth simulations.

The method is very similar to position based cloth simulation. The main difference is that the strain limiting pass described above does not influence the velocities. These are updated by the force-based solver. In contrast, position based cloth simulation derives the new velocities from the projections, making an additional solver pass obsolete. Therefore, every position based strain limiting method used in force based simulations can directly be used in a PBD solver.

The result of projecting along edges depends on the structure of the mesh. To reduce this artifact, Wang et al. [WOR10] propose to limit the principal strains of the 2D deformation field within each triangle. The 2D deformation field can be determined by considering the 2D coordinates of the vertices of a triangle within the planes of the rest and current triangle configurations. [WOR10] compute the principal strains of the 2D deformation gradient, clamp them and construct a new 2D transformation using the clamped strains. With this new transformation they correct the current positions of the triangle vertices. As before, to limit strain globally, they iterate through all triangles multiple times in a Gauss-Seidel fashion.




**Figure 4:** The image shows a mesh that is simulated using stretching and bending constraints. The top row shows  $(k_{stretching}, k_{bending}) = (1, 1)$ ,  $(\frac{1}{2}, 1)$  and  $(\frac{1}{100}, 1)$ . The bottom row shows  $(k_{stretching}, k_{bending}) = (1, 0)$ ,  $(\frac{1}{2}, 0)$  and  $(\frac{1}{100}, 0)$ .

Due to the relatively slow convergence rate of a Gauss-Seidel solver, high iteration counts are necessary to limit the strain globally which slows down the simulation. The two main methods to improve the convergence rate are the use of a global Newton-Raphson solver as proposed by Goldenthal et al. [GHF\*07b] or to perform Gauss-Seidel iterations on a hierarchy of meshes as proposed by [Mül08] and [WOR10]. However, these methods complicate the implementation and even though their convergence rate is higher, a single iteration can be significantly more expensive than a simple Gauss-Seidel iteration.

Recently, Kim et al. [KCM12] found a surprisingly simple and robust technique they call *Long Range Attachments* (LRA) to prevent cloth from getting stretched globally with low iteration counts. Their method exploits the fact that stretching artifacts almost only appear when cloth is attached. In this case, instead of only applying attachment constraints to the subset of the vertices near the region where the cloth is attached and relying on error propagation of the solver for all other vertices, they apply unilateral attachment constraints to *all* the vertices by attaching each vertex to one or more attachment point directly. The rest lengths of these long range attachments can either be set to the Euclidean distance in the rest state or via measuring geodesic lengths along the cloth. Figure 6 demonstrates the method on a single rope attached at one end. The method allows the simulation of a piece of cloth with 90K vertices at interactive rates as shown in Figure 7.

A similar approach was recently proposed by Müller et al. [MKC12] to guarantee zero stretch in a single pass for the case of attached ropes. This approach allows the simulation of thousands of hair strands in real time (Figure 9). Figure 8 visualizes the basic idea. Particle  $\mathbf{x}_1$  is attached. To



**Figure 6:** The Long Range Attachments (LRA) method used to simulate an inextensible rope attached at one end. Each particle is constrained or remain inside a sphere centered at the attachment point (red) whose radius is the initial distance from the particle to the attachment. For each configuration, target positions are shown in green when particles need to be projected. Particles inside the constraint spheres are allowed to move freely.

satisfy the first distance constraint, particle  $\mathbf{x}_2$  is moved towards  $\mathbf{x}_1$  such that their mutual distance is  $l_0$ . Particle  $\mathbf{x}_3$  is then moved towards the *new* position of  $\mathbf{x}_2$  and similarly along the chain until the last particle is reached. After this single pass, all the distance constraints are satisfied. This method is called *Follow The Leader* (FTL). While LRA guarantees zero stretch of all the particles w.r.t. the attachment points, the constraint between consecutive particles can still remain overstretched. On the other hand, in contrast to LRA which is momentum conserving, FTL introduces unphysical behavior. Not projecting distance constraints symmetrically means that a system is simulated for which each particle has






**Figure 7:** Simulation of a piece of cloth with 90K vertices at 20fps on a GPU using LRA.

infinitely more mass than its successor. To compensate for this behavior, the authors replace the PBD velocity update  $\mathbf{v}_i \leftarrow (\mathbf{p}_i - \mathbf{x}_i)/\Delta t$  by

$$\mathbf{v}_i \leftarrow \frac{\mathbf{p}_i - \mathbf{x}_i}{\Delta t} + s_{\text{damping}} \frac{-\mathbf{d}_{i+1}}{\Delta t}, \quad (13)$$


where  $\mathbf{d}_{i+1}$  is the position correction applied to particle  $i+1$  and  $s_{\text{damping}} \in [0, 1]$  a scaling factor do influence damping. While this modification of DFTL (dynamic FTL) hides the un-physical behavior of FTL, it introduces a certain amount of damping which is acceptable for the simulation of hair and fur as the author's results show.




**Figure 8:** Follow The Leader (FTL) projection. Starting from the attachment down, each particle is moved directly towards its predecessor such that their mutual distance constraint is satisfied.

### 3.6. Wrinkle Meshes

In cloth simulations, reducing the mesh resolution not only reduces the cost of a single solver iteration but also the number of iterations required to get visually pleasing results. In [MC10] the authors proposed a way to reduce the resolution of the dynamic mesh without losing too much visual detail. The most significant detail in cloth simulations are small wrinkles. The method is based on the observation that global dynamic behavior of the cloth and wrinkle formation can be separated. Therefore, expensive dynamic simulation including collision handling is performed on a low-resolution mesh. The wrinkle formation is handled on a high resolution mesh that is attached to the dynamic mesh (see Figures 10 and 11). Since wrinkles do not oscillate, it is sufficient to use a static solver with a low iteration count on the high-resolution mesh.



**Figure 10:** Basic idea of wrinkle meshes. The high resolution wrinkle mesh (white vertices) follows the low-resolution dynamic mesh (black vertices) by restricting the white vertices to remain within a certain distance (gray discs) to the dynamic mesh.




**Figure 11:** Visualization of the wrinkle mesh (solid) and the underlying dynamic mesh (wireframe).

Figure 12 shows the constraints defined on the high-resolution mesh to make it form wrinkles and follow the dynamic mesh. The attachment constraints makes sure that the vertices of the wrinkle mesh stay close to their attachment points on the dynamic mesh. If the dynamic mesh has outside/inside information, a one-sided constraint can be used



**Figure 9:** Dynamic FTL allows the simulation of every hair strand in real time. From left to right: 47k hair strands simulated at 25 fps including rendering and hair-hair repulsion. Long hair composed of 1.9m particles at 8 fps. Curly hair using visualization post-processing.

which makes sure that the wrinkle vertices stay on the outside of the dynamic mesh, thus avoiding penetrations with other objects. The stretching and bending constraints are responsible for wrinkle formation.



**Figure 12:** Static constraints on a wrinkle mesh: Attachment constraint (top left), one sided attachment constraint (top right), stretch constraint (bottom left) and bending constraint (bottom right).

### 3.7. Damping

The quality of dynamic simulations can generally be improved by the incorporation of an appropriate damping scheme. As a positive effect, damping can improve the stability by reducing temporal oscillations of the point positions of an object. This enables the use of larger time steps which increases the perceived performance of a dynamic simulation. On the other hand, damping changes the dynamic motion of the simulated objects. The resulting effects can be either desired, e.g. reduced oscillations of a deformable solid, or disturbing, e.g. changes of the linear or angular momentum of the entire object.

Generally, a damping term  $C\dot{\mathbf{X}}$  can be incorporated into the motion equation of an object where  $\dot{\mathbf{X}}$  denotes the vector of all first time derivatives of positions. If the user-defined matrix  $C$  is diagonal, absolute velocities of the points are

damped, which sometimes is referred to as point damping. If appropriately computed, such point damping forces result in an improved numerical stability by reducing the acceleration of a point. Such characteristics are desired in some settings, e.g. in the context of friction. In the general case, however, the overall slow-down of an object, caused by point damping forces, is not desired. Point damping forces are, e.g., used in [TF88] or in [PB88], where point damping is used for dynamic simulations with geometric constraints such as point-to-nail.

In order to preserve linear and angular momentum of deformable objects, symmetric damping forces, usually referred to as spring damping forces, can be used. Such forces can be represented by non-diagonal entries in the matrix  $C$ . Such damping forces are, e.g., described in [BW98] or [NMK\*06]. These approaches are independent from the deformation model. Therefore, they can also be applied to position-based methods. However, as [BW98] and [NMK\*06] rely on topological information of the object geometry, they cannot be applied to meshless techniques such as shape matching.

Point and spring damping can be used to reduce current velocities or relative velocities. However, it is generally more appropriate to consider predicted velocities or relative velocities for the next time step.

An interesting damping alternative has been presented in [SGT09]. Here, the idea of symmetric, momentum-conserving forces is extended to meshless representations. Global symmetric damping forces are computed with respect to the center of mass of an object. While such forces conserve the linear momentum, the preservation of the angular momentum is guaranteed by force projection onto relative positions or by torque elimination using Linear Programming. The approach presented in [SGT09] iteratively computes damping forces. The paper, however, also shows the convergence of the iterative process and how to explicitly compute the final solution. Therefore, the approach is an efficient alternative to compute damping forces for arbitrary position-based deformation models with or without connec-



tivity information. The approach can be used to damp oscillations globally or locally for user-defined clusters.

### 3.8. Parallelization

Parallelization of the PBD approach is an important topic since multi-core systems and massively parallel GPUs are ubiquitous today.


