Physically Based Modelling of Human Limbs

Lee Cooper *
Steve Maddock

Department of Computer Science
University of Sheffield
211 Portobello St.
Sheffield, S1 4DP, UK.

Abstract

This paper considers the modelling of human limbs featuring deformable tissue, which automatically exhibit realistic behaviour when simulated in a synthetic environment. A limb model is proposed that incorporates a simple skeleton joint, muscle tissue, non-muscle tissue (e.g. fat), tendons and a skin layer. The only input parameters to the model are for simple low-level control of muscle activation, with the resulting motion and deformation of the limb being generated automatically. To demonstrate the approach we have implemented a preliminary model of the limb which incorporates a new mass-spring-damper model based on octahedral elements for modelling the deformable tissue. The results show effects that occur in real limbs such as bulging and creasing at joints.
KEYWORDS: Animation, Modelling, Dynamics, Simulation, Deformation, Elastic, Human, Limb, Tissue.
CR CATEGORIES: I.3.7 - Three dimensional graphics and realism (Animation): I.6.3 - Simulation and modelling (Applications).

1 Introduction

Producing physically accurate computer based models of human limbs is a relatively new area of research. A number of applications exist for such models. One is in the entertainment area, to produce animated sequences showing realistic looking limbs or even entire bodies. Another important use is in medical applications. Here, physically accurate models could be used as a teaching aid, or to enable surgeons to perform simulated operations. A third application area is in the ergonomic design of environments. Here, virtual limb models or full human models can be used to test a proposed human environment for safety or comfort without the need for real physical prototypes. An example of this type of model is "JACK" [3].

Real limbs are composed of complex structures of deformable tissue (e.g. muscle, fat etc.) and non-deformable tissue (bone). Accurately modelling the way in which these structures behave and interact is a significant challenge. In order to appreciate the problem more fully, it is useful to consider the range of deformation that takes place for the limb surface and the internal tissue regions.

A primary cause of deformation is the position of the skeleton. Bending at joints generally causes bulging and creasing of flesh on the inside of the joint, and thinning of the flesh outside. Other effects are wrinkling of the skin and changes in surface texture due to stretching or squashing of the skin. A second cause is due to muscle activation. This is generally intended to cause changes in skeletal position. However, this may not be the case if the relevant limbs are restrained, in which case the muscle merely produces forces which are matched and opposed by the restraining forces. Similarly, a limb may be moved by external forces without the need for muscle activation. Muscles which do cause motion will generally contract resulting in bulging perpendicular to the direction of contraction. Muscle activation also often leads to increased prominence of the tendons associated with the muscle. This is particularly applicable when the tendon is being used to transmit force around a bent joint (e.g. at the back of a bent knee joint). The mass and inertia properties of flesh are another cause of deformation. This results in static effects e.g. sagging of flesh due to gravity, and dynamic effects e.g. the wobbling of flesh due to sudden changes in speed. Finally, external forces can cause deformation of the skin surface and internal tissue regions. They might include distributed forces e.g. gravity, and forces which act at the skin surface due to contact or collision with external objects.

Research into limb modelling has largely come from two distinct groups: those producing models for computer animation in entertainment, which generally need to be quick to simulate and easy to control, and those producing models for medical applications which need to be physically accurate and are often part of a simulation i.e. they respond realistically to the actions of a user rather than being controlled to perform some desired motion. An early model [17][11] aimed at use in animation, features a geometric skin surface model which deforms as the underlying skeleton moves. The method uses "Joint-dependent Local Deformation" (JLD) operators to transform a mesh of vertices which represent the skin surface. Models have been produced which create bulging and creasing around bent joints as well as mimicking the bulging due to muscle contraction. However, dynamic effects and deformation due to external forces is not included. Generally geometric models of this type are limited since they are only capable of exhibiting the behaviour that has been specifically designed into them, and will not automatically reproduce other effects.

Physically based models attempt to model the low level behaviour of material, and its interaction with other objects. If the model is sufficiently close to the real thing, and includes the relevant physical phenomena, then it should au-

*L.Cooper@dcs.shef.ac.uk

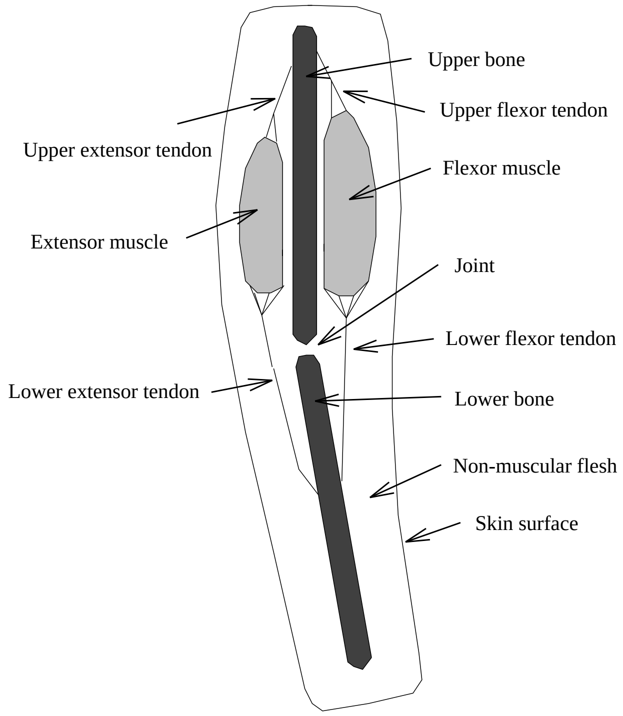

Figure 1: Cross section of the proposed limb model.

tomatically generate realistic behaviour when simulated. In [9] Gourret **et al.** present a physically based model of a hand consisting of a skeleton (driven by a "classical" control method, i.e. parametric, kinematic or dynamic), and a layer of deformable flesh modelled using the Finite Element Method (FEM). The model shows realistic static deformation of the hand as it grasps a deformable ball. In [7] Chen and Zelter produce a physically based FEM model of muscle whose properties are based on data from real muscle samples. A model is presented of the human arm, featuring a single muscle region, attached to a simple skeleton. In [10] Holton and Alexander use masses connected by damped springs to create "soft-cellular elements". These are small deformable elements which are combined to create larger deformable structures.

