Physically Based Deformable Models in Computer Graphics

Andrew Nealen[1], Matthias Müller[2,3], Richard Keiser[3], Eddy Boxerman[4] and Mark Carlson[5]

[1] Discrete Geometric Modeling Group, TU Darmstadt
[2] NovodeX / AGEIA
[3] Computer Graphics Lab, ETH Zürich
[4] Department of Computer Science, University of British Columbia
[5] DNA Productions, Inc.

Abstract

Physically based deformable models have been widely embraced by the Computer Graphics community. Many problems outlined in a previous survey by Gibson and Mirtich [GM97] have been addressed, thereby making these models interesting and useful for both offline and real-time applications, such as motion pictures and video games. In this paper, we present the most significant contributions of the past decade, which produce such impressive and perceivably realistic animations and simulations: finite element/difference/volume methods, mass-spring systems, meshfree methods, coupled particle systems and reduced deformable models based on modal analysis. For completeness, we also make a connection to the simulation of other continua, such as fluids, gases and melting objects. Since time integration is inherent to all simulated phenomena, the general notion of time discretization is treated separately, while specifics are left to the respective models. Finally, we discuss areas of application, such as elastoplastic deformation and fracture, cloth and hair animation, virtual surgery simulation, interactive entertainment and fluid/smoke animation, and also suggest areas for future research.

Categories and Subject Descriptors (according to ACM CCS): I.3.5 [Computer Graphics]: Physically Based Modeling I.3.7 [Computer Graphics]: Animation and Virtual Reality

1. Introduction

Physically based deformable models have two decades of history in Computer Graphics: since Lasseter's discussion of *squash and stretch* [Las87] and, concurrently, Terzopoulos et. al's seminal paper on elastically deformable models [TPBF87], numerous researchers have partaken in the quest for the visually and physically plausible animation of deformable objects and fluids. This inherently interdisciplinary field elegantly combines newtonian dynamics, continuum mechanics, numerical computation, differential geometry, vector calculus, approximation theory and Computer Graphics (to name a few) into a vast and powerful toolkit, which is being further explored and extended. The field is in constant flux and, thus, active and fruitful, with many visually stunning achievements to account for.

Since Gibson and Mirtich's survey paper [GM97], the field of physically based deformable models in Computer Graphics has expanded tremendously. Significant contributions were made in many key areas, e.g. object model-


Figure 1: *Hooke's Law, from* De Potentia Restitutiva *[1678].*

ing, fracture, plasticity, cloth animation, stable fluid simulation, time integration strategies, discretization and numer-

ical solution of PDEs, modal analysis, space-time adaptivity, multiresolution modeling and real-time simulation. Non-physical models, such as parametric curves and surfaces and free-form deformations, are not discussed in this report. The inclined reader is therefore encouraged to browse more recent literature on t-splines [SZBN03] [SCF*04], space-warping [MJBF02] [LKG*03] [BK05] and methods based on differential surface properties [SLCO*04] [YZX*04] [BK04] [LSLCO05] [NSACO05] [IMH05]. For advances in character skinning see e.g. [WP02] [KJP02] [JT05]. Since we are not able to cover basic elasticity theory and continuum mechanics in this report, we would like to point out that a nice review of the history of elasticity theory, starting with the discovery of Hooke's Law in 1660 (Fig. 1) and leading up to the general equations of Navier in 1821, is given in [Lov27]. Furthermore, great introductions to continuum mechanics and dynamics can be found in [WB97] and in general textbooks, such as [Chu96] [Coo95] [BW97] [Gdo93] [BLM00]. For application specific presentations, we refer the reader to a number of recent works. For cloth simulation, there is the text by House and Breen [HB00], as well as the recent, extensive tutorial by Thalmann et al. [MTCK*04]. For hair simulation, there is the (already slightly dated) overview by Thalmann et al. [MTHK00]; the paper by Volino and Thalmann [VMT04] gives a good, more recent overview. Collision detection and haptic force-feedback rendering for deformable objects are other challenging and active areas of research. For a summary of recent work in these fields, we refer the reader to the report by Teschner et al. [TKH*05] and the course notes of Lin and Otaduy [LO05].

In this report we take a *model based* point of view, motivated by the fact that there are many readily available physical models for very similar applications, i.e. we can animate an elastically or plastically deforming solid with many different underlying models, such as mass-spring systems, finite elements or meshfree methods. We furthermore make a distinction between *Lagrangian* methods, where the model consists of a set of points with varying locations and properties, and *Eulerian* methods, where model properties are computed for a set of stationary points. To give a coarse overview, we describe recent developments for

- Lagrangian Mesh Based Methods
  - Continuum Mechanics Based Methods
  - Mass-Spring Systems
- Lagrangian Mesh Free Methods
  - Loosely Coupled Particle Systems
  - Smoothed Particle Hydrodynamics (SPH)
  - Mesh Free Methods for the solution of PDEs
- Reduced Deformation Models and Modal Analysis
- Eulerian and Semi-Lagrangian Methods
  - Fluids and Gases
  - Melting Objects

In each section we present the basic model formulation, recent contributions, benefits and drawbacks, and various areas of application. The section on fluids, gases and melting objects contains an overview of recent work and establishes the connection to the field of physically based deformable models. A complete survey on the animation of fluids and gases would easily fill its own report and is therefore beyond our scope.

Our goal is to provide an up-to-date report to the Computer Graphics community, as an entry point for researchers and developers who are new to the field, thereby complementing the existing survey paper [GM97].

2. Background

2.1. Continuum Elasticity

A deformable object is typically defined by its undeformed shape (also called equilibrium configuration, rest or initial shape) and by a set of material parameters that define how it deforms under applied forces. If we think of the rest shape as a continuous connected subset M of $\mathbb{R}^3$, then the coordinates $\mathbf{m} \in M$ of a point in the object are called *material coordinates* of that point. In the discrete case M is a discrete set of points that sample the rest shape of the object.

When forces are applied, the object deforms and a point originally at location $\mathbf{m}$ (i.e. with material coordinates $\mathbf{m}$) moves to a new location $\mathbf{x}(\mathbf{m})$, the *spatial* or *world coordinates* of that point. Since new locations are defined for all material coordinates $\mathbf{m}$, $\mathbf{x}$ is a vector field defined on M. Alternatively, the deformation can also be specified by the *displacement* vector field $\mathbf{u}(\mathbf{m}) = \mathbf{x}(\mathbf{m}) - \mathbf{m}$ defined on M (see Fig. 2). From $\mathbf{u}(\mathbf{m})$ the elastic strain ε is computed (ε is a dimensionless quantity which, in the (linear) 1D case, is simply $\Delta l / l$). A spatially constant displacement field represents a translation of the object with no strain. Therefore, it becomes clear that strain must be measured in terms of spatial variations of the displacement field $\mathbf{u} = \mathbf{u}(\mathbf{m}) = (u, v, w)^T$. Popular choices in Computer Graphics are

$$\varepsilon_G = \frac{1}{2}(\nabla \mathbf{u} + [\nabla \mathbf{u}]^T + [\nabla \mathbf{u}]^T \nabla \mathbf{u}) \tag{1}$$

$$\varepsilon_C = \frac{1}{2}(\nabla \mathbf{u} + [\nabla \mathbf{u}]^T), \tag{2}$$

where the symmetric tensor $\varepsilon_G \in \mathbb{R}^{3x3}$ is Green's nonlinear strain tensor and $\varepsilon_C \in \mathbb{R}^{3x3}$ its linearization, Cauchy's linear strain tensor. The gradient of the displacement field is a 3 by 3 matrix

$$\nabla \mathbf{u} = \begin{bmatrix} u_{,x} & u_{,y} & u_{,z} \\ v_{,x} & v_{,y} & v_{,z} \\ w_{,x} & w_{,y} & w_{,z} \end{bmatrix}, \tag{3}$$

where the index after the comma represents a spatial derivative.

The material law (or *constitutive law*) is used for the computation of the symmetric internal stress tensor $\sigma \in \mathbb{R}^{3x3}$ for

each material point **m** based on the strain ε at that point (σ is measured as force per unit area, where $1 Pascal = 1 Pa = 1 N/m^2$). Most Computer Graphics papers use Hooke's linear material law

$$\sigma = \mathbf{E} \cdot \varepsilon, \tag{4}$$

where **E** is a rank four tensor which relates the coefficients of the stress tensor linearly to the coefficients of the strain tensor. For isotropic materials, the coefficients of **E** only depend on Young's modulus and Poisson's ratio.

2.2. Time Integration

In order to simulate dynamic deformable solids, we need to know the time dependent world coordinates $\mathbf{x}(\mathbf{m}, t)$ of all points in M. Given $\mathbf{x}(\mathbf{m}, t)$, we can subsequently display the configurations $\mathbf{x}(0), \mathbf{x}(\Delta t), \mathbf{x}(2\Delta t), ..$ resulting in an animation of the object. Here Δt is a fixed time step of the simulation and $\mathbf{x}(t)$ represents the entire vector field at time t.

The unknown vector fields $\mathbf{x}(t)$ are not given directly but implicitly as the solution of a differential equation, namely Newton's second law of motion of the form

$$\ddot{\mathbf{x}} = F(\dot{\mathbf{x}}, \mathbf{x}, t), \tag{5}$$

where $\ddot{\mathbf{x}}$ and $\dot{\mathbf{x}}$ are the second and first time derivatives of **x**, respectively and $F()$ a general function given by the physical model of the deformable object. In order to find the solution $\mathbf{x}(t)$, this second order differential equation is often rewritten as a coupled set of two first order equations

$$\dot{\mathbf{x}} = \mathbf{v} \tag{6}$$
$$\dot{\mathbf{v}} = F(\mathbf{v}, \mathbf{x}, t), \tag{7}$$

where the new quantity **v** represents $\dot{\mathbf{x}}$. A discrete set of values $\mathbf{x}(0), \mathbf{x}(\Delta t), \mathbf{x}(2\Delta t), ..$ of the unknown vector field **x** which is needed for the animation can now be obtained by numerically solving (i.e. integrating) this system of equations. Numerical integration of ordinary differential equations is the subject of many textbooks (e.g. [PTVF92, AP98]). See [HES02] for an excellent overview in the context of deformable modeling in computer graphics. We give a few examples here which appear in subsequent sections.

The simplest scheme is explicit (or forward) Euler integration, where the time derivatives are replaced by finite differences $\dot{\mathbf{v}}(t) = [\mathbf{v}(t + \Delta t) - \mathbf{v}(t)]/\Delta t$ and $\dot{\mathbf{x}}(t) = [\mathbf{x}(t + \Delta t) - \mathbf{x}(t)]/\Delta t$. Substituting these into the above equations and solving for the quantities at the next time step $t + \Delta t$ yields

$$\mathbf{x}(t + \Delta t) = \mathbf{x}(t) + \Delta t \mathbf{v}(t) \tag{8}$$
$$\mathbf{v}(t + \Delta t) = \mathbf{v}(t) + \Delta t F(\mathbf{v}(t), \mathbf{x}(t), t). \tag{9}$$

Time integration schemes are evaluated by two main criteria, their stability and their accuracy. Their accuracy is measured by their convergence with respect to the time step size Δt, i.e.

first order $O(\Delta t)$, second order $O(\Delta t^2)$, etc. In the field of physically based animation in Computer Graphics, stability is often much more important than accuracy.

The integration scheme presented above is called *explicit* because it provides explicit formulas for the quantities at the next time step. Explicit methods are easy to implement but they are only conditionally stable, i.e. stable only if Δt is smaller than a stability threshold (see [MHTG05] for a formalization). For stiff objects this threshold can be very small. The instability is due to the fact that explicit methods extrapolate a constant right hand side blindly into the future as the above equations show. For a simple spring and a large Δt, the scheme can overshoot the equilibrium position arbitrarily. At the next time step the restoring forces get even larger resulting in an exponential gain of energy and finally an explosion. This problem can be solved by using an *implicit* scheme that uses quantities at the next time step $t + \Delta t$ on both sides of the equation

$$\mathbf{x}(t + \Delta t) = \mathbf{x}(t) + \Delta t \mathbf{v}(t + \Delta t) \tag{10}$$
$$\mathbf{v}(t + \Delta t) = \mathbf{v}(t) + \Delta t F(\mathbf{v}(t + \Delta t), \mathbf{x}(t + \Delta t), t). \tag{11}$$

The scheme is now called implicit because the unknown quantities are *implicitly* given as the solution of a system of equations. Now, instead of extrapolating a constant right hand side blindly into the future, the right hand side is part of the solution process. Remarkably, the implicit (or backward) Euler scheme is stable for arbitrarily large time steps Δt (There is, however, a lower time step limit which, for practical purposes, poses no problem). This gain comes with the price of having to solve an algebraic system of equations at each time step (linear if $F()$ is linear, non-linear otherwise).

A simple improvement to the forward Euler scheme is to swap the order of the equations and use a forward-backward scheme

$$\mathbf{v}(t + \Delta t) = \mathbf{v}(t) + \Delta t F(\mathbf{v}(t), \mathbf{x}(t), t) \tag{12}$$
$$\mathbf{x}(t + \Delta t) = \mathbf{x}(t) + \Delta t \mathbf{v}(t + \Delta t). \tag{13}$$

The update to **v** uses forward Euler, while the update to **x** uses backward Euler. Note that the method is still explicit; $\mathbf{v}(t + \Delta t)$ is simply evaluated first. For non-dissipative systems (ie. when forces are independent of velocities), this reduces to the second order accurate Stoermer-Verlet scheme. The forward-backward Euler scheme is more stable than standard forward Euler integration, without any additional computational overhead.

3. Lagrangian Mesh Based Methods

3.1. The Finite Element Method

The Finite Element Method (FEM) is one of the most popular methods in Computational Sciences to solve Partial Differential Equations (PDE's) on irregular grids. In order to use the method for the simulation of deformable objects, the object is viewed as a continuous connected volume as

Figure 2: *In the Finite Element method, a continuous deformation (a) is approximated by a sum of (linear) basis functions defined inside a set of finite elements (b).*

in Section 2.1 which is discretized using an irregular mesh. Continuum mechanics, then, provides the PDE to be solved.

The PDE governing dynamic elastic materials is given by

$$\rho \ddot{\mathbf{x}} = \nabla \cdot \sigma + \mathbf{f}, \qquad (14)$$

where ρ is the density of the material and $\mathbf{f}$ externally applied forces such as gravity or collision forces. The divergence operator turns the 3 by 3 stress tensor back into a 3 vector

$$\nabla \cdot \sigma = \begin{bmatrix} \sigma_{xx,x} + \sigma_{xy,y} + \sigma_{xz,z} \\ \sigma_{yx,x} + \sigma_{yy,y} + \sigma_{yz,z} \\ \sigma_{zx,x} + \sigma_{zy,y} + \sigma_{zz,z} \end{bmatrix}, \qquad (15)$$

representing the internal force resulting from a deformed infinitesimal volume. Eq. 14 shows the equation of motion in *differential form* in contrast to the *integral form* which is used in the Finite Volume method.

The Finite Element Method is used to turn a PDE into a set of algebraic equations which are then solved numerically. To this end, the domain M is discretized into a finite number of disjoint elements (i.e. a mesh). Instead of solving for the spatially continuous function $\mathbf{x}(\mathbf{m}, t)$, one only solves for the discrete set of unknown positions $\mathbf{x}_i(t)$ of the nodes of the mesh. First, the function $\mathbf{x}(\mathbf{m}, t)$ is approximated using the nodal values by

$$\tilde{\mathbf{x}}(\mathbf{m}, t) = \sum_i \mathbf{x}_i(t) \mathbf{b}_i(\mathbf{m}), \qquad (16)$$

where $\mathbf{b}_i()$ are fixed nodal basis functions which are 1 at node i and zero at all other nodes, also known as the Kronecker Delta property (see Fig. 2). In the most general case of the Finite Element Method, the basis functions do not have this property. In that case, the unknowns are general parameters which can not be interpreted as nodal values. Substituting $\tilde{\mathbf{x}}(\mathbf{m}, t)$ into Eq. 14 results in algebraic equations for the $\mathbf{x}_i(t)$. In the Galerkin approach [Hun05], finding the unknowns $\mathbf{x}_i(t)$ is viewed as an optimization process. When substituting $\mathbf{x}(\mathbf{m}, t)$ by the approximation $\tilde{\mathbf{x}}(\mathbf{m}, t)$, the infinitely dimensional search space of possible solutions is reduced to a finite dimensional subspace. In general, no function in that subspace can solve the original PDE. The approximation will generate a deviation or residue when substituted

into the PDE. In the Galerkin method, the approximation which *minimizes the residue* is sought. In other words, we look for an approximation whose residue is perpendicular to the subspace of functions defined by Eq. 16.