In a single CPU implementation, the solver processes the constraints one by one in a Gauss-Seidel-type fashion. Thereby, after each constraint projection, the positions of affected particles are immediately updated. In a parallel implementation, the constraints are processed in parallel by multiple threads. If two constraints affecting the same particle are handled by two different threads simultaneously, they are not allowed to immediately update the particle's position because writing to the same position simultaneously leads to race conditions making the process unpredictable. A solution to circumvent this problem is to use atomic operations. Such operations are guaranteed not to be interrupted. However, atomics can slow down parallel execution significantly.

To avoid these issues, a parallel implementation of PBD needs to split the constraints into groups or phases. In each phase, none of the constraints are allowed to share a common particle. With this restriction, the constraints in the first phase can be processed in parallel without conflicts. Then, after a global synchronization, the next phase can be processed. This cycle is repeated until all constraints are processed.

As an example, if  $N$  particles are connected in a serial chain, the constraints  $1-2, 3-4, 5-6, 7-8, \dots$  can be processed in phase 1 and the constraints  $2-3, 4-5, 6-7, \dots$  in phase 2. This specific example corresponds to the Red-Black Gauss Seidel scheme, where there are two sets (colors) of constraints. For more general types of constraint such as the stretch, shear and bending constraints of cloth, more phases are needed. In this general case, splitting constraints into phases corresponds to the graph coloring problem, where each constraint corresponds to a node of the graph and two constraints are connected by an edge if they affect one or more common particles. The minimum number of colors determines how many phases are needed in the parallel execution of PBD. Keeping the number of phases small is not the only optimization criterion. The sets also need to have similar sizes for good load balancing.

## 4. Shape Matching

The geometrically motivated concept of shape matching to simulate deformable objects was introduced by Müller et al. [MHTG05]. Shape matching is a meshless approach which is able to simulate visually plausible elastic and plastic deformations (see Figure 13). This approach is easy to implement, very efficient and unconditionally stable.



**Figure 14:** The initial shape with the vertex positions  $\bar{\mathbf{x}}_i$  is matched to the deformed configuration  $\mathbf{x}_i$  to obtain goal positions  $\mathbf{g}_i$ . The deformed shape is pulled towards these goal positions to simulate elastic behavior.

The basic idea of simulating elastic behavior with shape matching is shown in Figure 14. For the simulation the initial configuration of the deformable object must be stored. Since no connectivity information is needed, this configuration is defined by the initial positions  $\bar{\mathbf{x}}_i$ . In each time step the positions and velocities of the particles are updated without considering any internal constraints between the particles. Only external forces and collision response are taken into account. Instead of using internal constraints, goal positions are determined by matching the initial shape with the deformed configuration. Then, each particle is pulled towards its goal position.

In the following we first describe how the goal positions are determined. Then we show how large deformations can be simulated using region-based shape matching and introduce fast summation techniques for this approach. In the end the concept of oriented particles and different extensions of the shape matching method are presented.

### 4.1. Goal Positions

In order to obtain goal positions for the deformed shape the best rigid transformation is determined which matches the set of initial positions  $\bar{\mathbf{x}}$  and the set of deformed positions  $\mathbf{x}$ . The corresponding rotation matrix  $\mathbf{R}$  and the translational vectors  $\mathbf{c}$  and  $\bar{\mathbf{c}}$  are determined by minimizing

$$\sum_i w_i (\mathbf{R}(\bar{\mathbf{x}}_i - \bar{\mathbf{c}}) + \mathbf{c} - \mathbf{x}_i)^2$$

where  $w_i$  are the weights of the individual points. The optimal translation vectors are given by the center of mass of the initial shape and the center of mass of the deformed shape:

$$\bar{\mathbf{c}} = \frac{1}{M} \sum_i m_i \bar{\mathbf{x}}_i, \quad \mathbf{c} = \frac{1}{M} \sum_i m_i \mathbf{x}_i, \quad M = \sum_i m_i. \quad (14)$$

If we minimize the term  $\sum_i (\mathbf{A}\bar{\mathbf{r}}_i - \mathbf{r}_i)^2$  with  $\mathbf{r}_i = \mathbf{x}_i - \mathbf{c}$  and  $\bar{\mathbf{r}}_i = \bar{\mathbf{x}}_i - \bar{\mathbf{c}}$ , we get the optimal linear transformation  $\mathbf{A}$  of the initial and the deformed shape. This transformation is determined by:

$$\mathbf{A} = \left( \sum_i m_i \mathbf{r}_i \mathbf{r}_i^T \right) \left( \sum_i m_i \bar{\mathbf{r}}_i \bar{\mathbf{r}}_i^T \right)^{-1} = \mathbf{A}_r \mathbf{A}_s. \quad (15)$$



**Figure 13:** Robust and volume-conserving deformations using shape matching. Armadillos (32442 particles total), 20 ducks and 20 tori (21280 particles total) and 20 balls (7640 particles total) were simulated in real-time on a GPU.

In our case we are only interested in the rotational part of this transformation. Since  $\mathbf{A}_s$  is symmetric, it contains no rotation. Therefore, we only need to extract the rotational part of  $\mathbf{A}_r$  to get the optimal rotation  $\mathbf{R}$  for shape matching. This can be done by a polar decomposition  $\mathbf{A}_r = \mathbf{R}\mathbf{S}$  of the transformation matrix where  $\mathbf{S}$  is a symmetric matrix.

Finally, the goal positions are determined by

$$\mathbf{g}_i = \mathbf{T} \begin{bmatrix} \bar{\mathbf{x}}_i \\ 1 \end{bmatrix}$$

where  $\mathbf{T} = [\mathbf{R} \quad (\mathbf{c} - \mathbf{R}\bar{\mathbf{c}})]$ . These goal positions are used to define the following modified integration scheme:


$$\begin{aligned} \mathbf{v}_i(t + \Delta t) &= \mathbf{v}_i(t) + \alpha \frac{\mathbf{g}_i(t) - \mathbf{x}_i(t)}{\Delta t} + \Delta t \frac{1}{m_i} \mathbf{f}_{\text{ext}}(t) \\ \mathbf{x}_i(t + \Delta t) &= \mathbf{x}_i(t) + \Delta t \mathbf{v}_i(t + \Delta t) \end{aligned}$$

where  $\alpha \in [0, 1]$  is a user defined stiffness parameter which defines how far the particles are pulled to their goal positions.

#### 4.2. Region-Based Shape Matching

The shape matching algorithm described above allows only for small deviations from the initial shape. For the simulation of large deformations the concept of region-based shape matching became popular, see e.g. [MHTG05, RJ07, DBB11]. The idea is to perform shape matching on several overlapping regions of the original shape. In each region we can have a small deviation from the corresponding part of the initial shape which results in a large deformation over all regions.

Diziol et al. [DBB11] propose to define a region for each particle of the model where the  $i$ -th region contains all particles in the  $\omega$ -ring of the  $i$ -th particle in the original mesh of the model. Shape matching is a meshless method but Diziol et al. require a mesh to define the shape matching regions.



**Figure 15:** The stiffness of the model depends on the region size. Smaller regions (top) allow larger deformations than larger regions (bottom). The hexagons in the left images represent the overlapping regions of the model. The right images show the goal positions after one particle is moved away.

Rivers and James [RJ07] use a regular lattice instead to define their regions. No matter which kind of regions are used, the stiffness of the model depends on the size of the overlapping regions (see Figure 15). Enlarging the regions results in a more global shape matching and therefore the stiffness of the simulated model is increased.

In region-based shape matching a particle is part of multiple regions. In the following we denote the set of regions to which a particle  $i$  belongs by  $\mathfrak{R}_i$ . Since particles can belong to more than one region, Rivers and James [RJ07] proposed to use modified particle masses  $\tilde{m}_i = m_i/|\mathfrak{R}_i|$  for shape matching. This ensures that a particle which is part of many regions has not more influence than others. The optimal translation vectors for a region  $i$  are determined by

$$\bar{\mathbf{c}}_i = \frac{1}{\tilde{M}_i} \sum_{j \in \mathfrak{R}_i} \tilde{m}_j \bar{\mathbf{x}}_j, \quad \mathbf{c}_i = \frac{1}{\tilde{M}_i} \sum_{j \in \mathfrak{R}_i} \tilde{m}_j \mathbf{x}_j \quad (16)$$

where  $\tilde{M}_i = \sum_{j \in \mathcal{R}_i} \tilde{m}_j$  is the effective region mass which can be precomputed. The optimal rotation matrix  $\mathbf{R}$  is computed by extracting the rotational part of the following matrix:

$$\mathbf{A}_{r,i} = \sum_{j \in \mathcal{R}_i} \tilde{m}_j \mathbf{x}_j \bar{\mathbf{x}}_j^T - \tilde{M}_i \mathbf{c}_i \bar{\mathbf{c}}_i^T. \quad (17)$$

In this form the first term depends on the particles  $j$  of the region while the second term depends on the region  $i$  if the translation vectors are already available. This isolation of the dependencies is required for fast summation techniques (see below).

After performing shape matching for all regions, we get multiple goal positions for each particle. The final goal position for a particle is determined by blending the goal positions of the corresponding regions:

$$\mathbf{g}_i = \frac{1}{|\mathcal{R}_i|} \sum_{j \in \mathcal{R}_i} \mathbf{T}_j \begin{bmatrix} \bar{\mathbf{x}}_j \\ 1 \end{bmatrix}.$$

### 4.3. Fast Summation Techniques

In the case of region-based shape matching the stiffness increases with growing region size  $\omega$ . However, at the same time the computation of the optimal translation  $\mathbf{c}$  and the transformation matrix  $\mathbf{A}_r$  becomes a bottleneck since large sums have to be computed for each region. For a mesh with the dimension  $d$  and  $n$  regions,  $O(\omega^d n)$  operations are required with the naive approach.

