

Invertible Finite Elements For Robust Simulation of Large Deformation

G. Irving[†], J. Teran[‡], and R. Fedkiw[§]

Stanford University

Abstract

We present an algorithm for the finite element simulation of elastoplastic solids which is capable of robustly and efficiently handling arbitrarily large deformation. In fact, our model remains valid even when large parts of the mesh are inverted. The algorithm is straightforward to implement and can be used with any material constitutive model, and for both volumetric solids and thin shells such as cloth. We also provide a mechanism for controlling plastic deformation, which allows a deformable object to be guided towards a desired final shape without sacrificing realistic behavior. Finally, we present an improved method for rigid body collision handling in the context of mixed explicit/implicit time-stepping.

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

1. Introduction

Significant effort has been placed into making finite element simulation robust in the regime of large deformations, including the arbitrary Lagrangian-Eulerian (ALE) formulations pioneered by [HAC74], continuous remeshing (see e.g. [CO97, EZE98] and the references therein), etc. However, as noted in [BBTL03], these approaches are often computationally intensive and difficult to implement, which has primarily limited their use to two spatial dimensions. Moreover, [BBTL03] points out that these difficulties often lead authors to less optimal techniques such as element deletion. This not only degrades the accuracy of the simulation, but is unsuitable for graphics applications where disappearing tetrahedra cause visual artifacts.

Since constitutive models for real materials are meaningful only for uninverted material, standard finite element simulation algorithms fail as soon as a single tetrahedron inverts. For this reason, various authors have proposed techniques for untangling inverted meshes. For example, [VGS04] used

feasible set methods and optimization to untangle two-dimensional meshes, [ERM*03] extended various tetrahedron quality metrics intended for mesh smoothing to the case of inverted tetrahedra allowing untangling to occur simultaneously with optimization, etc. However, none of these techniques are *guaranteed* to work, and fail quite often in practice especially if the boundary of the mesh is also tangled. And the failure to untangle a single tetrahedron forces the simulation to fail for most real world constitutive models.

As pointed out by [EZE98], element inversion can occur even if the vertex position of the mesh are identical to their true continuum values. A common case of this is illustrated in Figure 1 where a triangle with three nodes on the boundary is forced to invert during a collision with another object. Even if an object as a whole deforms by only a small amount, say 10%, an individual tetrahedron may undergo severe deformation due to errors in the discrete representation of the continuous material. In fact, large deformation and inversion can arise even when simulating incompressible material, since one typically cannot conserve volume for each individual tetrahedron. Given that it is difficult or impossible to prevent inversion in all cases, we propose a simpler approach that allows elements to invert gracefully and recover.

[†] e-mail: irving@cs.stanford.edu

[‡] e-mail: jteran@stanford.edu

[§] e-mail: fedkiw@cs.stanford.edu

If the particular material behavior and underlying physics is unimportant, there are several techniques for treating inversion. For mass-spring systems, altitude springs work surprisingly well [MBTF03]. In fact, one can even add altitude springs to finite element models, but this changes the underlying partial differential equations and thus the behavior of the material (losing one of the key benefits of the finite element method as compared to mass-spring systems). Moreover, one cannot use constitutive models that lose meaning for inverting elements (which is most of them). Of course, one could switch from finite element methods to mass-spring methods for flat, degenerate, and inverted elements but this leads to force discontinuities (detrimental to implicit methods), visual artifacts such as popping, etc. Various authors have proposed methods similar in spirit to altitude springs. [THMG04] changes the underlying partial differential equation by adding a volume preservation term to penalize inversion. [MG04] first computes and removes the rotation from material to spatial coordinates, and then applies a linear model in the unrotated space. Although the linear model works well for inverted elements, it is severely limited in the types of materials it can approximate. In fact, we stress that a spring is the most common linear model that correctly accounts for rotation.

In this paper, we take a different approach motivated by the desire to continuously (or even smoothly) extend the finite element method so that it behaves gracefully for both degenerate and inverted elements, even for arbitrary constitutive models. This is especially important as a number of recent graphics publications have advocated the use of more advanced constitutive models to capture more realistic physical behavior. Moreover, many materials found in nature exhibit complex nonlinearities under large deformation, such as the biphasic nature of biological tissue and the anisotropic behavior of muscle. Our approach begins by computing a diagonalization of the deformation mapping in order to determine the “direction” along which a given tetrahedron is inverted. We then extend the constitutive model to the inverted regime using C^0 , C^1 or higher continuity around the flat state, resulting in smooth behavior even in extreme situations. The resulting forces *always* act to restore the tetrahedron to its original shape, allowing objects to recover cleanly from flat or inverted configurations, as shown in Figure 2. We illustrate the generality and robustness of our approach with a number of simulations of objects undergoing large deformations, including nonlinear and anisotropic constitutive models, plasticity with and without control, both volumetric objects and thin shells, and fracture.

2. Related Work

[TPBF87, TF88b, TF88a] pioneered deformable models in computer graphics including early work on plasticity and fracture. Finite element simulations have been used to model a hand grasping a ball [GMTT89], to simulate muscles




Figure 1: A highly curved region of an object is pushed inwards during a collision with a grey object (left), and a triangle used to represent this deformation is forced to invert (right).

[CZ92], for virtual surgery [PDA01], and to simulate data from the NIH visible human data set [ZCK98, HFS*01]. [OH99] and [OBH02] simulated brittle and ductile fracture respectively, while [YOH00] coupled this work to explosions. Finite elements were also used for fracture in [MMDJ01]. Other work includes the adaptive framework of [DDCB01], the rotation based approach in [MDM*02] and [MG04], the hybrid finite element free form deformation approaches in [CGC*02a, CGC*02b], and the finite volume muscle models of [TSBNF03]. Other interesting approaches to the simulation of deformable objects include [JP02, JF03].

Many authors have worked to improve the robustness of mass-spring systems. [Pal94] used a pseudopressure term in addition to edge springs, [BC00] used springs emanating from the barycenter of each tetrahedron to preserve volume, and [CM97] introduced altitude springs to prevent triangles from collapsing. [MBTF03] later improved this model and applied it in three spatial dimensions to the case of tetrahedral mesh generation. If altitude springs are used correctly, not only is inversion not a problem, but the elements will work to un-invert. Unfortunately, spring systems do not allow the modeling of arbitrary constitutive models.