Many papers in Computer Graphics use a simple form of the Finite Element method for the simulation of deformable objects, sometimes called the *explicit* Finite Element Method, which is quite easy to understand and to implement (e.g. [OH99], [DDCB01], [MDM*02]). The explicit Finite Element Method is not to be confused with the standard Finite Element Method being *integrated* explicitly. The explicit Finite Element Method can be integrated either explicitly or implicitly.

In the explicit Finite Element approach, both, the masses and the internal and external forces are lumped to the vertices. The nodes in the mesh are treated like mass points in a mass-spring system while each element acts like a generalized spring connecting all adjacent mass points. The forces acting on the nodes of an element due to its deformation are computed as follows (see for instance [OH99]): given the positions of the vertices of an element and the fixed basis functions, the continuous deformation field $\mathbf{u}(\mathbf{m})$ inside the element can be computed using Eq. 16. From $\mathbf{u}(\mathbf{m})$, the strain field $\varepsilon(\mathbf{m})$ and stress field $\sigma(\mathbf{m})$ are computed. The deformation energy of the element is then given by

$$E = \int_V \varepsilon(\mathbf{m}) \cdot \sigma(\mathbf{m}) d\mathbf{m}, \qquad (17)$$

where the dot represents the componentwise scalar product of the two tensors. The forces can then be computed as the derivatives of the energy with respect to the nodal positions. In general, the relationship between nodal forces and nodal positions is nonlinear. When linearized, the relationship for an element e connecting n_e nodes can simply be expressed as

$$\mathbf{f}_e = \mathbf{K}_e \mathbf{u}_e, \qquad (18)$$

where $\mathbf{f}_e \in \mathbb{R}^{3n_e}$ contains the n_e nodal forces and $\mathbf{u}_e \in \mathbb{R}^{3n_e}$ the n_e nodal displacements of an element. The matrix $\mathbf{K}_e \in \mathbb{R}^{3n_e x 3n_e}$ is called the *stiffness matrix* of the element. Because elastic forces coming from adjacent elements add up at a node, a stiffness matrix $\mathbf{K} \in \mathbb{R}^{3n x 3n}$ for an enire mesh with n nodes can be formed by assembling the element's stiffness matrices

$$\mathbf{K} = \sum_e \mathbf{K}_e. \qquad (19)$$

In this sum, the element's stiffness matrices are expanded to the dimension of $\mathbf{K}$ by filling in zeros at positions related to nodes not adjacent to the element. Using the linearized elastic forces, the linear algebraic equation of motion for an entire mesh becomes ($\mathbf{u} = \mathbf{x} - \mathbf{x}_0$)

$$\mathbf{M}\ddot{\mathbf{u}} + \mathbf{D}\dot{\mathbf{u}} + \mathbf{K}\mathbf{u} = \mathbf{f}_{ext}, \qquad (20)$$

where $\mathbf{M} \in \mathbb{R}^{nxn}$ is the mass matrix, $\mathbf{D} \in \mathbb{R}^{nxn}$ the damping

matrix and $\mathbf{f}_{ext} \in \mathbb{R}^n$ externally applied forces. Often, diagonal matrices are used for $\mathbf{M}$ and $\mathbf{D}$, a technique called *mass lumping*. In this case, $\mathbf{M}$ just contains the point masses of the nodes of the mesh on its diagonal. The vectors $\mathbf{x}$ and $\mathbf{x}_0$ contain, respectively, the actual and the rest positions of the nodes.

The Finite Element method only produces a *linear* system of algebraic equations if applied to a *linear* PDE. If a linear strain measure is used and Hooke's law for isotropic materials is substituted into 14, Lamé's linear PDE results:

$$\rho\ddot{\mathbf{x}} = \mu\Delta\mathbf{u} + (\lambda + \mu)\nabla(\nabla \cdot \mathbf{u}), \qquad (21)$$

where the material constants λ and μ can be computed directly from Young's modulus and Poisson's ratio. This equation is solved in [DDBC99] in a multiresolution fashion. Using discretized versions of the Laplacian ($\Delta = \nabla^2$) and gradient-of-divergence ($\nabla(\nabla \cdot)$) operators, they solve the Lamé equation on an irregular, multiresolution grid. The system is optimized for limited deformations (linearized strain) and does not support topological changes. Based on Gauss' Divergence Theorem, the discrete operators are further evolved in [DDCB00], which leads to greater accuracy through defined error bounds. Furthermore, the cubic octree hierarchy employed in [DDBC99] is succeeded by a nonnested hierarchy of meshes, in conformance with the redefined operators, which leads to improved shape sampling. In [DDCB01] the previous linearized physical model is replaced by local explicit finite elements and Green's nonlinear strain tensor. To increase stability the simulation is integrated semi-implicitly in time [DSB99].

O'Brien et al. [OH99] and [OBH02] present a FEM based technique for simulating brittle and ductile fracture in connection with elastoplastic materials. They use tetrahedral meshes in connection with linear basis functions $\mathbf{b}_i()$ and Green's strain tensor. The resulting nonlinear equations are solved explicitly and integrated explicitly. The method produces realistic and visually convincing results, but it is not designed for interactive or real-time use. In addition to the strain tensor, they use the so-called *strain rate tensor* (the time derivative of the strain tensor), to compute damping forces. Other studies on the visual simulation of brittle fracture are [SWB00] and [MMDJ01].

Bro-Nielsen and Cotin [BNC96] use linearized finite elements for surgery simulation. They achieve significant speedup by simulating only the visible surface nodes (condensation), similar to the BEM. Many other studies successfully apply the FEM to surgery simulation, such as (but surely not limited to) [GTT89], [CEO*93], [SBMH94], [KGC*96], [CDA99], [CDA00], [PDA00] and [PDA01].

As long as the equation of motion is integrated explicitly in time, nonlinear elastic forces resulting from Green's strain tensor pose no computational problems. The nonlinear formulas for the forces are simply evaluated and used directly to integrate velocities and positions as in [OH99]. As

Figure 3: *The pitbull with its inflated head (left) shows the artifact of linear FEM under large rotational deformations. The correct deformation is shown on the right.*

mentioned earlier (Section 2.2), explicit integration schemes are stable only for small time steps while implicit integration schemes allow arbitrarily large time steps. However, in the latter case, a system of algebraic equations needs to be solved at every time step. Linear PDE's yield linear algebraic systems which can be solved more efficiently and more stably than nonlinear ones. Unfortunately, linearized elastic forces are only valid for small deformations. Large rotational deformations yield highly inaccurate restoring forces (see Fig. 3).

To eliminate these artifacts, Müller et al. extract the rotational part of the deformation for each finite element and compute the forces with respect to the non-rotated reference frame [MG04]. The linear equation 18 for the elastic forces of an element (in this case a tetrahedron) is replaced by

$$\mathbf{f} = \mathbf{R}\mathbf{K}(\mathbf{R}^T\mathbf{x} - \mathbf{x}_0), \qquad (22)$$

where $\mathbf{R} \in \mathbb{R}^{12x12}$ is a matrix that contains four 3 by 3 identical rotation matrices along its diagonal. The vector $\mathbf{x}$ contains the actual positions of the four adjacent nodes of the tetrahedron while $\mathbf{x}_0$ contains their rest positions. The rotation of the element used in $\mathbf{R}$ is computed by performing a polar decomposition of the matrix that describes the transformation of the tetrahedron from the configuration $\mathbf{x}_0$ to the configuration $\mathbf{x}$. This yields stable, fast and visually pleasing results. In an earlier approach, they extract the rotational part not per element but per node [MDM*02]. In this case, the global stiffness matrix does not need to be reassembled at each time step but ghost forces are introduced.

Another solution to this problem is proposed in [CGC*02]: each region of the finite element mesh is associated with the bone of a simple skeleton and then locally linearized. The regions are blended in each time step, leading to results which are visually indistinguishable from the nonlinear solution, yet much faster.

An adaptive nonlinear FEM simulation is described by Wu et al. [WDGT01]. Distance, surface and volume preservation for triangular and tetrahedral meshes is outlined in [THMG04]. Grinspun et al. introduce conforming, hierarchical, adaptive refinement methods (CHARMS) for general finite elements [GKS02], refining basis functions instead of elements. Irving et al. [ITF04] present a method which robustly handles large deformation and element inversion by

computing a problem-specific diagonalization of the deformation gradient, from which the forces are derived. Müller et al. [MTG04] propose simulating and fracturing objects represented by surface meshes, by embedding the surface in a cuboid element voxelization (Fig. 4). This domain embedding strategy is also used by James et al. in their *squashing cubes* simulator [JBT04]. To support topological changes while maintaining well-shaped elements, Molino et al. create duplicates of the original elements in their virtual node algorithm [MBF04].

Figure 4: *Embedding a topologically inconsistent surface mesh in hexahedral finite elements for deformation simulation and fracturing [MTG04].*

3.2. The Method of Finite Differences

If the object M is sampled using a *regular* spatial grid, the equation of motion 14 can be discretized using finite differences. The method of Finite Differences is easier to implement than the general Finite Element Method. However, it is difficult to approximate the boundary of an arbitrary object with a regular mesh. Also, local adaptations are only possible with irregular meshes.

Pioneering work in the field of physically based animation was carried out by Terzopoulos and his co-workers. In their seminal paper [TPBF87] the dynamics of deformable models are computed from the potential energy stored in the elastically deformed body. For volumetric objects, they define the deformation energy as the weighted matrix norm of the difference between the metric tensors of the deformed and original shape, integrated over the entire continuum (two and one dimensional objects involve further weighted norms of second and third order tensors). The continuous variational (or directional) derivative of this energy functional (the elastic force) is discretized using the Method of Finite Differences (FD), and the simulation is moved forward through time by semi-implicit integration. This work is further evolved in [TF88] and [TW88] where the model is extended to cover plasticity and fracture.

3.3. The Finite Volume Method

In the explicit Finite Element Method, the forces acting on the nodes of an element are computed as the derivatives of the deformation energy with respect to the nodal positions, where the deformation energy is computed via 17.

There is a more direct way to get to the nodal forces of an element, using the Method of Finite Volumes. The stress tensor σ can be used to compute the internal force $\mathbf{f}$ per unit area with respect to a certain plane orientation as

$$\mathbf{f} = \sigma \mathbf{n}, \tag{23}$$

where the normalized vector $\mathbf{n}$ is the normal on that plane. Thus, the total force acting on the face A of a finite element is given by the surface integral

$$\mathbf{f}_A = \int_A \sigma d\mathbf{A}. \tag{24}$$

When linear basis functions are used, the stress tensor is constant within an element. Then, for planar element faces, the surface integral reduces to the simple product

$$\mathbf{f}_A = A\sigma \mathbf{n}, \tag{25}$$

where the scalar A is the area of the face and $\mathbf{n}$ its normal. To get the nodal forces, one loops over all the faces A_i of the element and distributes the force $\mathbf{f}_{A_i}$ evenly among the nodes adjacent to that face. Teran et al. [TBHF03] use this method to simulate skeletal muscle. They also use a geometrically motivated way to compute strain which leads to an intuitive way of integrating the equations of motion.

3.4. The Boundary Element Method

The Boundary Element Method (BEM) is an interesting alternative to the standard Finite Element approach because all computations are done on the surface (boundary) of the elastic body instead of on its volume (interior). For a very good introduction to the subject see [Hun05]. Roughly speaking, the integral form of the equation of motion is transformed into a surface integral by applying the Green-Gauss theorem. The method achieves substantial speedup because the three dimensional problem is reduced to two dimensions. However, the approach only works for objects whose interior is composed of a homogenous material. Also, topological changes (e.g. fractures) are more difficult to handle than in the explicit FEM approach where only local changes to the stiffness matrix or the connectivity of the elements need to be done.

In the ArtDefo System [JP99] surface nodes are simulated using the Boundary Element Method and a database of precomputed reference boundary value problems (RBVPs). By employing a fast update process, which exploits coherence, accurate real-time deformation simulation is achieved. In [JP02b] the RBVPs are expressed in terms of linear elastostatic Green's functions (LEGFMs) and multiple RBVP's are linked via interface boundary conditions, resulting in a multizone elastokinematic model, which properly simulates large nonlinear relative strains. The system is further augmented by multiresolution Green's functions (wavelet Green's functions) in [JP03], with which large-scale objects can be simulated.

3.5. Mass-Spring Systems

Mass-spring systems are arguably the simplest and most intuitive of all deformable models. Instead of beginning with a PDE such as Eq. 14 and subsequently discretizing in space, we begin directly with a discrete model. As the name implies, these models simply consist of point masses connected together by a network of massless springs.

The state of the system at a given time t is defined by the positions $\mathbf{x}_i$ and velocities $\mathbf{v}_i$ of the masses $i = 1..n$. The force $\mathbf{f}_i$ on each mass is computed due to its spring connections with its neighbors, along with external forces such as gravity, friction, etc. The motion of each particle is then governed by Newton's second law $\mathbf{f}_i = m_i\ddot{\mathbf{x}}_i$, which for the entire particle system can be expressed as

$$\mathbf{M}\ddot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \mathbf{v}) \qquad (26)$$

where $\mathbf{M}$ is a $3n \times 3n$ diagonal mass matrix. Thus, mass-spring systems "only" require the solution of a system of coupled ordinary differential equations (ODEs). This is done via a numerical integration scheme as described in Section 2.2.


Figure 5: *A mass-spring system.*

Depicted in Fig. 5 is a simple example of a volumetric mass-spring system (similar to that in [CZKM98]). The masses are regularly spaced in a lattice. They are connected together along the twelve edges of each hexahedron by "structural" springs. Masses on opposite corners of each hexahedron are connected together by "shear" springs. These springs cause the solid to resist longitudinal and shear deformations respectively, and their rest lengths define the undeformed state of the body. Typically, the spring types have different properties; and for anisotropic materials each springs' properties also depend on its orientation.

Springs are commonly modeled as being linearly elastic; the force acting on mass i, generated by a spring connecting i and j together is

$$\mathbf{f}_i = k_s(|\mathbf{x}_{ij}| - l_{ij})\frac{\mathbf{x}_{ij}}{|\mathbf{x}_{ij}|} \qquad (27)$$

where $\mathbf{x}_{ij}$ is the difference between the two masses' position vectors $(\mathbf{x}_j - \mathbf{x}_i)$, k_s is the spring's stiffness and l_{ij} is its rest length.

Physical bodies are not perfectly elastic; they dissipate energy during deformation. To account for this, viscoelastic springs are used to damp out relative motion. Thus, in addition to (27), each spring exerts a viscous force. It is common to model this as

$$\mathbf{f}_i = k_d(\mathbf{v}_j - \mathbf{v}_i) \qquad (28)$$

where k_d is the spring's damping constant. This is unfortunate, as it damps rigid body rotations. Worse still, when modeling cloth it damps bending and wrinkling — a key feature of these objects. It is preferable to use

$$\mathbf{f}_i = k_d(\frac{\mathbf{v}_{ij}^T\mathbf{x}_{ij}}{\mathbf{x}_{ij}^T\mathbf{x}_{ij}})\mathbf{x}_{ij} \qquad (29)$$

where $\mathbf{v}_{ij} = \mathbf{v}_j - \mathbf{v}_i$. This *projects* the velocity difference onto the vector separating the masses, and only allows a force along that direction.

In the literature, the concept of a mass-spring system is often more general than the canonical example given above. These models are still represented by point masses and a fixed connectivity, but the notion of a spring is generalized. Often the name "particle system" is used instead, which is perhaps a more fitting name. (Though they should be distinguished from the models found in Section 4.1, which have no fixed connectivity.) In any case, the term "mass-spring system" is very suggestive, and will most likely persist.

In more general particle systems, such as those found in [BHW94, BW98, THMG04], deformation energies are defined. These are often derived from soft constraints that are to be maintained in the model; given some constraint of the form $\mathbf{C}(\mathbf{x}) = 0$, an associated energy is defined as $\frac{k_s}{2}\mathbf{C}^T(\mathbf{x})\mathbf{C}(\mathbf{x})$. These energies are minimized at the model's rest state, and are used to enforce the preservation of mesh distances, angles, areas, volumes, etc. Particle forces are then computed as the derivatives of the energies with respect to the particle positions

$$\mathbf{f}_i = -\frac{\partial E}{\partial \mathbf{x}_i}. \qquad (30)$$

Each energy term typically involves only a few particles. In the case of a distance constraint, we are back to our simple, linear spring model described above. However, for other constraints and energies, we must imagine more general spring types: angular, bending, areal, volumetric, etc.