Rivers and James demonstrated in [RJ07] how the number of operations for computing the sums can be reduced to  $O(n)$  for regular lattices ( $d = 3$ ). In their approach they compute the summation for a set of particles just once and reuse it for all regions that contain this set. This reduces redundant computations significantly for a system with large overlapping regions. The fast summation of Rivers and James is based on the usage of cubical regions. These cubical regions can be subdivided in two-dimensional plate regions which can again be subdivided in one-dimensional bar regions. The region summation is performed in three passes. In the first pass the sum for each bar is determined. The results are used to compute the sums for the plates which are again used to obtain the final region sum. Each pass requires  $O(\omega)$  operations. However, the region sum can even be determined in constant time if we take into account that the sum of two neighboring bars, plates or cubes only differs by one element. Lattice shape matching can be performed in linear time if the sums in Equations (16) and (17) are evaluated using the fast summation technique described above.

The FastLSM method of Rivers and James has several limitations. To handle regions where the lattice is not regular, e.g. on the boundary, several sums are defined in a pre-processing step for the corresponding node. In the case of fracturing the definition of these sums must be performed at run-time which is expensive to compute. Small features

need a fine sampling to obtain realistic results. Since a regular lattice is used, a fine sampling yields an explosion of the computational costs. FastLSM does not support a varying region size to simulate inhomogeneous material.


Steinemann et al. [SOG08] introduce an adaptive shape matching method which is based on lattice shape matching to overcome these limitations. A fast summation is realized by an octree-based sampling and an interval-based definition of the shape matching regions. The hierarchical simulation model is created by starting with a coarse cubic lattice and then performing an octree subdivision. The subdivision process can be controlled by a user-defined criterion. At the end of the process a simulation node is placed at the center of each leaf cell and a virtual node at the center of each non-leaf cell. A virtual node stores the sum of all its descendant simulation nodes.

The fast summation for the hierarchical model is performed by an interval-based method which requires  $O(1)$  operations per region. For each simulation node  $n_i$  a shape matching region is defined by a region width  $\omega_i$ . To perform a fast summation, all summation nodes of a region are determined in a pre-processing step by a top-down traversal. During the traversal each node  $n_j$  where all descendant leaf nodes are closer than  $\omega_i$  is added to region  $i$ . If the descendant leaf nodes are contained only partially in region  $i$ , the current node must be refined. Only in this case the traversal continues.

The top-down traversal assigns  $O(1)$  summation nodes to each region. A fast summation can now be performed in two steps. In the first step the sums of all nodes in the hierarchy are determined. This is done by first computing the sums for the simulation nodes which are the leaf nodes of the hierarchy, and then updating the sums of the virtual nodes in a bottom-up fashion. The second step sums up the values of the summation nodes for each region. For a roughly balanced octree the computation of the sums takes  $O(n)$  time where  $n$  is the number of simulation nodes. Hence, the adaptive shape matching method requires linear time when using the described fast summation technique to evaluate Equations (16) and (17).

In contrast to Rivers and James, Dziol et al. [DBB11] only use the surface mesh of a volumetric model to simulate its deformation. Therefore, no interior elements are required for the simulation which reduces the computational costs. Dziol et al. introduce a fast summation technique for arbitrary triangle meshes ( $d = 2$ ) to compute the large sums of the region-based approach efficiently. This technique only requires  $O(\omega n)$  operations instead of  $O(\omega^2 n)$  and can be performed very efficiently in parallel.

The fast summation technique of Dziol et al. is based on a subdivision of all particles of the mesh in disjoint paths. A path  $i$  is a set of vertices  $\mathbf{x}_{i_1}, \dots, \mathbf{x}_{i_n}$  which are connected by edges. The paths are determined in a precomputation step.



**Figure 16:** Fast summation technique for arbitrary triangle meshes [DBB11]. First the prefix sums for the disjoint paths are determined. Then the region sum is computed by adding the difference of the intersection interval for each path.

The goal of the path construction algorithm is that each region is intersected by a minimum number of paths. To determine the optimal path layout is computational expensive. Therefore, a heuristic is used to find a good path layout.

The fast summation is split in two phases (see Figure 16). In the first phase the prefix sum for each path  $i$  is computed with  $j \in [1, n_i]$ :

$$\mathbf{c}_{i_j}^p = \sum_{k=1}^j \tilde{m}_{i_k} \mathbf{x}_{i_k}, \quad \mathbf{A}_{i_j}^p = \sum_{k=1}^j \tilde{m}_{i_k} \mathbf{x}_{i_k} \tilde{\mathbf{x}}_{i_k}^T.$$

Since the prefix sums for all paths are independent from each other, they can be computed in parallel. The sums for a region  $r$  are computed by first setting  $\mathbf{c}_r := \mathbf{0}$  and  $\mathbf{A}_r := \mathbf{0}$ . Then for each path  $i$  which intersects the region in the interval  $[i_k, \dots, i_l]$ , the following terms are added:

$$\mathbf{c}_r := \mathbf{c}_r + \mathbf{c}_{i_l}^p - \mathbf{c}_{i_{k-1}}^p, \quad \mathbf{A}_r := \mathbf{A}_r + \mathbf{A}_{i_l}^p - \mathbf{A}_{i_{k-1}}^p. \quad (18)$$

The final translational vector and the affine matrix are determined by  $\mathbf{c}_r := (1/\tilde{M}_r)\mathbf{c}_r$  and  $\mathbf{A}_r := \mathbf{A}_r - \tilde{M}_r \mathbf{c}_r \tilde{\mathbf{c}}_r^T$  respectively.

#### 4.4. Oriented Particles

For a small number of particles or particles that are close to co-linear or co-planar (as in Figure 17), the matrix  $\mathbf{A}_r$  in Equation (15) becomes ill-conditioned and the polar decomposition needed to obtain the optimal rotation tends to be numerically unstable.

To solve this problem, Müller et al. [MC11b] proposed to use oriented particles. By adding orientation information to particles, the polar decomposition becomes stable even for single particles. The moment matrix of a single spherical particle with orientation  $\mathbf{R} \in \mathbb{R}^{3 \times 3}$  and finite radius  $r$  at the origin is well defined and can be computed via an integral

over its volume as

$$\begin{aligned} \mathbf{A}_{\text{sphere}} &= \int_{V_r} \rho(\mathbf{R}\mathbf{x}) \mathbf{x}^T dV = \rho \mathbf{R} \int_{V_r} \mathbf{x} \mathbf{x}^T dV \quad (19) \\ &= \frac{4}{15} \pi r^5 \rho \mathbf{R} = \frac{4}{15} \pi r^5 \frac{m}{V_r} \mathbf{R} \\ &= \frac{1}{5} m r^2 \mathbf{R}, \end{aligned}$$

where  $V_r$  is the volume of a sphere of radius  $r$ . Since  $\mathbf{R}$  is an orthonormal matrix,  $\mathbf{A}_i$  always has full rank and an optimal condition number of 1. For an ellipsoid with radii  $a, b$  and  $c$  we get

$$\mathbf{A}_{\text{ellipsoid}} = \frac{1}{5} m \begin{bmatrix} a^2 & 0 & 0 \\ 0 & b^2 & 0 \\ 0 & 0 & c^2 \end{bmatrix} \mathbf{R}. \quad (20)$$

However, the moment matrices of the individual particles cannot simply be added because each one is computed relative to the origin. We need the moment matrix of particle  $i$  relative to the position  $\mathbf{x}_i - \mathbf{c}$ .

Fortunately, this problem can be fixed easily. As we saw above, the equation for computing the moment matrix

$$\mathbf{A} = \sum_i m_i (\mathbf{x}_i - \mathbf{c})(\bar{\mathbf{x}}_i - \bar{\mathbf{c}})^T, \quad (21)$$

can be re-written as

$$\mathbf{A} = \sum_i m_i \mathbf{x}_i \bar{\mathbf{x}}_i^T - M \mathbf{c} \bar{\mathbf{c}}^T, \quad (22)$$

where  $\bar{\mathbf{c}}$  and  $\mathbf{c}$  are the centers of mass of the initial and the deformed shape, respectively (see Equation (14)).

Therefore, shifting the evaluation from the origin to the position  $\mathbf{x}_i - \mathbf{c}$  yields

$$\mathbf{A}_i^{\text{global}} = \mathbf{A}_i + m_i \mathbf{x}_i \bar{\mathbf{x}}_i^T - m_i \mathbf{c} \bar{\mathbf{c}}^T. \quad (23)$$

Equation (21) now generalizes to

$$\mathbf{A} = \sum_i \left( \mathbf{A}_i + m_i \mathbf{x}_i \bar{\mathbf{x}}_i^T \right) - M \mathbf{c} \bar{\mathbf{c}}^T \quad (24)$$

$$= \sum_i \left( \mathbf{A}_i + m_i (\mathbf{x}_i - \mathbf{c})(\bar{\mathbf{x}}_i - \bar{\mathbf{c}})^T \right). \quad (25)$$

As you can see, the last form looks like Equation (21) but with all the individual particle moment matrices added in the sum.