We also show examples of our method at work for the in-plane deformations of lower dimensional manifolds such as cloth and shells. Here, since triangles cannot invert in three spatial dimensions, our method is similar to the work of [EKS03], except that they do not consider degenerate elements. For out-of-plane forces, we use the bending model of [BMF03] (see also [GHDS03]), and for self-collisions we use the method of [BFA02]. Moreover, for volumetric collisions we also use the method in [BFA02] simply applied to the triangulated boundary surface of the tetrahedron mesh. Other interesting work on cloth and shells includes the implicit time stepping of [BW98], the bending model of [CK02], the adaptive simulation work of [GKS02], and the self-collision untangling strategy of [BWK03].

3. Measuring deformation

A deformable object is characterized by a time dependent map ϕ from material coordinates \mathbf{X} to world coordinates \mathbf{x} . The stress at a given point \mathbf{X} in the material depends only on the deformation gradient $\mathbf{F}(\mathbf{X}) = \partial\mathbf{x}/\partial\mathbf{X}$ of this mapping. Since we are using a purely Lagrangian framework, all mappings are based in material space. We restrict ourselves to constant strain tetrahedral elements, and therefore work with a fixed tetrahedron mesh in material space. In this case, ϕ becomes a piecewise linear map completely determined by its values at the vertices of the mesh, and \mathbf{F} is a constant 3×3 matrix in each tetrahedron.

Since [TSBNF03] proved that their constant strain tetrahedra finite volume method was *exactly* the finite element method for this case (see also [TBC03]), we adopt their geometrically intuitive notation. For simplicity, consider two spatial dimensions where each element is a triangle. We define edge vectors for each triangle as $\mathbf{d}_{m_1} = \mathbf{X}_1 - \mathbf{X}_0$, $\mathbf{d}_{m_2} = \mathbf{X}_2 - \mathbf{X}_0$, $\mathbf{d}_{s_1} = \mathbf{x}_1 - \mathbf{x}_0$, $\mathbf{d}_{s_2} = \mathbf{x}_2 - \mathbf{x}_0$, and construct 2×2 matrices \mathbf{D}_m with columns \mathbf{d}_{m_1} and \mathbf{d}_{m_2} , and \mathbf{D}_s with columns \mathbf{d}_{s_1} and \mathbf{d}_{s_2} . Then we have $\mathbf{F} = \mathbf{D}_s \mathbf{D}_m^{-1}$ within this triangle. In three dimensions, \mathbf{D}_m and \mathbf{D}_s become 3×3 matrices, and $\mathbf{F} = \mathbf{D}_s \mathbf{D}_m^{-1}$ still holds. Since the tetrahedron mesh is fixed in material coordinates, \mathbf{D}_m^{-1} is constant and can be precomputed for efficiency. More importantly, as long as the initial tetrahedron mesh is reasonable, \mathbf{D}_m is well-conditioned, and therefore $\mathbf{F} = \mathbf{D}_s \mathbf{D}_m^{-1}$ is well-defined and finite regardless of the current state of the object. Furthermore, \mathbf{F} contains all the information about the deformation of each tetrahedron, since we can recover \mathbf{D}_s via $\mathbf{D}_s = \mathbf{F} \mathbf{D}_m$. In particular, we can use \mathbf{F} to detect whether elements are inverted by checking the sign of $\det \mathbf{F}$.

If the material is isotropic, we can save storage space by performing a QR-decomposition of \mathbf{D}_m and storing only the upper triangular part, as noted in [MMDJ01]. This corresponds to rotating material space, and therefore has no effect on an isotropic material. This optimization can be performed for an anisotropic model by rotating the anisotropic terms via the rotation from the QR-decomposition.

The next step is usually to define the Green strain $\mathbf{G} = 1/2(\mathbf{F}^T \mathbf{F} - \mathbf{I})$, and compute stress and forces based on \mathbf{G} . We do not do this, however, since \mathbf{G} is invariant with respect to all orthogonal transformations, including reflection, and is therefore incapable of detecting element inversion. Furthermore, \mathbf{G} is already nonlinear in the deformation, and it is therefore more difficult to interpret the large deformation behavior of a constitutive model based on \mathbf{G} than one based on \mathbf{F} , which is linearly related to deformation. Thus, for the remainder of this paper, we make the (nonrestrictive) assumption that the constitutive model is written explicitly in terms of \mathbf{F} .

4. Force computation

Suppose the constitutive model is given as a first Piola-Kirchhoff stress \mathbf{P} , i.e. a mapping from area-weighted normals in material space to traction vectors in world space. The force on a node i due to a single tetrahedron incident to it is $\mathbf{g}_i = -\mathbf{P}(A_1 \mathbf{N}_1 + A_2 \mathbf{N}_2 + A_3 \mathbf{N}_3)/3$, where $A_j \mathbf{N}_j$ are the area weighted normals (in material coordinates) of the faces of the tetrahedron incident to node i . Since these do not change during the simulation, we can precompute a vector \mathbf{b}_i such that $\mathbf{g}_i = \mathbf{P} \mathbf{b}_i$. We optimize the computation by calculating \mathbf{g}_0 as $\mathbf{g}_0 = -(\mathbf{g}_1 + \mathbf{g}_2 + \mathbf{g}_3)$, and compactly express the other \mathbf{g}_i as $\mathbf{G} = \mathbf{P} \mathbf{B}_m$, where $\mathbf{G} = (\mathbf{g}_1, \mathbf{g}_2, \mathbf{g}_3)$ and $\mathbf{B}_m = (\mathbf{b}_1, \mathbf{b}_2, \mathbf{b}_3)$. Note that since \mathbf{B}_m is constant, the nodal forces are linearly related to \mathbf{P} . Therefore, the key to obtaining robust forces in face of large deformation is an accurate calculation of \mathbf{P} .

If a constitutive model is given in terms of a Cauchy stress σ or second Piola-Kirchhoff stress \mathbf{S} , we can easily convert to a first Piola-Kirchhoff stress via the formulas $\mathbf{P} = \mathbf{F} \mathbf{S}$ and $\mathbf{P} = J \sigma \mathbf{F}^{-T}$ where $J = \det \mathbf{F}$. Alternatively, one could rewrite the force formula $\mathbf{G} = \mathbf{P} \mathbf{B}_m$ directly in terms of the other stresses as $\mathbf{G} = \sigma \mathbf{B}_s$ or $\mathbf{G} = \mathbf{F} \mathbf{S} \mathbf{B}_m$. Unlike the first Piola-Kirchhoff case where obtaining a valid \mathbf{P} is sufficient to obtain robust forces, computing a valid σ or \mathbf{S} is not enough. For example, if the tetrahedron is a single point, \mathbf{F} and thus $\mathbf{G} = \mathbf{F} \mathbf{S} \mathbf{B}_m$ are identically zero. Moreover, \mathbf{B}_s (the analogue to \mathbf{B}_m in world or spatial coordinates) depends on the area weighted normals of the deformed, possibly degenerate tetrahedra. And thus \mathbf{B}_s and $\mathbf{G} = \sigma \mathbf{B}_s$ are identically zero as well. That is, there are no restorative forces in either instance. Therefore, we write all constitutive models in terms of \mathbf{P} before force computation.