These models largely concentrate on single tissue regions. A significant step would be to produce models incorporating distinct tissue regions in a single integrated model. A new model based on this idea is presented next.

2 Proposed limb model

The proposed model attempts to maintain simplicity whilst incorporating the main features required to produce a working model capable of representing the mechanics of real limbs (e.g. muscles produce contractile forces which are transfered to bones using tendons). It is not based on a particular limb, e.g. the arm or leg, but on the general structure found in all limbs. Anatomically correct models featuring more complex arrangements of tissue regions can be considered later after the basic methods have been developed.

A cross section of the model is shown in figure 1. The model is based around a single degree of freedom hinge joint with a limited range of motion, of the type found in the knee or interphalangeal joints (joints in the fingers and toes). Two muscle regions are included, one for flexing, and the other for extending the limb. These muscles are attached to the bones at suitable points using tendons. The above elements are surrounded by a region of non-muscular tissue, which represents a merging of all the other deformable tissue types found in a typical limb. Finally, the surface of the limb is represented by a skin surface.

The bones are simply modelled as two non-deformable cylinders with rounded ends which approximate the geometry of the bones found in the arms and legs. The bones rotate about an axis at the centre of the joint, but do not necessarily need to come into contact with one another. The physics of the bones and joint are modelled, i.e. their mass and moments of inertia. Various constraints can be used, for example the upper bone could be held in a fixed vertical position whilst the lower bone is free to swing. The non-muscular flesh is modelled as an elastic homogeneous material with damping, and a certain degree of incompressibility. The muscle regions have similar properties with the addition that they can generate internal contractile forces. The tendons are modelled as inextensible "string-like" objects which are capable of carrying tension. The tendons are constrained to follow a fixed channel within the deformable flesh, but are free to slide through this channel. At the point of attachment with the muscle, a number of sub-tendons could be used to ensure an even distribution of force over the end of the muscle region. These sub-tendons would then combine into a single tendon. Interaction between the flesh and the tendon should be included, e.g. forces and possibly friction, in order to recreate the deformation caused by tendons. Finally, a purely geometric skin surface model is used which can be rendered using standard techniques. The distinct physical properties of skin are neglected at this stage.

Central to the proposed model is a suitable deformable material model, and this is discussed next.

3 Deformable material models

Modelling the various types of tissue that are present in human limbs presents difficulties. This is partly due to the large scale of deformation that normally occurs. Modelling of deformable objects is used in engineering to simulate components which deform during use, however, the methods used generally apply only to small amounts of deformation. As a result, they are not directly applicable, and new ones have to be developed. Before discussing models which have been implemented, it is useful to consider the properties that are required by this application.

- The general class of material is an **elastic**, isotropic material with internal damping.

- **Variable incompressibility** is present in real tissue and is necessary to recreate many of the required effects, e.g. bulging and thinning around joints, and bulging at the limb surface due to muscle contraction.

- Generally the computation required increases as the resolution (the smallest feature which can be represented) increases. Because certain areas of a limb experience lower amounts of deformation that other areas (for example flesh at the middle of bones compared to flesh around joints) then it is desirable to be able to vary the resolution in different regions to improve efficiency.

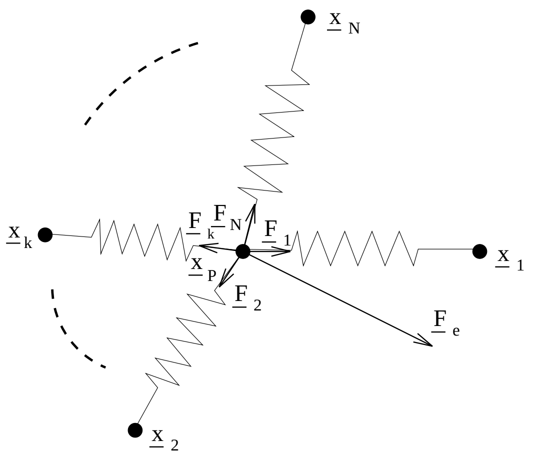

Figure 2: Forces acting on a single mass.

- When creases appear on the inside of joints, different regions of the surface of the limb may come into contact with one another. **Surface self-intersection prevention** is required to produce interaction between the surfaces in order to prevent them from merging.

- Interaction with external objects should be modelled to produce deformation due to contact and collision between the limb surface and the external objects.

- The method should allow arbitrary structures of deformable material to be easily modelled.

- The model should be **robust** enough to reproduce the range of stresses and strains which occur in real limbs.

3.1 Existing models

A number of researchers have considered models for this type of application. Terzopoulos **et al.** [15][13] [14] create models based on the theory of elastic materials, and the Finite Difference Method (FDM). Other researchers have created FEM models [18][8][4]. An alternative method [19][10] [6] uses networks of mass-spring-damper elements. This is an intuitive approach and is relatively easy to implement. The networks consist of point masses connected together by springs and dampers. These elements are "ideal" meaning that the spring elements have no mass or damping, and the masses are point masses with no volume, etc. If the material behaviour required is principally that of an elastic material with damping (as is the case here), then the masses are usually connected by a spring and damper in parallel. However, it is possible to achieve more complex non-elastic behaviour using other arrangements, and introducing slip and fracture elements (see [14]).