3.5.1. Early Work

Mass-spring networks first saw use in Computer Graphics for facial modeling [PB81, Wat87]. These early works solve static problems of the form (18). Soon after, dynamic models were introduced to model skin, fat and muscle [CHP89, TW90, WT91].

The locomotion of simple creatures such as snakes,

worms and fish [Mil88, TT94] was simulated using mass-spring systems. In these systems, spring rest-lengths vary over time to simulate muscle actuation.

In work by Terzopoulos et al. [TPF89] the equations of thermal conductivity are used to simulate heat transfer in a volumetric mass-spring system. Each spring's stiffness is determined by its temperature, which is set to zero once the melting threshold is exceeded. A discrete fluid model (similar to those presented in Section 4.1) is applied to particles for which all connecting springs have melted.

Breen et al. [BHW94] presented a particle system to model cloth. They posited that cloth is a mechanism of warp and weft fibres — not a continuum — and that mass-spring systems are thus *more* appropriate than Finite Element techniques for modeling cloth. Using measured data from the *Kawabata Evaluation System* [Kaw80], they took care in formulating their energy functions for stretching, bending and trellising (shearing), and predicted the static drape of real materials quite accurately. Since [BHW94], particle systems have dominated the cloth simulation literature — although new FEM formulations (such as the one described in [EKS03]) continue to be proposed.

3.5.2. Drawbacks and Improvements

Particle systems tend to be intuitive and simple to implement. Moreover, they are computationally efficient and handle large deformations with ease. However, unlike the Finite Element and finite difference methods, which are built upon elasticity theory, mass-spring systems are not necessarily accurate. Primarily, most such systems are not convergent; that is, as the mesh is refined, the simulation does not converge on the true solution. Instead, the behavior of the model is dependent on the mesh resolution and topology. In practice, spring constants k_s are often chosen arbitrarily, and one can say little quantitatively about the material being modeled. In addition, there is often coupling between the various spring types. For instance, the "shear" springs of the model in Fig. 5 also resist stretching. For medical applications, as well as for virtual garment simulation in the textile industry, greater accuracy is required. For applications such as film and games, this lack of accuracy may be acceptable; convincing animations have been produced using these systems. However, a judicious choice of model is still advised, as the behavior of the material being modeled can be highly mesh dependent.

Several researchers have explored and have tried to mitigate the accuracy issues in mass-spring systems. Kass [Kas95] presents a simple equivalence in one dimension between a mass-spring system and a corresponding finite difference spatial discretization. Eischen and Bigliani [HB00] perform a comparison between a Finite Element model and a carefully tuned particle system, which gave similar experimental results for small deformations. Etzmuss et al. [EGS03] carefully derive a cloth particle system from a continuous formulation by a finite difference dis-

Figure 6: *Cloth carefully modeled using a mass-spring system. Image courtesy of Robert Bridson, UBC.*

cretization. Their model is convergent, and they show how to choose spring constants based on continuous material parameters.

These more accurate particle systems, however, only apply to rectangular meshes. Like finite difference techniques, they do not generalize easily to triangular (or tetrahedral) meshes. Volino and Thalmann [VMT97] present a triangular mesh which attempts to model the physical properties of cloth regardless of edge orientations. However, the accuracy of their method is unclear. Baraff and Witkin [BW98] model cloth using a triangular mesh, deriving in-plane particle forces from a continuum formulation. Their approach supports anisotropic behavior, but is not convergent. In [BC00], the authors present a novel mass-spring system which gives consistent results (though it is not convergent) for tetrahedral and hexahedral meshes, independent of the orientation of the elements. In their system, springs emanate from the barycenter of each element along principle coordinate axes and are attached to element faces; forces acting on each particle are then computed via interpolation. Their approach also allows the user to control anisotropic material behavior. Teschner et al. [THMG04] model volumes and surfaces using tetrahedral and triangular meshes respectively. They employ generalized springs which preserve distances, areas and volumes. Their model is efficient, convergent (though it does not incorporate continuous material properties), and supports plastic deformation.

Bending models for surfaces tend to be particularly ad-hoc; Bridson et al. [BMF03] present a physically correct bending model for triangle meshes by isolating the bending mode from all other modes of deformation (Fig. 7). Grinspun et al. present a similar bending model for the simulation of discrete shells [GHDS03].

An interesting approach is to use learning algorithms to *search* for mesh parameters. (See [BTH*03, BSSH04] and references therein.) Bhat et al. use simulated annealing to estimate the spring constants in a cloth mesh. They employ Baraff and Witkin's model [BW98], and use video of real cloth as their reference for comparison. Bianchi et al. use a genetic algorithm to identify spring constants as well as mesh topology in a volumetric mass-spring system. They

Figure 7: *A bending model for triangle meshes with nonzero rest angles [BMF03]. The model on the right has significantly stronger bending resistance, giving it a "water bottle" effect. Image courtesy of Robert Bridson, UBC.*

compare their results to FEM reference solutions, and obtain close correspondence. While such methods are general, they are computationally expensive and require reference solutions; moreover, the results of a specific search may not hold for different mesh resolutions.

Adaptive meshing techniques are particularly complicated by the non-convergent behavior of many mass-spring systems, as they *require* modeling at multiple resolutions. (Another difficulty is the handling of so called T-junctions between regions of differing resolutions, though this problem must also be addressed in FEM and other methods.) A number of authors have addressed this issue for cloth simulation. In these systems, splitting/merging criterion are often based on bending angles between triangles. Villard and Borouchaki [VB02] apply adaptive refinement to a simple, convergent model based on quadrilateral meshes. They show consistent results between meshes of multiple resolutions, maintaining rich detail at reduced computation times. Earlier work by Hutchinson et al. [HPH96] takes a similar approach, though the convergence of their model is not clear as they couple their spatial and time discretizations together in an ad-hoc manner. Li and Volkov [LV05] present an adaptive framework for triangular cloth meshes, avoiding the problematic T-junctions. They present attractive results, though their physical model is not convergent.

Finally, a number of researchers have sought to improve the realism of particle systems by modeling nonlinear material properties. For instance, Eberhardt et al. [EWS96] base their spring model on measured cloth data to model hysteresis effects; Choi and Ko [CK02] approximate cloth's buckling response using a fifth order polynomial.

3.5.3. Time Integration

Once a force model is chosen, the particle system is stepped forward in time via a numerical integration scheme as in Section 2.2. Time integration schemes have received particular attention in the cloth simulation literature (see [HES02]). As cloth is often modeled by mass-spring systems, we further discuss the subject of time integration here.

Cloth is a *highly* anisotropic material due to its structure: it resists bending weakly, but has a relatively strong resistance to stretching. When simulating a highly stiff volume

of material, one can employ a reduced deformation model, such as modal analysis (see Section 5) — and in the limit, use a rigid body model. Clearly, this cannot be done with cloth. We are interested in visualizing the large scale, out-of-plane folding and wrinkling of cloth, however we must still deal with these planar energies in our model. Numerically speaking, the resulting ODE system (26) is *stiff*; that is, it possesses a wide range of eigenvalues. This forces us to take excessively small time steps when using an explicit scheme.

Provot [Pro95] tackles the issue of stiffness by using much weaker stretching energies and then post-processing the cloth mesh at each time step, iteratively enforcing constraints. Springs that are stretched by more than 10% are relaxed (shortened). This in turn stretches nearby neighboring springs which are then relaxed, and so on until convergence is obtained. This, in effect, also models the biphasic nature of cloth — small deformations are resisted weakly until a threshold is reached, whereupon stiffness dramatically increases. In practice, this method gives efficient, attractive results, though convergence is not guaranteed.

Baraff and Witkin [BW98] present a linear implicit scheme which allows for large time steps while maintaining stability. This has proven to be a robust and efficient solution to the stiffness problem, and has become a popular technique. Specifically, they apply a backward Euler scheme to Eq. 26, giving

$$\begin{bmatrix} \Delta \mathbf{x} \\ \Delta \mathbf{v} \end{bmatrix} = \Delta t \begin{bmatrix} \mathbf{v}_n + \Delta \mathbf{v} \\ M^{-1} \mathbf{f}(\mathbf{x}_n + \Delta \mathbf{x}, \mathbf{v}_n + \Delta \mathbf{v}) \end{bmatrix} \quad (31)$$

where $\mathbf{x}_n = \mathbf{x}(t)$ and $\mathbf{v}_n = \mathbf{v}(t)$. This is a nonlinear equation in $\Delta \mathbf{x}$ and $\Delta \mathbf{v}$. A linear implicit version of (31) is obtained by using a first order Taylor series expansion of $\mathbf{f}$

$$\mathbf{f}(\mathbf{x}_n + \Delta \mathbf{x}, \mathbf{v}_n + \Delta \mathbf{v}) = \mathbf{f}_n + \frac{\partial \mathbf{f}}{\partial \mathbf{x}} \Delta \mathbf{x} + \frac{\partial \mathbf{f}}{\partial \mathbf{v}} \Delta \mathbf{v} \quad (32)$$

where $\frac{\partial \mathbf{f}}{\partial \mathbf{x}}$ and $\frac{\partial \mathbf{f}}{\partial \mathbf{v}}$ are the Jacobian matrices of the particle forces with respect to position and velocity respectively. In their paper, Baraff and Witkin derive expressions for the Jacobians of the various internal forces in their model. Due to the local connectivity structure of the mesh, these are sparse matrices. They then solve the resulting linear system at each time step

$$A \Delta \mathbf{v} = \mathbf{b}, \quad (33)$$

where

$$A \equiv (I - \Delta t M^{-1} \frac{\partial \mathbf{f}}{\partial \mathbf{v}} - \Delta t^2 M^{-1} \frac{\partial \mathbf{f}}{\partial \mathbf{x}}) \quad (34)$$

$$\text{and} \quad \mathbf{b} \equiv \Delta t M^{-1} (\mathbf{f}_n + \Delta t \frac{\partial \mathbf{f}}{\partial \mathbf{x}} \mathbf{v}_n). \quad (35)$$

They solve this system, in the presence of constraints, using a modified preconditioned conjugate gradient solver (MPCG). This is equivalent to applying one Newton iteration to Eq. 31. Although solving this system increases the

Figure 8: *Decomposing cloth into parts which can be solved more efficiently [BA04].*

computational cost at each time step, this is more than offset by the ability to take large steps when desired. However, as illustrated by Desbrun et al. [DSB99], the larger the time step size, the more artificial damping is added to the system.

Others have since attempted to improve upon Baraff and Witkin's approach. Desbrun et al. [DSB99] precompute the linear part of A, achieving an $O(n)$, unconditionally stable scheme. Their technique, however, is inaccurate and does not generalize well to large systems. Kang et al. [KCC*00] improve upon this approximation, but ultimately are only replacing the CG iterations of [BW98] with a single, Jacobi-like iteration. Volino and Magnenat-Thalmann [VMT00] use a weighted implicit-midpoint method that appeared to give attractive results, but which is less stable and may be difficult to tune in practice. Choi and Ko [CK02] use the more accurate second-order backward difference formula (BDF2); moreover, their coupled compression/bending model improves the stability of the system, eliminating the need for fictitious spring damping.

Instead of applying an implicit scheme to the entire system, a number of researchers have employed implicit-explicit (IMEX) schemes [ARW95]. The essential idea is to separately treat the stiff and non-stiff parts of the ODE, handling the stiff parts with an implicit method and the non-stiff parts with an explicit method. This combines the stability of an implicit scheme where needed with the simplicity and efficiency of an explicit scheme where possible. Hauth et al. [HES02] base their IMEX splitting on connection type: stretch springs are handled implicitly, whereas shear and bend "springs" are handled explicitly. This simplifies implementation and improves the sparsity of A — and thus the cost of solving Eq. 33. The authors also use BDF2, and embed their linear solver within a Newton solver, making theirs a fully (as opposed to linear) implicit technique. Boxerman and Ascher [BA04] take a similar approach, optimizing their IMEX splitting adaptively based on a local stability criterion. Each spring is handled either explicitly or implicitly, based on time step, local mesh resolution and material properties. They then take further advantage of the improved system sparsity to decompose the mesh into components which can be solved more efficiently (Fig. 8).

Bridson et al. [BMF03] apply a similar IMEX splitting to cloth as that used for advection-diffusion equations [ARW95], applying an implicit method to the damping term and an explicit method to the elastic term. This is appropriate when damping plays a significant role, as damping imposes quadratic time step restrictions with respect to mesh size (as opposed to a linear one for elastic terms). They incorporate this within a second order accurate integration scheme, and include a strain limiting step. Coupled with their careful handling of bending and collisions, they produced visually stunning cloth animations.

4. Lagrangian Mesh Free Methods

4.1. Loosely Coupled Particle Systems

Particle systems were developed by Reeves [Ree83] for explosion and subsequent expanding fire simulation in the movie "Star Trek II: The Wrath of Khan". The same technique can also be used for modeling other fuzzy objects such as clouds and water, i.e. objects that do not have a well-defined surface. Particles are usually graphical primitives such as points or spheres, however, they might also represent complex group dynamics such as a herd of animals [Rey87]. Each particle stores a set of attributes, e.g. position, velocity, temperature, shape, age and lifetime. These attributes define the dynamical behavior of the particles over time and are subject to change due to procedural stochastic processes. Particles pass through three different phases during their lifetime: generation, dynamics and death.

In [Ree83], particles are points in 3D space which represent the volume of an object. A stochastic process generates particles in a predetermined *generation shape* which defines a region about its origin into which the new particles are randomly placed. Attribute values are either fixed or may be determined stochastically. Initially, particles move outward away from the origin with a random speed. During the dynamics phase, particle attributes might change as functions of both time and attributes of other particles. Position is updated by simply adding the velocity. Finally, a particle dies if its lifetime reaches zero or if it does not contribute to the animation anymore, e.g. if it is outside of a region of interest. A particle is rendered as a point light source which adds an amount of light depending on its color and transparency attribute. Motion blurring is very easy to achieve by rendering a particle as a streak using its position and velocity. An advantage of particles is their simplicity, which enables the animation of a huge number of particles for complex scenes. The procedural definition of the model and its stochastic control simplifies the human design of the system. Furthermore, with particle hierarchies, complicated fuzzy objects such as clouds can be assembled and controlled. Although the particles are simulated omitting inter-particle forces, the resulting animations are convincing and fast for inelastic phenomena. Therefore, the technique has been widely employed in movies and video games. Examples of modeling waterfalls, ship wakes, break-

ing waves and splashes using particle systems can be found in [FR86, Pea86, Gos90, Sim90, OH95].

Particles which interact with each other depending on their spatial relationship are referred to as spatially coupled particle systems [Ton92]. The interactions between particles evolve dynamically over time, thus, complex geometry and topological changes can be easily modeled using this approach. Tonnesen presented a framework for physically based animation of solids and liquids based on dynamically coupled particles which represent the volume of an object [Ton91, Ton98]. Each particle p_i has a potential energy ϕ_i which is the sum of the pairwise potential energies ϕ_{ij} between p_i and all other other particles p_j, i.e.

$$\phi_i = \sum_{j \neq i} \phi_{ij}. \tag{36}$$

The force $\mathbf{f}_i$ exerted on p_i with position $\mathbf{x}_i$ is then

$$\mathbf{f}_i = -\nabla_{\mathbf{x}_i} \phi_i = -\sum_{j \neq i} \nabla_{\mathbf{x}_i} \phi_{ij}. \tag{37}$$

So far, all particles interact with each other, resulting in $O(n^2)$ complexity where n is the number of particles. The computational costs for computing the force can be reduced to $O(n)$ when restricting the interaction to a neighborhood within a certain distance, and $O(n \log n)$ for neighbor searching. To avoid discontinuities at the neighborhood boundary, the potentials are weighted with a continuous, smooth and monotonically decreasing weighting function which depends on the distance to the particles and ranges from one at the particle position to zero at the neighborhood boundary.