5. Diagonalization

Since rigid body rotations do not change the physics of a deformable object, the stress \mathbf{P} satisfies $\mathbf{P}(\mathbf{U}\mathbf{F}) = \mathbf{U}\mathbf{P}(\mathbf{F})$ for any rotation \mathbf{U} . (Here $\mathbf{P}(\mathbf{F})$ denotes function application.) Furthermore, if we temporarily assume an isotropic constitutive model, \mathbf{P} is invariant under rotations of material space, i.e. $\mathbf{P}(\mathbf{F}\mathbf{V}^T) = \mathbf{P}(\mathbf{F})\mathbf{V}^T$. Therefore, if we diagonalize \mathbf{F} via rotations \mathbf{U} and \mathbf{V} to obtain $\mathbf{F} = \mathbf{U}\hat{\mathbf{F}}\mathbf{V}^T$, \mathbf{P} becomes

$$\mathbf{P} = \mathbf{P}(\mathbf{F}) = \mathbf{U}\mathbf{P}(\hat{\mathbf{F}})\mathbf{V}^T = \hat{\mathbf{U}}\hat{\mathbf{P}}\hat{\mathbf{V}}^T \quad (1)$$

where a hat superscript denotes the corresponding rotated quantity. Since the elastic energy of an isotropic material is invariant under world and material rotations, it can depend only on the invariants of \mathbf{F} , or equivalently on the entries of the diagonalization $\hat{\mathbf{F}}$ (see e.g. [BW97]). Therefore, the gradient of the energy, $\hat{\sigma}$, will also be diagonal. Moreover, since the three stresses are related via $\hat{\sigma} = (1/J)\hat{\mathbf{P}}\hat{\mathbf{F}}^T$ and $\hat{\mathbf{P}} = \hat{\mathbf{F}}\hat{\mathbf{S}}$, the diagonalization of \mathbf{F} actually results in the simultaneous diagonalization of all three stresses. In particular, $\hat{\mathbf{P}}$ in equation 1 is diagonal for an isotropic constitutive model. For an anisotropic constitutive model, a diagonal $\hat{\mathbf{F}}$ does not result

in a diagonal $\hat{\mathbf{P}}$. However, this is not restrictive, and we show examples of anisotropic constitutive models in section 6.1.

The diagonalization of \mathbf{F} is not unique, however. While the ordering of the entries of the diagonal matrix $\hat{\mathbf{F}}$ is unimportant, the signs of the entries determine whether the tetrahedron is inverted in a particular direction. The standard SVD convention of choosing all nonnegative entries works only when $\det \mathbf{F} \geq 0$. When $\det \mathbf{F} < 0$, the signs of the entries must be chosen carefully in order to guarantee that the forces act to uninvert the tetrahedron. In this case, $\hat{\mathbf{F}}$ has either one or three negative entries. We heuristically assume that each tetrahedron is as uninverted as possible, and thus we assume that only one entry (not three) is negative. Moreover, the entry with the smallest magnitude is chosen to be negative. This is motivated by the geometric fact that an inverted tetrahedron can be uninverted by moving any one node across the plane of the opposite face, and it is most efficient to choose the node that is closest to the opposite face.

We compute the correct diagonalization by finding any diagonalization and correcting the signs. When doing this, we must be careful to ensure that the final \mathbf{U} and \mathbf{V} are pure rotations, i.e., $\det \mathbf{U} = \det \mathbf{V} = 1$. This is because deformable objects are not invariant under reflections of material or world space, and equation 1 does not hold if either \mathbf{U} or \mathbf{V} is a reflection. We compute the SVD of $\mathbf{F} = \mathbf{U}\hat{\mathbf{F}}\mathbf{V}^T$ as follows. First we form the normal equations $\mathbf{F}^T\mathbf{F} = \mathbf{V}\hat{\mathbf{F}}\mathbf{U}^T\mathbf{U}\hat{\mathbf{F}}\mathbf{V}^T = \mathbf{V}\hat{\mathbf{F}}^2\mathbf{V}^T$. Then we rearrange to obtain an eigenproblem, $\mathbf{F}^T\mathbf{F}\mathbf{V} = \mathbf{V}\hat{\mathbf{F}}^2$ for the symmetric positive semidefinite $\mathbf{F}^T\mathbf{F}$. Here, \mathbf{V} is an orthogonal matrix of eigenvectors and $\hat{\mathbf{F}}^2$ is a diagonal matrix with nonnegative entries. Robust computation of eigensystems for 3×3 matrices (even with repeated or zero eigenvalues) is a solved problem. And since it is relevant to rigid body simulations, it has received a lot of attention. Note that if \mathbf{V} is a reflection with $\det \mathbf{V} = -1$ we can simply multiply a column of \mathbf{V} by -1 to make \mathbf{V} a rotation with $\det \mathbf{V} = 1$. The entries of $\hat{\mathbf{F}}$ are then determined by taking the square root of the diagonal elements of $\hat{\mathbf{F}}^2$, and \mathbf{U} can be found via $\mathbf{U} = \mathbf{F}\mathbf{V}\hat{\mathbf{F}}^{-1}$ for well shaped elements. However, if a diagonal entry of $\hat{\mathbf{F}}$ is near zero, which is the case for flat tetrahedra, we do not use this formula for the corresponding column of \mathbf{U} , but instead take it to be orthogonal to the other columns. For example, in the extreme case where the tetrahedron is a single point and $\mathbf{F} = \mathbf{0}$, we choose $\mathbf{U} = \mathbf{I}$. Finally, for inverted tetrahedra with $\det \mathbf{F} < 0$, we have $\det \mathbf{U} = -1$ implying that \mathbf{U} is a reflection. This is removed by negating the minimal element of $\hat{\mathbf{F}}$ and the corresponding column of \mathbf{U} . Figure 2 illustrates degeneracy and inversion handling for a torus mesh. Moreover, we have tested our approach for a variety of degenerate configurations, such as when a tetrahedron collapses to a single point or line, and the method always leads to robust recovery from inversion.




Figure 2: A torus with zero strength collapses into a puddle. When the strength is increased, the torus recovers.