The spring and damper elements exert forces on the masses, and this, along with external forces, causes motion of the masses, and hence distortion of the material. Figure 2 shows a single mass with position vector $\underline{x}_p$ connected to N other point masses with position vectors $\underline{x}_1 \ldots \underline{x}_N$ by damped springs. The force exerted on particle p by the spring connecting it to particle k is given by $\underline{F}_k$. $\underline{F}_e$ also acts on p and represents the resultant of all the external

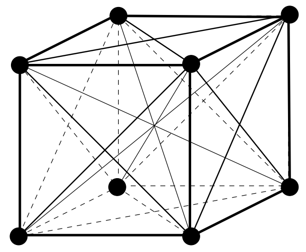

Figure 3: Standard cubic element.

forces. Applying Newton's second law to the particle gives,

$$\underline{F}_e + \sum_{k=1}^{N} \underline{F}_k = m_p \underline{\ddot{x}}_p$$

where m_p is the mass of the particle. $\underline{F}_k$ is given by,

$$\underline{F}_k = -\left[k_s(|\underline{l}| - r) + k_d \frac{\dot{\underline{l}} \cdot \underline{l}}{|\underline{l}|} \right] \frac{\underline{l}}{|\underline{l}|}$$

where $\underline{l} = \underline{x}_p - \underline{x}_k$, $\dot{\underline{l}} = \underline{\dot{x}}_p - \underline{\dot{x}}_k$, r is the rest length of the damped spring, and k_s and k_d are the spring and damping constants respectively. By using a suitable method of numerical integration the motion of the masses can be computed. In order to model large blocks of material, basic small deformable elements are usually created which are then combined to form the required shape. The most common element used [6][19] is a cubic element consisting of eight masses (positioned at the corners of the cube) with springs connecting each mass to every other mass as seen in figure 3. Models composed of cubic elements give behaviour approximating that of real continuous elastic bodies over a range of conditions, however we have performed tests which show that they become unstable at certain levels of compression. This is discussed next.

4 Analysing mass-spring-damper model behaviour

A simple test used to analyse real materials is to compress or stretch them whilst measuring the force required to do so. This has been carried out on samples of the simulated material based on cubic elements composed of mass-spring-damper elements. Consider the rectangular sample shown in figure 4 with unloaded length l, and cross sectional area A. As the sample is loaded under tension (stretched) or compression (squashed) by applying a force F to the end of the sample, its length changes by an amount u. Quantities called *nominal stress* and *nominal strain* can be used to give crude measures of the loading on the sample and its resulting deformation. More refined measures exist, e.g. *true stress* and *true strain* (see [2]), but the former will suffice for this discussion. Nominal stress is given by $\sigma = F/A$, and nominal strain by $\epsilon_n = u/l$. To simulate the tests, one end of the sample was effectively fixed to the floor, and the other end fixed to a moving block. This was achieved by constraining the vertical position and velocity of the point masses lying at the ends of the sample to be the vertical position and velocity of the floor or the block. The block was then moved with a constant velocity, and the resultant reaction forces acting on

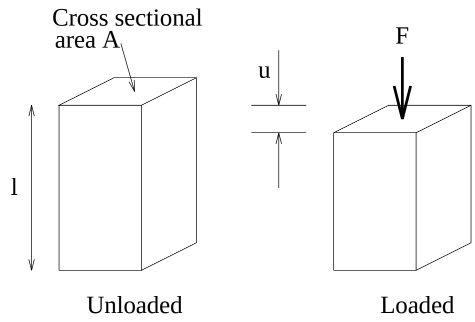

Figure 4: Testing a material sample.

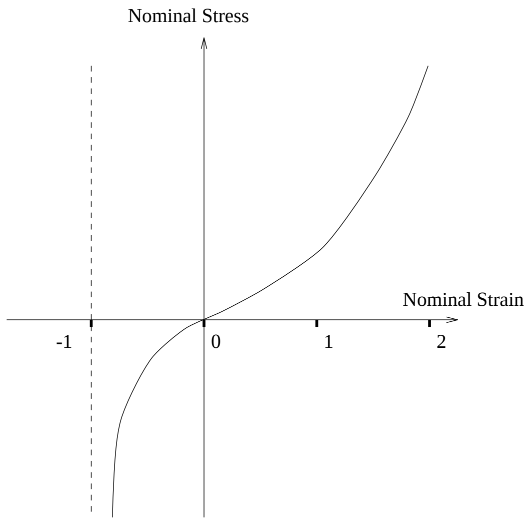

Figure 5: A typical stress-strain curve.

it by the sample were periodically computed. Figure 5 shows a typical stress-strain (nominal stress-nominal strain) curve that might be produced by a *real* sample of certain elastic materials. Initially, the stress-strain relationship is approximately linear, however as the material becomes stretched (or compressed) further, the forces begins to increase (or decrease) non-linearly. Due to the fact that nominal strain is being used, a strain of -1 represents total compression of the sample (which is not possible), and hence this line becomes an asymptote, with the curve tending towards it as it is compressed further.