For deriving inter-particle forces, the Lennard-Jones potential ϕ_{LJ} is used:

$$\phi_{LJ}(d) = \frac{B}{d^n} - \frac{A}{d^m}, \tag{38}$$

where d is the distance between two particles, and n, m, B and A are constants. The Lennard-Jones potential is well known in molecular dynamics for modeling the interaction potential between pairs of atoms. It creates long-range attractive and short-range repulsive forces, yielding particles arranged into hexagonally ordered 2D layers in absence of external forces. A more convenient formulation, called the Lennard-Jones bi-reciprocal function, is written as

$$\phi(d) = \frac{-e_o}{m-n} \left(m (\frac{d_0}{d})^n - n (\frac{d_0}{d})^m \right), \tag{39}$$

where d_0 is the *equilibrium separation distance*, and $-e_0$ is the minimal potential (the magnitude is called the *dissociation energy*). Increasing the dissociation energy increases the stiffness of the model, while with the exponents n and m the width of the potential well can be varied. Therefore, large dissociation energy and high exponents yield rigid and brittle material, while low dissociation energy and small exponents result in soft and elastic behavior of the object. This allows the modeling of a wide variety of physical properties

ranging from stiff to fluid-like behavior. By coupling the dissociation energy with thermal energy such that the total system energy is conserved, objects can be melted and frozen. Furthermore, thermal expansion and contraction can be simulated by adapting the equilibrium separation distance d_0 to the temperature.

One problem of particle systems is that the surface is not explicitly defined. Blinn [Bli82] proposed using an algorithm which was developed to model electron density maps of molecular structures. A Gaussian potential $\varphi_i(\mathbf{x}) = be^{-ad^2}$ (which is often called a *blob*) is assigned to each particle, where a and b are constants and $d = \|\mathbf{x} - \mathbf{x_i}\|$ is the distance from an arbitrary point $\mathbf{x}$ in space to the particle's position $\mathbf{x}_i$. A continuously defined potential field $\varphi(\mathbf{x})$ in space is obtained by summing the contribution from each particle

$$\varphi(\mathbf{x}) = \sum_i \varphi_i(\mathbf{x}). \tag{40}$$

The surface is then defined as an iso-value S of φ. This yields an implicit coating of the particle which handles topological changes such as splitting and merging by construction. For a more intuitive control of the surface, the constants a and b can be computed as $a = -B/R^2$ and $b = Se^{-B}$, where R is the radius in isolation and B the *blobiness*.

The implicit coating of particles for soft inelastic substances undergoing topological changes pose challenging problems such as volume preservation, avoiding unwanted blending and contact modeling. These were addressed by Desbrun and Cani in a series of papers [Can93, DC94, DC95]. A hybrid model is used which is composed of two layers: Particles are used to simulate soft inelastic substances as described above, while an elastic implicit layer defines the surface of an object and is locally deformed during collisions. A problem of the implicit coating is that the volume may change significantly during deformation, especially for splitting and merging. However, efficiently computing the volume of a soft substance is not trivial. A *territory* of a particle p_i is defined as the (volumetric) part V_i of the object where the field contribution of p_i is the highest. Note that territories form a partition of the implicit volume of an object. Each particle samples its territory boundary by sending a fixed number of points called *seeds* in a set of distributed directions until they reach the boundary. The volume of a particle is approximated by simply summing up the distances from the particle to the seeds. The local volume variation can then be easily approximated for each particle, and the field function is changed accordingly such that the volume is preserved. Another problem is that split object parts might blend with each other when they come close. To avoid this, an *influence graph* is built at each animation step by recursively adding the neighbors of a particle which are in its sphere of influence to the same influence connected component. Only the particles of the same component can interact and their field functions are

blended. However, the problem arises that two or more separated components might collide. For detecting a collision, the seed points on the iso-surface S are tested against the field function of another component. For resolving interpenetrations between two components with potential functions φ_1 and φ_2, exact contact surfaces are computed by applying negative compressing potentials $g_{2,1}$ and $g_{1,2}$ such that $\varphi_1 + g_{2,1} = \varphi_2 + g_{1,2} = S$, resulting in a local compression. To compensate the compression and ensure C^1 continuity of the deformed surfaces, positive dilating potentials are applied in areas defined around the interpenetration zone [Can93, OC97]. For collision response, the compression potentials $g_{i,j}$ are computed for each colliding seed and then transmitted as response force to the corresponding particle. Additionally, the two components might be merged locally where the collision force exceeds a threshold.

Szeliski and Tonnesen introduced oriented particles for deformable surface modeling [ST92, Ton98], where each particle represents a small surface element (called surfel). Each surfel has a local coordinate frame, given by the position of the particle, a normal vector and a local tangent plane to the surface. To arrange the particles into surface-like arrangements, interaction potentials are defined which favor locally planar or locally spherical arrangements. The Lennard-Jones potential described above is used to control the average inter-particle spacing. The weighted sum of all potentials yields the energy of a particle, where the weights control the bending and stiffness of the surface. Variation of the particle energy with respect to its position and orientation yields forces acting on the positions and torques acting on the orientations, respectively. Using these forces and torques, the Newtonian equations for linear and angular motion are solved using explicit integration as described in Section 2.2.

Recently, Bell et al. presented a method for simulating granular materials, such as sand and grains, using a particle system [BYM05]. A (non-spherical) grain is composed of several round particles, which move together as a single rigid body. Therefore, stick-slip behavior naturally occurs. Molecular dynamics (MD) is used to compute contact forces for overlapping particles. The same contact model is used for collision of granular materials with rigid bodies, or even between rigid-bodies, by simply sampling the rigid body surface with particles.

4.2. Smoothed Particle Hydrodynamics

Smoothed Particle Hydrodynamics (SPH) is a Lagrangian technique where discrete, smoothed particles are used to compute approximate values of needed physical quantities and their spatial derivatives. Forces can be easily derived directly from the state equations. Furthermore, as a particle-based Lagrangian approach it has the advantage that mass is trivially conserved and convection is dispensable. This reduces both the programming and computational complexity

Figure 9: *Simulating interface tension forces between fluids of different polarity, temperature diffusion and buoyancy [MSKG05].*

and is therefore suitable for interactive applications. A drawback of SPH is that it is not possible so far to exactly maintain the incompressibility of material.

Following, we will give a short introduction of the SPH method and discuss applications in Computer Graphics. For a more detailed introduction we refer the reader to the excellent paper of Monaghan [Mon92].

A function A is interpolated at a position $\mathbf{x}$ from its neighboring particles using a summation interpolant:

$$A(\mathbf{x}) = \sum_j m_j \frac{A_j}{\rho_j} W(\mathbf{r}, h), \qquad (41)$$

where A_j denotes the function value of a particle p_j at $\mathbf{x}_j$ (the interpolation point), m_j and ρ_j are the mass and density of a particle p_j, respectively, and $\mathbf{r} = \mathbf{x} - \mathbf{x_j}$. $W(\mathbf{r}, h)$ is a smoothing kernel with the properties $\int W(\mathbf{r}, h) d\mathbf{r} = 1$ and $\lim_{h \to 0} W(\mathbf{x}, h) = \delta(\mathbf{x})$, where h is the support radius and δ the Dirac-function. For efficiency reasons, h is usually chosen such that the kernel $W(\mathbf{r}, h)$ has compact support. The choice of the smoothing kernel $W(\mathbf{r}, h)$ is very important. Often spline Gaussian kernels are used, see e.g. [Mon92] for a discussion of kernels. If $W(\mathbf{r}, h)$ is differentiable, a differentiable interpolant of a function can be derived by simply computing the gradient of $W(\mathbf{r}, h)$, i.e.,

$$\nabla A(\mathbf{x}) = \sum_j m_j \frac{A_j}{\rho_j} \nabla W(\mathbf{r}, h). \qquad (42)$$

Therefore, there is no need for finite differences or a grid. To obtain higher accuracy the interpolant can be computed as

$$\rho \nabla A(\mathbf{x}) = \nabla(\rho A(\mathbf{x})) - A(\mathbf{x}) \nabla \rho. \qquad (43)$$

The density is estimated at an arbitrary point as

$$\rho(\mathbf{x}) = \sum_j m_j W(\mathbf{r}, h). \qquad (44)$$

The particle densities could be computed for each time step. In practice, this is often not appropriate because the density drops close to the object boundary. Furthermore, difficulties arise when adapting the spatial resolution of the particles [DC99]. Instead, we can assign an initial density to each particle. The continuity equation (conservation of mass) is then used for computing the variation of density over time

$$\dot{\rho}_i = -\rho_i \nabla \mathbf{v}_i, \qquad (45)$$

Figure 10: *An elastic solid is dropped onto a heated box and melts to a fluid [KAG*05].*

where the divergence of velocity $\nabla \mathbf{v}_i$ is approximated by

$$\nabla \mathbf{v}_i = \frac{1}{\rho_i} \sum_{j \neq i} m_j (\mathbf{v}_j - \mathbf{v}_i) \cdot \nabla W(\mathbf{r}, h). \qquad (46)$$

The disadvantage of this update of the density relative to the motion of the neighboring particles is that exact mass conservation is not guaranteed.

The equations of motion are solved by deriving forces. E.g. the pressure force can be estimated using Eq. 43

$$\mathbf{f}_i^{pressure} = -m_i \sum_{j} m_j \left(\frac{P_i}{\rho_i^2} + \frac{P_j}{\rho_j^2} \right) \nabla W(\mathbf{r}, h), \qquad (47)$$

where the force is symmetrized to fulfill the action-reaction principle of Newton's third law. For keeping a constant rest density ρ_0, the pressure P_i is computed by a variation of the ideal gas state equation [DC96]

$$P_i = k(\rho - \rho_0), \qquad (48)$$

where k is a gas constant.

For smoothed particles, local stability criteria have been found which, in conjunction with time-adaptive integration, yield both stable and efficient animations. The Courant-Friedrichs-Lewy (CFL) criterion requires that each particle i must not be passed by, giving a limit for the time step $\Delta t \leq \lambda h/c$, where λ is the Courant number and c is the speed of sound of the material, which is the maximum velocity of a deformation wave inside the material. Pressure waves propagate at speed $c = \sqrt{\partial P/\partial \rho}$, therefore using Eq. 48 induces $c = \sqrt{k}$. Considering viscosity further decreases the maximum time step, see [Mon92] and [DC96] for details.

Generally, the animation quality and accuracy increases with a higher number of particles. In [DC99], an adaptive framework is proposed where space and time resolutions are chosen automatically. Individual particles are refined where pressure differences are above a user-defined threshold. A refined particle p_i with mass m_i and volume m_i/ρ_0 is replaced by n particles with smaller volume and mass m_i/n such that they occupy the same volume as p_i. Considering the particles as spheres with radius proportional to their support radius, an individual support radius h_i can be computed as $h_i = \xi \sqrt[3]{m_i/\rho_0}$, where the constant ξ is chosen according to the desired average number of neighboring particles. Neighboring particles can be merged in stable areas where pressure is almost uniform. However, since a particle is isotropic, the neighboring particles should approximately

fill a sphere. The new particle is positioned at the center of gravity of the replaced ones. Its mass and velocity are computed such that mass and linear momentum are conserved.

Considering the object as a set of smeared-out masses, the surface can be defined as an iso-surface of the mass density function [DC96]. This yields an implicit representation coherent with the physical model. Desbrun and Cani [DC98] improve this implicit coating of particles by using an active surface which evolves depending on a velocity field, similar to snakes [KWT88]. Therefore, surface tension and other characteristics such as constant volume can be added to the surface model.

4.2.1. Applications

SPH was introduced independently by Gingold and Monaghan [GM77] and Lucy [Luc77], for the simulation of astrophysical problems such as fission of stars. Stam and Fiume introduced SPH to the Computer Graphics community to depict fire and other gaseous phenomena [SF95]. They solve the advection-diffusion equation for densities composed of "warped blobs". Desbrun and Cani solve the state equations for the animation of highly deformable bodies using SPH [DC96]. They achieve the animation of inelastic bodies with a wide range of stiffness and viscosity. Lava flows are animated by coupling viscosity with a temperature field and simulated heat transfer between the particles [SAC*99]. By considering hair as a fluid-like continuum, Hadap and Magnenat-Thalmann [HMT01] used a modified formulation of SPH to simulate hair-hair interactions. Premože et al. introduced the use of the moving particle semi-implicit method (MPS) for simulating incompressible multiphase fluids [PTB*03]. Impressive visual results were produced by coupling the physical particles with level sets for surface reconstruction. Müller et al. presented a method based on SPH and new smoothing kernels, with which fluids with free surfaces can be simulated at interactive rates with up to 5000 particles [MCG03]. Furthermore, they proposed a term to model surface tension effects. In [MST*04], the interaction of Lagrangian fluids and mesh-based deformable solids are modeled by placing virtual boundary particles, so-called ghost particles [Mon94], on the surface of the solid objects according to Gaussian quadrature. Recently, the SPH method was extended in [MSKG05] so that the simulation of phenomena such as boiling water, trapped air and the dynamics of a lava lamp are possible (Fig. 9). Liquids with different polarities are simulated by

Figure 11: *Highly plastic deformations and ductile fracture [PKA*05].*

computing a body force which acts perpendicular to the interface of two liquids. The force is proportional to the curvature of the interface and the interface tension. Additionally, air particles are generated and deleted dynamically where air pockets are likely to be formed, making it possible to simulate trapped air.

4.3. Mesh Free Methods

Interesting and fairly new approaches to deformable modeling are the so-called *meshfree methods for the solution of partial differential equations*, which originated in the FEM community approximately a decade ago ([NTVR92], [Suk03]). For an extensive and up-to-date classification and overview of mesh-free methods, see [FM03], [Liu02] and also the excellent review paper [BKO*96]. A nice introduction to the element-free Galerkin method is given in [Ask97].

4.3.1. Point-based Animations

Recently, the combination of mesh-free physics with point-sampled surfaces [PZvBG00] in so-called point-based animations [pba04] have become popular. Müller et al. introduced to the graphics community a mesh-free continuum-mechanics-based framework for the animation of elastic, plastic and melting objects [MKN*04]. Elastic body forces are derived via the strain energy density

$$U = \frac{1}{2}(\varepsilon \cdot \sigma), \qquad (49)$$

where ε is the strain and σ the stress tensor as described in Section 3.1. The elastic force per unit volume at a particle p_i with material coordinates $\mathbf{m}_i$ is computed using the directional derivative $\nabla_{\mathbf{u}_i}$. For a Hookean material, i.e. $\sigma = E \cdot \varepsilon$ (see Eq. 4), this yields

$$\mathbf{f}_i = -\nabla_{\mathbf{u}_i} U = -\frac{1}{2}\nabla_{\mathbf{u}_i}(\varepsilon_{\mathbf{i}} \cdot \mathbf{E} \cdot \varepsilon_{\mathbf{i}}) = -\sigma \cdot \nabla_{\mathbf{u}_i}\varepsilon_{\mathbf{i}}, \qquad (50)$$

where Green's nonlinear strain tensor (see Eq. 1) is used for ε_i. For computing ε_i, the spatial derivatives of the displacement field $\nabla\mathbf{u}_i$ at $\mathbf{m}_i$ is needed. To guarantee zero elastic forces for rigid body modes, the approximation of $\nabla\mathbf{u}_i$ from the displacement vectors $\mathbf{u}_j$ of the neighboring particles must be at least first order accurate. Therefore, the Moving Least Squares method [LS81] with a linear basis is used which yields first order accurate interpolation of point-sampled functions (contrary to the SPH approximation described in Section 4.2 which is not even zeroth-order

accurate). For the following, we will only consider the x-component u of the displacement field $\mathbf{u} = (u, v, w)^T$. The basic idea is to approximate a continuous scalar field $u(\mathbf{m})$ using a Taylor approximation. For particles p_j close to a particle p_i we get a first order approximation $\tilde{u}_j$ of the values u_j as

$$\tilde{u}_j = u_i + \nabla u|_{\mathbf{m}_i} \cdot \mathbf{m}_{ij}, \qquad (51)$$

where $\mathbf{m}_{ij} = \mathbf{m}_j - \mathbf{m}_i$ and $\nabla u|_{\mathbf{m}_i} = (u_{,x}, u_{,y}, u_{,z})^T$ at $\mathbf{m}_i$. The index after the coma denotes a spatial derivative. The weighted least-squares error e_i of the approximation $\tilde{u}_j$ is given by

$$e_i = \sum_j (\tilde{u}_j - u_j)^2 w_{ij}, \qquad (52)$$

where w_{ij} is a normalized, continuously defined and monotonically decreasing weighting function. We want to find the unknowns $u_{,x}$, $u_{,y}$ and $u_{,z}$ such that the error e is minimized. Therefore, we set the derivatives of e with respect to $u_{,x}$, $u_{,y}$ and $u_{,z}$ to zero, yielding three equations for the three unknowns

$$\left(\sum_j \mathbf{m}_{ij}\mathbf{m}_{ij}^T w_{ij}\right) \nabla u|_{\mathbf{m}_i} = \sum_j (u_j - u_i)\mathbf{m}_{ij}w_{ij}. \qquad (53)$$

Finally, we obtain the spatial derivatives of $u(\mathbf{m})$ at $\mathbf{m}_i$ as

$$\nabla u|_{\mathbf{m}_i} = \mathbf{A}^{-1}\left(\sum_j (u_j - u_i)\mathbf{m}_{ij}w_{ij}\right), \qquad (54)$$

where the inverse of $A = \sum_j \mathbf{m}_{ij}\mathbf{m}_{ij}^T w_{ij}$ needs to be computed only once per particle and can be used for computing the derivatives of v and w as well.