5.1. Other Rotations

Typically, authors use a polar decomposition to remove the world rotation of a tetrahedron producing a symmetric \mathbf{F}_s with $\mathbf{F} = \mathbf{Q}\mathbf{F}_s$. To recover from inversion, one must be careful to control the signs of the eigenvalues of \mathbf{F}_s as in the diagonalization case. However, we know of no way to do this without first computing the full diagonalization $\mathbf{F} = \mathbf{U}\hat{\mathbf{F}}\mathbf{V}^T$, and forming the polar decomposition via $\mathbf{Q} = \mathbf{U}\mathbf{V}^T$, $\mathbf{F}_s = \mathbf{V}\hat{\mathbf{F}}\mathbf{V}^T$. Polar decomposition was used by [EKS03] for cloth simulation and [MG04] for volumetric solids, but neither showed how to correctly handle inverting elements. Since inversion does not occur for triangles in three spatial dimensions, this is only a difficulty in the volumetric case.

Alternatively, one could attempt to remove the world rotation with a QR-decomposition, i.e. $\mathbf{F} = \mathbf{Q}\mathbf{F}_r$ with \mathbf{F}_r an upper triangular matrix. However, any stress which depends linearly on \mathbf{F}_r will be anisotropic in a mesh dependent way, since it is not invariant under rotations of material space. To see this, it suffices to note that if \mathbf{V} is a rotation of material space, then $\mathbf{F}\mathbf{V}^T = \mathbf{Q}\mathbf{F}_r\mathbf{V}^T$, and $\mathbf{F}_r\mathbf{V}^T$ is not upper triangular. Therefore, QR-decomposition is inadvisable even if physical accuracy is not a requirement. Moreover, this is a problem with any method for removing rotations, such as [MDM*02], that does not use $\mathbf{Q} = \mathbf{U}\mathbf{V}^T$ for the world rotation.

Note that our approach departs from the typical goal of determining (or approximating) the rotation from material space to world space, i.e. \mathbf{Q} from the polar decomposition. Instead, we look for *two* rotations \mathbf{U} and \mathbf{V} such that \mathbf{U}^T and \mathbf{V}^T rotate the world and material spaces, respectively, to a space where the deformation gradient is a diagonal matrix. This is preferable to the space obtained using \mathbf{Q} in which the deformation gradient is a more complex symmetric matrix.

6. Constitutive Models

Once we have carefully diagonalized \mathbf{F} , we can extend our constitutive models to behave reasonably under inversion. The diagonal setting makes this quite simple.

If St. Venant-Kirchhoff material is compressed beyond a certain point, it gets weaker and weaker as the compression increases, and the stress drops to zero as the object becomes flat. Moreover, if an element inverts, the forces act to keep the element inverted. See Figure 3 (upper left). As noted in [BW97], this makes the St. Venant-Kirchhoff model completely useless for modeling large deformations. [OH99] noted these difficulties, but dismissed them since they were simulating rigid materials. However, as discussed previously, stiff or incompressible objects may still have inverted tetrahedra due to discretization error, especially on the coarse grids common in the computer graphics community.

In order to alleviate the problems with the St. Venant-Kirchhoff model, various authors (e.g. [PDA01]) have proposed adding a pseudo-pressure term to prevent element inversion. In fact, the classical neo-Hookean constitutive model already does this as shown in Figure 3 (upper right). The singularity at the origin means that infinite energy is required to completely flatten a tetrahedron, and as long as the equations for this constitutive model are accurately simulated, inversion is prevented. However, preventing inversion also prevents the handling of situations where inversion is the desired, correct response, as in Figure 1. Moreover, since the forces become arbitrarily large, the system can become arbitrarily stiff and difficult to integrate, making it difficult to handle situations such as that shown in Figure 6 where a volumetric Buddha model is pulled through rigid, interlocking gears.

To avoid the unnecessary stiffness associated with the neo-Hookean constitutive model, we modify the constitutive model near the origin to remove the singularity by either linearizing at a given compression limit or simply clamping the stress at some maximum value. Moreover, as shown in Figure 3 (lower right), we extend the model past the origin into the inverted regime in order to obtain valid forces for inverted elements. These forces act to uninvert the element. Note that since we have removed both spatial and material rotations by diagonalizing, the modified model is automatically rotation invariant and isotropic.

The major strength of the diagonal setting is that these modifications can be applied to arbitrary constitutive models. This is quite natural, since the diagonal setting is also commonly used in the experimental determination of material parameters. The resulting model is identical to the physical model most of the time, and allows the simulation to continue if a few tetrahedrons invert. Furthermore, our extensions provide C^0 or C^1 continuity around the flat case, which avoids sudden jumps or oscillations which might effect neighboring elements.




Figure 3: The relationship between the first Piola-Kirchhoff stress $\hat{\mathbf{P}}$ and the deformation gradient $\hat{\mathbf{F}}$ for various constitutive models.

While it may seem nonphysical to modify a constitutive model for inversion handling, most constitutive models lose accuracy long before inversion occurs. It is exceedingly difficult to measure material response in situations of extreme compression, so constitutive models are often measured for moderate deformation and continued heuristically down to the flat cases. Given that some accuracy loss is unavoidable when tetrahedrons are nearly flat, it is preferable to provide smooth, consistent handling of inversion in order to avoid unnecessary corruption of the more meaningful parts of the simulation.

If a specific qualitative material behavior is desired but the exact quantitative model is less important, we can use the diagonal setting to construct a suitable constitutive model. For example, most biological material is soft under small deformation, but becomes stiffer as the deformation increases. A simple model capturing this behavior is given by choosing threshold values for compression and elongation, specifying the slope of the stress curve outside these threshold values and at the undeformed state, and using a cubic spline to interpolate between them. This model, equipped with a linear pressure component, was used for the simulations of the volumetric Buddha model shown in Figures 5 and 6. Note that any isotropic constitutive model expressed in diagonal form will automatically preserve angular momentum, since if $\hat{\mathbf{P}}$ is diagonal, $\hat{\mathbf{S}} = \hat{\mathbf{F}}^{-1}\hat{\mathbf{P}}$ is symmetric (see e.g. [BW97]). For the torus puddle and plastic sphere simulations (see Figures 2 and 7), where the focus is on degeneracy and plasticity, respectively, we used the simple rotated linear model $\hat{\mathbf{P}} = 2\mu(\hat{\mathbf{F}} - \mathbf{I}) + \lambda tr(\hat{\mathbf{F}} - \mathbf{I})$ depicted in Figure 3 (lower left).




Figure 4: A simulation of muscles driven by a key-framed skeleton. The muscle is represented with a transversely isotropic constitutive model, and the strength along the fiber direction in the muscle is based on activation levels.