The first simulated test to be carried out was on a sample of size $2cm * 2cm * 4cm$, composed of regular cubic elements with sides of length $1cm$. The cubic elements were composed of masses connected by damped springs with the following parameters: normalised spring constant (spring constant/length) $= 125.0 N/m^2$, damping factor $= 0.01$, particle mass $= 0.000036 Kg$ (corresponding to a material density of $100 Kg/m^3$). The model was composed of 45 masses and 296 damped springs. The dampers were included to prevent errors in the simulation leading to instability, however the tests were performed sufficiently slowly to prevent them having an effect on the measured stress. Figure 6 shows the stress-strain graph from the test. The graph shows the sample being stretched from rest to a positive strain of 1 (twice its

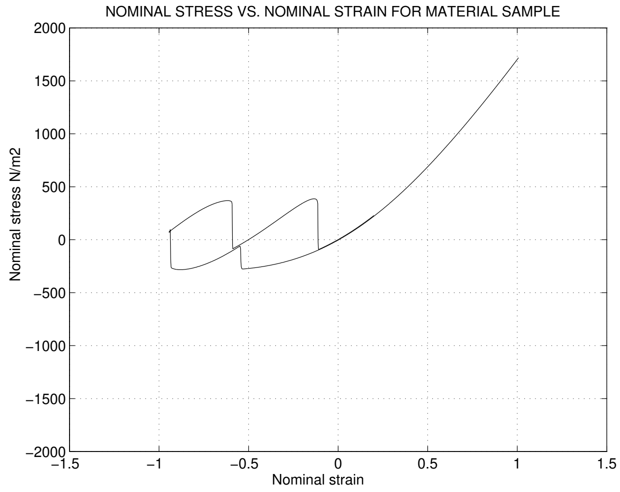

Figure 6: Stress-strain relationship for first sample.

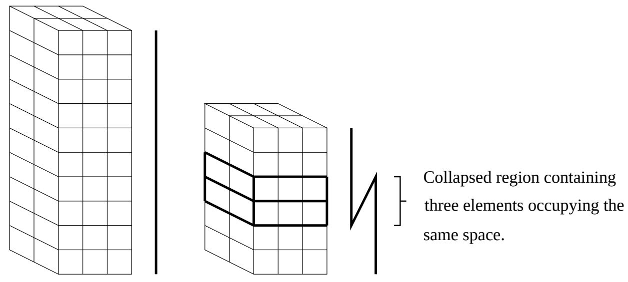

Figure 7: Collapse of the deformable material model.

rest length), squashed to a negative strain of approximately -0.8 (one fifth its rest length), and then brought back to 0 strain. When stretched, the sample can be seen to behave as expected. However, under compression, at approximately -0.6 nominal strain (40% original length) significant divergence from the "correct" behaviour occurs. Analysis shows that this occurs due to layers of elements within the structure folding internally, to adopt a less strained (lower potential energy) configuration. This is illustrated in figure 7. Here the elements are represented by cubes, and samples of the material are shown before (left) and after collapse (right). The line at the side of each block shows a one dimensional analogy of the configuration of the elements. In the collapsed state, the elements are folded once, and then back on themselves a second time. This results in elements merging such that three elements occupy the same space. This effect will be referred to as *internal self-intersection*. The sudden collapse along entire planes of the material can be explained by the fact that regions undergoing collapse exert extra strain on surrounding elements causing them to collapse also.

In figure 6 we see two instances of collapse along different planes during compression at strains of -0.6 and -0.8. This is accompanied by a sudden decrease in the compressive stress that has built up. As the strain is removed, the stress now become positive, and the sample literally has to be pulled back into its undeformed state. This implies that the collapsed state is stable, which makes the problem worse since elements experiencing collapse will not automatically

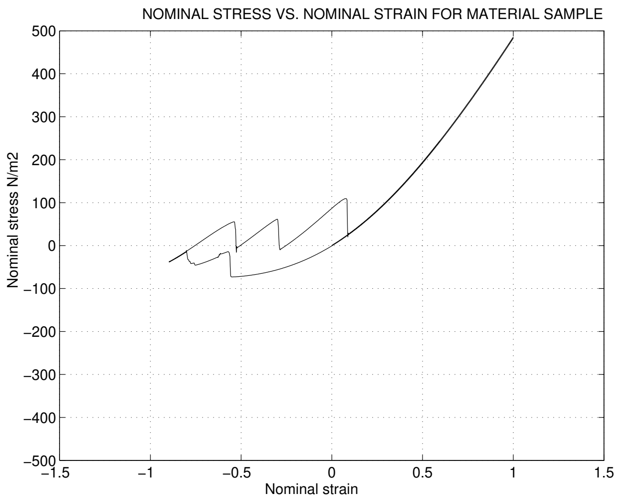

Figure 8: Stress-strain relationship for the second sample.

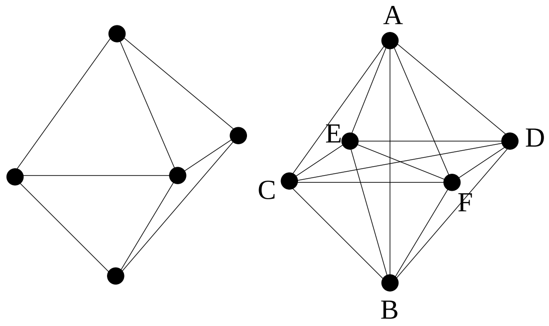

Figure 9: Octahedral element.

resume their correct state when the applied forces have been removed. The sudden drop in tensile stress at strains of -0.6 and -0.2 corresponds to the collapsed layers flipping back to their normal state.