With this approach, material properties ranging from stiff elastic to highly plastic can be simulated. In [KAG*05], the authors describe how to extend this range to viscous materials such as fluids by merging the solid mechanics equation with the Navier-Stokes equations. Viscosity, pressure and surface tension forces are computed using SPH as described in Section 4.2 and added to the elastic force. Note that this allows the simulation of non-realistic materials such as elastic fluids, similar to [GBO04]. The material properties such as stiffness, viscosity, plasticity and cohesion can be defined per particle. By coupling these properties to the temperature of a particle, local melting and freezing of objects can be achieved (Fig. 10).

In the case of elastic materials which do not undergo

any topological changes, the surface can be simply dragged along with the particles [MKN*04]. This is done by computing the displacement for each surface element (surfel) from the displacements of neighboring particles p_j. The first order approximation of $\nabla \mathbf{u}$ at $\mathbf{m}_j$ can be used to obtain a smooth displacement vector field which is invariant under linear transformations:

$$\mathbf{x}(\mathbf{m}_{s_i}) = \mathbf{m}_{s_i} + \sum_j \overline{\omega}(\mathbf{m}_{s_i}, \mathbf{m}_j)(\mathbf{u}_j + \nabla \mathbf{u}|_{\mathbf{m}_j}(\mathbf{m}_j - \mathbf{m}_{s_i})),$$

(55)

where $\mathbf{m}_{s_i}$ and $\mathbf{x}(\mathbf{m}_{s_i})$ are the material and world coordinates of a surfel s_i, respectively, and $\overline{\omega}$ is a normalized weighting kernel. The deformation is computed and applied to both the center and the axes of a surfel. To adapt the sampling of the surface in case of stretching or compression, surfels can be split or merged similar to [PKKG03]. The animation of the physics particles combined with this surface skin allows the simulation of elastic and plastic materials of detailed models in real-time. Furthermore, Adams et al. [AKP*05] showed that the surface can be efficiently raytraced by exploiting the deformation field for bounding hierarchy updates.

However, dragging the surface along with the particles fails if topological changes occur. In this case, an implicit representation can be used as discussed in the Section 4.2. In [MKN*04], this implicit representation is sampled with surfels. After an animation step the new position of the surfels is estimated using the displacement computation in Eq. 55. The deformed surfels are then projected onto the implicit surface. Finally, a resampling and relaxation scheme ensures that the surface is completely covered and uniformly sampled with surfels. In [KAG*05], the implicit representation is improved by applying energy potentials and geometric constraints which give better control over the surface. However, as the particle representation of the volume is usually much coarser than the surface representation, the implicit coating of the particles cannot represent fine detail of the surface. A possible solution is to use the explicit deformation approach of (55) for solids (with low temperature), the implicit representation for fluids (with high temperature), and blend between the explicit and the implicit representation depending on the temperature for melting object parts. This way, surface detail smoothly fades out and the melting surface approaches the implicit representation. Freezing fluids are handled analogously.

The framework is extended in [PKA*05] to cope with fracturing (Fig. 11). Cracks are created at surfels where the main principal stress exceeds a threshold. The crack propagates in a plane perpendicular to the main principal stress, where the crack surface is dynamically sampled with surfels. Visibility tests detect neighboring particles, which are separated by a crack surface. A splitting, merging and termination operator handles the topological events which occur when cracks merge or branch. Furthermore, a resampling scheme adapts the particles resolution to handle small fractured object pieces. With the suggested methods a wide range of materials can be fractured, from stiff elastic to highly plastic objects that exhibit brittle and/or ductile fracture.

Collision detection of point-based models has been discussed for rigid [KZ04], quasi-rigid [PPG04] and deformable objects [KMH*04]. For rigid collision detection, a bounding sphere hierarchy is built with the surfels as leaves, which is suitable for time-critical collision detection [KZ04]. Pauly et. al [PPG04] compute a consistent contact surface and the traction distribution by solving a Linear Complementarity Problem (LCP). Collision response forces are computed from the local tractions and integration yields the total wrench that acts on the quasi-rigid bodies. Keiser et al. propose to decouple the collision handling and deformation to achieve stable and efficient contact handling for deformable objects [KMH*04]. Collisions are detected using a high resolution surface sampled with surfels and a simple scheme is used to compute a consistent contact surface. A penalty force depending on its displacement onto the contact surface is computed for each surfel. These forces are then distributed to the neighboring particles where they are applied as external collision response forces in the next animation step.

Wicke et al. propose a method for animating point-sampled thin shells [WSG05]. The geometric surface properties are approximated using splines (so-called *fibres*) embedded in the surface. Physical effects such as plasticity and fracturing can be easily modeled. The authors also describe a skinning technique for point-sampled surfaces, with which a high-resolution surface can be animated using a sparse set of samples.

4.3.2. Meshless Deformations Based on Shape Matching

Müller et al. [MHTG05] propose a meshless method for animating deformable objects which does not fit into any of the preceding categories. The nodes of a volumetric mesh are treated as point masses and animated as a simple particle system without connectivity. Then, at every time step, the original configuration of the points (the rest state mesh) is fitted to the actual point cloud in the least squares sense, using shape matching techniques for point clouds with correspondence. The fitted rest shape yields goal positions for all point masses. Each point mass is then pulled towards its goal position while the displacement divided by the time step is added to the velocity of the point. This geometric integration scheme, although explicit, is shown to be unconditionally stable (Fig. 12).

Figure 12: *Stability and the ability to recover from highly deformed or inverted configurations, shown in [MHTG05].*

5. Reduced Deformation Models and Modal Analysis

5.1. Basic Formulation

As described by James and Pai [JP04], given N undeformed point locations $\mathbf{p} = [\mathbf{p}_1, \ldots, \mathbf{p}_N]$, a reduced deformation model approximates deformed point locations $\mathbf{p}'$ by a linear superposition of M displacement fields (the columns $\mathbf{U}_i$ of $\mathbf{U}$ in Eq. 56). The amplitude of each displacement field is given by the reduced coordinates $\mathbf{q}$ such that

$$\mathbf{p}' = \mathbf{p} + \mathbf{U}\mathbf{q}, \tag{56}$$

as shown in Fig. 13. In [JP04] it is stated that the columns of $\mathbf{U}$ could arise from any possibly nonlinear black box process, such as an interpolation process, multiresolution modeling or nonlinear/linear modal analysis. The latter will be the focus of this section.

Figure 13: *Reduced deformation models: (a) reference shape* $\mathbf{p}$, *(b) displacement field* $\mathbf{U}_1$, *(c) displacement field* $\mathbf{U}_2$, *and (d) one possible deformed shape* $\mathbf{p}' = \mathbf{p} + \mathbf{U}_1 + 0.5\mathbf{U}_2$.

5.2. Linear Modal Analysis

Pentland and Williams [PW89] pioneered the use of reduced deformable models in Computer Graphics, using modal analysis. Given the mass, damping and stiffness matrices $\mathbf{M}$, $\mathbf{C}$ and $\mathbf{K}$, Eq. 20 can be decoupled into $N = 3n$ linearly independent ordinary differential equations (ODEs) by solving the generalized eigenvalue problem

$$\mathbf{M}\mathbf{\Phi}\mathbf{\Lambda} = \mathbf{K}\mathbf{\Phi}, \tag{57}$$

such that $\mathbf{\Phi}^T\mathbf{M}\mathbf{\Phi} = \mathbf{I}$ and $\mathbf{\Phi}^T\mathbf{K}\mathbf{\Phi} = \mathbf{\Lambda}$ are diagonal matrices. The entries of $\mathbf{\Lambda}$ contain the eigenvalues, and the columns of $\mathbf{\Phi} = [\mathbf{\Phi}_1\mathbf{\Phi}_2\ldots\mathbf{\Phi}_N]$ contain the eigenvectors of $\mathbf{M}^{-1}\mathbf{K}$. $\mathbf{\Phi}$ is often termed the *modal matrix* or *modal displacement matrix*, where the i-th column $\mathbf{\Phi}_i$ represents the i-th mode shape, and the i-th entry of $\mathbf{\Lambda}$ is proportional to the resonant frequency of the $\mathbf{\Phi}_i$. The columns of $\mathbf{\Phi}$ form a basis (the modal basis, or *eigen*basis) of $3n$ dimensional space, so any displacement $\mathbf{u}(t) = \mathbf{x}(t) - \mathbf{x}_0$ can be written as a linear combination of the columns

$$\mathbf{u}(t) = \mathbf{\Phi}\mathbf{q}(t). \tag{58}$$

The vector $\mathbf{q}(t)$ contains the modal coordinates (or modal amplitudes). Substituting Eq. 58 into Eq. 20 and premultiplying by $\mathbf{\Phi}^T$ yields

$$\mathbf{\Phi}^T\mathbf{M}\mathbf{\Phi}\,\ddot{\mathbf{q}} + \mathbf{\Phi}^T\mathbf{C}\mathbf{\Phi}\,\dot{\mathbf{q}} + \mathbf{\Phi}^T\mathbf{K}\mathbf{\Phi}\,\mathbf{q} = \mathbf{\Phi}^T\mathbf{f}_{ext}. \tag{59}$$

Note that for general damping, $\mathbf{\Phi}^T\mathbf{C}\mathbf{\Phi}$ is dense, but if we assume proportional (Raleigh) damping $\mathbf{C} = \alpha\mathbf{M} + \beta\mathbf{K}$, then $\mathbf{\Phi}^T\mathbf{C}\mathbf{\Phi} = \alpha\mathbf{I} + \beta\mathbf{\Lambda}$ is also diagonal (taking into consideration

that $\mathbf{M}$, $\mathbf{C}$ and $\mathbf{K}$ are normally symmetric positive definite). With $\mathbf{g} = \mathbf{\Phi}^T\mathbf{f}_{ext}$ we can now express Eq. 59 as $3n$ independent scalar 2nd order differential equations

$$\ddot{q}_i + (\alpha + \beta\lambda_i)\dot{q}_i + \lambda_i q_i = g_i, \tag{60}$$

where λ_i is the i-th diagonal element of $\mathbf{\Lambda}$.

With this decoupling, one can examine the eigenvalues, discard high frequency mode shapes, and thereby only use dominant modes (i.e. dominant columns $\mathbf{\Phi}_i$). This can significantly reduce computational cost. Furthermore, the motion components q_i of the individual modes can now be computed independently and combined by linear superposition.

In detail, each of these equations (60) has an analytical solution of the form

$$q_i = c_1 e^{r_1 t} + c_2 e^{r_2 t}, \tag{61}$$

where the constants c_1 and c_2 depend on the initial conditions (q_i and $\dot{q}_i$ at a time $t = t_0$), and the roots of the characteristic equation $r^2 + (\alpha + \beta\lambda_i)r + \lambda_i = 0$ are

$$r_{1,2} = \frac{-(\alpha + \beta\lambda_i) \pm \sqrt{(\alpha + \beta\lambda_i)^2 - 4\lambda_i}}{2}. \tag{62}$$

The solution depends on the sign of $R = (\alpha + \beta\lambda_i)^2 - 4\lambda_i$: $R > 0$, $R = 0$ and $R < 0$ produces the overdamped (r_1, r_2 real and different), critically damped (r_1, r_2 real but repeated) and underdamped case (r_1, r_2 are complex conjugates) respectively. If r_1, r_2 are real then c_1, c_2 are also real, if r_1, r_2 are complex conjugate pairs then so will be c_1, c_2. In any case, q_i in Eq. 61 will be a real value.

Figure 14: *Underdamped (solid, blue) and critically damped (dashed, red) vibration. The black (dash-point) line is the plot of* $\pm e^{-(\alpha + \beta\lambda_i)t/2}$.

Most interest is given to the underdamped case, where damped oscillations occur, and for which Eq. 61 can be written as

$$q_i = e^{-(\alpha + \beta\lambda_i)t/2}(c_1 \cos \mu t + c_2 \sin \mu t), \tag{63}$$

$$\mu = \frac{\sqrt{4\lambda_i - (\alpha + \beta\lambda_i)^2}}{2}, \tag{64}$$

see Fig. 14. This can be used to compute the response of a mode to an external force at some time $t = t_0$ (see [HSO03]): first, the generalized force $\mathbf{f}_{ext}$ must be transformed to modal

Figure 15: *The first four dominant low frequency mode shapes $\Phi_1, \ldots, \Phi_4$ of a constrained bar. The distortion inherent to linear modal analysis is most evident in the bottom right figure [BJ05].*

coordinates by $\Delta t \mathbf{g} = \Delta t \Phi^T \mathbf{f}_{ext}$, and then c_1, c_2 can be computed by substituting Eq. 63 into Eq. 60 and setting $q_i = 0$ and $\dot{q}_i = \Delta t g$ [HSO03]. Naturally, the system of decoupled ODEs (Eq. 60) can also be numerically integrated through time.

5.3. Applications and Other Reduced Models

Interestingly, modal analysis was introduced to Computer Graphics in 1989, but only relatively few significant contributions have been presented since then. The first work of note is Stam's application of modal analysis to the simulation of tree branches subjected to turbulence [Sta97].

James and Pai [JP02a] map runtime dynamics to graphics hardware: in a precomputation stage they build a so-called dynamic response texture (DyRT), where mode shapes Φ_i and other quantities are stored (see Fig. 15). At runtime, the modal coordinates $\mathbf{q}$ are computed from rigid bone transforms or external excitations. Finally, displacements are applied to the undeformed mesh by computing $\mathbf{u} = \sum_i \Phi_i q_i$ for a few dominant low frequency modes using a vertex program running on graphics hardware.

Whereas enforcing direct manipulation and collision constraints is relatively straightforward with node positions in Euclidean space, applying these in a modal basis can be somewhat unintuitive. Hauser et al. [HSO03] provide a solution, where generalized forces are computed for constrained nodes using the modal basis. Since the force computation involves evaluating a pseudoinverse using the singular value decomposition (SVD), only few constraints (up to ten in their examples) can be applied in a real-time simulation environment. Furthermore, they couple the deformable model with a rigid body reference frame and simulate dynamics, collisions and friction.

In the above methods, a linear Cauchy strain model is

Figure 16: *Evolution of one mode shape in linear modal analysis (top row) and modal warping (bottom row). Image courtesy of Min Gyu Choi, Kwangwoon University.*

employed to obtain the stiffness matrix K, resulting in well known artifacts for large rotational deformations away from the rest shape (see Section 3.1 and Fig. 15). To suppress these artifacts, Choi and Ko [CK05] identify per-node rotations and extend the basic modal analysis formulation (Section 5.2) to accommodate these rotations, similar in spirit to the warped stiffness approach [MDM*02]. Although their modal warping procedure is not guaranteed to perform well for large deformations, their results show visually convincing improvements over the standard linear model (Fig.16). The decoupled ODEs are solved using semi-implicit numerical integration, instead of applying the analytical solution of Eq. 61.

Barbič and James [BJ05] take a different approach: instead of employing Cauchy's linearized strain (Eq. 2) they use full quadratic Green strain (Eq. 1) throughout the entire computation, thereby necessitating the solution of a nonlinear version of Eq. 20 per time step. The otherwise computationally burdening implicit Newmark integration is greatly accelerated by carrying it out in reduced coordinates (also known as *subspace integration*), and thereafter performing the matrix multiplication of Eq. 56 in the fragment shader of current graphics hardware. For the generation of an appropriate reduced deformation basis $\mathbf{U}$, which contains sufficient nonlinear deformation, they provide a fully automatic method based on modal derivatives (Fig. 17), and also show how user-defined (force) sketches can assist a full unreduced offline static solver in precomputing a suitable $\mathbf{U}$. Note that due to the nonlinear coupling of modes, using independent harmonic oscillators as previously described (Eq. 61) does not suffice.