Once we have computed the diagonalized stress $\hat{\mathbf{P}}$, the force computation becomes

$$\mathbf{G} = \mathbf{U}\hat{\mathbf{P}}\mathbf{V}^T \mathbf{B}_m = \mathbf{U}\hat{\mathbf{P}}\hat{\mathbf{B}}_m \quad (2)$$

where $\hat{\mathbf{B}}_m = \mathbf{V}^T \mathbf{B}_m$ can be computed and stored if the rotation is fixed for multiple force computations, as in some versions of Newmark time integration (see section 9).

6.1. Anisotropy

If the constitutive model includes anisotropic components, it is no longer invariant under rotations of material space. However, we can continue to fully diagonalize \mathbf{F} , and rotate the anisotropic terms using \mathbf{V} . Since we still work with a diagonal $\hat{\mathbf{F}}$, the large deformation behavior of the constitutive model is still apparent and easy to modify to handle inversion. For example, if the material is stronger in a certain material direction \mathbf{a} , we diagonalize \mathbf{F} and use $\mathbf{V}^T \mathbf{a}$ in the computation of $\hat{\mathbf{P}}$. $\hat{\mathbf{P}}$ is no longer a diagonal matrix, but we can still compute forces using equation 2. When constructing anisotropic constitutive models that allow inversion, we write $\hat{\mathbf{P}}$ as a diagonal matrix plus $\hat{\mathbf{F}}$ times a symmetric matrix for the anisotropic terms. Then $\hat{\mathbf{S}} = \hat{\mathbf{F}}^{-1} \hat{\mathbf{P}}$ is symmetric (preserving angular momentum) as required.

We illustrate this with an example simulation of skeletal muscle in the upper limb (see Figure 4). We use a nonlinear transversely-isotropic quasi-incompressible constitutive model. See [TSBNTH*04] for more details. This is an intricate region of the body articulated with complex joints in the shoulder, elbow and wrist. Inaccuracy in the joint models and motion data leads to skeletal configurations that are incompatible with the musculature creating boundary conditions that degenerately deform muscles and tendons leading to spurious element inversion. However, these configurations often only occur in limited regions of the mesh and only for brief moments during a given motion. Our algorithm allows simulations to progress past these temporary problems by letting elements invert and then later recover.

6.2. Damping

Damping forces can be implemented by rotating the velocity gradient $\hat{\mathbf{F}}$ by the same \mathbf{U} and \mathbf{V} used to diagonalize \mathbf{F} , computing the damping stress $\hat{\mathbf{P}}$ in the rotated frame, and

computing the force exactly as for the elastic case. Note that the rotated velocity gradient will in general not be diagonal.

As in the case of anisotropic elastic forces, a damping model will only preserve angular momentum if $\hat{\mathbf{P}}$ can be expressed as $\hat{\mathbf{F}}\hat{\mathbf{S}}$, with $\hat{\mathbf{S}}$ symmetric. This was not a problem for anisotropy since the anisotropic terms are usually not important for flat or inverted elements. However, in order to prevent visually unpleasant oscillations, we do not want the damping forces to disappear for flat tetrahedra. For example, the analogous damping model to the rotated linear constitutive model, $\hat{\mathbf{P}} = \beta(\hat{\mathbf{F}} + \hat{\mathbf{F}}^T) + \alpha \text{tr}(\hat{\mathbf{F}})$, does not preserve angular momentum unless $\hat{\mathbf{F}}$ is a uniform scaling. However, since the angular momentum errors are small around the undeformed state, and highly deformed tetrahedra are usually interacting with other objects, we have not found this lack of conservation to be visually noticeable. In simulations where more physical accuracy is desired, we use a correct damping model for moderate deformations and a more robust but nonphysical model for the few flat or inverted tetrahedra.

6.3. Plasticity

We represent plastic deformation with a multiplicative decomposition of the deformation $\mathbf{F} = \mathbf{F}_e \mathbf{F}_p$, where \mathbf{F}_p represents the permanent plastic deformation and \mathbf{F}_e the elastic deformation, see e.g. [BW97] or [AL03]. The multiplicative formulation allows a complete separation between plastic flow and elastic forces, and makes constraints such as volume preservation simple to enforce. In contrast, the additive plasticity formulation of [OBH02] does not support true incompressibility, though this might not be a significant problem for graphics applications. Note that if the elastic constitutive model is isotropic, the rotational part of \mathbf{F}_p is arbitrary, e.g. we can choose \mathbf{F}_p to be symmetric.

We restrict ourselves to rate-independent plasticity models, and use the return mapping algorithm to transfer deformation from the elastic part \mathbf{F}_e to the plastic part \mathbf{F}_p whenever a yield criterion on \mathbf{F}_e is exceeded. The details of the computation of plastic flow are as follows. Compute the trial elastic deformation $\mathbf{F}_{e,trial} = \mathbf{F}\mathbf{F}_p^{-1}$, and find the diagonalization $\mathbf{F}_{e,trial} = \mathbf{U}\hat{\mathbf{F}}_{e,trial}\mathbf{V}^T$. If a yield criterion on $\hat{\mathbf{F}}_{e,trial}$ is exceeded, project back onto the yield surface producing a new diagonal matrix $\hat{\mathbf{F}}_{e,proj}$. Compute the trial plastic deformation $\mathbf{F}_{p,trial} = \hat{\mathbf{F}}_{e,proj}^{-1} \mathbf{U}^T \mathbf{F}$ (dropping \mathbf{V} since rotations

of \mathbf{F}_p are unimportant). If $\mathbf{F}_{p,trial}$ exceeds a separate limit criterion, project it back onto the limit surface producing the final \mathbf{F}_p . Compute and store \mathbf{F}_p^{-1} for future use.

This structure supports an arbitrary plastic yield criterion while still ensuring that the plastic deformation does not become too extreme. This is important, since the time step required for stability depends on the conditioning of \mathbf{F}_p . In particular, \mathbf{F}_p should never invert. This can be implemented in the diagonal framework by ensuring that the projection of $\hat{\mathbf{F}}_e$ always results in an $\hat{\mathbf{F}}_{e,proj}$ with positive entries. The final limiting of $\mathbf{F}_{p,trial}$ can then be adjusted to ensure a well-conditioned \mathbf{F}_p . Figure 9 shows a simulation of ductile fracture using this technique. See [MBF04] for more details.