In order to determine whether or not the robustness could be improved by varying the model parameters, the same test was performed with the element density doubled. This involved reducing the element size to $0.5cm$ cubes with the sample size kept constant. Figure 8 shows the new stress-strain curve. Here we see that again collapse occurs initially around a compressive strain of -0.6. This is then followed by a succession of sudden collapses as compression is increased, and also a succession of sudden de-collapses as the elements return to their normal state. The fact that there are fewer apparent collapses than de-collapses can be explained by the fact that collapse occurs simultaneously on multiple planes.

5 A new deformable material model

An alternative element design is presented which attempts to make the model more robust. This uses an irregular octahedral shaped element shown in figure 9. Table 1 compares this octahedral element with the commonly used square element, and shows that it has fewer masses and connections between masses. It can therefore be considered to be the

Number of:	masses	springs
Cube	8	12+12+4=28
Octahedron	6	8+4+3=15

Table 1: Comparing the cubic and octahedral elements.

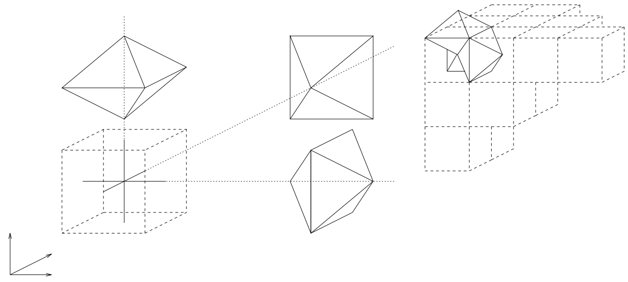

Figure 10: Tessellating the octahedral elements.

more "elementary" of the two. Figure 10 illustrates how the octahedral elements are tessellated in three dimensions. Three elements are brought together at a point such that each is aligned with one of the orthogonal axes. This then gives a shape which can be stacked like a cube to fill space.

5.1 Modelling incompressibility

A new method of modelling incompressibility within octahedral elements has been developed. The element has three cross springs AB, CD, and EF (see figure 9) which pass through the centre of the element, and have lengths given by l_{AB}, l_{CD}, l_{EF}. When the element is in equilibrium (and the cross springs are at right angles to one another) then the relative lengths are given by:

$$l^o_{CD} = l^o_{EF} = \sqrt{2} l^o_{AB}$$

The volume of an element in the rest state is given by:

$$V^o = \frac{(l^o_{AB} * l^o_{CD} * l^o_{EF})}{3}$$

which is deduced from the fact that three elements in equilibrium combine to produce a shape with the same volume as a cube with sides l^o_{AB}, l^o_{CD} and l^o_{EF}.

Now consider the element in a distorted state, but with the cross springs still lying orthogonal to one another. It is possible to suggest new lengths for each of the cross springs (denoted by superscript s) which if true would lead to the volume of the element being maintained at the original equilibrium volume. These are given by:

$$l^S_{AB} = \frac{3 * V^O}{(l_{CD} * l_{EF})}$$

$$l^S_{CD} = \frac{3 * V^O}{(l_{AB} * l_{EF})}$$

$$l^S_{EF} = \frac{3 * V^O}{(l_{AB} * l_{CD})}$$

The method used is to try and force the cross springs to their suggested lengths by adding extra springs (called volume springs) which act in parallel with the existing cross

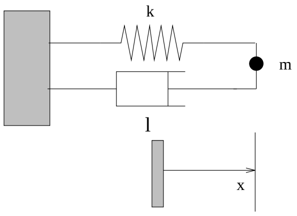

Figure 11: Simple mass-spring-damper arrangement.

springs. The equilibrium length of the volume springs is constantly changed to the suggested length, resulting in forces being applied whenever the volume of the element changes from the rest volume. Since the volume conservation is enforced by springs, then by altering the strength of the springs, we have some control over the degree of incompressibility of the elements. This method is extremely simple and requires little additional computation. However it is also crude mainly due to the fact that the volume calculation is only exact when the cross springs are orthogonal.

5.2 Creating independent material parameters

Varying the mass, spring and damper parameters directly does not give intuitive control over the behaviour of the material. A more independent set of parameters can be derived by considering the simple mass-spring-damper system shown in figure 11. The differential equation governing the behaviour of this system is given by:

$$\frac{m\ddot{x}}{k} + \frac{l\dot{x}}{k} + x = 0$$

Analysis shows that this can also be expressed using the alternative form:

$$\frac{\ddot{x}}{\omega_n^2} + \frac{2c\dot{x}}{\omega_n} + x = 0$$

where ω_n is the *undamped natural frequency* of oscillation of the system, and c is a parameter called the *damping factor*, which determines how quickly the system returns to rest after being disturbed. Equating coefficients in the above equations gives:

$$k = \omega_n^2 m$$

$$l = 2\omega_n mc$$

This change of independent variables from k and l to ω_n and c gives a more intuitive method of determining the behaviour of this simple system, and tests have shown that it can also be used to control the behaviour of blocks of material composed of the octahedral elements described above. The value of m used is the mass of each point mass within the mesh, given by:

$$m = \frac{\text{density} * \text{volume of block}}{\text{number of masses in block}}$$

5.3 Evaluating the model

A number of other element shapes have been used by other researchers, for example in [10] six irregular tetrahedrons are combined to form a cube, and in [16] tetrahedral and hexahedral elements are used. Full analysis of the octahedral element and these alternative designs is an area for future research.

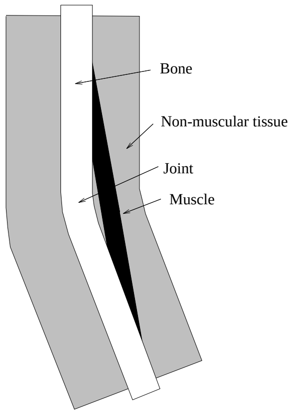

Figure 12: A cross section through the simplified model.

