

Elasticity-based Clustering for Haptic Interaction with Heterogeneous Deformable Objects

Benoît Le Gouis¹ Maud Marchal¹ Anatole Lécuyer² Bruno Arnaldi¹

¹INSA Rennes/IRISA, Campus de Beaulieu, Rennes, France

²Inria Rennes, France




Figure 1: Overview of our elasticity-based clustering approach: (left) Based on an heterogeneous object composed of elements with different elasticity values, (center) we propose to build elasticity clusters to improve the computation time performances, (right) thus allowing haptic interaction while keeping similar perceptual sensations.

Abstract

Physically-based simulation of heterogeneous objects remains computationally-demanding for many applications, especially when involving haptic interaction with virtual environments. In this paper, we introduce a novel multiresolution approach for haptic interaction with heterogeneous deformable objects. Our method called "Elasticity-based Clustering" is based on the clustering and aggregation of elasticity inside an object, in order to create large homogeneous volumes preserving important features of the initial distribution. The design of such large and homogeneous volumes improves the attribution of elasticity to the elements of the coarser geometry. We could successfully implement and test our approach within a complete and real-time haptic interaction pipeline compatible with consumer-grade haptic devices. We evaluated the performance of our approach on a large set of elasticity configurations using a perception-based quality criterion. Our results show that for 90% of studied cases our method can achieve a 6 times speedup in the simulation time with no theoretical perceptual difference.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Virtual reality I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Physically based modeling

1. Introduction

As almost all materials of our daily life are non-homogeneous, the ability of interacting with such materials is crucial when designing physically-based simulations of virtual environments. The simulation of non-homogeneous objects remains however computationally demanding and often leads to simplification in the material simulation. Approaches such as embedded models [NKJF09] or

frame-based methods [GFBP13] for instance have been proposed to address the heterogeneity of the object.

If we aim at simulations allowing haptic interaction, the computation time performances become a bottleneck that prevent from simulating complex heterogeneous objects that could be used in many potential applications such as medical gestures or virtual prototyping. Haptic rendering of complex objects requires therefore a trade-off between accuracy and interactivity performances. Within

this context, methods relying on a geometric multiresolution of the object have been for instance proposed in the literature [KS04]. However, there is currently very little work performed on the haptic perception of object heterogeneity [TOT13], and multiresolution methods seldom take the heterogeneity of the object into account.

In order to simulate a heterogeneous deformable object, the Finite Element Method (FEM) could be used to provide accurate results. It relies on the decomposition of the object into elements, often tetrahedra. The properties of an object, such as elasticity, are embedded in the stiffness matrix of the object. As explained by [KMOD09], changing the resolution of the object implies to change the stiffness matrix, and attributing an elasticity to each element is not trivial for heterogeneous objects.

In this paper, we propose a novel approach that facilitates the computation of elasticity parameters of multiresolution methods in the context of haptic interaction with a heterogeneous deformable object, as illustrated in Figure 1. Our method creates clustered volumes with homogeneous elasticity inside the object, while focusing on the distribution pattern of elasticity inside the object and its effect on haptic perception. Our method coarsens the heterogeneity of our object, while keeping important features of the elasticity of the object. In particular, the overall elasticity of the simulated object remains similar. The method is evaluated using a perception-based quality criterion of the obtained distribution.

In the remainder of this paper, we first present the related work on physically-based simulation of heterogeneous objects and multiresolution methods for haptic interaction (Section 2). The overall pipeline of our method with its four main phases is then presented in Section 3. The evaluation protocol of the method using a perception-based quality criterion and the corresponding perceptual results and performance gains are shown in Section 4. A use case illustrating our method, as well as the haptic rendering scheme are then presented in Section 5. The paper concludes with a discussion on the method, and possible future work.

2. Related Work

In this section we first present the related work on the physically-based simulation of heterogeneous deformable objects. Level-of-Details (LoD) methods for haptic interaction are then presented.

Physically-based simulation of heterogeneous deformable objects

Many physically-based approaches propose nowadays the simulation of deformable objects with various applications in different domains such as medicine, computer animation or virtual prototyping. An extensive review of most main simulation methods can be found in [NMK*06]. The simulation of various deformable objects such as the elasticity heterogeneity remains computationally-demanding and only a few methods allow the interactive simulation of such objects. Thus, the use of Finite Element Method (FEM) and Frame-based method have both been proposed to deal with volumetric elasticity heterogeneity. For instance, [NKJF09] embedded the object into a regular simulated grid, with the deformation of the object being an interpolation of the deformation of the key frames in the grid. A superposition of different boxes inside the

