COMPUTER ANIMATION AND VIRTUAL WORLDS
Comp. Anim. Virtual Worlds 2010; **21**: 245–254
Published online 28 May 2010 in Wiley InterScience
(www.interscience.wiley.com) DOI: 10.1002/cav.371

Perceptually validated global/local deformations

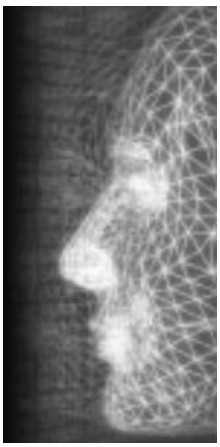

By Marcos García, Miguel Anguel Otaduy and Carol O'Sullivan*

Modal analysis techniques are often used to animate deformable objects in real time. In order to achieve performance improvements, the degrees of freedom of the problem may be reduced by using the main global deformations alone. However, this can lead to a reduction in quality and realism due to the lack of local behaviors. To solve this problem, we present a new method to add local deformations to modal analysis simulations. We perceptually evaluate our method with a set of experiments, thereby deriving guidelines for when and how local deformations can best be used. Copyright © 2010 John Wiley & Sons, Ltd.

KEY WORDS: local deformations; perceptual validation; physically based animation

Introduction

Nowadays, many computer graphics applications use physically based simulations to generate automatic animations. Rigid body dynamics have been extensively researched and many effective solutions have been proposed in this area. Nevertheless, the simulation of deformable models (e.g., elastic, plastic, fluid, and fabric objects) remains an open problem, with one of the main challenges being to find a technique that provides both good performance and high quality results. *Reduced coordinate* techniques,[1] such as modal analysis, provide for particularly good performance. However, while these methods capture an object's global behavior very well, they cannot handle local deformations. We propose a hybrid method that combines modal analysis with local deformations.

Our method can be considered as belonging to a set of techniques that try to add local deformations to fast global models. All of these models are based on the hypothesis that adding local deformation to a global model will increase the visual quality and appeal of an animation. While this is a plausible hypothesis, its validity has not been formally tested to date. Furthermore, it may be the case that sometimes the global model alone can provide results that are just as effective for a given situation. Our method can therefore be used to simulate only global deformations if required, thus allowing for level of detail (LOD) simulation. To investigate these issues and to evaluate our method, we ran a set of perceptual experiments. We found that the use of local deformations does indeed help to increase realism in certain situations, especially for soft objects, and can in fact ameliorate the effects of significantly reducing the quality of the global simulation. The local model can also help to make collisions look more realistic. Our results can be used to tune systems depicting deformable object simulations in order to achieve the optimal speed/quality balance.

*Correspondence to: M. García, Grupo de Modelado y Realidad Virtual (GMRV)—Rey Juan Carlos University, Despacho 0051, Ampliación de Rectorado, Calle Tulipán S/N, Móstoles, Madrid, Spain. E-mail: marcos.garcía@urjc.es

Previous Work

Since the use of physically based simulation to animate deformable objects was first proposed by Terzopoulos *et al.*,[2] many other methods have appeared. For a thorough overview, we refer the reader to Reference [3] and concentrate here on those methods most closely related to our own.

The finite element method (FEM)[4] is commonly used to animate elastic objects, due to the realism of the resulting simulation. However, such methods are generally very slow, so several simplifications have been proposed. For example, models based on the co-rotational formulation of FEM[5,6] rely on linear equations to describe object behavior.

From a performance point of view, reduced coordinates can deliver some of the best results. Reduced

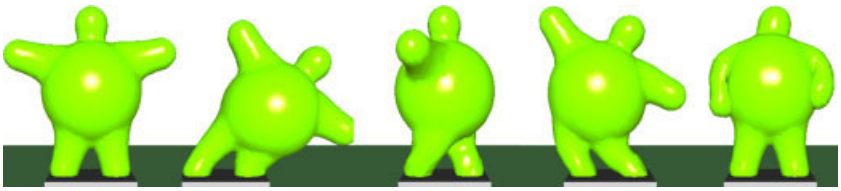

Figure 1. Object main modes. From left to right: rest position, 2nd, 3rd, 5th, and 7th mode.

coordinates represent displacements of the nodes $\mathbf{u}$ using another basis:

$$\mathbf{u} = \mathbf{\Phi q} \tag{1}$$

Each degree of freedom of $\mathbf{q}$ represents an object deformation (see Figure 1) and the dimension of $\mathbf{u}$ is much higher than the dimension of $\mathbf{q}$. One of the most commons ways of obtaining $\mathbf{\Phi}$ is using modal analysis.[1] One of the many advantages of such methods is that they can be easily accelerated using programmable graphics hardware.[7] Choi and Ko[8] handle the problems that arise when a linear strain tensor is used, by combining modal analysis and the co-rotational formulation, whereas Barbič and James[9] use a linear constitutive equation and a nonlinear strain tensor. In a later work,[10] they apply their technique to a haptic system.

As in this paper, combining global deformation methods with local deformations has been proposed to speed up simulations. Terzopoulos and Witkin[11] were amongst the first to separate global rigid motions from elastic deformations. A similar approach was stated by Galoppo *et al.*,[12] who later combined skeleton driven global deformation with elastic surface deformation.[13] Multiresolution methods also try to reduce the degrees of freedom of the problem, by using a coarse mesh to capture the global behavior and a finer mesh to model the local behavior.[14–16] Steinemann *et al.*[17] add a multiresolution scheme to shape-matching deformation method proposed by Müller *et al.*[18] and improved by Rivers and James.[19] Finally, we would like to point out that very few works treat influence of perception on physical based simulations. Among them, we want to point out the work.[20]

Model Description

As stated in the background section, FEM models are often favored when simulating elastic objects. These methods discretize the continuum mechanics equations and can easily model a variety of physical phenomena such as volume conservation, due to their physical basis. However, their main drawback is speed, so a common acceleration strategy is to use a linear model characterized by:

- the linear constitutive equation $\sigma = \mathbf{E}\varepsilon$, where σ is the stress tensor, $\mathbf{E}$ is the elasticity matrix and ε is the strain tensor, and
- the Cauchy strain tensor $\varepsilon = 1/2(\nabla\mathbf{u}^T + \nabla\mathbf{u})$

The Cauchy strain tensor is not invariant to rotations and the co-rotational formulation is used to solve this problem. The rotational component of each element's deformation gradient is extracted. Then, the forces in the un-rotated configuration are computed and subsequently rotated: $\mathbf{f_e} = \mathbf{R_e}\mathbf{K_e}\mathbf{R_e}^{-1}\mathbf{u_e}$, where $\mathbf{f_e}$ are the internal forces on the nodes of the element e, $\mathbf{R_e}$ is the rotation matrix of that element and $\mathbf{u_e}$ are the displacements of the element nodes.

The performance of linear FEM can be significantly increased using modal analysis, in which the vibration modes of an object are pre-computed (Figure 1). The lower frequency modes are used as a basis to represent all possible deformations of the model (see Equation (1)). Good results have been achieved using a reduced number of modes, less than 10 in some cases.[8,9]

As a consequence of using only the lower frequency modes, these methods capture the global movement of an object but fail to adequately simulate local behaviors. In this work we propose a new approach to add local deformations to modal analysis methods. As a starting point, we have used the model proposed in Reference [8], but our method could be adapted to use other modal analysis techniques, even nonlinear ones as Reference [9].

Model Overview

As mentioned above, our algorithm consists of two main steps: first, following collision detection, we compute global displacements using the global model and then select the elements to be used by the local model; finally, we compute the local displacements using the local

DOI: 10.1002/cav

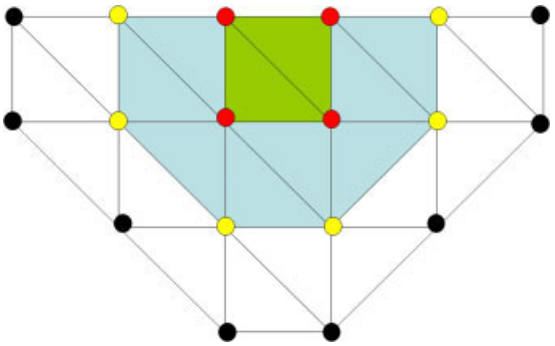

Figure 2. The adjacency level is defined by the designer. The green triangles are the e_c elements, the blue ones are the elements in the first level of adjacency, the red dots are the nodes $p \in L_p$ and the yellow dots are the nodes $p \in B_p$.

model, taking into account collision information and the previously computed global displacements.

Global Model

Starting from Choi and Ko's[8] model, we need to solve the dynamic equation, which can be written as:

$$\mathbf{M}\ddot{\mathbf{u}} + \mathbf{C}\dot{\mathbf{u}} + \mathbf{K}\mathbf{u} = \mathbf{F} \tag{2}$$

where $\mathbf{M}$ is the mass matrix, $\mathbf{C}$ is the damping matrix, $\mathbf{K}$ is the stiffness matrix, $\mathbf{u}$ is the displacement field and $\mathbf{F}$ is the vector that contains the forces applied over the mesh nodes. To compute the vibration modes, the generalized eigenvalue problem $\mathbf{K}\mathbf{\Phi} = \lambda\mathbf{M}\mathbf{\Phi}$ is solved at the pre-processing stage. The columns of $\mathbf{\Phi}$ contain the vibration modes and the elements in the diagonal matrix λ contain the associated frequencies. To reduce the degrees of freedom of the problem, only the lower frequency vibration modes are used. From now on we refer to the matrix containing the reduced set of eigenvectors as $\mathbf{\Phi}$ and to the diagonal matrix containing their associated eigenvalues as λ. Using Rayleigh damping ($\mathbf{C} = c_1\mathbf{M} + c_2\mathbf{K}$, where c_1 and c_2 are constants), $\mathbf{\Phi}$, λ and Equation (1), the dynamic equation can be rewritten as: $\ddot{\mathbf{q}} + c_1\dot{\mathbf{q}} + c_2\lambda\dot{\mathbf{q}} + \lambda\mathbf{q} = \mathbf{\Phi}^t\mathbf{F}$. It should be noted that there are significantly fewer equations in this system and they are independent.

The model presented above uses the Cauchy strain tensor, which means that rotations will cause unrealistic deformations. This problem can be resolved by computing the internal forces $\mathbf{F_{int}}$ in the un-rotated configuration and then rotating the forces again: $\mathbf{F_{int}} = \mathbf{R}\mathbf{K}\mathbf{u_l}$, where $\mathbf{u_l}$ is the linear displacement field and $\mathbf{R}$ its associated rotation. Now the displacement field and their derivatives can be written as:

$$\mathbf{u} = \int_0^t \mathbf{R}\dot{\mathbf{u}}_\mathbf{l} d\tau,\ \dot{\mathbf{u}} = \mathbf{R}\dot{\mathbf{u}}_\mathbf{l},\ \ddot{\mathbf{u}} = \mathbf{R}\ddot{\mathbf{u}}_\mathbf{l} + \dot{\mathbf{R}}\dot{\mathbf{u}}_\mathbf{l} \tag{3}$$

For simplicity the coriolis $\dot{\mathbf{R}}_\mathbf{l}\dot{\mathbf{u}}_\mathbf{l}$ effect term is neglected.