6 A simplified limb model

A model has been implemented which is a simplified version of the proposed model discussed earlier. The structure is shown in figure 12. The bone and joint model are based on the proposed model with the alteration that the joint is modelled as a smooth bent section between the two bones. This was adopted to try and keep compression around the inside of the joint to a minimum since it was feared that the octahedral elements may also be susceptible to collapse. A further simplification is to have only a single long thin muscle region which is connected to the bones either side of the joint. The non-muscular tissue regions and skin model are consistent with the proposed model.

6.1 Adding the deformable flesh

The method used to add the two deformable tissue regions to the bone and joint model involves the following stages, with figure 13 showing cross sections of the model during this process:

1. The bone and joint model is placed in position with a slight bend in the joint, and a large block of material is used to surround the bone (see fig. 13a).

2. The deformable tissue is attached to the bone using the following method (see fig. 13b):

   - Springs wholly inside the bone are discarded.
   - Springs intersecting the bone surface are labelled as *bone anchor springs*. The point of intersection with the surface becomes the new end of the spring, and this point is attached to the surface of the bone at a *bone anchor point*.

3. The muscle region is determined by selecting deformable elements lying within a cylinder of material. The orientation of the mesh is chosen such that entire rows of deformable elements can be selected as muscle elements (see fig. 13c).

4. A surface representing the outer surface of the limb is used to carve out the required shape, using the following method (see fig. 13d):

   - Deformable elements wholly outside the surface are discarded.

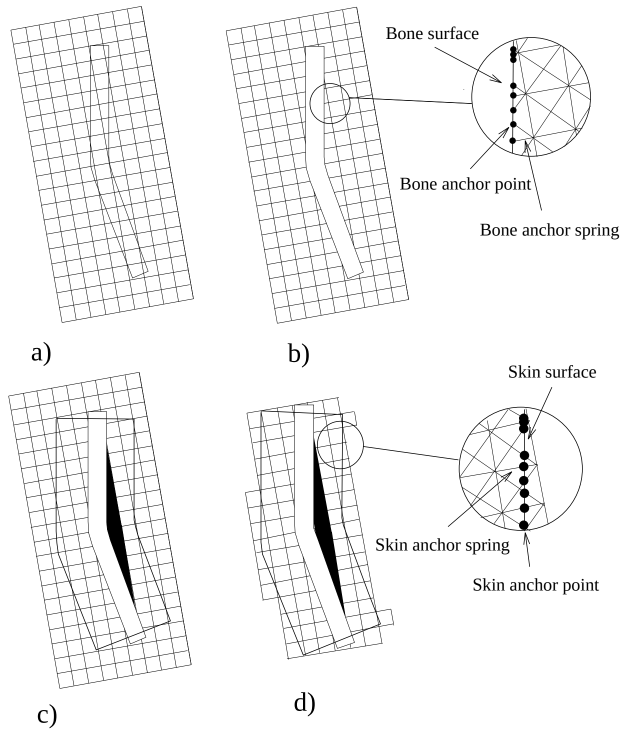

Figure 13: Stages involved in producing the simplified limb model.

- Deformable elements partially outside the surface are kept intact.
- Springs which intersect the surface are labelled *skin anchor springs*, and the point of intersection is recorded as a *skin anchor point*. The position of intersection is stored as a ratio of the distance along the spring divided by the spring length, so that during simulation, as the spring is stretched or compressed, the new position of the skin anchor point can be easily computed.

Figure 14 shows the deformable material mesh around the bone. Figure 15 shows the mesh and the surface used to carve out the desired limb shape. The particular type of surface geometry used to determine the limb shape is not restricted, and could include polygon meshes, parametric surface patches or implicit surfaces. The minimum requirement is that a test can be performed to determine whether a point is inside or outside the surface. This then enables us to test whether a spring is outside or partially outside the surface, and in the latter case to find the point of intersection using a binary search method along the spring.

6.2 The muscle elements

Contraction of the muscle elements is achieved by substituting an alternative set of equilibrium spring lengths corresponding to the elements being scaled by a certain factor in the required direction of contraction of the muscle. The new set of equilibrium spring lengths is switched with the original set during simulation to represent tensing of the muscle elements. All the other properties of the element are left unchanged, including the volume of the element, which

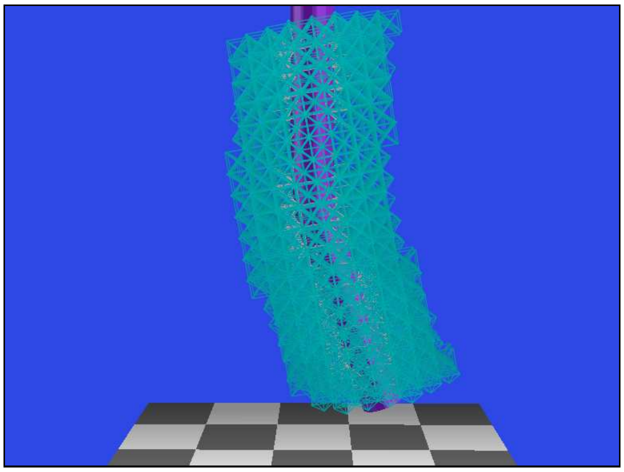

Figure 14: The deformable mesh attached to the bone.