In addition to position  $\mathbf{x}$  and velocity  $\mathbf{v}$ , oriented particles carry a rotation which can be defined as an orthonormal matrix  $\mathbf{R}$  as above or a unit quaternion  $\mathbf{q}$ . They also carry the angular velocity  $\boldsymbol{\omega}$ . In the prediction step of position-based dynamics, these two quantities have to be integrated as well:

$$\mathbf{x}_p \leftarrow \mathbf{x} + \mathbf{v} \Delta t \quad (26)$$

$$\mathbf{q}_p \leftarrow \left[ \frac{\boldsymbol{\omega}}{|\boldsymbol{\omega}|} \sin\left(\frac{|\boldsymbol{\omega}| \Delta t}{2}\right), \cos\left(\frac{|\boldsymbol{\omega}| \Delta t}{2}\right) \right] \mathbf{q}, \quad (27)$$

For stability reasons,  $\mathbf{q}_p$  should directly be set to  $\mathbf{q}$  if  $|\boldsymbol{\omega}| < \epsilon$ .

After the prediction step, the solver iterates multiple times through all shape match constraints in a Gauss-Seidel type fashion as before. To simulate objects represented by a mesh of linked particles, [MC11b] define one shape matching group per particle. A group contains the corresponding particle and all the particles connected to it via a single edge. The positions of the particles in a group are updated as in regular shape matching by pulling them towards the goal positions while the orientation of the center particle only is replaced by the optimal rotation of shape matching.

After the solver has modified the predicted state  $(\mathbf{x}_p, \mathbf{q}_p)$ , the current state is updated using the integration scheme

$$\mathbf{v} \leftarrow (\mathbf{x}_p - \mathbf{x}) / \Delta t \quad (28)$$

$$\mathbf{x} \leftarrow \mathbf{x}_p \quad (29)$$

$$\boldsymbol{\omega} \leftarrow \text{axis}(\mathbf{q}_p \mathbf{q}^{-1}) \cdot \text{angle}(\mathbf{q}_p \mathbf{q}^{-1}) / \Delta t \quad (30)$$

$$\mathbf{q} \leftarrow \mathbf{q}_p, \quad (31)$$

where  $\text{axis}()$  returns the normalized direction of a quaternion and  $\text{angle}()$  its angle. Again, for stability reasons,  $\boldsymbol{\omega}$  should be set to zero directly if  $|\text{angle}(\mathbf{q}_p \mathbf{q}^{-1})| < \epsilon$ . There are two rotations,  $\mathbf{r} = \mathbf{q}_p \mathbf{q}^{-1}$  and  $-\mathbf{r}$  transforming  $\mathbf{q}$  into  $\mathbf{q}_p$ . It is important to always choose the shorter one, i.e. if  $r.w < 0$  use  $-\mathbf{r}$ . As in traditional PBD for translation, changing the rotational quantity  $\mathbf{q}_p$  in the solver also affects its time derivative  $\boldsymbol{\omega}$  through the integration step creating the required second order effect.

The orientation information of particles cannot only be used to stabilize shape matching but also to move a visual mesh along with the physical mesh. With position and orientation, each particle defines a full rigid transformation at every point in time. This allows the use of traditional linear blend skinning with particles replacing skeletal bones.

An additional advantage of having orientation information is that ellipsoids can be used as collision volumes for particles. This allows a more accurate approximation of the object geometry than with the same number of spherical primitives (see Figure 18).

#### 4.5. Extensions

There exist several extensions for shape matching. In the following we will introduce volume conservation and plastic deformation.


##### 4.5.1. Volume Conservation

The conservation of volume plays an important role in the dynamic simulation of deformable bodies. Since most soft biological tissues are incompressible, this is an essential extension in the field of medical simulation. However, it is also used in the field of shape modeling [vFTS06] since volume conserving deformations appear more realistic.

In the following we introduce the position-based approach



**Figure 17:** This underwater scene demonstrates the ability of the oriented particle approach to handle sparse meshes such as the one-dimensional branches of the plants or the fins of the lion fish.



**Figure 18:** The rotation information of oriented particles cannot only be used to stabilize shape matching, it also allows the use of ellipsoids as collision primitives. The figure shows how the same mesh is approximated much more accurately with ellipsoids (right) than with the same number of spheres (left).

for volume conservation of Diziol et al. [DBB11]. This method considers only the surface of a simulated object and does not require interior particles which reduces the computational effort. The volume  $V$  of a volumetric 3D shape  $\mathcal{V}$  can be determined by using the divergence theorem as proposed by [HJCW06]:

$$\iiint_{\mathcal{V}} \nabla \cdot \mathbf{x} \, d\mathbf{x} = \iint_{\partial\mathcal{V}} \mathbf{x}^T \mathbf{n} \, d\mathbf{x} = 3V \quad (32)$$

where  $\partial\mathcal{V}$  is the boundary of the shape and  $\mathbf{n}$  is the surface normal. If the boundary is given as triangle mesh, the inte-



gral can be written as sum over all triangles  $i$ :

$$V(\mathbf{X}) := \frac{1}{3} \iint_{\partial \mathcal{V}} \mathbf{x}^T \mathbf{n} d\mathbf{x} = \frac{1}{9} \sum_i A_i (\mathbf{x}_{i_1} + \mathbf{x}_{i_2} + \mathbf{x}_{i_3})^T \mathbf{n}_i \quad (33)$$

where  $A_i$  is the area and  $i_1$ ,  $i_2$  and  $i_3$  are the vertex indices of the  $i$ -th triangle. Now we can define a volume constraint  $C := V(\mathbf{X}) - V(\bar{\mathbf{X}}) = 0$  and compute a position correction (see Section 3):

$$\Delta \mathbf{x}_i = - \frac{w_i C(\mathbf{X})}{\sum_j w_j \|\nabla_{\mathbf{x}_j} C(\mathbf{X})\|^2} \nabla_{\mathbf{x}_i} C(\mathbf{X}). \quad (34)$$

The weights  $w_i$  are used to realize a local volume conservation (see below). The gradient can be approximated by

$$\nabla C(\mathbf{X}) \approx \frac{1}{3} [\bar{\mathbf{n}}_1^T, \dots, \bar{\mathbf{n}}_n^T]^T$$

where  $\bar{\mathbf{n}}_i = \sum A_j \mathbf{n}_j$  is the sum of the area weighted normals of all triangles which contain particle  $i$ .

The weights in Equation 34 are chosen as follows:

$$w_i = (1 - \alpha) w_i^l + \alpha w_i^g, \quad w_i^l = \frac{\|\Delta \mathbf{x}_i\|}{\sum_j \|\Delta \mathbf{x}_j\|}, \quad w_i^g = \frac{1}{n}$$

where  $w_i^l$  and  $w_i^g$  are the weights for local and global volume conservation, respectively, and the user-defined value  $\alpha \in [0, 1]$  is used to blend between both. The vector  $\Delta \mathbf{x}_i$  contains the position change of the  $i$ -th particle in the shape matching step. Hence, strongly deformed particles participate more in volume correction. The weight of a colliding particle is set to zero in order to ensure that a collision constraint is not violated during the position correction for the volume conservation. Finally, the weights are smoothed by a Laplacian filter.


Diziol et al. also propose another definition for the local weights  $w_i^l$ . To propagate volume changes through the object, they first determine pairs of opposing particles in a pre-processing step by intersecting the geometry with multiple rays. Then they choose a local weight which does not only depend on the position change of a particle but also on the distance change of the corresponding particle pairs.

Analogous to the positions correction we perform a velocity correction to fulfill the constraint  $\partial C / \partial t = 0$ . This leads to a divergence free velocity field.

In Figure 19 different configurations for the presented volume conservation method are compared with each other.

#### 4.5.2. Plastic Deformation

Shape matching can be extended in order to simulate plastic deformations [MHTG05]. If we perform a polar decomposition  $\mathbf{A}_r = \mathbf{R}\mathbf{S}$  for the linear transformation matrix  $\mathbf{A}_r$  (see Equation 15), we get a rotational part  $\mathbf{R}$  and a symmetric part  $\mathbf{S} = \mathbf{R}^T \mathbf{A}_r$ . The matrix  $\mathbf{S}$  represents a deformation in the unrotated reference frame. Hence, for each region we can store the plastic deformation state in a matrix  $\mathbf{S}^p$  which is initialized with the identity matrix  $\mathbf{I}$ . As proposed by O'Brien



**Figure 19:** Four spheres with different volume conservation squeezed by a plate. Left to right: global conservation, local conservation with distance constraints, local conservation without distance constraints and no volume conservation. The maximum volume loss was 0.6%, 0.7%, 0.7% and 40% respectively.

et al. [OBH02], we use two parameters  $c_{\text{yield}}$  and  $c_{\text{creep}}$  to control the plastic behavior of the material. If the condition  $\|\mathbf{S} - \mathbf{I}\|_2 > c_{\text{yield}}$  is fulfilled for the deformation matrix  $\mathbf{S}$  of the current time step, the plastic deformation state is updated as follows:

$$\mathbf{S}^p \leftarrow [\mathbf{I} + \Delta t c_{\text{creep}} (\mathbf{S} - \mathbf{I})] \mathbf{S}^p.$$

After this update,  $\mathbf{S}^p$  is divided by  $\sqrt[3]{\det(\mathbf{S}^p)}$  in order to conserve the volume. The plastic state  $\mathbf{S}^p$  is integrated in the shape matching process by deforming the reference shape in Equation 15. This is done by replacing the definition of  $\bar{\mathbf{r}}_i$  (see Section 4.1) with

$$\bar{\mathbf{r}}_i = \mathbf{S}^p (\bar{\mathbf{x}}_i - \bar{\mathbf{c}}).$$

Note that the plasticity can be bound by the condition  $\|\mathbf{S}^p - \mathbf{I}\|_2 > c_{\text{max}}$  where  $c_{\text{max}}$  is the threshold for the maximum plastic deformation. If this condition is fulfilled, we use  $\mathbf{S}^p \leftarrow \mathbf{I} + c_{\text{max}} (\mathbf{S}^p - \mathbf{I}) / \|\mathbf{S}^p - \mathbf{I}\|_2$ .