Instead of Equation (1) we use:

$$\mathbf{u_l} = \mathbf{\Phi}\mathbf{q_l} \tag{4}$$

Multiplying both sides of Equation (2) by $\mathbf{R}^{-1}$, Equations (3) and (4) can be used to rewrite the dynamic equation: $\ddot{\mathbf{q}}_\mathbf{l} + c_1\dot{\mathbf{q}}_\mathbf{l} + c_2\lambda\dot{\mathbf{q}}_\mathbf{l} + \lambda\mathbf{q}_\mathbf{l} = \mathbf{\Phi}^t\mathbf{R}^{-1}\mathbf{F}$. This equation is solved using an implicit integration scheme.

The rotation vector $\mathbf{w}(\mathbf{x})$ of each point $\mathbf{x}$ in the material can be computed using the curl of the displacement field $\mathbf{w}(\mathbf{x}) = \frac{1}{2}\nabla \times \mathbf{u}(\mathbf{x})$. The following matrix is used to compute the rotation of a material point $\mathbf{x}$ inside the element e, from the displacement of the element nodes: $\mathbf{W_e}(\mathbf{x})$, where $\mathbf{w}(\mathbf{x}) = \mathbf{W_e}(\mathbf{x})\mathbf{u_e}$. When using tetrahedral elements and barycentric coordinates as shape function, $\mathbf{W_e}$ is constant and can be computed as $\mathbf{W_e} = \frac{1}{2}\nabla \times \mathbf{H_e}$, where $\mathbf{H_e}(\mathbf{x})$ is the linear shape function of the element e. The rotations of each node are calculated using the average rotation of all nodes' elements. If all the $\mathbf{W_e}$ matrices are assembled in $\mathbf{W}$, which is constant over time, the rotations of all nodes $\mathbf{w}$ can be computed directly from the reduced coordinates, $\mathbf{w} = \mathbf{W}\mathbf{u} = \mathbf{W}\mathbf{\Phi}\mathbf{q} = \mathbf{\Psi}\mathbf{q}$.

Local Model

In this section we explain how the local deformations are computed and integrated with the global model. This local model simulates the type of behavior that is not captured by the global model. The idea is to solve most of the mesh nodes with the global technique (using the reduced coordinates basis) and to solve the nodes in contact using linear displacements.

In order to add the local behavior we mantein a set of local elements L_e. These elements will be solved using the traditional co-rotational formulation. After the collision detection stage, new elements e_c in contact are added to the local element set L_e. After the simulation, we remove from L_e those elements whose local solution is close to the global one. Some of the elements adjacent to e_c can also be added to L_e. The level of adjacency used is a design choice and depends on the mesh and the material properties (see Figure 2). Furthermore, our implementation allows the user to define elements that cannot be part of L_e, even if they are in direct contact with a scene object. This lets

DOI: 10.1002/cav

the user reduce the number of elements L_e can contain, selecting nodes where local deformations are not important, for example interior nodes. The boundary node set B_p consists of all nodes p which belong both to elements of L_e and elements of the complement of L_e, while the local node set L_p consists of nodes that belong only to elements of L_e. It should be noted that L_e may contain elements from previous simulation steps.

As stated above, the nodes $p \in L_p$ will be solved using the linear displacements rather than the reduced coordinates. The displacement vector $\mathbf{u_l}$ can be arranged in such a way that $\mathbf{u_l} = [\mathbf{u_L}, \mathbf{u_B}, \mathbf{u_Q}]$, where $\mathbf{u_L}$ are the displacements of the nodes $p \in L_p$, $\mathbf{u_B}$ are the displacements of the nodes $p \in B_p$ and $\mathbf{u_Q}$ are the displacements of the nodes $p \notin L_p$ and $p \notin B_p$. The values of $\mathbf{u_B}$ and $\mathbf{u_Q}$ are calculated using the global algorithm described in the previous section. Therefore, the equations that contain $\mathbf{u_Q}$ can be removed from the co-rotational dynamic equation system:

$$\mathbf{M}\ddot{\mathbf{u}}_\mathbf{l} + c_1\mathbf{M}\dot{\mathbf{u}}_\mathbf{l} + c_2\mathbf{K}\dot{\mathbf{u}}_\mathbf{l} + \mathbf{K}\mathbf{u}_\mathbf{l} = \mathbf{R}^{-1}\mathbf{F} \tag{5}$$

Now $\mathbf{M}$ and $\mathbf{K}$ for the local nodes in L_p can be written as:

$$\mathbf{M} = \begin{bmatrix} \mathbf{M_L} & \mathbf{M_{LB}} \\ \mathbf{M_{BL}} & \mathbf{M_B} \end{bmatrix},\ \mathbf{K} = \begin{bmatrix} \mathbf{K_L} & \mathbf{K_{LB}} \\ \mathbf{K_{BL}} & \mathbf{K_B} \end{bmatrix} \text{ and } \mathbf{F} = \begin{bmatrix} \mathbf{F_L} \\ \mathbf{F_B} \end{bmatrix} \tag{6}$$

The state of nodes $p \in B_p$ is known, so any equations relating to those nodes with forces $\mathbf{F_B}$ applied to them can be removed from Equation (5), which can be rewritten as:

$$\mathbf{F_C} = \mathbf{M_{LB}}\ddot{\mathbf{u}}_\mathbf{B} + c_1\mathbf{M_{LB}}\dot{\mathbf{u}}_\mathbf{B} + c_2\mathbf{K_{LB}}\dot{\mathbf{u}}_\mathbf{B} + \mathbf{K_{LB}}\mathbf{u_B} \tag{7}$$

$$\mathbf{M_L}\ddot{\mathbf{u}}_\mathbf{L} + c_1\mathbf{M_L}\dot{\mathbf{u}}_\mathbf{L} + c_2\mathbf{K_L}\dot{\mathbf{u}}_\mathbf{L} + \mathbf{K_L}\mathbf{u_L} + \mathbf{F_C} = \mathbf{R}^{-1}\mathbf{F_L} \tag{8}$$

where $\mathbf{F_C}$ is a constraint force, produced by the elements in B_p.