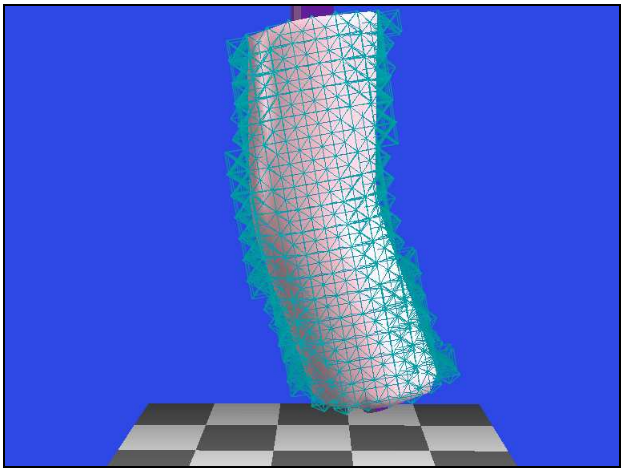

Figure 15: Limb model with deformable mesh still visible.

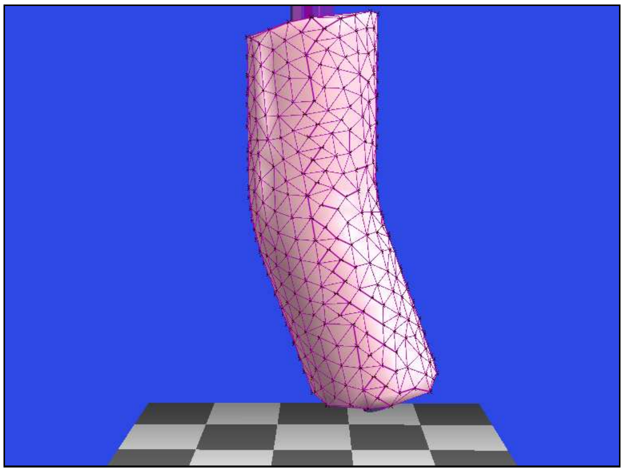

Figure 16: Triangulation of the skin surface.

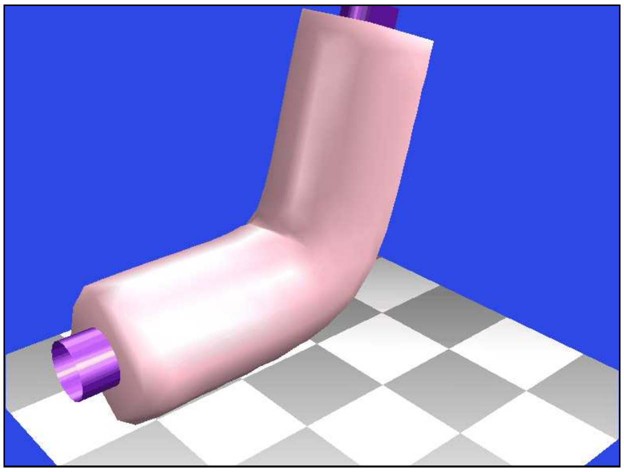

Figure 17: Limb model showing the onset of creasing.

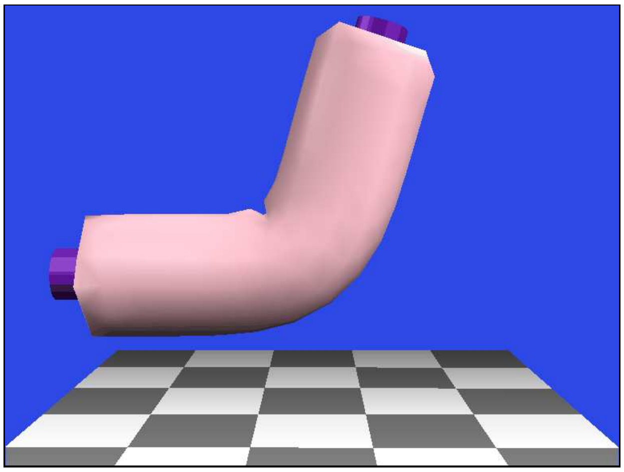

Figure 18: Creasing of the skin surface.

leads to bulging of the muscle perpendicular to the direction of contraction, adding to the realism. It is also possible to achieve intermediate levels of muscle activation by linearly interpolating between the sets of equilibrium lengths.

6.3 Creating the skin surface

A geometric skin model is produced by determining a triangular surface mesh based on the skin anchor points. In order to produce this the points are mapped along the normal to the surface onto a cylinder which surrounds the limb. The position of the points on the cylinder are then transformed into cylindrical coordinates producing a two dimensional set of points. These points are then triangulated using a 2D Delauney Triangulation algorithm [5]. Finally, the triangulation is applied to the original points to produce a surface mesh on the limb. The surface is rendered using standard Gouraud shading (other techniques could be used to give improved realism, e.g. bump mapping and texture mapping). Figure 16 shows the completed model and the triangulation of the skin surface. The final model is composed of 3570 octahedral elements which are in turn composed of 53550 springs and dampers, and 3132 masses. However, the authors believe that this could be significantly reduced by lowering the resolution of the mesh in areas of low deformation, and using other similar techniques.

6.4 Dynamic simulation of the model

Simulation of the model has been carried out so far using a basic Euler method with manual control over the size of the time steps. This is the simplest method to implement, and allows the basic behaviour of the model to be tested, however it is inefficient since it is prone to instability unless the time steps are kept sufficiently small. Improved methods are the midpoint method or a fourth order Runga-Kutta (see [1]). Using automatically varied time steps is also desirable otherwise the simulation has to be monitored to ensure it remains stable. Adaptive time step methods are included in [1] and [12]. Another feature of the current system is that the simulation stepsize is constant throughout the model. This is also inefficient since regions experiencing little deformation can be updated less often that regions of high deformation whilst still preserving stability.