Figure 17: *Using subspaces based on the first few dominant vibration modes and their derivatives is sufficient to describe large nonlinear deformations [BJ05].*

Figure 18: *A silver block catapulting some wooden blocks into an oncoming wall of water [CMT04].*

A somewhat different model reduction strategy is devised by James and Fatahalian [JF03]: they entirely precompute collision, contact and dynamics (as well as co-parameterized global illumination) of deformable objects by limiting the range of interactions to a finite set of user impulses. The dynamic responses, termed *impulse response functions* (IRFs), each a chronological set of generalized deformation vectors, are dimensionality reduced using principal component analysis (PCA). At runtime, IRFs are blended based on user interaction.

6. Eulerian and Semi-Lagrangian Methods

There are two general points of view we can take when describing an object that we wish to physically simulate. The main distinctions of the two view points is how they discretize an object in order to work with it numerically, and how they define the boundary of the object itself. So far, this report has covered the *Lagrangian* point of view. As Section 2 explains, the Lagrangian point of view describes an object as a set of moving points (material coordinates) that travel around and change position over time; these points carry their material properties with them as they move through the world. Lagrangian techniques are a convenient way to define an object as a connected mesh of points (Section 3) or simply a cloud of points (Section 4). The *Eulerian* point of view, on the other hand, looks at a stationary set of points and calculates how the material properties stored at those stationary grid points change over time. One of the drawbacks of changing to an Eulerian perspective is that the boundary of the object is no longer explicitly defined. This section covers Eulerian techniques as they are used in Computer Graphics, and focuses heavily on fluids because fluids are often defined in an Eulerian framework. This section will also explore how object boundaries are represented in the Eulerian framework. This section is not meant to give an exhaustive reference for Eulerian techniques; it is meant to give the reader an idea of what types of deformable and fuzzy objects have been simulated in an Eulerian framework.

The Eulerian approach to fluids was popularized by a series of papers by Foster and Metaxas [FM96], [FM97b], and [FM97a]. Their formulation solves the fluid (Navier-Stokes) equations on a regular voxel grid and uses a finite differ-

ence formulation for all the differential equations. To understand the Eulerian framework, and how it differs from a Lagrangian one, we will look briefly at how the Navier-Stokes equations (equations of motion for a fluid) are formulated and solved for a single time step. We refer the reader to [Car04] for an in-depth discussion of liquids, and to [Sta03] for a practical implementation of a smoke simulator.

There are two parts to the Navier-Stokes equations:

$$\nabla \cdot \mathbf{u} = 0 \qquad (65)$$

$$\mathbf{u}_t = -(\mathbf{u} \cdot \nabla)\mathbf{u} + \nabla \cdot (\nu \nabla \mathbf{u}) - \nabla p + \mathbf{f}. \qquad (66)$$

These two equations represent the conservation of mass and momentum for an incompressible fluid. The vector field $\mathbf{u}_t$ is the time derivative of the fluid velocity. The scalar pressure field is $p^\dagger$, and ν is the kinematic viscosity. The vector field $\mathbf{f}$ represents the body force per unit mass (usually just gravity). In this Eulerian formulation the quantities of the fluid are stored in a grid of cells (a regular grid of cubes or voxels if you like). The velocity is stored on the cell faces and the pressure is stored at the center of the cells. This is commonly called a staggered (or MAC) grid. Notice that the position, $\mathbf{x}$, of the fluid is not defined; in this Eulerian frame work, the grid positions remain fixed.

Solving equations 65 and 66 is usually done in several steps (breaking up a PDE into simpler pieces and solving the pieces separately is known as *operator splitting*). The $\mathbf{f}$ term is simply scaled by the time step and added to the current velocity. The advection term, $-(\mathbf{u} \cdot \nabla)\mathbf{u}$, is then solved. One popular way to solve the advection term is with the *semi-Lagrangian* technique [Sta99] because it is stable for large time step sizes. The semi-Lagrangian technique continues as follows–to solve for the advection of any quantity through a velocity field (in this case we are solving for the advection of the velocity field through itself), the velocity at the grid point we wish to update is found, and a path is traced from the grid point backwards in time with that velocity to a new location.

---

† In this formulation we assume a constant density and absorb it into the pressure term for simplicity.

The quantity that we desire to update is then interpolated from the new location and that interpolated value is used as the updated quantity. It is interesting to note that if the time step is scaled so that the path traverses only a single cell, then the semi-Lagrangian technique is a first order upwind technique. After advection, the viscosity term is solved. If the viscosity is expected to be large then an implicit technique should be used (see Section 2.2). At this point we have what can be called a *best guess velocity*:

$$\tilde{\mathbf{u}} = \mathbf{u} + \Delta t[-(\mathbf{u} \cdot \nabla)\mathbf{u} + \nabla \cdot (\nu \nabla \mathbf{u}) + \mathbf{f}], \qquad (67)$$

which does not take into account pressure or mass conservation (Eq. 65). The final step, commonly called a *pressure-projection*, is in fact one of the reasons that Eulerian techniques are popular for fluids. The trick of the pressure-projection step is to solve for a pressure such that subtracting that pressure's gradient from the best guess velocity will give a final velocity that conserves mass. So, we know two things: first, the final velocity, $\mathbf{u}^{new}$, must contain the pressure term which is missing from Eq. 67, so we add it back in to get

$$\mathbf{u}^{new} = \tilde{\mathbf{u}} - \Delta t \nabla p. \qquad (68)$$

Second, we know the final velocity has to conserve mass, so we plug Eq. 68 into Eq. 65 to get

$$\nabla \cdot \mathbf{u}^{new} = \nabla \cdot \tilde{\mathbf{u}} - \Delta t \nabla \cdot (\nabla p) = 0. \qquad (69)$$

Rearranging Eq. 69 gives us the following Poisson equation

$$\Delta t \nabla^2 p = \nabla \cdot \tilde{\mathbf{u}} \qquad (70)$$

with which we must solve for p. We then plug p back into Eq. 68 to complete the pressure projection and find our final velocity. The final velocity is incompressible (divergence free), and for a constant density fluid this ensures mass conservation. Compressible fluids can also conserve mass, but their density must change to do so (see [YOH00]).

The system of equations formed by Eq. 70 is symmetric and positive definite (as long as there is at least one known Dirichlet boundary condition), and can be solved with a conjugate gradient method. Although the above explanation assumes a regular grid, it can also be carried out on an adaptively refined grid, such as an octree [SY02, LGF04].

In the rest of this section we will explore the numerous types of deformable objects that are modeled within an Eulerian framework. As the techniques are described we will also describe how the standard Navier-Stokes equations are changed in each case, and how the objects are represented.

In the early work of Foster and Metaxas [FM96] liquid is displayed either as a height field or a collection of massless particles. These massless particles are different from the material coordinate particles used in a Lagrangian simulation because they are passively moved in the velocity field defined by the fluid, and the particles' velocity is interpolated from the grid. Foster and Metaxas render liquid particles as smoothed ellipsoids [FM97a] with an orientation


Figure 19: *Melting bunny [CMVT02].*

based on the velocity of the particle and a normal computed from the position of the other nearby particles. Foster and Metaxas also define smoke as massless particles [FM97b]: they advect and diffuse an added temperature field on the grid, and use that temperature to define thermal buoyancy in the smoke to create turbulent motion.

Stam [Sta99] uses a scalar density field to define quantities of smoke. These smoke densities are diffused on the grid and advected with the semi-Lagrangian technique (which Stam did not know about at the time, so when he published his seminal work he called the semi-Lagrangian technique the *method of characteristics*). Stam also adds buoyancy forces based on the local smoke density, and Fedkiw et al. [FSJ01] add a *vorticity confinement* term to artificially increase small scale swirling motion which is lost when using semi-Lagrangian advection on coarse grids. Kim et al. [KLLR05] used the semi-Lagrangian technique with *back and fourth error compensation and correction* to keep both small and large scale swirling motion in the fluid while maintaining stability. Their technique greatly reduces the dissipation present in the semi-Lagrangian technique, while maintaining its stability.

Carlson et al. [CMVT02] add a temperature and variable viscosity field to model solid objects that can melt into liquid (Fig. 19). In their model, a solid object is simply a fluid with very high viscosity, and as the temperature rises the viscosity decreases, and the object melts into a liquid state. Like [FM96], the object in their simulation is defined by a set of massless particles, but instead of just rendering the particles they use a splatting technique to extract a mesh from the particles for rendering. Their method also animates the building of a sand castle, but only high viscosity is used to model the mud.

More recently, Zhu and Bridson [ZB05] model sand as a

Figure 20: *Bunny made of sand [ZB05]. Image courtesy of Robert Bridson, UBC.*

continuum (fluid) on an Eulerian grid (Fig. 20). They use a Lagrangian technique to move particles, then interpolate the velocity from the nearby particles on to the Eulerian grid. They then use the grid velocity to solve the pressure projection step. Thus, their technique combines the strengths of both the Lagrangian and Eulerian techniques. After solving the Navier-Stokes equations, they add in a model for sand. At each grid cell they identify the strain rate (Eq. 2) and use it to calculate stress forces in that cell. All connected groups of cells which can resist the inertia trying to make them slide have their velocity projected to rigid body motion. The velocity in all other cells get forces created by sliding frictional stress added to them. Their method also describes a way to build an implicit function from the particles for rendering.

Carlson et al. [CMT04] project rigid bodies on an Eulerian grid to model their interaction with a fluid (Fig. 18). A rigid body solver is run in tandem with a fluid solver: at each frame, the rigid bodies that are in contact with fluid are projected onto the fluid, where their equations of motion are solved with a Navier-Stokes equation that is modified with buoyancy and collision forces. After solving the fluid equation they project the velocity in the rigid body cells to rigid body motion (effectively setting the strain rate to zero inside the cells that contain rigid bodies) and use that velocity for the next step of the rigid body solver.

Foster and Fedkiw [FF01] were the first in graphics to model the boundary of a fluid as a level set. A level set, in this context, is an extra scalar field, ϕ, that stores the signed distance to the fluid surface at each grid point. The change of ϕ is computed each step with the level set equation,

$$\phi_t + \mathbf{u} \cdot \nabla \phi = 0, \tag{71}$$

which effectively updates the position of the iso-contour that delineates the fluid's surface. The level set method has become one of the most popular ways to define a fluid's surface. It has the drawback of defining the surface only at Eulerian grid points, and thus rounding off high resolution details, but this drawback can be reduced by using the particle level set technique [EMF02], which patches the level set values with a group of particles that is passively advected with the surface. The level set can be directly ray traced with a root finder, or a mesh can be extracted from it with marching cubes. We would like to take note of a common confu-

sion when discussing level sets and fluids. Level sets are not used to solve the fluid equations; they are used to define the boundary of the fluid (often called *surface tracking*), so that we know where to solve the fluid equations.

Another way to define a fluid surface is with an explicit polygonal mesh. A mesh representation not only allows a direct way to render the surface, but can also keep track of surface properties like texture coordinates. It can be difficult to deal with topology changes, warping, and self intersection in a mesh if a single mesh is used. The contouring method described by Bargteil et al. [BGOS05] defeats these difficulties by re-defining a mesh each frame of the fluid simulation. They also maintain a signed distance field and show volume conservation on par with the particle level set method with octree refinement.

Figure 21: *A dripping viscoelastic fluid [GBO04].*

Goktekin et al. [GBO04] animate viscoelastic substances (that display characteristics of both a liquid and a solid) by adding elastic terms to the basic Navier-Stokes equations (Fig. 21). The elastic terms are controlled by von Mises's yield condition and a quasi-linear plasticity model. They keep track of a strain tensor field that is advected throughout the fluid grid, and its components are stored at edge centers (off-diagonal components) and cell centers (diagonal components). The surface of their fluid is defined with a particle level set.

Feldman et al. [FOK05] use an Eulerian grid based technique to model smoke, but they use an unstructured mesh, tetrahedra instead of cubes, near the boundary of complicated objects. Their *hybrid mesh* formulation allows them to model the boundary with curved or highly detailed objects in a more accurate fashion. In their paper they render smoke as a set of massless particles.

7. Conclusion

Physically based deformable models have seen wide application in many fields of Computer Graphics, and research and development efforts are as active and fruitful as ever. As pointed out throughout this paper, no one model is suited best for any given application. Instead, many parameters and considerations need to be taken into account, such as model representations, the range of physical parameters, topological changes, realtime or interactive simulation, and so on.

And even though such a seemingly exhaustive toolkit exists, with which very impressive results have been achieved, we still have quite a way to go until we can plausibly and interactively simulate the natural phenomena in our everyday lives. To quote Richard Feynman [FLS63]

> "The things with which we concern ourselves in science appear in myriad forms, and with a multitude of attributes. For example, if we stand on the shore and look at the sea, we see the water, the waves breaking, the foam, the sloshing motion of the water, the sound, the air, the winds and the clouds, the sun and the blue sky, and light; there is sand and there are rocks of various hardness and permanence, color and texture. There are animals and seaweed, hunger and disease, and the observer on the beach; there may be even happiness and thought. Any other spot in nature has a similar variety of things and influences. It is always as complicated as that, no matter where it is. Curiosity demands that we ask questions, that we try to put things together and try to understand this multitude of aspects as perhaps resulting from the action of a relatively small number of elemental things and forces acting in an infinite variety of combinations."

We think it is safe to say that this statement still holds, and probably will hold for years to come, a burden and a blessing alike to researchers in the field.

Yet even with the current methodology, the algorithms and models have seen somewhat limited application in production environments and video games. One reason for this is the lack of computational power: many of the presented techniques are inherently offline, and can take hours or days to produce results. In the field of interactive entertainment, small physically based deformable models are already being implemented, and thanks to developments such as the upcoming physics processing unit (PPU), this trend is likely to carry on. A second reason we see for a reluctancy to adopt these new methods is the limited control the users have over the resulting animations. Although this has been previously addressed, we see great potential in coupling more intuitive user interfaces with physically based simulations.