We apply the following implicit integration scheme to Equation (8):

$$\mathbf{u}(t) = \mathbf{u}(t-h) + h\dot{\mathbf{u}}(t) \tag{9}$$

$$\dot{\mathbf{u}}(t) = \dot{\mathbf{u}}(t-h) + h\ddot{\mathbf{u}}(t) \tag{10}$$

which leads to the following system:

$$\mathbf{S_L}\dot{\mathbf{u}}_\mathbf{L}(t) = \mathbf{M_L}\dot{\mathbf{u}}_\mathbf{L}(t-h) + h\mathbf{R}^{-1}\mathbf{F_L} - h\mathbf{K_L}\mathbf{u}(t-h) - h\mathbf{F_C}(t) \tag{11}$$

where

$$\mathbf{S_L} = \mathbf{M_L} + c_1h\mathbf{M_L} + c_2h\mathbf{K_L} + h^2\mathbf{K_L} \tag{12}$$

$$\mathbf{F_C}(t) = \mathbf{M_{LB}}\ddot{\mathbf{u}}_\mathbf{B}(t) + \mathbf{C_{LB}}\dot{\mathbf{u}}_\mathbf{B}(t) + \mathbf{K_{LB}}\mathbf{u_B}(t) \tag{13}$$

$$\mathbf{C_{LB}} = c_1\mathbf{M_{LB}} + c_2\mathbf{K_{LB}} \tag{14}$$

$$\ddot{\mathbf{u}}(t) = \frac{\dot{\mathbf{u}}(t) - \dot{\mathbf{u}}(t-h)}{h} \tag{15}$$

To solve Equation (11) we use the Jacobi preconditioned conjugate gradient method. Once $\mathbf{u_L}$ is obtained, any node $p \in L_p$ is erased if it satisfies: $\|\mathbf{u_p} - \mathbf{u'_p}\|^2 < \epsilon_u$ and $\|\dot{\mathbf{u}}_\mathbf{p} - \dot{\mathbf{u}}'_\mathbf{p}\|^2 < \epsilon_{\dot{u}}$ where $\mathbf{u_p}$ and $\dot{\mathbf{u}}_\mathbf{p}$ are the position and linear velocity respectively, as computed by the local method, $\mathbf{u'_p}$ and $\dot{\mathbf{u}}'_\mathbf{p}$ are those computed by the global method and ϵ_u and $\epsilon_{\dot{u}}$ are user defined constants. If all the nodes p of an element e are removed from the set L_p, that element e is also removed from the set L_e.

Perceptual Validation

We ran a series of perceptual experiments, which examined how the local deformations affects how much participants enjoy an animation or how well they perform a task. Three experiments were designed to test the *appeal* of a simulation, while a fourth tested whether local deformations aided in the perception of contacts.

Appeal: Active Task

In the first experiment, 15 naïve participants (6F/9M, aged 20–40) performed an active task and indicated their preferences between pairs of simulations. We hypothesized that the use of local deformation would improve the realism of the animation, especially in the more obvious cases. The simulations depicted two objects: a cannon and a deformable cylinder placed on a square base. The cylinder was simulated as an elastic object either using only its main three vibration modes, i.e., the *Global* model, or combining its three main vibration modes with local deformations, i.e., the *Global/Local* model. The cylinder was used because of its neutral shape, thus ensuring that the participants did not make assumptions about the material it was made of. Rigid spheres were shot toward

DOI: 10.1002/cav

the cylinder at random time intervals and in random directions

We allowed the user to control the cylinder, and the task was to prevent collisions between the spheres and the cylinder by grabbing the latter with the mouse pointer and pulling it. The number of hits was displayed at the bottom of the screen. For each trial, participants were asked to play two versions of this game (in randomized order), where the only difference between them was the deformation model used, i.e., Global or Global/Local. After playing each pair, the participants had to answer which of the two they found more realistic.

Three different scenarios were tested: near camera with a soft cylinder; near camera with a stiffer cylinder; and far camera with a soft cylinder. There were two repetitions of each scenario, so the participants played six different pairs of games in random order.

Results. We performed a three-way, within subjects analysis of the variance (ANOVA) on the data, where the factors were model and scenario. No statistically significant preference for the Global model or the Global/Local model was found. There was no clear preference between the different scenarios either. The lowest levels of appeal were recorded for the Global model in the near-camera soft-stiffness scene. For the stiff object, the Global/Local model was clearly preferred 66% of the time ($p < 0.05$).

It appears that *the addition of local deformations did not significantly affect the appeal of the simulation for the users when they were engaged in an active task*.

Appeal: Passive Task

In the second experiment, 16 naïve participants (7F/9M, aged 20–40) were asked to watch videos created from pre-recording a user performing the task described in Experiment 1. We hypothesized that in the passive test the effect of the local deformations would be higher because the participant was not immersed in performing a task and would therefore pay more attention to the quality of the simulation. The simulations and display were almost identical, except that the score was not displayed at the bottom of the screen. The videos were displayed as before in Global/Global–Local pairs (in randomized order), after which the participant indicated which one they found more realistic. Again, the three scenarios (near soft, near stiff, far soft) were repeated twice, so participants viewed six different pairs of videos in random order.

Results. We found that there was a main effect of model ($F(1, 16) = 19.105$, $p < 0.0005$), where the Global/Local model was preferred 72% of the time. As before, the scenario did not have a significant effect. From these results we can conclude that *the addition of local deformations significantly enhanced the appeal of the simulation when the user was passively viewing it*.

Appeal: Effect of Material Properties

The third experiment examined the effect of material properties on participants' preferences for the Global or Global/Local deformation models. Our hypothesis was that for the soft material, the Global/Local model would be preferred more often. We also tested the impact of reducing the number of modes used for global deformations. We hypothesized that knowing the material properties would make participants more inclined to prefer the Global/Local model for softer materials. We also predicted that local deformations would improve the realism of the object even when the number of modes was reduced.