#### 4.6. Cloth Simulation

Stumpp et al. [SSBT08] presented a region-based shape matching approach for the simulation of cloth. In their work they define a region for each triangle in the model. But instead of using the triangles directly as regions for shape matching, overlapping regions are defined. The region of a triangle is defined by the outer corners of its adjacent tri-



angles. These overlapping regions enable the bending resistance of the cloth model.

To perform shape matching the optimal translation vectors of the regions are determined by evaluating Equation 16. Since the triangular regions are two-dimensional, the optimal rotation can be extracted by a 2D polar decomposition [SD92]. For each region the current positions of the vertices  $\mathbf{x}$  and their rest state positions  $\bar{\mathbf{x}}$  are projected in the two-dimensional space of the corresponding triangle. Then a 2D version of the matrix  $\mathbf{A}_r$  is determined by evaluating Equation 17 for the projected position vectors. By a polar decomposition of the resulting matrix the 2D rotation matrix is obtained. This rotation matrix is used to compute 2D goal positions for the particles. Finally, these goal positions are transformed to world space and the particles are pulled towards the resulting positions.

The size of the shape matching regions directly influences the stiffness of the model. Since the simulation model of Stump et al. has triangular regions, the stiffness of high resolution models is too low for realistic results. Therefore, they introduce so-called fiber clusters which are used to increase the stretching stiffness. These one-dimensional regions are determined in a pre-processing step by subdividing the mesh into multiple edge strips with an angle of approximately 180 degree between adjacent edges. During the simulation each strip is traversed in both directions to obtain goal positions for the particles. In each step of a traversal the next particle in the strip is pulled towards the current one in order to restore the initial length of the corresponding edge. After performing both traversals the resulting goal positions are averaged. Each goal position defines a displacement. These displacements are translated so that they sum up to 0 to preserve the momentum of the model. The final goal positions are blended with the goal positions of the triangular regions. The usage of fiber clusters is not physically correct but the results are visually plausible and the increased stiffness yields more realistic results.

#### 4.7. Parallelization

In Section 4.3 we presented different fast summation techniques for shape matching. The one of Diziol et al. [DBB11] is best suited for a parallel implementation on the GPU. In the following the GPU implementation of this technique with CUDA is described in detail. For such an implementation memory access and memory layouts play an important role as well as the number of kernel calls.

Since each kernel call introduces a computational overhead, the particles of all objects in a simulation are packed into one single array. This array is ordered according to the path layout which is used for the fast summation (see Section 4.3). Since the array contains the paths one after another, a segmented prefix sum [SHZO07] can be used to determine the prefix sums of all paths at once. To avoid numerical problems due to the 32 bit floating-point arithmetics

on the GPU, the path length is limited to 512. The resulting prefix sums are stored in texture memory to benefit from the texture cache when the translational vectors and the affine matrices are determined (see Equation 18).

The volume conservation introduced in Section 4.5.1 can also be performed efficiently on the GPU. This is done by evaluating the volume integral (see Equation 32), the integral which is required to obtain a divergence free velocity field and the weights for the local volume conservation in parallel. Both integrals can be written as sums over the vertices (see Equation 33). Hence, the integrals as well as the weights can be computed by a segmented sum reduction. Finally, the smoothing of the weights by a Laplacian filter can be performed in parallel using the fast summation technique as described above.

## 5. Geometric Data-Driven Methods

The behavior of solid objects can be accurately described using well-known mechanical models, but real-world objects display other inherent sources of complexity that largely limit the results of traditional models in computer animation. Complexity is produced, for example, by nonlinear or anisotropic behaviors, by heterogeneous properties, or by a high dynamic range. These sources of complexity are typically addressed by designing complex nonlinear constitutive models to describe the mechanical behavior of solid objects. However, these models require computationally expensive simulation algorithms, and their parameters are difficult and tedious to tune, particularly if the properties are heterogeneous. All in all, the animation of solid objects is limited by the domain of effects captured by the underlying physical models, but also by their parameterization accuracy.


Data-driven methods offer an alternative to complex constitutive models, as they turn the modeling metaphor into the knowledge of a system's response under several example conditions. This section describes geometric data-driven methods in computer animation. It formulates a two-scale representation of geometry and dynamics, describes the computation of detailed geometry as the interpolation of example data, and discusses several successful examples.

### 5.1. Two-Scale Geometry and Dynamics

Let us consider the surface of a solid object (e.g., the cloth in Figure 20), with vertex positions  $\mathbf{x} \in \mathbb{R}^3$ . These positions can be decomposed into a low-resolution position  $\mathbf{x}_0$  and a fine-scale displacement  $\Delta\mathbf{x}$ , expressed in a local reference system for each vertex (i.e., with orientation  $\mathbf{R}$ ):

$$\mathbf{x} = \mathbf{x}_0 + \mathbf{R}\Delta\mathbf{x}. \quad (35)$$

This definition of vertex positions essentially decomposes large-scale geometry (i.e., the overall shape of the object) from the small-scale deformation (i.e., wrinkles).



**Figure 20:** Data-driven animation of cloth wrinkles. On the left, a low-resolution feature mesh, a smooth mesh obtained by subdivision, and a high-resolution detail mesh. On the right, a schematic depiction of the data-driven wrinkle computation. The feature mesh is subdivided to obtain the smooth mesh, and then detailed wrinkles are added in a local reference frame by combining detail from examples.

The large-scale and fine-scale geometry can be represented at different resolutions, connected through subdivision schemes. As shown in Figure 20 for a cloth object, a low-resolution *feature mesh* defines the large-scale deformation. A high-resolution *smooth mesh* is obtained by subdividing the feature mesh a user-defined number of times, and can be regarded as an upsampled version of the feature mesh. Finally, a high-resolution *detail mesh* is obtained by adding local displacements onto the smooth mesh. The smooth and detail meshes fully share the connectivity, thus trivially defining their correspondence.

Once large-scale and fine-scale geometry are separated, they can be computed using different models. The choice of models can be made based on the following observations. First, for many objects such as the face or cloth, the most salient dynamic effects can be captured at a large scale. For example, Bickel et al. [BLB\*08] and Ma et al. [MJC\*08] compute the large-scale face geometry from an actor’s performance, using a single face scan and sparse mocap markers as input, and a linear deformation model. Wang et al. [WRO11] compute the large-scale geometry of tight cloth as a function of a character’s pose. Zurdo et al. [ZBO12] and Kavan et al. [KGBS11] compute the large-scale geometry of loose cloth using a low-resolution dynamics model with contact handling.

A second important observation is that plausible high-resolution wrinkles can often be defined as a quasi-static function in a reduced low-resolution domain  $\mathbf{u}$ . For the face, and due to the repetitive nature of facial expressions, tissue becomes weaker at certain locations, and expressive wrin-

kles appear in a deterministic fashion as a function of large-scale deformation, and indirectly as a function of muscle activations and facial bone configurations. For cloth, even though real wrinkles require a large number of degrees of freedom to capture their true diversity, plausible wrinkles can be defined in the reduced domain of large-scale deformation. We can formally write the dependency between high-resolution detail  $\Delta\mathbf{x}$  and the low-resolution configuration  $\mathbf{u}$  as a position-based model  $\Delta\mathbf{x} = f(\mathbf{u})$ .

The reduced-domain definition presents some limitations, which should be mentioned upfront. Fine-scale wrinkle dynamics cannot be captured, as wrinkles are defined quasi-statically. And the model captures only a limited set out of all the possible wrinkles that a solid object might present. However, the power of data-driven methods is that the generic function  $f$  makes use of data from real deformation examples; therefore, data-driven wrinkles preserve natural characteristics such as length, width, and consistency over time, and they appear plausible despite their limitations.

## 5.2. Data-Driven Geometric Detail

At this point, we have a suitable setting to define a position-based data-driven model. For data collection, we need to record example deformations (denoted by the subscript  $i$ ), with vertex displacements  $\{\Delta\mathbf{x}_i\}$  and low-resolution configurations  $\{\mathbf{u}_i\}$  in correspondence. Then, we apply learning methods to design a data-driven approximation of the function  $f$ , which can be formally defined as  $\Delta\mathbf{x} \approx \hat{f}(\mathbf{u}, \{\Delta\mathbf{x}_i\}, \{\mathbf{u}_i\})$ .

A general and successful approach to define a data-driven approximation  $\hat{f}$  is through a linear combination of example-based basis functions,

$$\Delta \mathbf{x} = \sum_j w_j(\mathbf{u}) \mathbf{b}_j, \quad (36)$$

with weights  $w_j$  computed as a function of the low-resolution configuration  $\mathbf{u}$ . In the expression above,  $\mathbf{b}_j \in \mathbb{R}^3$  constitutes an example-based detail displacement, and represents the values associated with one vertex in the  $j^{\text{th}}$  basis function.

In the rest of this section, we discuss several examples of position-based data-driven deformation models, which differ in terms of the choice of low-resolution configuration, basis functions, or interpolation method.

### 5.2.1. Weighted Pose-Space Deformation

Bickel et al. [BLB\*08] proposed a data-driven method to compute facial wrinkles as a function of large-scale deformation and a small set of example deformations, and later Zurdo et al. [ZBO12] followed a similar strategy for cloth animation. Figure 21 shows an example facial animation by Bickel et al., and Figure 22 an example cloth animation by Zurdo et al.