7 Conclusion

The simplified limb model successfully reproduces a number of effects that occur in real limbs. Figures 17 and 18 show slight creasing on the inside of a bent joint, thinning of the flesh on the outside of the joint, and a small amount of bulging around this joint. This bulging can be made more prominent by increasing the incompressibility of the material. Muscle activation leads to motion of the limb, although very little surface bulging is seen. This may be partly due to the small size of the muscle region, and again the muscle material being too compressible. The model displays excellent inertial effects, and animation produced shows behaviour which looks very similar to real limbs which have a thick layer of fatty tissue surrounding the bone. A number of problems occured due to collapse of the material model. To prevent this occuring the limb model had to be prevented from bending too far. Bends larger than 90 degrees (from the straight position) resulted in the material on the inside of the joint (where compression is highest) collapsing, leading to strange effects. Further work is required to develop a more robust model.

Current work includes implementing the full proposed limb model. Future research is planned which involves making more anatomically correct models featuring accurate bone and joint models, and a greater number and range of tissue regions. Possibilities include the use of geometry data (possibly obtained via MRI or CAT scans) to produce anatomically accurate models of real limbs.

Acknowledgements

Thanks to the members of the Computer Graphics Unit, University of Manchester for their assistance during this research, particularly Martin Preston and Terry Hewitt. This research was funded by the EPSRC.

References

[1] *Siggraph '93 Course Notes: An Introduction to Physically based Modelling*, 1993.

[2] Michael F. Ashby and David R.H. Jones. *Engineering Materials 1*. Pergamon Press, 1980.

[3] Norman I. Badler, Dimitri Metaxas, Bonnie Webber, and Mark Steedman. The center for human modeling and simulation. Technical report, Department of Computer and Information Science, University of Pennsylvania, 1994.

[4] David Baraff and Andrew Witkin. Dynamic simulation of non-penetrating flexible bodies. In Edwin E. Catmull, editor, *Computer Graphics (SIGGRAPH '92 Proceedings)*, volume 26, pages 303–308, July 1992.

[5] C. Bradford Barber, David P. Dobkin, and Hannu Huhdanpaa. The quickhull algorithm. Technical Report Technical Report GCG53, The Geometry Center, University of Minnesota, Minneapolis.

[6] John E. Chadwick, David R. Haumann, and Richard E. Parent. Layered construction for deformable animated characters. In Jeffrey Lane, editor, *Computer Graphics (SIGGRAPH '89 Proceedings)*, volume 23, pages 243–252, July 1989.

[7] David T. Chen and David Zeltzer. Pump it up: Computer animation of a biomechanically based model of muscle using the finite element method. In Edwin E. Catmull, editor, *Computer Graphics (SIGGRAPH '92 Proceedings)*, volume 26, pages 89–98, July 1992.

[8] Marie-Paule Gascuel and Claude Puech. Dynamic Animation of Deformable Bodies. In S. Coquillart, W. Straßer, and P. Stucki, editors, *From Object Modelling to Advance Visual Communication*. Springer–Verlag, 1994.

[9] Jean-Paul Gourret, Nadia Magnenat Thalmann, and Daniel Thalmann. Simulation of object and human skin deformations in a grasping task. In Jeffrey Lane, editor, *Computer Graphics (SIGGRAPH '89 Proceedings)*, volume 23, pages 21–30, July 1989.

[10] Matthew Holton and Simon Alexander. Soft cellular modelling: A technique for the simulation of non-rigid materials. *Computer Graphics: Developments in Virtual Environments (Proceedings of Computer Graphics International '95*, 1995.

[11] Nadia Magnenat-Thalmann, Richard Laperriere, and Daniel Thalmann. Joint-dependent local deformations for hand animation and object grasping. In *Proceedings of Graphics Interface '88*, pages 26–33, June 1988.

[12] W.H. Press, B.P. Flannery, Saul A. Teukolsky, and William T. Vetterling. *Numerical Recipes in C*. Cambridge University Press, 1986.

[13] Demetri Terzopoulos and Kurt Fleischer. Deformable models. *The Visual Computer*, 4(6):306–331, December 1988.

[14] Demetri Terzopoulos and Kurt Fleischer. Modeling inelastic deformation: Viscoelasticity, plasticity, fracture. *Computer Graphics*, 22(4):269–278, August 1988.

[15] Demetri Terzopoulos, John Platt, Alan Barr, and Kurt Fleischer. Elastically deformable models. In Maureen C. Stone, editor, *Computer Graphics (SIGGRAPH '87 Proceedings)*, volume 21, pages 205–214, July 1987.

[16] Demetri Terzopoulos and Keith Waters. Techniques for realistic facial modelling and animation. In D. Thalmann and N.M. Thalmann, editors, *Computer Animation '91*. Springer-Verlag, 1991.

[17] Nadia Magnenat Thalmann and Daniel Thalmann. *Computer Animation Theory and Practice*. Springer-Verlag, 1990.

[18] Pascal Volino, Martin Courchesne, and Nadia Magnenat-Thalmann. Versatile and efficient techniques for simulating cloth and other deformable objects. *Computer Graphics Proceedings, SIGGRAPH'95*, 1995.

[19] Shin ya Miyazaki, Masao Ishiguro, Takami Yasuda, Shigeki Yokoi, and Jun ichiro Toriwaki. A study of virtual manipulation of elastic objects. *Computer Graphics: Developments in Virtual Environments (Proceedings of Comue Graphics International '95)*, 1995.