Twenty-one naïve participants (8F/13M, aged 20–40) viewed sets of videos depicting a physically simulated anthropomorphic blobby object and were asked to sort them based on realism. This stimulus was chosen to give the impression of a flexible toy (see Figure 1). Again, multiple spheres were fired at the object from many different directions, while the impacts made it bend and deform. The camera viewpoint was rotated around the object following an elliptical trajectory.

Three factors were varied: the deformation model (Global or Global/Local as before); the object elasticity (soft, medium, stiff); and the quality of the global deformation (4 or 10 modes).

In this test, four videos were shown simultaneously on a large screen, placed at the four screen corners. The position of each video was selected randomly. Then, the participant was asked to sort them from the more realistic to the less realistic. The videos were rated by the participant, putting the randomly assigned name of each video on a list. Each participant viewed three sets of videos in random order and each set was animated with a different material, but the same material was used within each set and the user was told what the material properties were. At the beginning of the experiment all the materials were briefly described to the users. In each set, four modes were used for the global deformations in two videos and

DOI: 10.1002/cav

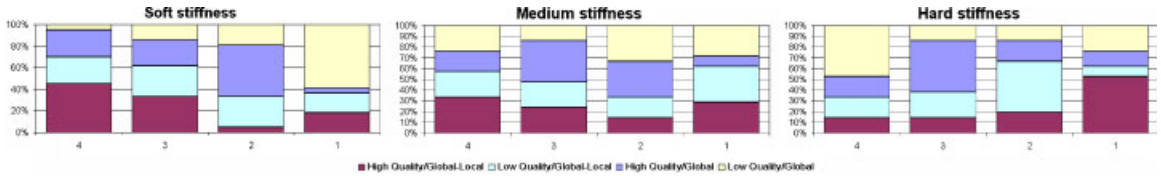

Figure 3. Results from the third experiment (material properties). The categories on the x-axis show the participants' preferences in terms of realism. 4 is the most realistic and 1 is the least realistic.

10 modes for the other two. Similarly, two videos were animated using the Global model and the others using the Global/Local method in such a way that all possible factor combinations were shown.

Results. It should first be noted that objects simulated using reduced coordinate methods tend to look stiffer, because most of the vibration modes are neglected. A log-linear analysis of frequencies showed that there were no main effects, but there were some significant interactions (see Figure 3). The material interacted with the quality of the global deformation ($\chi^2(6, N = 21) = 14.17$, $p < 0.05$), in that the softer the object, the higher the global quality that is needed to simulate it. There was also an interaction between the model used and the material ($\chi^2(6, N = 21) = 16.77$, $p < 0.05$), due to the Global/Local model being preferred for simulating soft objects. Finally there was an interaction between the model and the global deformation quality ($\chi^2(3, N = 21) = 20.63$, $p < 0.05$). We found that when local deformations are active, the difference between the two global qualities is smaller.

One conclusion from this experiment is that local deformations are particularly important when simulating soft objects. Particularly important is the result that *the number of vibration modes used can be reduced when using local deformations, even for softer objects, while still maintaining visual quality*.

Collision Perception

In the final experiment, we tested whether local deformations helped participants to detect contacts between the balls and the blobby object as shown in the previous experiment.

Fourteen naïve participants (6F/8M, aged 20–40) viewed videos depicting a similar scenario to the one used in the third experiment. The participants were asked to click the right mouse button each time they saw a collision. They were asked not to anticipate the collision based on the ball trajectory or the click would not count. To ensure that all contacts were visible, the balls were fired so that they never hit the back of the object. Furthermore, the camera was either static or moved back and forth along a trajectory that was restricted to a half-ellipse.

Four factors were tested: deformation model, object elasticity, quality of global deformation, and camera trajectory (static or moving). All 24 possible scenarios were shown twice, giving rise to 48 videos which were displayed in random order. We hypothesized that the use of local deformations can help to detect collisions, thus assisting the user to perform some tasks. The camera trajectory factor was added to see if the Global/Local model can be especially useful when ball trajectory cues cannot be used to detect collisions.

Results. The data were analyzed using a within-subjects four-factor ANOVA, with post-hoc analysis using the Newman–Keuls test for comparison of means. We found main effects for all factors except global quality. The participants found it easier to count the collisions when the object was stiffer ($F(2, 54) = 8, 12$, $p < 0.05$). The results for stiffer objects were significantly different ($p < 0.05$) from the other two elasticities. As expected, there was a main effect of the model used ($F(1, 27) = 59$, $p < 0.05$), where participants could detect 90% of the collisions when the Global/Local model was used and 84% otherwise. The camera trajectory affected the perception of the collisions ($F(1, 27) = 69.8$, $p < 0.05$), with the static camera allowing for higher accuracy, as hypothesized.

Elasticity interacted with camera trajectory ($F(2, 54) = 11.1$, $p < 0.05$) with participants' performance especially bad in the lower stiffness cases with a moving camera. There was also an interaction between elasticity and

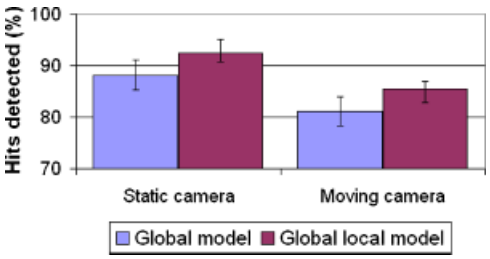

Figure 4. Interaction between the elasticity and scene.