In their methods, the basis function combination in Equation (36) follows the weighted pose-space deformation (WPSD) approach [KM04]. In a nutshell, the basis functions constitute local vertex displacements in example deformations, called *poses*, and the weights  $w$  are convex weights for each of the poses. The weights of the poses are computed using scattered-data interpolation based on *radial basis functions* (RBFs) in *pose space*. In this case, the pose space is given by the large-scale deformation  $\mathbf{u}$ .

As shown by Zurdo et al., pose weights can be computed through WPSD on the sparse vertices of the feature mesh, and then simply interpolated to the vertices of the detail mesh using subdivision weights. For the computation of pose weights on a given vertex, WPSD requires a local definition of the pose space  $\mathbf{u}$ . Both Bickel et al. and Zurdo et al. use a metric of local low-resolution strain, formed by concatenating the deformation of the 16 closest edges of the feature mesh, multiplied by fast-decaying weights. With RBF interpolation, the weight of each pose is computed as

$$w_j = \sum_i \omega_{i,j} \phi(\|\mathbf{u} - \mathbf{u}_i\|). \quad (37)$$

The RBF weights  $\omega_{k,j}$  are precomputed such that pose weights fulfill the Kronecker delta for the database of poses, i.e.,  $w_j(\mathbf{u}_i)$  is 0 if  $j \neq i$  and 1 if  $j = i$ , for poses  $j$  and  $i$  in the database. The function  $\phi$  represents a specific type of RBF. Both Bickel et al. and Zurdo et al. use RBFs with global support  $\phi(r) = r$ , as they avoid complex tuning of support radii for unevenly sampled data [CBC\*01].

The selection of poses starts with the generation of training data, which requires a database of synchronized feature and detail meshes. Zurdo et al. employ the TRACKS method [BMWG07] to precompute cloth simulations where the detail mesh tracks the motion of the feature mesh. In the training simulations, contact is solved both on the feature mesh and the detail mesh, and then the high-resolution information in the poses captures the response to contact. From all mesh pairs in the database, Zurdo et al. select only a small number of poses (typically 6), until the  $L^2$  error between synthesized animations and the training dataset is below a certain threshold. The set of poses is grown in a greedy manner, adding each time the mesh pair with largest  $L^2$  error between the training and synthesized configurations.

### 5.2.2. Polynomial Displacement Maps

Ma et al. [MJC\*08] designed Polynomial Displacement Maps (PDM), a method for data-driven computation of wrinkles in facial animation, following the computational strategy of Polynomial Texture Maps [MGW01]. Despite the computational differences, the method shares many similarities with the WPSD approach of Bickel et al. [BLB\*08].

In the general data-driven framework described by Equation (36), the low-resolution configuration of PDM is defined as a two-dimensional local strain metric of the feature mesh,  $\mathbf{u} = (u_1, u_2)$ . In particular, this metric is obtained through principal component analysis of a five-dimensional vector formed by the local vertex offset and the in-plane strain of the feature mesh. Based on this low-resolution configuration, the weights in Equation (36) are defined based on biquadratic polynomials,  $\mathbf{w}(\mathbf{u}) = (u_1^2 u_2^2 u_1 u_2 u_1 u_2 1)$ . Finally, the detail position of each vertex in the detail mesh,  $\Delta \mathbf{x}$ , is defined as a scalar displacement in the local normal direction.

Due to the choice of biquadratic PDMs, the per-vertex displacements are computed through linear combination of six basis functions, which are stored in texture maps. Ma et al. compute the coefficients in these texture maps as the result of a least-squares problem that minimizes the fitting error over the training data.

### 5.2.3. Physics-Inspired Upsampling

Kavan et al. [KGBS11] proposed a data-driven method to compute cloth animations with detailed wrinkles at videogame frame rates. In their approach, high-resolution cloth positions are computed following a position-based data-driven method, but the data is learned from dynamic simulations with tracking of high-resolution and low-resolution geometry, similar to the approach by Zurdo et al. [ZBO12] described earlier.

In physics-inspired upsampling (PIU), the weights  $w$  in Equation (36) correspond directly to the vertex positions of the feature mesh. The basis functions coefficients  $\mathbf{b}$  extend subdivision schemes, and are learned from a database of example deformations.



**Figure 21:** Example of facial animation with the data-driven method of Bickel et al. [BLB\*08]. From left to right: large-scale deformation interpolating mocap markers, full result data-driven detail computation, the same result with full shading, and comparison to the real actor’s face.



**Figure 22:** Top row: a low-resolution model of a dress, with only 381 triangles, defines the dynamics, the large-scale deformation, and response to contact. The dress is then subdivided to 24384 triangles for rendering. Bottom row: high-resolution wrinkles are interpolated from 6 example poses based on the large-scale deformation of the dress, using the data-driven method of Zurdo et al. [ZBO12].

For the description of the method in more detail, let us define the low-resolution deformation  $\mathbf{u}$  as a vector that concatenates the vertex positions of the low-resolution feature mesh. We also define a matrix  $\mathbf{B}$  of upsampling basis functions, being each row of  $\mathbf{B}$  the basis function for one vertex coordinate. Finally, we define a vector  $\mathbf{X}$  that concatenates all vertex positions of the high-resolution detail mesh. Note that the method of Kavan et al. computes high-resolution positions directly, not local displacements as discussed earlier. With these definitions, the data-driven detail from Equation (36) can be rewritten as  $\mathbf{X} = \mathbf{B}\mathbf{u}$ .

Given the data captured in training simulations, expressed as example deformations of the feature mesh,  $\{\mathbf{u}_i\}$ , and the detail mesh,  $\{\mathbf{X}_i\}$ , Kavan et al. compute basis functions

through an optimization problem

$$\mathbf{B} = \arg \min \sum_i \|\mathbf{B}\mathbf{u}_i - \mathbf{X}_i\|. \quad (38)$$

The optimization must satisfy additional constraints. In particular, the basis functions  $\mathbf{B}$  must constitute a partition of unity, and the objective function includes a regularization term to prevent overfitting.

As mentioned earlier, one of the limitations of pure position-based data-driven methods is that wrinkles are quasi-static functions of low-resolution deformations. In PIU, Kavan et al. add oscillatory modes to allow the simulation of traveling waves.

### 5.3. Parallelization

One of the most attractive features of position-based data-driven methods is that they can be massively parallelized. Taking as reference Equation (36), the multiplication of basis weights times the basis values follows the same procedure at all vertices of the detail mesh, but it can be executed in a completely independent manner on each vertex. Therefore, this multiplication is very well suited for implementation on SIMD architectures. In fact, it can be handled as a dense matrix-vector product.

Moreover, the evaluation of weights  $w$  can also be done independently for each basis. In the PIU method described above, the weights are trivially defined as the positions of vertices in the feature mesh, hence they do not need further evaluation. In the WPSD and PDM methods, on the other hand, these weights are computed as a function of low-resolution deformation. As shown by Zurdo et al. [ZBO12], the nonlinear weights at feature vertices can be evaluated easily on the CPU, and then interpolated to detail vertices using subdivision weights. This interpolation step is trivial to parallelize on SIMD architectures.

Thanks to the highly parallel nature of data-driven algorithms, existing implementations exhibit very high performance. For example, Zurdo et al. [ZBO12] simulate cloth with 25.6K triangles at 125 fps including contact handling on the feature mesh, Kavan et al. [KGBS11] simulate a cape with 10K vertices at 1kHz, and Bickel et al. [BLB\*08] animate a face with more than 1M triangles at 30 fps.

### 5.4. Discussion

The particular data-driven methods discussed here present different pros and cons. Among all methods, PIU is probably the fastest method, as it relies on simple matrix-vector multiplication. However, the result is limited to a linear operator, and wrinkles may appear smoothed. WPSD and PDM, on the other hand, allow nonlinear operators, and the combination of the input data is decided in a local manner.

With WPSD and PDM, the representation is more compact, thanks to the choice of low-resolution strain as interpolation domain. Low-resolution strain has also been used as a parameter space for cloth wrinkles in procedural models [RPC\*10], or in the Wrinkle Meshes method [MC10]. In the position-based Wrinkles Meshes method, described in Section 3.6, the low-resolution strain affects wrinkle modeling indirectly, as it influences the deformation of the high-resolution cloth.

One limitation of most data-driven methods is that wrinkles are defined as a quasi-static function of the feature mesh, hence they exhibit no dynamics. Kavan et al. [KGBS11] somewhat alleviate this problem by adding oscillatory modes, but the solution does not support all dynamic effects. The work of de Aguiar et al. [dASTH10]

offers an interesting approach to data-driven simulation of dynamic effects, but it models low-resolution motion, not high-resolution details.

## 6. Applications

In this section we introduce different application areas of position-based methods. These methods are mainly used in interactive applications where performance, controllability and stability is more important than accuracy, like e.g. in [SGdA\*10]. But there exist also other works which use a position-based approach for stabilization.

One application area for position-based methods is interactive surgical simulation. In this area Wang et al. [WXX\*06] introduce a mass-spring model based on a surface mesh to simulate deformable bodies in real-time. Since such a model can neither preserve its volume nor resume its rest shape in the absence of external forces, the authors propose to couple the surface model with a rigid core by using spring forces. This rigid core is simulated using shape matching [MHTG05] which results in a fast and stable simulation. Kubiak et al. [KPGF07] present a simulation method for surgical threads which is based on the position-based dynamics approach of Müller et al. [MHHR07]. Their method simulates the stiffness, bending and torsion of a thread and also provides feedback for a haptic device. For the simulation Kubiak et al. define distance constraints for stiffness and bending, torsion constraints, contact constraints and friction constraints. The presented method allows for an interactive and robust simulation of knots.