7. Controlling Plasticity

Various authors, such as [FM97, PSE*00, TMPS03], have considered controlling physics based simulations. The ability to control a simulation alleviates the need for laborious parameter tuning to achieve a desired effect, and makes possible animations which could not be achieved through physical accuracy alone. In the context of plasticity, we can use the plastic limiting projection step to control the plastic deformation toward any desired state without sacrificing realism. To do this, we compute a goal deformation $\mathbf{F}_{p,goal}$ at the beginning of the simulation. In the plastic projection step, we are given a tentative plastic flow from the old deformation \mathbf{F}_p to the trial deformation $\mathbf{F}_{p,trial}$. In order to always move towards $\mathbf{F}_{p,goal}$, we choose the final plastic deformation to be the point on the segment from \mathbf{F}_p to $\mathbf{F}_{p,trial}$ which is closest to $\mathbf{F}_{p,goal}$. This computation is actually performed on the logarithms of each \mathbf{F}_p after removing the world rotation. Since the mapping from rest to goal state will rarely preserve volume locally, volume preservation should not be used during the elastic projection step.

Allowing some flexibility in the plastic flow allows the deformation to pick up additional fine detail not present in the goal state. For example, Figure 7 shows a plastic sphere pulled through interlocking gears. The sphere is controlled towards the flattened disk shape. In particular, the goal state does not include the teeth marks present in the final state of the sphere. A more obvious example of control is shown in Figure 7. Both of these examples used the simple yield criterion $\|\log \mathbf{F}_e\| \leq \gamma$.

8. Thin Shells and Cloth

The diagonalized framework is readily extended to handle the in-plane behavior of thin shells and cloth (see Figure 5). Here, \mathbf{F} is a 3×2 matrix decomposed as $\mathbf{F} = \mathbf{U}\hat{\mathbf{F}}\mathbf{V}^T$ where \mathbf{U} is a 3×2 matrix with orthonormal columns, $\hat{\mathbf{F}}$ is a 2×2 diagonal matrix, and \mathbf{V} is a 2×2 rotation matrix. Everything else follows in a straightforward manner.

Inversion does not occur for freely moving thin shells and

cloth, since an “inverted” triangle is indistinguishable from a triangle that has been rotated 180° out of plane. However, when triangles degenerate to lines or points special care is needed. Moreover, when a shell approximates a two dimensional surface such as during surface mesh generation (see e.g. [MBTF03]), “inversion” can occur. That is, a triangle can be tested for inversion by considering the sign of the dot product between its face normal and a known approximation to the surface normal at the center of the triangle. If this sign is negative, the triangle can be considered inverted, and the signs of the entries of $\hat{\mathbf{F}}$ can be corrected as before. Thus, the triangle acts to uninvert by flipping the direction of its face normal.

For bending forces, we use the formulation of [BMF03], which is similar to that of [GHDS03]. The bending model is independent of the in-plane model, and in-plane plasticity is analogous to the three dimensional case. To allow plastic bending, we apply the plastic flow algorithm to the rest angles between each pair of adjacent triangles. An example of a shell simulation showing both in-plane and bending plasticity is shown in Figure 8.

9. Time-Stepping and Collision Handling

We use the Newmark time-stepping scheme of [BMF03] with explicit integration for the elastic forces and implicit integration for the damping forces. Treating only the damping forces implicitly removes the strict quadratic time step restriction required by fully explicit schemes without introducing the extra artificial damping characteristic of fully implicit schemes. As most damping models are linear in the velocities with a positive definite, symmetric Jacobian, the implicit integration can be implemented using a fast conjugate gradient solver.

We modify [BMF03]’s scheme slightly to improve the handling of rigid body collisions. Specifically, we use the velocity from the last implicit update as input to the rigid body collision algorithm, and use constraints in the velocity update to prevent motion in the direction normal to the rigid body for points experiencing a collision. The resulting algorithm to move from step n to $n+1$ is as follows:

- $\tilde{v}^{n+1/2} = v^n + \frac{\Delta t}{2} a(t^n, x^n, \tilde{v}^{n+1/2})$
- $\tilde{x}^{n+1} = x^n + \Delta t \tilde{v}^{n+1/2}$
- Process rigid body collisions using \tilde{x}^{n+1} and v^n , producing final positions x^{n+1} and modified velocities \tilde{v}^n .
- $v^{n+1} = \tilde{v}^n + \Delta t (a(t^n, x^n, \tilde{v}^n) + a(t^{n+1}, x^{n+1}, v^{n+1})) / 2$

Note that the last line is exactly the trapezoidal rule applied to the velocities. This algorithm supports a variable time step with second order accuracy and monotone behavior. Since the positions change only in lines 2 and 3 of the algorithm, we can compute \mathbf{F} and its diagonalization only once per time step after step 3. Plastic flow is also computed at this time.

With this optimization, the cost of diagonalization becomes *negligible* compared to the cost of the implicit velocity updates.

The rigid body collision processing is based on the algorithm of [BMF03]. We represent rigid bodies as implicit surfaces, which simplifies collision detection. Each such node is projected to the surface of the object, and its normal velocity is set to that of the object if it is not already moving away from it. We incorporate friction by changing the relative tangential velocity $v_{T,rel}$ to

$$v_{T,rel}^{new} = \max\left(0, 1 - \mu \frac{\Delta v_N + \Delta x_N / \Delta t}{|v_{T,rel}|}\right) v_{T,rel}$$

where Δx_N and Δv_N are the changes in position and normal velocity from the projection step. The $\Delta x_N / \Delta t$ term ensures that the particle will experience the correct friction for the change in position imparted by the object. This term was not considered in [BMF03].

Any node involved in a collision is flagged, and its normal velocity is held fixed during the final trapezoidal rule step. Enforcing normal velocities of colliding particles via constraints during the velocity update further increases the stability of the collision scheme, since it allows a nonlocal response to collision. This strategy is similar to that proposed in [BW98], who implemented rigid body collisions in their fully implicit scheme via constraints in the conjugate gradient solver.

Since projecting points to the surface of an object tends to crush tetrahedra, the ability to handle flat or inverted tetrahedra is essential to enable the use of reasonable time steps. Also, since the rigid body collision algorithm is applied to vertices only, not triangles, it is useful to apply the algorithm to the interior points as well as the surface points, to prevent small rigid bodies from slipping between surface points into the interior of the object. The importance of this increases for very soft objects, as very soft surface triangles can easily expand and pass around even moderately sized obstacles.

For self-collisions we extract the boundary surface and apply the cloth collision algorithm of [BFA02]. This algorithm is applied “outside” of the time-stepping algorithm outlined above. Although a surface-only collision algorithm does not prevent the interior of the object from extending outside its boundary, our method has no difficulty with this inversion and only the surface is needed for rendering.