DOI: 10.1002/cav

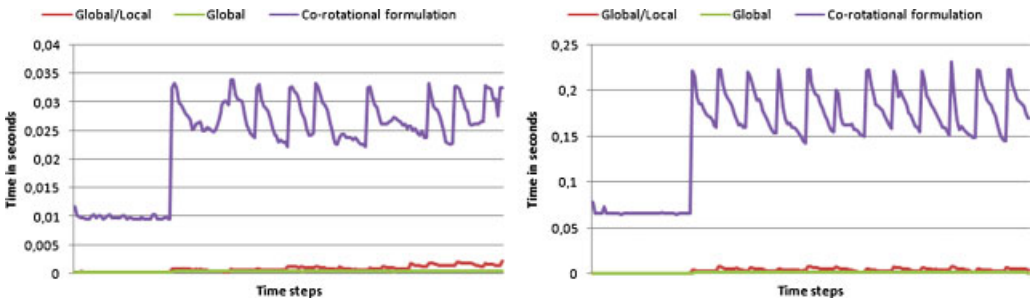

Figure 5. Performance comparison between the Global method, the CF (Stiffness Warping), and our model (Global/Local). The plot on the left shows the results for the Flubber model and the plot on the right shows the results for the high complexity cylinder.

model ($F(2, 54) = 3.4$, $p < 0.05$), as the Global model in the two softer elasticity cases performed significantly worse than in the other cases. We did not find any interaction between the model and the camera trajectory. The Global/Local model always obtained better results than the Global model: (Global/Static 88.1% of the hits, Global/Moving 81.2%, Global–Local/Static 92.5% of the hits and Global–Local/Moving 85.5%; see Figure 4).

Clearly, *adding local deformations increases the realism of contacts* and thereby helps the user to detect collisions.

Performance Evaluation

In this section, we analyze our model from the performance point of view. We have used three different models: a Flubber model (324 nodes, 970 elements), a low complexity Cylinder (253 nodes, 990 elements) and a high complexity Cylinder (861 nodes, 3600 elements). All models were simulated in the same scenario as the one used for the collision perception experiment. The models were placed at the center of the scene and several balls were fired at them. We compared our Global/Local model with the Global model[8] and the co-rotational formulation (CF).[6] In the CF setting, the state update was solved using the Conjugate Gradient method (CGM)[21] with a Jacobi preconditioner. We stop the CGM after 60 iterations or when an error threshold is reached. All our implementations run in a single thread and we have not used any GPU optimization in a Intel(R) Core(TM)2 QUAD Q6600 at 2.4 GHz.

Figure 5 shows the cost per simulation step of the three techniques. The plot on the left displays the results for the Flubber object and the plot on the right displays the results for the high quality Cylinder. The Flubber model was animated using 10 modes and one level of adjacency. The high quality Cylinder was simulated using 10 modes and two levels of adjacency. We have shown the results of both models to ilustrate that the behavior of the different techniques is similar despite of the model used. When the simulation starts there are no contacts and the model is in its rest position, therefore the CGM exits in just one iteration. In that case, the cost of the CF setting is dominated by building the system matrix and the computation of per-element rotations. Figure 5 demonstrates that both the Global method and our Global/Local method run orders of magnitude faster than the CF method.

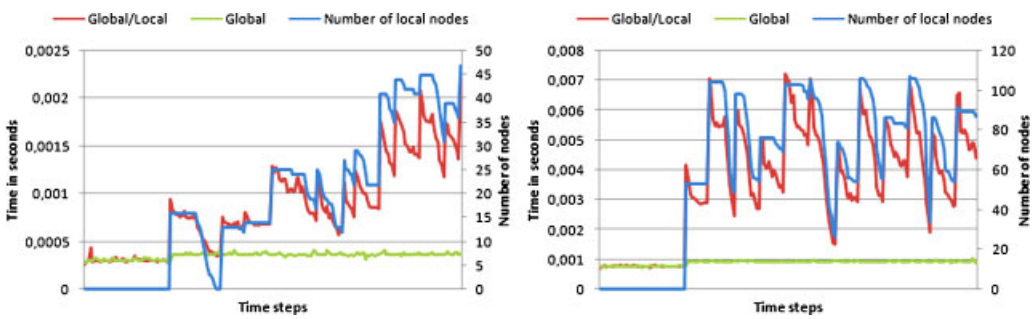

Figure 6. Performance comparison between the Global method, our model (Global/Local), and number of local nodes considered in the Global/Local technique. The plot on the left shows the results for the Flubber model and the plot on the right shows the results for the high complexity cylinder.

DOI: 10.1002/cav

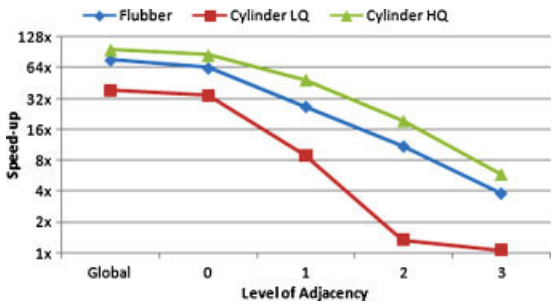

Figure 7. Speed-up of the Global model and Global/Local model (using different adjacency levels) with respect to the CF model.

Figure 6 compares the Global method with ours, using the Flubber model (left) and the high quality Cylinder model (right). The number of local elements used by the Global/Local model is also shown. In this figure, we can see that the cost of the Global/Local model is roughly linear in the number of active local nodes. At the beginning of the simulation, when no local elements are active, the Global/Local model and the Global have almost the same performance. In our test scenario the number of active local nodes never exceeds 20% of the model total nodes. It can also be pointed out that the behavior of both techniques is similar regardless of the model complexity. In our experiments we also found that computing the local node set took only a few microseconds, a negligible cost in comparison to the total solve time.