The simulation of complex hairstyles using a shape matching approach is presented by Rungjiratananon et al. [RKN10]. Their approach is based on Lattice Shape Matching which was originally introduced by Rivers and James [RJ07]. For the simulation each hair strand is represented by a chain of particles which is subdivided in overlapping chain regions. After shape matching an additional position-based strain limiting is applied to each strand which moves the particles in direction of their root. Different hair styles are realized by using appropriate initial configurations and by modifying the region sizes of a chain.

O'Brien et al. [ODC11] use position-based dynamics for the physically plausible adaptation of motion-captured animations. In their work they use a vertex-based character skeleton and different constraints to preserve the skeleton structure, to define joint limits and to implement a center of mass control. In addition to the kinematic constraints, they define a couple of dynamics constraints which consider vertices in multiple frames. Dynamics constraints are used to enforce smooth acceleration and dynamical correctness.

Fierz et al. [FSAH12] introduce a position-based approach to stabilize a finite element simulation. When using an explicit time integration for a finite element simulation, the time step size is typically limited by the stiffness

of the model and its spatial discretization. In each simulation step Fierz et al. use the Courant-Friedrichs-Lewy (CFL) condition to determine the maximum allowed time step size for each tetrahedral element in their volumetric simulation model. However, instead of using the time step size given by the CFL condition to perform a stable simulation step with an explicit integration scheme, they use a fixed size and mark all elements where the condition is not met. The marked elements are then simulated using a shape matching approach while for all other elements a linear finite element method is used for the simulation.

## 7. Conclusion

In this survey, we focused on a popular and practically relevant subset of approaches for dynamically deforming solids, namely on position-based approaches. Such geometrically motivated techniques are not force-driven and particularly appropriate in interactive applications due to their versatility, robustness, controllability and efficiency. We explained general ideas of position-based methods, shape-matching approaches and data-driven techniques. Various deformation aspects for 2D and 3D solids and efficient solution strategies were discussed with a particular focus on the benefits of position-based approaches compared to force-driven techniques.

## References

- [BETC12] BENDER J., ERLBEN K., TRINKLE J., COUMANS E.: Interactive Simulation of Rigid Body Dynamics in Computer Graphics. In *EG 2012 - State of the Art Reports* (Cagliari, Sardinia, Italy, 2012), Cini M.-P., Ganovelli F., (Eds.), Eurographics Association, pp. 95–134. 1
- [BLB\*08] BICKEL B., LANG M., BOTSCH M., OTADUY M. A., GROSS M.: Pose-space animation and transfer of facial details. In *Proc. of the ACM SIGGRAPH / Eurographics Symposium on Computer Animation* (2008), pp. 57–66. 16, 17, 18, 19
- [BMR03] BRIDSON R., MARINO S., REDKIW R.: Simulation of clothing with folds and wrinkles. In *Proc. ACM/Eurographics Symposium on Computer Animation* (2003), pp. 28–36. 5
- [BMWG07] BERGOU M., MATHUR S., WARDETZKY M., GRINSPUN E.: TRACKS: Toward directable thin shells. *Proc. of ACM SIGGRAPH* (2007). 17
- [BW98] BARAFF D., WITKIN A.: Large steps in cloth simulation. In *Proceedings of Computer graphics and interactive techniques* (New York, NY, USA, 1998), SIGGRAPH '98, ACM, pp. 43–54. 8
- [CBC\*01] CARR J. C., BEATSON R. K., CHERRIE J. B., MITCHELL T. J., FRIGHT W. R., MCCALLUM B. C., EVANS T. R.: Reconstruction and representation of 3D objects with radial basis functions. In *Proc. of ACM SIGGRAPH* (2001), pp. 67–76. 17
- [CK02] CHOI K.-J., KO H.-S.: Stable but responsive cloth. *ACM Transactions on Graphics* 21, 3 (2002), 604–611.
- [CK05] CHOI M. G., KO H.-S.: Modal warping: Real-time simulation of large rotational deformation and manipulation. *IEEE Transactions on Visualization and Computer Graphics* 11, 1 (Jan. 2005), 91–101.
- [dASTH10] DE AGUIAR E., SIGAL L., TREUILLE A., HODGINS J. K.: Stable spaces for real-time clothing. *ACM Transactions on Graphics* 29, 4 (July 2010), 106:1–106:9. 19
- [DBB11] DIZIOL R., BENDER J., BAYER D.: Robust real-time deformation of incompressible surface meshes. In *Proceedings of the 2011 ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (2011), SCA '11, Eurographics Association. 10, 11, 12, 13, 15
- [DGW11] DICK C., GEORGII J., WESTERMANN R.: A hexahedral multigrid approach for simulating cuts in deformable objects. *IEEE Transactions on Visualization and Computer Graphics* 17, 11 (2011), 1663–1675.
- [DSB99] DESBRUN M., SCHRÖDER P., BARR A.: Interactive animation of structured deformable objects. In *Proc. of SIGGRAPH 99* (1999), ACM, pp. 1–8. 1, 5
- [FSAH12] FIERZ B., SPILLMANN J., AGUINAGA I., HARDERS M.: Maintaining large time steps in explicit finite element simulations using shape matching. *Visualization and Computer Graphics, IEEE Transactions on* 18, 5 (may 2012), 717–728. 19
- [GHDS03] GRINSPUN E., HIRANI A. N., DESBRUN M., SCHRÖDER P.: Discrete shells. In *Proceedings of the 2003 ACM SIGGRAPH/Eurographics symposium on Computer animation* (Aire-la-Ville, Switzerland, Switzerland, 2003), SCA '03, Eurographics Association, pp. 62–67. 4
- [GHF\*07a] GOLDENTHAL R., HARMON D., FATTAL R., BERCOVIER M., GRINSPUN E.: Efficient simulation of inextensible cloth. *ACM Transactions on Graphics* 26, 3 (2007), 49. 3
- [GHF\*07b] GOLDENTHAL R., HARMON D., FATTAL R., BERCOVIER M., GRINSPUN E.: Efficient Simulation of Inextensible Cloth. *SIGGRAPH (ACM Transactions on Graphics)* 26, 3 (Jul 2007). 6
- [GM97] GIBSON S. F., MIRTICH B.: *A survey of deformable modeling in computer graphics*. Tech. Rep. TR-97-19, Mitsubishi Electric Research Lab., Cambridge, MA, 1997. 1
- [HES03] HAUTH M., ETZMUSS O., STRASSER W.: Analysis of numerical methods for the simulation of deformable models. *The Visual Computer* 19, 7-8 (2003), 581–600.
- [HJCW06] HONG M., JUNG S., CHOI M., WELCH S.: Fast volume preservation for a mass-spring system. *IEEE Comput. Graph. Appl.* 26 (2006), 83–91. 13
- [ISF07] IRVING G., SCHROEDER C., FEDKIW R.: Volume conserving finite element simulations of deformable models. *ACM Trans. on Graphics* 26, 3 (July 2007), 13:1–13:6.
- [ITF04] IRVING G., TERAN J., FEDKIW R.: Invertible finite elements for robust simulation of large deformation. In *Proc. of the 2004 ACM SIGGRAPH/Eurographics Symp. on Comput. Anim.* (2004), Eurographics Association, pp. 131–140.
- [Jak01] JAKOBSEN T.: Advanced character physics. In *Proceedings, Game Developer's Conference 2001* (San Jose, 2001). 4
- [JP99] JAMES D. L., PAI D. K.: Artdefo: accurate real time deformable objects. In *Proc. of SIGGRAPH 99* (1999), ACM, pp. 65–72. 1
- [KCM12] KIM T.-Y., CHENTANEZ N., MÜLLER M.: Long Range Attachments - A Method to Simulate Inextensible Clothing in Computer Games. In *Eurographics/ ACM SIGGRAPH Symposium on Computer Animation* (Lausanne, Switzerland, 2012), Lee J., Kry P., (Eds.), Eurographics Association, pp. 305–310. 6
- [KGBS11] KAVAN L., GERSZEWSKI D., BARGTEIL A. W., SLOAN P.-P.: Physics-inspired upsampling for cloth simulation in games. *Proc. of ACM SIGGRAPH* (2011). 16, 17, 19