10. Examples

We used the algorithm of [MBTF03] to generate the meshes used in this paper. Even without preconditioning in the CG solver, computation times were generally under 20 minutes per frame for the largest meshes. Of course, coarser meshes can be simulated in just a few minutes a frame. For example, the torus simulation in Figure 2 ran at around 10 to 20 seconds per frame with the 115K element mesh, and .5

to 1 second per frame with an 11K element mesh. All the simulations involved large numbers of inverted elements: a typical frame from the Buddha simulation in Figure 6 had about 29K inverted tetrahedrons out of a total of 357K tetrahedrons, or about 8% of the mesh.

The Buddha with cape example in Figure 5 was simulated in two layers, with one-way coupling from the Buddha to the cloth using the collision processing algorithm from section 9. We used the exact triangulated surface geometry of the Buddha in order for the cloth to resolve the many features of the Buddha mesh. To evaluate the Buddha as an implicit surface at a cloth vertex \mathbf{v} , we find the closest point \mathbf{p} to \mathbf{v} on the Buddha surface (which may lie on a vertex, edge, or face) and define the “normal” at \mathbf{v} to be in the direction from \mathbf{v} to \mathbf{p} or \mathbf{p} to \mathbf{v} , whichever points outwards.

11. Conclusions

We have presented a new method for modifying an elastic constitutive model to behave robustly for inverted elements, which works by carefully diagonalizing the deformation mapping prior to computing forces. Examples were presented to demonstrate that this algorithm works well for volumetric and thin shell simulations involving degeneracy, complex geometries, anisotropic constitutive models, plasticity with and without control, ductile fracture, and coupling between different types of deformable objects.

12. Acknowledgement

Research supported in part by an ONR YIP award and a PECASE award (ONR N00014-01-1-0620), a Packard Foundation Fellowship, a Sloan Research Fellowship, ONR N00014-03-1-0071, ARO DAAD19-03-1-0331, NSF IIS-0326388, NSF ACI-0323866 and NSF ITR-0205671. In addition, G. I. and J. T. were supported in part by NSF Graduate Research Fellowships.

We would also like to thank Mike Houston for providing computing resources used for rendering.

References

- [AL03] ARMERO F., LOVE E.: An arbitrary lagrangian-eulerian finite element method for finite strain plasticity. *Int. J. Num. Meth. Eng.* 57 (2003), 471–508.
- [BBTL03] BESSETTE G., BECKER E., TAYLOR L., LITTLEFIELD D.: Modeling of impact problems using an h-adaptive, explicit lagrangian finite element method in three dimensions. *Comput. Meth. in Appl. Mech. and Eng.* 192 (2003), 1649–1679.
- [BC00] BOURGUIGNON D., CANI M. P.: Controlling anisotropy in mass-spring systems. In

- Eurographics* (2000), Eurographics Assoc., pp. 113–123.
- [BFA02] BRIDSON R., FEDKIW R., ANDERSON J.: Robust treatment of collisions, contact and friction for cloth animation. *ACM Trans. Graph. (SIGGRAPH Proc.)* 21 (2002), 594–603.
- [BMF03] BRIDSON R., MARINO S., FEDKIW R.: Simulation of clothing with folds and wrinkles. In *Proc. of the 2003 ACM SIGGRAPH/Eurographics Symp. on Comput. Anim.* (2003), pp. 28–36.
- [BW97] BONET J., WOOD R.: *Nonlinear continuum mechanics for finite element analysis*. Cambridge University Press, Cambridge, 1997.
- [BW98] BARAFF D., WITKIN A.: Large steps in cloth simulation. In *Proc. SIGGRAPH 98* (1998), pp. 1–12.
- [BWK03] BARAFF D., WITKIN A., KASS M.: Untangling cloth. *ACM Trans. Graph. (SIGGRAPH Proc.)* 22 (2003), 862–870.
- [CGC*02a] CAPELL S., GREEN S., CURLESS B., DUCHAMP T., POPOVIĆ Z.: Interactive skeleton-driven dynamic deformations. *ACM Trans. Graph. (SIGGRAPH Proc.)* 21 (2002), 586–593.
- [CGC*02b] CAPELL S., GREEN S., CURLESS B., DUCHAMP T., POPOVIĆ Z.: A multiresolution framework for dynamic deformations. In *ACM SIGGRAPH Symp. on Comput. Anim.* (2002), ACM Press, pp. 41–48.
- [CK02] CHOI K.-J., KO H.-S.: Stable but responsive cloth. *ACM Trans. Graph. (SIGGRAPH Proc.)* 21 (2002), 604–611.
- [CM97] COOPER L., MADDOCK S.: Preventing collapse within mass-spring-damper models of deformable objects. In *The 5th Int. Conf. in Central Europe on Comput. Graphics and Vis.* (1997).
- [CO97] CAMACHO G., ORTIZ M.: Adaptive Lagrangian modelling of ballistic penetration of metallic targets. *Comput. Meth. in Appl. Mech. and Eng.* 142 (1997), 269–301.
- [CZ92] CHEN D., ZELTZER D.: Pump it up: Computer animation of a biomechanically based model of muscle using the finite element method. *Comput. Graph. (SIGGRAPH Proc.)* (1992), 89–98.
- [DDCB01] DEBUNNE G., DESBRUN M., CANI M., BARR A.: Dynamic real-time deformations using space & time adaptive sampling. In *Proc. SIGGRAPH 2001* (2001), vol. 20, pp. 31–36.
- [EKS03] ETZMUSS O., KECKEISEN M., STRASSER W.: A fast finite element solution for cloth modelling. In *Pacific Graph.* (2003), pp. 244–251.
- [ERM*03] ESCOBAR J., RODRÍGUEZ E., MONTENEGRO R., MONTERO G., GONZÁLEZ-YUSTE J.: Simultaneous untangling and smoothing of tetrahedral meshes. *Comput. Meth. in Appl. Mech. and Eng.* 192 (2003), 2775–2787.
- [EZE98] ESPINOSA H., ZAVATTIERI P., EMORE G.: Adaptive FEM computation of geometric and material nonlinearities with application to brittle failure. *Mech. Materials* 29 (1998), 275–305.
- [FM97] FOSTER N., METAXAS D.: Controlling fluid animation. In *Computer Graphics International 1997* (1997), pp. 178–188.
- [GHDS03] GRINSPUN E., HIRANI A., DESBRUN M., SCHRODER P.: Discrete shells. In *Proc. of the 2003 ACM SIGGRAPH/Eurographics Symp. on Comput. Anim.* (2003), pp. 62–67.
- [GKS02] GRINSPUN E., KRYSL P., SCHRODER P.: CHARMS: A simple framework for adaptive simulation. *ACM Trans. Graph. (SIGGRAPH Proc.)* 21 (2002), 281–290.
- [GMTT89] GOURRET J.-P., MAGNENAT-THALMANN N., THALMANN D.: Simulation of object and human skin deformations in a grasping task. *Comput. Graph. (SIGGRAPH Proc.)* (1989), 21–30.
- [HAC74] HIRT C., AMSDEN A., COOK J.: An arbitrary Lagrangian-Eulerian computing method for all flow speeds. *J. Comput. Phys.* 135 (1974), 227–253.
- [HFS*01] HIROTA G., FISHER S., STATE A., LEE C., FUCHS H.: An implicit finite element method for elastic solids in contact. In *Comput. Anim.* (2001).
- [JF03] JAMES D., FATAHALIAN K.: Precomputing interactive dynamic deformable scenes. *ACM Trans. Graph. (SIGGRAPH Proc.)* 22 (2003), 879–887.
- [JP02] JAMES D., PAI D.: DyRT: Dynamic response textures for real time deformation simulation with graphics hardware.