In Figure 7 we evaluate the impact of the number of adjacency levels on the speed-up gained with our Global/Local model. Specifically, the plots show the computational cost ratio between the CF setting and our Global/Local model or the full Global model. We use as reference technique the co-rotational formulation. The results are shown using a logarithmic scale. The data gathered were obtained from the Flubber model, low quality Cylinder and the high quality Cylinder. All the models were simulated using 10 modes. Can be pointed out that the plot lines fit in a logarithmic curve. From these plots, one could derive an optimal adjacency level for each particular model in a particular application. For example, the low quality Cylinder has the lowest number of elements, and using an adjacency level higher than one has no contribution in the performance because almost all its elements are added into L_e. Experiments using higher number of modes were also carried out and similar results were obtained.

Conclusions and Future Work

In this paper, we have described a technique to add local deformations to global models based on modal analysis. The results produced by our technique can be shown in Figure 8. Futhermore, the data analysis of our perceptual experiments provides guidance as to when it is most effective to use local deformations and when it is not so useful.

Our proposed approach allows for local deformations to be disabled if they are not necessary. In some circumstances it may not be worthwhile to use local deformations, for example, when processor units (CPUs or GPUs) are overloaded; when the user may not be paying attention to the collisions or interactions due to an absorbing task; or when simulating stiff objects. When local deformations are disabled, our technique is as fast as the modal analysis alone. Our approach is similar to many multiresolution methods but it has the added advantage of not having to compute and store other volumetric meshes.

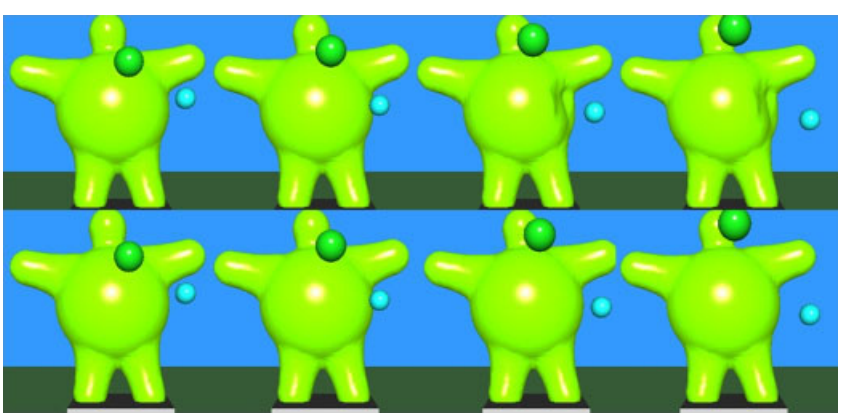

Figure 8. These images compare our Global/Local method (on the top) with the Modal Warping method (on the bottom).

Comp. Anim. Virtual Worlds 2010; **21**: 245–254
DOI: 10.1002/cav

From our experiments, we learned that local deformations did not have a significant effect when the user was performing an absorbing task. However, we were using a linear material, which means that the deformation is linearly related to the stress. This is not a common situation in nature. The forces that the participant applied to the object had a linear effect on the object deformation. This was especially noticeable when the user was pulling the soft object and the camera was near. This may have had a disturbing effect on the user. In the following experiments we tried to prevent these effects. No grabbing actions were permitted so the possibly unrealistic pulling deformations described above were avoided. From our second experiment, it is clear that the use of local deformations helps to increase realism for passive tasks, while our third experiment demonstrates that the use of local deformation is especially useful when animating soft objects. We have also found that, when using local deformations, the number of vibration modes used can be reduced even for the softer objects. Finally, we found that adding local deformations helps the user to detect collisions. We also discovered that softer elasticities make collision perception more difficult, especially when just global deformations are used. Our conclusion is that the large deformations can be disturbing for the user, thus distracting him from his task. However, the use of local deformation can also ameliorate this effect.

From the performance evaluation experiments we can conclude that using our Global/Local technique can drastically reduce the computational cost in comparison to a full-resolution solver based on co-rotational methods , with little overhead compared to a Global modal warping approach.

In this work we have presented some guidelines that can be used to determine when to apply local deformations. Our plan for future work is to find perceptual thresholds to guide the automatic activation and deactivation of our local deformation model. It will also be interesting to apply our technique to other model analysis techniques to solve some of the limitations of our global model.

ACKNOWLEDGMENTS

This work has been partially funded by the Trinity College Dublin, Science Foundation Ireland, URJC - Comunidad de Madrid (project CCG08-URJC/DPI-3647), and the Spanish Ministry of Science and Innovation (projects TIN2009-07942, PSE-300000-2009-5, and TIN2007-67188). Authors thank professor Angel Rodríguez for his support.