- [KM04] KURIHARA T., MIYATA N.: Modeling deformable human hands from medical images. In *Proc. of the ACM SIGGRAPH / Eurographics Symposium on Computer Animation* (2004), pp. 357–366. 17
- [KNE10] KELAGER M., NIEBE S., ERLEBEN K.: A Triangle Bending Constraint Model for Position-Based Dynamics. In *VRIPHYS 10: 7th Workshop on Virtual Reality Interactions and Physical Simulations* (Copenhagen, Denmark, 2010), Erleben K., Bender J., Teschner M., (Eds.), Eurographics Association, pp. 31–37. 4
- [KPGF07] KUBIAK B., PIETRONI N., GANOVELLI F., FRATAR-CANGELI M.: A robust method for real-time thread simulation. In *Proceedings of the 2007 ACM symposium on Virtual reality software and technology* (New York, NY, USA, 2007), VRST '07, ACM, pp. 85–88. 19
- [LG98] LIN M. C., GOTTSCHALK S.: Collision detection between geometric models: A survey. In *In Proc. of IMA Conference on Mathematics of Surfaces* (1998), pp. 37–56. 2
- [MC10] MÜLLER M., CHENTANEZ N.: Wrinkle meshes. In *Proceedings of the 2010 ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (Aire-la-Ville, Switzerland, Switzerland, 2010), SCA '10, Eurographics Association, pp. 85–92. 7, 19
- [MC11a] MÜLLER M., CHENTANEZ N.: Adding Physics to Animated Characters with Oriented Particles. In *VRIPHYS 11: 8th Workshop on Virtual Reality Interactions and Physical Simulations* (Lyon, France, 2011), Bender J., Erleben K., Galin E., (Eds.), Eurographics Association, pp. 83–91.
- [MC11b] MÜLLER M., CHENTANEZ N.: Solid simulation with oriented particles. *ACM Trans. Graph.* 30, 4 (July 2011), 92:1–92:10. 12, 13
- [MG04] MÜLLER M., GROSS M.: Interactive virtual materials. In *Proceedings of Graphics Interface 2004* (School of Computer Science, University of Waterloo, Waterloo, Ontario, Canada, 2004), GI '04, Canadian Human-Computer Communications Society, pp. 239–246.
- [MGW01] MALZBENDER T., GELB D., WOLTERS H.: Polynomial texture maps. In *Proceedings of ACM SIGGRAPH 2001* (Aug. 2001), Computer Graphics Proceedings, Annual Conference Series, pp. 519–528. 17
- [MHHR07] MÜLLER M., HEIDELBERGER B., HENNIX M., RATCLIFF J.: Position based dynamics. *Journal of Visual Communication and Image Representation* 18, 2 (2007), 109–118. 2, 19
- [MHTG05] MÜLLER M., HEIDELBERGER B., TESCHNER M., GROSS M.: Meshless deformations based on shape matching. *ACM Transactions on Graphics* 24, 3 (2005), 471–478. 9, 10, 14, 19
- [MJC\*08] MA W.-C., JONES A., CHIANG J.-Y., HAWKINS T., FREDERIKSEN S., PEERS P., VUKOVIC M., OUHYOUNG M., DEBEVEC P.: Facial performance synthesis using deformation-driven polynomial displacement maps. *ACM Trans. Graph. (Proc. of ACM SIGGRAPH Asia)* 27, 5 (2008). 16, 17
- [MKC12] MÜLLER M., KIM T.-Y., CHENTANEZ N.: Fast Simulation of Inextensible Hair and Fur. In *in Proceedings of Virtual Reality Interactions and Physical Simulations (VRIPhys)* (2012), Eurographics Association. 6
- [MSJT08] MÜLLER M., STAM J., JAMES D., THÜREY N.: Real time physics: class notes. In *ACM SIGGRAPH 2008 classes* (New York, NY, USA, 2008), SIGGRAPH '08, ACM, pp. 88:1–88:90. 1
- [MTV05] MAGNENAT-THALMANN N., VOLINO P.: From early draping to haute couture models: 20 years of research. *The Visual Computer* 21 (2005), 506–519. 1
- [Mül08] MÜLLER M.: Hierarchical Position Based Dynamics. In *VRIPHYS 08: Fifth Workshop in Virtual Reality Interactions and Physical Simulations* (Grenoble, France, 2008), Faure F., Teschner M., (Eds.), Eurographics Association, pp. 1–10. 6
- [NMK\*06] NEALEN A., MÜLLER M., KEISER R., BOXERMAN E., CARLSON M.: Physically based deformable models in computer graphics. *Computer Graphics Forum* 25, 4 (December 2006), 809–836. 1, 8
- [OBH02] O'BRIEN J. F., BARGTEIL A. W., HODGINS J. K.: Graphical modeling and animation of ductile fracture. *ACM Trans. Graph.* 21, 3 (July 2002), 291–294. 14
- [ODC11] O'BRIEN C., DINGLIANA J., COLLINS S.: Space-time vertex constraints for dynamically-based adaptation of motion-captured animation. In *Proceedings of the 2011 ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (New York, NY, USA, 2011), SCA '11, ACM, pp. 277–286. 19
- [OH99] O'BRIEN J. F., HODGINS J. K.: Graphical modeling and animation of brittle fracture. In *SIGGRAPH '99: Proceedings of the 26th annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 1999), ACM Press/Addison-Wesley Publishing Co., pp. 137–146. 1
- [PB88] PLATT J. C., BARR A. H.: Constraints methods for flexible objects. In *Proceedings of the 15th annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 1988), SIGGRAPH '88, ACM, pp. 279–288. 8
- [Pro95] PROVOT X.: Deformation constraints in a mass-spring model to describe rigid cloth behavior. In *In Graphics Interface* (1995), Davis W. A., Prusinkiewicz P., (Eds.), Canadian Human-Computer Communications Society, pp. 147–154. 5
- [RJ07] RIVERS A. R., JAMES D. L.: FastLSM: fast lattice shape matching for robust real-time deformation. In *SIGGRAPH '07: ACM SIGGRAPH 2007 papers* (New York, NY, USA, 2007), ACM, p. 82. 10, 11, 19
- [RKN10] RUNGJIRATANANON W., KANAMORI Y., NISHITA T.: Chain shape matching for simulating complex hairstyles. *Computer Graphics Forum* 29, 8 (2010), 2438–2446. 19
- [RPC\*10] ROHMER D., POPA T., CANI M.-P., HAHMANN S., SHEFFER A.: Animation wrinkling: Augmenting coarse cloth simulations with realistic-looking wrinkles. *ACM Transactions on Graphics* 29, 5 (2010), 157:1–157:8. 19
- [SD92] SHOEMAKE K., DUFF T.: Matrix animation and polar decomposition. In *Proceedings of the conference on Graphics interface '92* (San Francisco, CA, USA, 1992), Morgan Kaufmann Publishers Inc., pp. 258–264. 15
- [SGdA\*10] STOLL C., GALL J., DE AGUIAR E., THRUN S., THEOBALT C.: Video-based reconstruction of animatable human characters. *ACM Trans. Graph.* 29, 6 (Dec. 2010), 139:1–139:10. 19
- [SGT09] SCHMEDDING R., GISSLER M., TESCHNER M.: Optimized damping for dynamic simulations. In *Spring Conference on Computer Graphics* (2009), pp. 205–212. 8
- [SHZO07] SENGUPTA S., HARRIS M., ZHANG Y., OWENS J. D.: Scan primitives for GPU computing. In *Proc. of the 22nd ACM SIGGRAPH/Eurographics Symp. on Graph. Hardware* (2007), Eurographics Association, pp. 97–106. 15
- [SOG08] STEINEMANN D., OTADUY M. A., GROSS M.: Fast adaptive shape matching deformations. In *Proceedings of the 2008 ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (Aire-la-Ville, Switzerland, Switzerland, 2008), SCA '08, Eurographics Association, pp. 87–94. 11

- [SSBT08] STUMPP T., SPILLMANN J., BECKER M., TESCHNER M.: A Geometric Deformation Model for Stable Cloth Simulation. In *VRIPHYS 08: Fifth Workshop in Virtual Reality Interactions and Physical Simulations* (Grenoble, France, 2008), Faure F., Teschner M., (Eds.), Eurographics Association, pp. 39–46. 14
- [TBHF03] TERAN J., BLEMKER S., HING V. N. T., FEDKIW R.: Finite volume methods for the simulation of skeletal muscle. In *Proc. of the 2003 ACM SIGGRAPH/Eurographics Symp. on Comput. Anim.* (2003), Eurographics Association, pp. 68–74. 1
- [TF88] TERZOPOULOS D., FLEISCHER K.: Deformable models. *The Visual Computer* 4 (1988), 306–331. 8
- [THMG04] TESCHNER M., HEIDELBERGER B., MULLER M., GROSS M.: A versatile and robust model for geometrically complex deformable solids. In *Proceedings of the Computer Graphics International* (Washington, DC, USA, 2004), CGI '04, IEEE Computer Society, pp. 312–319. 1
- [TKH\*05] TESCHNER M., KIMMERLE S., HEIDELBERGER B., ZACHMANN G., RAGHUPATHI L., FUHRMANN A., CANI M.-P., FAURE F., MAGNENAT-THALMANN N., STRASSER W., VOLINO P.: Collision detection for deformable objects. *Computer Graphics Forum* 24, 1 (Mar. 2005), 61–81. 2
- [TPBF87a] TERZOPOULOS D., PLATT J., BARR A., FLEISCHER K.: Elastically deformable models. In *Proceedings of the 14th annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 1987), SIGGRAPH '87, ACM, pp. 205–214. 1
- [TPBF87b] TERZOPOULOS D., PLATT J., BARR A., FLEISCHER K.: Elastically deformable models. In *Computer Graphics (Proceedings of SIGGRAPH 87)* (1987), vol. 21, ACM, pp. 205–214. 1
- [vFTS06] VON FUNCK W., THEISEL H., SEIDEL H.-P.: Vector field based shape deformations. *ACM Trans. on Graphics* 25, 3 (July 2006), 1118–1125. 13
- [WOR10] WANG H., O'BRIEN J., RAMAMOORTHI R.: Multi-resolution isotropic strain limiting. *ACM Trans. Graph.* 29, 6 (Dec. 2010), 156:1–156:10. 5, 6
- [WRO11] WANG H., RAMAMOORTHI R., O'BRIEN J.: Data-driven elastic models for cloth: Modeling and measurement. *ACM Trans. Graph. (Proc. SIGGRAPH)* 30, 4 (2011), 71. 16
- [WXX\*06] WANG Y., XIONG Y., XU K., TAN K., GUO G.: A mass-spring model for surface mesh deformation based on shape matching. In *Proceedings of the 4th international conference on Computer graphics and interactive techniques in Australasia and Southeast Asia* (New York, NY, USA, 2006), GRAPHITE '06, ACM, pp. 375–380. 19
- [ZBO12] ZURDO J. S., BRITO J. P., OTADUY M. A.: Animating wrinkles by example on non-skinned cloth. *IEEE Transactions on Visualization and Computer Graphics* 99, PrePrints (2012). 16, 17, 18, 19