References

[AKP*05] ADAMS B., KEISER R., PAULY M., GUIBAS L. J., GROSS M., DUTRÉ P.: Efficient raytracing of deforming point-sampled surfaces. In *Proceedings of Eurographics* (2005). To appear. 15

[AP98] ASCHER U., PETZOLD L.: *Computer Methods for Ordinary Differential Equations and Differential-Algebraic Equations.* Society for Industrial & Applied Mathematics, 1998. 3

[ARW95] ASCHER U., RUUTH S., WETTON B.: Implicit-explicit methods for time-dependent pde's. *SIAM J. Numer. Anal.*, 32 (1995), 797–823. 10

[Ask97] ASKES H.: *Everything you always wanted to know about the Element-Free Galerkin method, and more.* Tech. rep., TU Delft nr. 03.21.1.31.29, 1997. 14

[BA04] BOXERMAN E., ASCHER U.: Decomposing cloth. In *SCA '04: Proceedings of the 2004 ACM SIGGRAPH/Eurographics symposium on Computer animation* (2004), pp. 153–161. 10

[BC00] BOURGUIGNON D., CANI M.-P.: Controlling anisotropy in mass-spring systems. In *Computer Animation and Simulation '00* (2000), pp. 113–123. 8

[BGOS05] BARGTEIL A. W., GOKTEKIN T. G., O'BRIEN J. F., STRAIN J. A.: A semi-lagrangian contouring method for fluid simulation. *ACM Transactions on Graphics (in press)* (2005). 20

[BHW94] BREEN D., HOUSE D., WOZNY M.: Predicting the drape of woven cloth using interacting particles. In *SIGGRAPH '94* (1994), pp. 365–372. 7, 8

[BJ05] BARBIČ J., JAMES D. L.: Real-time subspace integration for St.Venant-Kirchhoff deformable models. *ACM Transactions on Computer Graphics (ACM SIGGRAPH 2005) 24*, 3 (2005). 17

[BK04] BOTSCH M., KOBBELT L.: An intuitive framework for real-time freeform modeling. *ACM Trans. Graph. 23*, 3 (2004), 630–634. 2

[BK05] BOTSCH M., KOBBELT L.: Real-time shape editing using radial basis functions. In *Eurographics 2005, to appear* (2005). 2

[BKO*96] BELYTSCHKO T., KRONGAUZ Y., ORGAN D., FLEMING M., KRYSL P.: Meshless methods: An overview and recent developments. *Computer Methods in Applied Mechanics and Engineering 139*, 3 (1996), 3–47. 14

[Bli82] BLINN J. F.: A generalization of algebraic surface drawing. *ACM Trans. Graph. 1*, 3 (1982), 235–256. 11

[BLM00] BELYTSCHKO T., LIU W. K., MORAN B.: *Nonlinear Finite Elements for Continua and Structures.* John Wiley & Sons Ltd., 2000. 2

[BMF03] BRIDSON R., MARINO S., FEDKIW R.: Simulation of clothing with folds and wrinkles. In *ACM SIGGRAPH/Eurographics Symposium Computer Animation* (2003), pp. 28–36. 8, 9, 10

[BNC96] BRO-NIELSEN M., COTIN S.: Real-time volumetric deformable models for surgery simulation using finite elements and condensation. *Computer Graphics Forum 15*, 3 (1996), 57–66. 5

[BSSH04] BIANCHI G., SOLENTHALER B., SZÉKELY G., HARDERS M.: Simultaneous topology and stiffness identification for mass-spring models based on fem reference deformations. In *MICCAI (2)* (2004), pp. 293–301. 8

[BTH*03] BHAT K., TWIGG C., HODGINS J. K., KHOSLA P., POPOVIC Z., SEITZ S.: Estimating cloth simulation parameters from video. In *ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (July 2003). 8

[BW97] BONET J., WOOD R. D.: *Nonlinear continuum mechanics for finite element analysis.* Cambridge Univ. Press, NY, 1997. 2

[BW98] BARAFF D., WITKIN A.: Large steps in cloth simulation. In *Proceedings of SIGGRAPH 1998* (1998), pp. 43–54. 7, 8, 9, 10

[BYM05] BELL N., YOU Y., MUCHA P. J.: Particle-based simulation of granular materials. ACM. to appear. 12

[Can93] CANI M.-P.: An implicit formulation for precise contact modeling between flexible solids. In *SIGGRAPH '93* (1993), pp. 313–320. 11, 12

[Car04] CARLSON M. T.: *Rigid, Melting, and Flowing Fluid.* PhD thesis, Georgia Institute of Technology, 2004. 18

[CDA99] COTIN S., DELINGETTE H., AYACHE N.: Real-time elastic deformations of soft tissues for surgery simulation. In *IEEE Transactions on Visualization and Computer Graphics*, vol. 5 (1). 1999. pp. 62–73. 5

[CDA00] COTIN S., DELINGETTE H., AYACHE N.: A hybrid elastic model allowing real-time cutting, deformations and force-feedback for surgery training and simulation. *The Visual Computer 16*, 8 (2000), 437–452. 5

[CEO*93] COVER S., EZQUERRA N., O'BRIEN J., ROWE R., GADACZ T., PALM E.: Interactively deformable models for surgery simulation. *IEEE Computer Graphics and Applications 13*, 6 (1993), 68–75. 5

[CGC*02] CAPELL S., GREEN S., CURLESS B., DUCHAMP T., POPOVIC Z.: Interactive skeleton-driven dynamic deformations. In *Proceedings of SIGGRAPH 2002* (2002), pp. 586–593. 5

[CHP89] CHADWICK J., HAUMANN D., PARENT R.: Layered construction for deformable animated characters. In *SIGGRAPH '89* (1989), pp. 243–252. 7

[Chu96] CHUNG T. J.: *Applied Continuum Mechanics*. Cambridge Univ. Press, NY, 1996. 2

[CK02] CHOI K.-J., KO H.-S.: Stable but responsive cloth. In *SIGGRAPH '02* (2002), pp. 604–611. 9, 10

[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 (2005), 91–101. 17

[CMT04] CARLSON M., MUCHA P. J., TURK G.: Rigid fluid: Animating the interplay between rigid bodies and fluid. *ACM Transactions on Graphics 23*, 3 (Aug. 2004), 377–384. 18, 20

[CMVT02] CARLSON M., MUCHA P., VAN HORN III R. B., TURK G.: Melting and flowing. In *Proceedings of the 2002 ACM SIGGRAPH Symposium on Computer Animation* (July 2002), pp. 167–174. 19

[Coo95] COOK R. D.: *Finite Element Modeling for Stress Analysis*. John Wiley & Sons, NY, 1995. 2

[CZKM98] CHEN Y., ZHU Q., KAUFMAN A., MURAKI S.: Physically-based animation of volumetric objects. In *CA '98: Proceedings of the Computer Animation* (1998), p. 154. 7

[DC94] DESBRUN M., CANI M.-P.: Highly deformable material for animation and collision processing. In *Fifth Eurographics Workshop on Animation, Simulation* (septembre 1994). 11

[DC95] DESBRUN M., CANI M.-P.: Animating soft substances with implicit surfaces. In *Computer Graphics Proceedings* (1995), ACM SIGGRAPH, pp. 287–290. 11

[DC96] DESBRUN M., CANI M.-P.: Smoothed particles: A new paradigm for animating highly deformable bodies. In *6th Eurographics Workshop on Computer Animation and Simulation '96* (1996), pp. 61–76. 13

[DC98] DESBRUN M., CANI M.-P.: Active implicit surface for animation. In *Proceedings of Graphics Interface* (1998), pp. 143–150. 13

[DC99] DESBRUN M., CANI M.-P.: *Space-Time Adaptive Simulation of Highly Deformable Substances*. Tech. rep., INRIA Nr. 3829, 1999. 12, 13

[DDBC99] DEBUNNE G., DESBRUN M., BARR A., CANI M.-P.: Interactive multiresolution animation of deformable models. In *Eurographics Workshop on Computer Animation and Simulation '99* (1999), pp. 133–144. 5

[DDCB00] DEBUNNE G., DESBRUN M., CANI M.-P., BARR A. H.: Adaptive simulation of soft bodies in real-time. In *Computer Animation '00* (2000), pp. 15–20. 5

[DDCB01] DEBUNNE G., DESBRUN M., CANI M.-P., BARR A.: Dynamic real-time deformations using space & time adaptive sampling. In *Computer Graphics Proceedings* (Aug. 2001), Annual Conference Series, ACM SIGGRAPH 2001, pp. 31–36. 4, 5

[DSB99] DESBRUN M., SCHRÖDER P., BARR A. H.: Interactive animation of structured deformable objects. In *Graphics Interface '99* (1999). 5, 10

[EGS03] ETZMUSS O., GROSS J., STRASSER W.: Deriving a particle system from continuum mechanics for the animation of deformable objects. *IEEE Transactions on Visualization and Computer Graphics 9*, 4 (2003), 538–550. 8

[EKS03] ETZMUSS O., KECKEISEN M., STRASSER W.: A fast finite element solution for cloth modelling. *Proceedings of Pacific Graphics 2003* (2003). 8

[EMF02] ENRIGHT D., MARSCHNER S., FEDKIW R.: Animation and rendering of complex water surfaces. In *SIGGRAPH '02* (2002), pp. 736–744. 20

[EWS96] EBERHARDT B., WEBER A., STRASSER W.: A fast, flexible, particle-system model for cloth draping. *IEEE Comput. Graph. Appl. 16*, 5 (1996), 52–59. 9

[FF01] FOSTER N., FEDKIW R.: Practical animation of liquids. In *Proceedings of ACM SIGGRAPH 2001* (2001), pp. 23–30. 20

[FLS63] FEYNMAN R. P., LEIGHTON R. B., SANDS M.: *The Feynman Lectures on Physics: Volume I*. Addison Wesley, 1963. 21

[FM96] FOSTER N., METAXAS D.: Realistic animation of liquids. *Graphical Models and Image Processing 58*, 5 (1996), 471–483. 18, 19

[FM97a] FOSTER N., METAXAS D.: Controlling fluid animation. In *Proceedings CGI '97* (1997), pp. 178–188. Winner of the Androme Award 1997. 18, 19

[FM97b] FOSTER N., METAXAS D.: Modeling the motion of a hot, turbulent gas. In *SIGGRAPH '97* (1997), pp. 181–188. 18, 19

[FM03] FRIES T.-P., MATTHIES H. G.: *Classification and Overview of Meshfree Methods*. Tech. rep., TU Brunswick, Germany Nr. 2003-03, 2003. 14

[FOK05] FELDMAN B. E., O'BRIEN J. F., KLINGNER B. M.: Animating gases with hybrid meshes. In *Proceedings of ACM SIGGRAPH 2005* (2005), p. (in press). 20

[FR86] FOURNIER A., REEVES W. T.: A simple model of ocean waves. In *SIGGRAPH '86* (1986), pp. 75–84. 11

[FSJ01] FEDKIW R., STAM J., JENSEN H. W.: Visual simulation of smoke. In *SIGGRAPH '01* (2001), pp. 15–22. 19

[GBO04] GOKTEKIN T. G., BARGTEIL A. W., O'BRIEN J. F.: A method for animating viscoelastic fluids. In *Proc. of ACM SIGGRAPH* (2004), vol. 23, pp. 463–468. 14, 20

[Gdo93] GDOUTOS E. E.: *Fracture Mechanics*. Kluwer Academic Publishers, Netherlands, 1993. 2

[GHDS03] GRINSPUN E., HIRANI A. N., DESBRUN M., SCHRÖDER P.: Discrete shells. In *Proceedings of the 2003 ACM SIGGRAPH/Eurographics Symposium on Computer animation* (2003), pp. 62–67. 8

[GKS02] GRINSPUN E., KRYSL P., SCHRÖDER P.: CHARMS: A simple framework for adaptive simulation. In *Proceedings of SIGGRAPH 2002* (2002), pp. 281–290. 5

[GM77] GINGOLD R. A., MONAGHAN J. J.: Smoothed particle hydrodynamics - theory and application to non-spherical stars. *Royal Astronomical Society, Monthly Notices 181* (1977), 375–389. 13

[GM97] GIBSON S. F., MIRTICH B.: *A survey of deformable models in computer graphics*. Technical Report TR-97-19, MERL, Cambridge, MA, 1997. 1, 2

[Gos90] GOSS M. E.: Motion simulation: A real time particle system for display of ship wakes. *IEEE Comput. Graph. Appl. 10*, 3 (1990), 30–35. 11

[GTT89] GOURRET J.-P., THALMANN N. M., THALMANN D.: Simulation of object and human skin formations in a grasping task. In *SIGGRAPH '89* (1989), pp. 21–30. 5

[HB00] HOUSE D., BREEN D.: *Cloth modeling and animation*. A. K. Peters, Ltd., 2000. 2, 8

[HES02] HAUTH M., ETZMUSS O., STRASSER W.: Analysis of numerical methods for the simulation of deformable models. *The Visual Computer* (2002). Accepted for publication. 3, 9, 10

[HMT01] HADAP S., MAGNENAT-THALMANN N.: Modeling dynamic hair as continuum. In *Eurographics Proceedings. Computer Graphics Forum* (2001), vol. 20. 13

[HPH96] HUTCHINSON D., PRESTON M., HEWITT T.: Adaptive refinement for mass/spring simulations. In *Proceedings of the Eurographics workshop on Computer animation and simulation '96* (1996), pp. 31–45. 9

[HSO03] HAUSER K. K., SHEN C., O'BRIEN J. F.: Interactive deformation using modal analysis with constraints. In *Graphics Interface '03* (2003). 16, 17

[Hun05] HUNTER P.: *FEM/BEM Notes*. University of Oakland, New Zealand, 2005. http://www.bioeng.auckland.ac.nz/cmiss/fembemnotes/fembemnotes.pdf. 4, 6

[IMH05] IGARASHI T., MOSCOVICH T., HUGHES J. F.: As-rigid-as-possible shape manipulation. *ACM Transactions on Graphics (ACM SIGGRAPH 2005) 24*, 3 (2005). 2

[ITF04] IRVING G., TERAN J., FEDKIW R.: Invertible finite elements for robust simulation of large deformation. In *Proceedings of the 2004 ACM SIGGRAPH/Eurographics symposium on Computer animation* (2004), pp. 131–140. 5

[JBT04] JAMES D. L., BARBIČ J., TWIGG C. D.: Squashing cubes: Automating deformable model construction for graphics. In *Proceedings of the SIGGRAPH 2004 Conference on Sketches & Applications* (2004). 6

[JF03] JAMES D. L., FATAHALIAN K.: Precomputing interactive dynamic deformable scenes. *ACM Transactions on Graphics (Proceedings of ACM SIGGRAPH 2003) 22*, 3 (2003), 879–887. 18

[JP99] JAMES D. L., PAI D. K.: ArtDefo: accurate real time deformable objects. In *SIGGRAPH '99* (1999), pp. 65–72. 6

[JP02a] JAMES D. L., PAI D. K.: DyRT: Dynamic response textures for real time deformation simulation with graphics hardware. In *Proceedings of SIGGRAPH 2002* (2002). 17

[JP02b] JAMES D. L., PAI D. K.: Real time simulation of multizone elastokinematic models. In *2002 IEEE Intl. Conference on Robotics and Automation* (2002). 6

[JP03] JAMES D. L., PAI D. K.: Multiresolution Green's function methods for interactive simulation of large-scale elastostatic objects. *ACM Transactions on Graphics 22*, 1 (2003), 47–82. 6

[JP04] JAMES D. L., PAI D. K.: BD-Tree: Output-sensitive collision detection for reduced deformable models. *ACM Transactions on Graphics (SIGGRAPH 2004) 23*, 3 (Aug. 2004). 16

[JT05] JAMES D. L., TWIGG C. D.: Skinning mesh animations. *ACM Transactions on Graphics (SIGGRAPH 2005) 24*, 3 (2005). 2

[KAG*05] KEISER R., ADAMS B., GASSER D., BAZZI P., DUTRÉ P., GROSS M.: A unified lagrangian approach to solid-fluid animation. In *Proceedings of the Eurographics Symposium on Point-Based Graphics* (2005). To appear. 13, 14, 15

[Kas95] KASS M.: An introduction to physically based modeling, chapter: Introduction to continuum dynamics for computer graphics. *In SIGGRAPH 95 Course Notes* (1995). 8

[Kaw80] KAWABATA S.: The standardization and analysis of hand evaluation. *The Textile Machinery Society of Japan* (1980). 8

[KCC*00] KANG Y., CHOI J., CHO H., LEE D., PARK C.: Real-time animation technique for flexible and thin objects. In *WSCG 2000* (2000), pp. 322–329. 10

[KGC*96] KOCH R. M., GROSS M. H., CARLS F. R., VON BÜREN D. F., FANKHAUSER G., PARISH Y. I. H.: Simulating facial surgery using finite element models. *Computer Graphics 30*, Annual Conference Series (1996), 421–428. 5

[KJP02] KRY P. G., JAMES D. L., PAI D. K.: Eigenskin: real time large deformation character skinning in hardware. In *Proceedings of the 2002 ACM SIGGRAPH/Eurographics symposium on Computer animation* (2002), pp. 153–159. 2

[KLLR05] KIM B., LIU Y., LLAMAS I., ROSSIGNAC J.: Flowfixer: Using bfecc for fluid simulation. (submitted). 19

[KMH*04] KEISER R., MÜLLER M., HEIDELBERGER B., TESCHNER M., GROSS M.: Contact handling for deformable point-based objects. In *Proceedings of Vision, Modeling, Visualization VMV'04* (Nov 2004), pp. 339–347. 15

[KWT88] KASS M., WITKIN A., TERZOPOULOS D.: Snakes: active contour models. *International Journal of Computer Vision 1*, 4 (1988), 321–331. 13

[KZ04] KLEIN J., ZACHMANN G.: Point cloud collision detection. In *Proceedings of EUROGRAPHICS 2004* (2004), pp. 567–576. 15

[Las87] LASSETER J.: Principles of traditional animation applied to 3D computer animation. In *SIGGRAPH '87* (1987), pp. 35–44. 1

[LGF04] LOSASSO F., GIBOU F., FEDKIW R.: Simulating water and smoke with an octree data structure. *ACM Transactions on Graphics 23*, 3 (Aug. 2004), 457–462. 19

[Liu02] LIU G. R.: *Mesh-Free Methods*. CRC Press, 2002. 14

[LKG*03] LLAMAS I., KIM B., GARGUS J., ROSSIGNAC J., SHAW C. D.: Twister: a space-warp operator for the two-handed editing of 3D shapes. *ACM Trans. Graph. 22*, 3 (2003), 663–668. 2