- ACM Trans. Graph. (SIGGRAPH Proc.)* 21 (2002), 582–585.
- [MBF04] MOLINO N., BAO J., FEDKIW R.: A virtual node algorithm for changing mesh topology during simulation. *ACM Trans. Graph. (SIGGRAPH Proc.)*, (in press) (2004).
- [MBTF03] MOLINO N., BRIDSON R., TERAN J., FEDKIW R.: A crystalline, red green strategy for meshing highly deformable objects with tetrahedra. In *12th Int. Meshing Roundtable* (2003), pp. 103–114.
- [MDM*02] MULLER M., DORSEY J., MCMILLAN L., JAGNOW R., CUTLER B.: Stable real-time deformations. In *ACM SIGGRAPH Symp. on Comput. Anim.* (2002), pp. 49–54.
- [MG04] MULLER M., GROSS M.: Interactive virtual materials. In *Graph. Interface* (May 2004).
- [MMDJ01] MULLER M., MCMILLAN L., DORSEY J., JAGNOW R.: Real-time simulation of deformation and fracture of stiff materials. In *Comput. Anim. and Sim. '01* (2001), Proc. Eurographics Workshop, Eurographics Assoc., pp. 99–111.
- [OBH02] O'BRIEN J., BARGTEIL A., HODGINS J.: Graphical modeling of ductile fracture. *ACM Trans. Graph. (SIGGRAPH Proc.)* 21 (2002), 291–294.
- [OH99] O'BRIEN J., HODGINS J.: Graphical modeling and animation of brittle fracture. In *Proc. SIGGRAPH 99* (1999), vol. 18, pp. 137–146.
- [Pal94] PALMERIO B.: An attraction-repulsion mesh adaption model for flow solution on unstructured grids. *Comput. and Fluids* 23, 3 (1994), 487–506.
- [PDA01] PICINBONO G., DELINGETTE H., AYACHE N.: Non-linear and anisotropic elastic soft tissue models for medical simulation. In *IEEE Int. Conf. Robot. and Automation* (2001).
- [PSE*00] POPOVIĆ J., SEITZ S., ERDMANN M., POPOVIĆ Z., WITKIN A.: Interactive manipulation of rigid body simulations. *ACM Trans. Graph. (SIGGRAPH Proc.)* 19 (2000).
- [TBC03] TAYLOR G., BAILEY C., CROSS M.: A vertex-based finite volume method applied to non-linear material problems in computational solid mechanics. *Int. J. Num. Meth. Eng.* 56 (2003), 507–529.
- [TF88a] TERZOPOULOS D., FLEISCHER K.: Deformable models. *The Visual Computer*, 4 (1988), 306–331.
- [TF88b] TERZOPOULOS D., FLEISCHER K.: Modeling inelastic deformation: viscoelasticity, plasticity, fracture. *Comput. Graph. (SIGGRAPH Proc.)* (1988), 269–278.
- [THMG04] TESCHNER M., HEIDELBERGER B., MÜLLER M., GROSS M.: A versatile and robust model for geometrically complex deformable solids. In *Proc. Computer Graphics International* (2004).
- [TMPS03] TREUILLE A., MCNAMARA A., POPOVIC Z., STAM J.: Keyframe control of smoke simulations. *ACM Trans. Graph. (SIGGRAPH Proc.)* 22 (2003), 716–723.
- [TPBF87] TERZOPOULOS D., PLATT J., BARR A., FLEISCHER K.: Elastically deformable models. *Comput. Graph. (Proc. SIGGRAPH 87)* 21, 4 (1987), 205–214.
- [TSBNF03] TERAN J., SALINAS-BLEMKER S., NG V., FEDKIW R.: Finite volume methods for the simulation of skeletal muscle. In *Proc. of the 2003 ACM SIGGRAPH/Eurographics Symp. on Comput. Anim.* (2003), pp. 68–74.
- [TSBNTH*04] TERAN J., SALINAS-BLEMKER S., NGTHOW-HING V., LAU C., FEDKIW R.: Robust simulation of muscle and connective tissues in the upper limb (in review). *IEEE Trans. on Vis. and Comput. Graph.* (2004).
- [VGS04] VACHAL P., GARIMELLA R., SHASHKOV M.: Untangling of 2D meshes in ALE simulation. *J. Comput. Phys.* 196 (2004), 627–644.
- [YOH00] YNGVE G., O'BRIEN J., HODGINS J.: Animating explosions. In *Proc. SIGGRAPH 2000* (2000), vol. 19, pp. 29–36.
- [ZCK98] ZHU Q., CHEN Y., KAUFMAN A.: Real-time biomechanically-based muscle volume deformation using FEM. *Comput. Graph. Forum* 190, 3 (1998), 275–284.




Figure 5: A deformable Buddha with a cape undergoing large deformation when hit by a ball (top). The same with the Buddha removed to illustrate the deformation (bottom). (Cape - 84K triangles, Buddha - 357K tetrahedrons)




Figure 6: A volumetric Buddha model is pushed down with a cylinder and pulled between rigid interlocking gears, then recovers its shape elastically. (300K elements)




Figure 7: A plastic sphere controlled towards a flattened disk shape is pulled through rigid interlocking gears (upper left, upper right, lower left). A more obvious example of plasticity control (lower right).




Figure 8: Half of a torus (shell) simulated with in-plane and bending plasticity. (3.5K triangles).




Figure 9: A Buddha undergoing ductile fracture. (300K elements)