grid allows the heterogeneity. However, as such, this method relies on a regular grid, not suited for multiresolution methods, and it does not discuss haptics. Another frame-based method for heterogeneous objects has been proposed by [GFBP13], with the deformation of the object being interpolated based on the deformation of key frames in the object. A specific selection of interpolation functions allows for heterogeneity, but haptic interaction or multiresolution features are not discussed. Tagawa et al. [TOT13] presented a FEM-based method for the simulation of heterogeneous objects, with the embedding of heterogeneity properties in the stiffness matrix. This method allows for haptic interaction, and multiresolution, however multiresolution was not based on the heterogeneity of the object, and possible perceptual loss induced by such a simplification was not investigated. Other FEM-based methods, such as proposed in [KMOD09], [IYCN13] and [CLSM15], compute the elasticity of coarse elements based on the deformation of the object at fine and coarse resolution. While [KMOD09] provide with a homogeneous coarse mesh, [IYCN13] create a heterogeneous adaptive mesh and [CLSM15] create a heterogeneous coarse mesh. However neither of these methods involve haptic, and the effect of such heterogeneity on the perceived force was therefore not investigated.

Level-of-Details methods for haptic interaction

Level-of-Details methods have been widely used in computer graphics, especially for deformable objects [MWN*16]. For haptic interaction, multiresolution methods can act on different aspects of the simulation. Most methods rely on a geometric multiresolution approach, such as [PDZ05], [KS04] and [TYT12]. The geometry of the object is either locally refined in order to ensure a more precise interaction, or globally coarsened in order to achieve better performances. Other methods, such as [OL05], simplify the collision model for haptic interaction with a multiresolution hierarchy of collision models. Talvas et al. [TMD*15] proposed to aggregate the constraints applied on the object to allow for interactive simulations. Some methods also act on the simulation rate of the object, depending on the stiffness of the objects, stiffer objects requiring a more important simulation rate [PDC11].

While this method focuses on an object property (the object stiffness) to modify the accuracy of the simulation, other criteria have also been investigated to drive the multiresolution. For instance, maintaining the quality of tetrahedra from the deformed object provides a criterion deciding when to change resolution [TYT12]: tetrahedra are refined when they undergo a too large deformation. The use of a perception criterion has also been proposed in the literature as a simplification criterion. The simulation is performed on the coarsest model possible with no perceptual loss in the interaction. This approach, used by [OL05] and [PDZ05], studies the perception that the users have from the simulation, and derive a perception-based quality criterion in order to have the best possible performances without the user noticing that the simulation has been simplified. However to the best of our knowledge, no method studied the perception of the heterogeneity of the object as a criterion for simplification.

3. Our approach: Elasticity-Based Clustering for Haptic Interaction with Heterogeneous Objects

3.1. Motivations

As explained in [IYCN13], most methods using multiresolution on heterogeneous objects make the assumption that the object is separated into homogeneous domains. However, objects with complex elasticity distribution cannot be as easily separated into clear homogeneous domains. Such objects require methods that take into account the heterogeneity of the object, whatever the initial distribution of elasticity is within the object. The coarsening of the elasticity of an object should be carefully handled as it could greatly influence the haptic rendering. In this paper, we propose a method that aims at clustering the elasticity of an object while keeping similar haptic perception for the user. This method is particularly well suited for coarse interaction with the deformable object, such as probing with a large tool, or having two deformable objects pressed against each other. Our method is based on the use of co-rotational FEM [MG04] and allows the setting of elasticity values of a coarse mesh based on the elasticity distribution of a finer mesh, as illustrated in Figure 2.




Figure 2: Simulation pipeline with elasticity-based clustering: our novel method (in grey) simplifies the elasticity distribution inside a fine mesh in order to compute the elasticity inside the corresponding coarse mesh.

3.2. Our Algorithm: Elasticity-based Clustering

We introduce a novel approach called "Elasticity-based Clustering". The main principle is to modify the elasticity distribution of an object in a way that creates large homogeneous domains inside the object, while preserving the resulting elastic feedback force, as illustrated in Figure 3. Our method comprises 4 main phases: (1) Initialization, (2) Homogenization, (3) Aggregation and (4) Geometric Coarsening. In the Initialization phase, the goal is to regroup tetrahedra of similar elasticity, in order to create important homogeneous regions inside the object, for an easier geometry coarsening. The first step consists in defining similar elasticities through elasticity binning. Tetrahedra with similar elasticity are then clustered. The Homogenization phase attributes an average elasticity value to elements of similar elasticities, reducing the number of different elasticities inside the object. The Aggregation phase modifies the elasticity of the tetrahedra, in order to reflect the clusters, creating more important homogeneous volumes. Last, the Geometric Coarsening phase attributes elasticity values to the tetrahedra from the coarse mesh.