References

1. Pentland A, Williams J. Good vibrations: model dynamics for graphics and animation. In *SIGGRAPH '89: Proceedings of the 16th Annual Conference on Computer Graphics and Interactive Techniques*, ACM Press, New York, NY, 1989; 215–222.
2. Terzopoulos D, Platt J, Barr A, Fleischer K. Elastically deformable models. In *SIGGRAPH '87: Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques*, ACM, New York, NY, 1987: 205–214.
3. Nealen A, Müller M, Keiser R, Boxerman E, Carlson M. Physically based deformable models in computer graphics. *Computer Graphics Forum* 2006; 25(4): 809–836.
4. Bathe L-J. *Finite Element Procedures*. Prentice Hall: NJ, USA, 1996.
5. Etzmuß O, Keckeisen M, Strasser W. A fast finite element solution for cloth modelling. In *PG '03: Proceedings of the 11th Pacific Conference on Computer Graphics and Applications*, IEEE Computer Society, Washington, DC, 2003; 244.
6. Müller M, Gross M. Interactive virtual materials. In *GI '04: Proceedings of the 2004 Conference on Graphics Interface*, Canadian Human-Computer Communications Society, School of Computer Science, University of Waterloo, Waterloo, Ontario, Canada, 2004; 239–246.
7. James DL, Pai DK. DyRT: dynamic response textures for real time deformation simulation with graphics hardware. In *SIGGRAPH '02: Proceedings of the 29th Annual Conference on Computer Graphics and Interactive Techniques*, ACM, New York, NY, 2002; 582–585.
8. Choi MG, Ko H-S. Modal warping: real-time simulation of large rotational deformation and manipulation. *IEEE Transactions on Visualization and Computer Graphics* 2005; **11**(1): 91–101.
9. Barbič J, James DL. Real-time subspace integration for St. Venant-Kirchhoff deformable models. *ACM Transactions on Graphics (SIGGRAPH 2005)* 2005; **24**(3): 982–990.
10. Barbič J, James DL. Six-DoF haptic rendering of contact between geometrically complex reduced deformable models. *IEEE Transactions on Haptics* 2008; **1**(1): 39–52.
11. Terzopoulos D, Witkin A. Physically based models with rigid and deformable components. *IEEE Computer Graphics and Applications* 1988; **8**(6): 41–51.
12. Galoppo N, Otaduy MA, Mecklenburg P, Gross M, Lin MC. Fast simulation of deformable models in contact using dynamic deformation textures. In *SCA '06: Proceedings of the 2006 ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 2006; 73–82.
13. Galoppo N, Otaduy MA, Tekin S, Gross MH, Lin MC. Soft articulated characters with fast contact handling. *Computer Graphics Forum* 2007; **26**(3): 243–253.
14. Debunne G, Desbrun M, Cani M-P, Barr AH. Dynamic real-time deformations using space & time adaptive sampling. In *SIGGRAPH '01: Proceedings of the 28th Annual Conference on Computer Graphics and Interactive Techniques*, ACM, New York, NY, 2001; 31–36.
15. Capell S, Green S, Curless B, Duchamp T, Popović Z. A multiresolution framework for dynamic deformations. In *SCA '02: Proceedings of the 2002 ACM SIGGRAPH/Eurographics*

DOI: 10.1002/cav

Symposium on Computer Animation, ACM, New York, NY, 2002; 41–47.
16. Grinspun E, Krysl P, Schröder P. Charms: a simple framework for adaptive simulation. *ACM Transactions on Graphics* 2002; **21**(3): 281–290.
17. Steinemann D, Otaduy MA, Gross M. Fast adaptive shape matching deformations. In *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, Jul 2008; 87–94.
18. Müller M, Heidelberger B, Teschner M, Gross M. Meshless deformations based on shape matching. In *SIGGRAPH '05: ACM SIGGRAPH 2005 Papers*, ACM Press, New York, NY, 2005; 471–478.
19. Rivers AR, James DL. Fastlsm: fast lattice shape matching for robust real-time deformation. *ACM Transactions on Graphics* 2007; **26**(3): 82.
20. Theetten A, Grisoni L, Duriez C, Merlhiot X. Quasi-dynamic splines. In *SPM '07: Proceedings of the 2007 ACM Symposium on Solid and Physical Modeling*, ACM, New York, NY, 2007; 409–414.
21. Krizek M, Neittaanmäki P, Glowinski R, Korotov S. *Conjugate Gradient Algorithms and Finite Element Methods*. Springer: Berlin, 2004.

Authors' biographies:

Marcos García received his degree in Computer Science and Engineering in 2003 and his Ph.D. in Computer Graphics in 2007, both from the Universidad Politécnica de Madrid. From December 2008 to February 2009, he works as research fellow at the Trinity College Dublin. Currently, he is an assistant professor in the University Rey Juan Carlos (Madrid, Spain). His Ph.D. was centered on the task of deformable object simulation and has published works in the fields of computer graphics and virtual reality.

Miguel A. Otaduy is an assistant professor at Universidad Rey Juan Carlos (URJC Madrid), where he works at the Modeling and Virtual Reality Group (GMRV), in the Department of Computer Science. He received his B.S. (2000) in Electrical Engineering from Mondragon Unibertsitatea (Spain), and his M.S. (2003) and Ph.D. (2004) in Computer Science from the University of North Carolina at Chapel Hill. He completed his Ph.D. thesis in the field of haptic rendering, and supported by fellowships from the Government of the Basque Country and the UNC Computer Science Alumni. Between 1995 and 2000, he was a research assistant at Ikerlan research lab, and between August 2000 and December 2004 he was a research assistant with the Gamma group at UNC. In the Summer of 2003 he worked at Immersion Medical. From February 2005 to February 2008, he worked as a research associate at the Computer Graphics Laboratory of ETH Zurich. Miguel is at URJC Madrid since February 2008. His main research areas are physically-based simulation, haptic rendering, collision detection, virtual reality, and geometric algorithms, and he is particularly interested in the simulation and interaction with virtual objects in contact, with application to virtual prototyping, computational medicine, animation, or videogames

Carol O'Sullivan leads the Graphics, Vision and Visualization group (GV2) in Trinity College Dublin. After receiving a B.A. in Mathematics from Trinity College in 1988, she worked for several years as a software engineer in industry (mainly in Germany), followed by a Masters degree from Dublin City University in 1996 and a Ph.D. in Computer Graphics from TCD in 1999. Her research interests include perception, animation, virtual humans, and crowds. She was elected a Fellow of Trinity College Dublin and of Eurographics in 2003 and 2007, respectively. She has been a member of many IPCs, including the Eurographics and SIGGRAPH papers committees, and has published over 100 peer-reviewed papers. She has organized and co-chaired several conferences and workshops, including Eurographics'05 in Dublin, the SIGGRAPH/EG Symposium on Computer Animation 2006, and the SIGGRAPH/EG Campfire on Perceptually Adaptive Graphics 2001. She is the programme co-chair of the SIGGRAPH Symposium on Applied Perception in Graphics and Visualization 2009, co-editor in chief of ACM Transactions on Applied Perception and an editorial board member of IEEE Computer Graphics & Applications and Graphical Models.

DOI: 10.1002/cav