[LO05] LIN M., OTADUY M.: Recent advances in haptic rendering and applications. *Siggraph Course Notes* (2005). 2

[Lov27] LOVE A.: *A Treatise on the Mathematical Theory of Elasticity*. Cambridge University Press, 1927. 2

[LS81] LANCASTER P., SALKAUSKAS K.: Surfaces generated by moving least squares methods. *Mathematics of Computation 87* (1981), 141–158. 14

[LSLCO05] LIPMAN Y., SORKINE O., LEVIN D., COHEN-OR D.: Linear rotation-invariant coordinates for meshes. *ACM Transactions on Graphics (ACM SIGGRAPH 2005) 24*, 3 (2005). 2

[Luc77] LUCY L. B.: A numerical approach to the testing of the fission hypothesis. *The Astronomical Journal 82*, 12 (1977), 1013–1024. 13

[LV05] LI L., VOLKOV V.: Cloth animation with adaptively refined meshes. In *ACSC* (2005), pp. 107–114. 9

[MBF04] MOLINO N., BAO Z., FEDKIW R.: A virtual node algorithm for changing mesh topology during simulation. *ACM Trans. Graph. 23*, 3 (2004), 385–392. 6

[MCG03] MÜLLER M., CHARYPAR D., GROSS M.: Particle-based fluid simulation for interactive applications. In *Proceedings of the 2003 ACM SIGGRAPH/Eurographics Symposium on Computer animation* (2003), pp. 154–159. 13

[MDM*02] MÜLLER M., DORSEY J., MCMILLAN L., JAGNOW R., CUTLER B.: Stable real-time deformations. In *Proceedings of the 2002 ACM SIGGRAPH/Eurographics symposium on Computer animation* (2002), pp. 49–54. 4, 5, 17

[MG04] MÜLLER M., GROSS M.: Interactive virtual materials. In *GI '04: Proceedings of Graphics Interface 2004* (2004), pp. 239–246. 5

[MHTG05] MÜLLER M., HEIDELBERGER B., TESCHNER M., GROSS M.: Meshless deformations based on shape matching. *ACM Transactions on Computer Graphics (ACM SIGGRAPH 2005) 24*, 3 (2005). 3, 15

[Mil88] MILLER G. S. P.: The motion dynamics of snakes and worms. In *SIGGRAPH '88* (1988), pp. 169–173. 8

[MJBF02] MILLIRON T., JENSEN R. J., BARZEL R., FINKELSTEIN A.: A framework for geometric warps and deformations. *ACM Trans. Graph. 21*, 1 (2002), 20–51. 2

[MKN*04] MÜLLER M., KEISER R., NEALEN A., PAULY M., GROSS M., ALEXA M.: Point based animation of elastic, plastic and melting objects. In *Proceedings of the 2004 ACM SIGGRAPH/Eurographics symposium on Computer animation* (2004), pp. 141–151. 14, 15

[MMDJ01] MÜLLER M., MCMILLAN L., DORSEY J., JAGNOW R.: Real-time simulation of deformation and fracture of stiff materials. *Proceedings of Eurographics Workshop on Animation and Simulation 2001* (2001). 5

[Mon92] MONAGHAN J.: Smoothed particle hydrodynamics. *Annu. Rev. Astron. Physics 30* (1992), 543. 12, 13

[Mon94] MONAGHAN J. J.: Simulating free surface flows with SPH. *J. Comput. Phys. 110*, 2 (1994), 399–406. 13

[MSKG05] MÜLLER M., SOLENTHALER B., KEISER R., GROSS M.: Particle-based fluid-fluid interaction. ACM. to appear. 12, 13

[MST*04] MÜLLER M., SCHIRM S., TESCHNER M., HEIDELBERGER B., GROSS M.: Interaction of fluids with deformable solids. In *Computer Animation and Social Agents CASA'04* (2004), pp. 159–171. 13

[MTCK*04] MAGNENAT-THALMANN N., CORDIER F., KECKEISEN M., KIMMERLE S., KLEIN R., MESETH J.: Simulation of clothes for real-time applications. In *Eurographics 2004, Tutorials 1: Simulation of Clothes for Real-time Applications* (2004). 2

[MTG04] MÜLLER M., TESCHNER M., GROSS M.: Physically-based simulation of objects represented by surface meshes. In *Proceedings of Computer Graphics International (CGI)* (Jun 2004), pp. 26–33. 6

[MTHK00] MAGNENAT-THALMANN N., HADAP S., KALRA P.: State of the art in hair simulation. In *International Workshop on Human Modeling and Animation* (June 2000), Korea Computer Graphics Society, pp. 3–9. 2

[NSACO05] NEALEN A., SORKINE O., ALEXA M., COHEN-OR D.: A sketch-based interface for detail-preserving mesh editing. *ACM Transactions on Computer Graphics (ACM SIGGRAPH 2005) 24*, 3 (2005). 2

[NTVR92] NAYROLES B., TOUZOT G., VILLON P., RICARD A.: Generalizing the finite element method: diffuse approximation and diffuse elements. *Computational Mechanics 10*, 5 (1992), 307–318. 14

[OBH02] O'BRIEN J. F., BARGTEIL A. W., HODGINS J. K.: Graphical modeling and animation of ductile fracture. In *Proceedings of SIGGRAPH 2002* (2002), pp. 291–294. 5

[OC97] OPALACH A., CANI M.: Local deformations for animation of implicit surfaces. In *13th Spring Conference on Computer Graphics* (1997), Straßer W., (Ed.), pp. 85–92. 12

[OH95] O'BRIEN J. F., HODGINS J. K.: Dynamic simulation of splashing fluids. In *CA '95: Proceedings of the Computer Animation* (1995), p. 198. 11

[OH99] O'BRIEN J. F., HODGINS J. K.: Graphical modeling and animation of brittle fracture. In *Proceedings of SIGGRAPH 1999* (1999), pp. 287–296. 4, 5

[PB81] PLATT S. M., BADLER N. I.: Animating facial expressions. In *SIGGRAPH '81* (1981), pp. 245–252. 7

[pba04] Point based animation: Resource collection on the world wide web. http://www.pointbasedanimation.org, 2004. 14

[PDA00] PICINBONO G., DELINGETTE H., AYACHE N.: Real-time large displacement elasticity for surgery simulation: Non-linear tensor-mass model. *Third International Conference on Medical Robotics, Imaging And Computer Assisted Surgery: MICCAI 2000* (Oct. 2000), 643–652. 5

[PDA01] PICINBONO G., DELINGETTE H., AYACHE N.: Nonlinear and anisotropic elastic soft tissue models for medical simulation. In *Proceedings of the IEEE International Conference on Robotics and Automation* (2001). 5

[Pea86] PEACHEY D. R.: Modeling waves and surf. In *SIGGRAPH '86* (1986), pp. 65–74. 11

[PKA*05] PAULY M., KEISER R., ADAMS B., DUTRÉ P., GROSS M., GUIBAS L. J.: Meshless animation of fracturing solids. In *Proceedings of ACM Siggraph* (2005). To appear. 14, 15

[PKKG03] PAULY M., KEISER R., KOBBELT L. P., GROSS M.: Shape modeling with point-sampled geometry. In *Proceedings of ACM Siggraph* (2003), pp. 641–650. 15

[PPG04] PAULY M., PAI D. K., GUIBAS L. J.: Quasi-rigid objects in contact. In *Proceedings of the 2004 ACM SIGGRAPH/Eurographics symposium on Computer animation* (2004), pp. 109–119. 15

[Pro95] PROVOT X.: Deformation constraints in a mass-spring model to describe rigid cloth behaviour. In *Proc. Graphics Interface* (1995), pp. 147–154. 9

[PTB*03] PREMOZE S., TASDIZEN T., BIGLER J., LEFOHN A., WHITAKER R. T.: Particle-based simulation of fluids. In *Proceedings of Eurographics 2003* (2003), pp. 401–410. 13

[PTVF92] PRESS W., TEUKOLSKY S., VETTERLING W., FLANNERY B.: *Numerical Recipes in C - The Art of Scientific Computing*, 2nd ed. Cambridge University Press, 1992. 3

[PW89] PENTLAND A., WILLIAMS J.: Good vibrations: Modal dynamics for graphics and animation. *Compter Graphics (Proceedings of SIGGRAPH 1989) 23*, 3 (1989), 215–222. 16

[PZvBG00] PFISTER H., ZWICKER M., VAN BAAR J., GROSS M.: Surfels: Surface elements as rendering primitives. In *Siggraph 2000, Computer Graphics Proceedings* (2000), pp. 335–342. 14

[Ree83] REEVES W. T.: Particle systems – a technique for modeling a class of fuzzy objects. *ACM Trans. Graph. 2*, 2 (1983), 91–108. 10

[Rey87] REYNOLDS C. W.: Flocks, herds and schools: A distributed behavioral model. In *SIGGRAPH '87* (1987), pp. 25–34. 10

[SAC*99] STORA D., AGLIATI P.-O., CANI M.-P., NEYRET F., GASCUEL J.-D.: Animating lava flows. In *Proceedings of Graphics Interface* (1999), pp. 203–210. 13

[SBMH94] SAGAR M. A., BULLIVANT D., MALLINSON G. D., HUNTER P. J.: A virtual environment and model of the eye for surgical simulation. In *SIGGRAPH '94* (1994), pp. 205–212. 5

[SCF*04] SEDERBERG T. W., CARDON D. L., FINNIGAN G. T., NORTH N. S., ZHENG J., LYCHE T.: T-spline simplification and local refinement. *ACM Trans. Graph. 23*, 3 (2004), 276–283. 2

[SF95] STAM J., FIUME E.: Depicting fire and other gaseous phenomena using diffusion processes. In *SIGGRAPH '95* (1995), pp. 129–136. 13

[Sim90] SIMS K.: Particle animation and rendering using data parallel computation. In *SIGGRAPH '90* (1990), pp. 405–413. 11

[SLCO*04] SORKINE O., LIPMAN Y., COHEN-OR D., ALEXA M., RÖSSL C., SEIDEL H.-P.: Laplacian surface editing. In *Proceedings of the Eurographics/ACM SIGGRAPH Symposium on Geometry processing* (2004), pp. 179–188. 2

[ST92] SZELISKI R., TONNESEN D.: Surface modeling with oriented particle systems. *Computer Graphics 26*, 2 (1992), 185–194. 12

[Sta97] STAM J.: Stochastic dynamics: Simulating the effects of turbulence on flexible structures. *Computer Graphics Forum 16*, 3 (1997), 159–164. 17

[Sta99] STAM J.: Stable fluids. In *SIGGRAPH '99* (1999), pp. 121–128. 18, 19

[Sta03] STAM J.: Real-time fluid dynamics for games. In *Proceedings of the Game Developer Conference* (March 2003). 18

[Suk03] SUKUMAR N.: Meshless methods and partition of unity finite elements. In *of the Sixth International ESAFORM Conference on Material Forming* (2003), pp. 603–606. 14

[SWB00] SMITH J., WITKIN A., BARAFF D.: Fast and controllable simulation of the shattering of brittle objects. In *Graphics Interface* (May 2000). 5

[SY02] SHI L., YU Y.: *Visual Smoke Simulation with Adaptive Octree Refinement*. Tech. Rep. UIUCDCS-R-2002-2311, University of Illinois at Urbana-Champaign, Dec. 2002. 19

[SZBN03] SEDERBERG T. W., ZHENG J., BAKENOV A., NASRI A.: T-splines and t-nurccs. *ACM Trans. Graph. 22*, 3 (2003), 477–484. 2

[TBHF03] TERAN J., BLEMKER S., HING V. N. T., FEDKIW R.: Finite volume methods for the simulation of skeletal muscle. In *Eurographics/SIGGRAPH Symposium on Computer Animation* (2003). 6

[TF88] TERZOPOULOS D., FLEISCHER K.: Modeling inelastic deformation: viscolelasticity, plasticity, fracture. In *SIGGRAPH '88* (1988), pp. 269–278. 6

[THMG04] TESCHNER M., HEIDELBERGER B., MÜLLER M., GROSS M.: A versatile and robust model for geometrically complex deformable solids. In *Proceedings of Computer Graphics International (CGI)* (Jun 2004), pp. 312–319. 5, 7, 8

[TKH*05] TESCHNER M., KIMMERLE S., HEIDELBERGER B., ZACHMANN G., RAGHUPATHI L., FUHRMANN A., CANI M.-P., FAURE F., MAGNETAT-THALMANN N., STRASSER W., VOLINO P.: Collision detection for deformable objects. *Computer Graphics Forum 24*, 1 (2005), 61–81. 2

[Ton91] TONNESEN D.: Modeling liquids and solids using thermal particles. In *Graphics Interface* (June 1991), pp. 255–262. 11

[Ton92] TONNESEN D.: Spatially coupled particle systems. 4.1 – 4.21. 11

[Ton98] TONNESEN D.: *Dynamically Coupled Particle Systems for Geometric Modeling, Reconstruction, and Animation*. PhD thesis, University of Toronto, November 1998. 11, 12

[TPBF87] TERZOPOULOS D., PLATT J., BARR A., FLEISCHER K.: Elastically deformable models. In *SIGGRAPH '87* (1987), pp. 205–214. 1, 6

[TPF89] TERZOPOULOS D., PLATT J., FLEISCHER K.: Heating and melting deformable models (from goop to glop). In *Graphics Interface '89* (1989), pp. 219–226. 8

[TT94] TU X., TERZOPOULOS D.: Artificial fishes: physics, locomotion, perception, behavior. In *SIGGRAPH '94* (1994), pp. 43–50. 8

[TW88] TERZOPOULOS D., WITKIN A.: Physically based models with rigid and deformable components. *IEEE Computer Graphics and Applications 8*, 6 (1988), 41–51. 6

[TW90] TERZOPOULOS D., WATERS K.: Physically-based facial modeling, analysis, and animation. *Journal of Visualization and Computer Animation 1*, 1 (1990), 73–80. 7

[VB02] VILLARD J., BOROUCHAKI H.: Adaptive meshing for cloth animation. In *11th International Meshing Roundtable* (Ithaca, New York, USA, 15–18 September 2002), Sandia National Laboratories, pp. 243–252. 9

[VMT97] VOLINO P., MAGNENAT-THALMANN N.: Developing simulation techniques for an interactive clothing system. In *Proceedings of the 1997 International Conference on Virtual Systems and MultiMedia* (1997), p. 109. 8

[VMT00] VOLINO P., MAGNENAT-THALMANN N.: Implementing fast cloth simulation with collision response. *IEEE Computer Society* (2000), 257–268. 10

[VMT04] VOLINO P., MAGNENAT-THALMANN N.: Animating complex hairstyles in real-time. In *VRST* (November 2004). 2

[Wat87] WATERS K.: A muscle model for animation three-dimensional facial expression. In *SIGGRAPH '87* (1987), pp. 17–24. 7

[WB97] WITKIN A., BARAFF D.: Physically based modeling: Principles and practice. *Siggraph Course Notes* (1995, 1997). 2

[WDGT01] WU X., DOWNES M. S., GOKTEKIN T., TENDICK F.: Adaptive nonlinear finite elements for deformable body simulation using dynamic progressive meshes. *Eurographics* (Sept. 2001), 349–358. 5

[WP02] WANG X. C., PHILLIPS C.: Multi-weight enveloping: least-squares approximation techniques for skin animation. In *Proceedings of the 2002 ACM SIGGRAPH/Eurographics symposium on Computer animation* (2002), pp. 129–138. 2

[WSG05] WICKE M., STEINEMANN D., GROSS M.: Efficient animation of point-based thin shells. In *Proceedings of Eurographics '05* (2005), p. to appear. 15

[WT91] WATERS K., TERZOPOULOS D.: Modeling and animating faces using scanned data. *Journal of Visualization and Computer Animation 2*, 2 (1991), 123–128. 7

[YOH00] YNGVE G. D., O'BRIEN J. F., HODGINS J. K.: Animating explosions. In *Proc. of SIGGRAPH* (2000), Computer Graphics Proceedings, Annual Conference Series, pp. 29–36. 19

[YZX*04] YU Y., ZHOU K., XU D., SHI X., BAO H., GUO B., SHUM H.-Y.: Mesh editing with poisson-based gradient field manipulation. *ACM Trans. Graph. 23*, 3 (2004), 644–651. 2

[ZB05] ZHU Y., BRIDSON R.: Animating sand as a fluid. *ACM Transactions on Graphics (Proc. of ACM SIGGRAPH 2005) (in press)* (2005). 19, 20