Initialization

The input of our method is a volumetric tetrahedral mesh representing the object, as illustrated in Figure 3, Cube 1. This object is heterogeneous, which means that to each tetrahedron t_i is associated a local Young Modulus $E(t_i)$. The goal of the method is to reorganize the elasticity inside the object in order to simplify its structure and to create large homogeneous volumes inside the object for more efficient simulation. The first step of this simplification is to reduce the number of different elasticities inside the object, by regrouping tetrahedra of similar elasticity together, for a later homogenization of similar tetrahedra. This process is similar to the data analysis method called *binning*, in which data is partitioned into a relative smaller number of categories based on similarity values, for easier analysis. We here adapt this method for the elasticity value of the tetrahedra. To do so, the elasticity range of the tetrahedra ($E_{max} - E_{min}$) is divided into *numberElasticity* elasticity ranges (later referred as *elasticity bins*) of same size $\frac{E_{max} - E_{min}}{\text{numberElasticity}}$. While *numberElasticity* = 1 represents the most extreme possible simplification, i.e. completely homogenizing the object, a too large number does not simplify the object enough, and a compromise between efficiency and accuracy must be found. To each tetrahedron is then associated an elasticity bin. It is important to note that, even if the elasticity ranges are of same size, the number of tetrahedra is not necessarily similar in each elasticity bin. Our method does not depend on a specific elasticity attribution method. Other methods such as an equal distribution of tetrahedra inside each elasticity bin could be further investigated. The chosen method however enforces an elasticity similarity between elements from the same elasticity bin.

The second step for mesh simplification is to regroup tetrahedra of similar elasticity together, in order to create important homogeneous volumes inside the object. To do so, the tetrahedra inside each elasticity bin are clustered using a spatial criterion, with standard *k-means* clustering using an euclidean distance. After the initialization, each tetrahedron has then an elasticity bin and belongs to a cluster of close tetrahedra within its elasticity bin. The result of this step can be seen in Figure 3, Cube 2.

This phase requires important computation time, up to a few seconds with an important number of elasticities and clusters on a large mesh, but is performed entirely offline. It is also important to note that in this phase, no modification to the object is performed. The object is only given more information required for an online execution of the elasticity simplification. Actual modifications to the elasticity are described in the two further phases of the algorithm.

Homogenization

In this phase, each tetrahedron is in a bin of elements with similar elasticity. The first simplification is to homogenize these tetrahedra, by giving them all the same mean elasticity. The mean elasticity is computed as follow, for each elasticity bin EB : $E_{mean} = \frac{\sum_{t \in EB} V_t * E(t)}{\sum_{t \in EB} V_t}$. This value for the elasticity corresponds to the resulting elasticity value for two beams in parallel and is similar to the computation performed in [NPF06]. While [KMOD09] have shown that this method is not accurate for the homogenization of elements with important stiffness difference, this averaging is performed with elements of close elasticity, and gives a force-wise

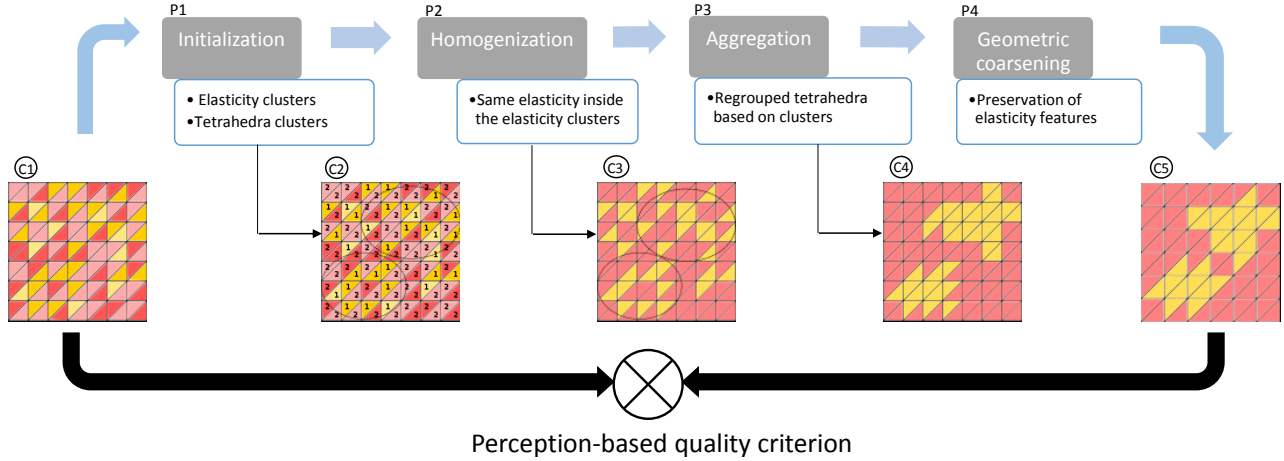


Figure 3: Schematic diagram of our elasticity-based clustering approach. It is composed of 4 phases: 1) Initialization: creates elasticity bins and tetrahedra clusters inside these elasticity bins. 2) Homogenization: unifies the elasticity inside the elasticity bins. 3) Aggregation: modifies the elasticity inside the object according to the tetrahedra clusters. 4) Geometric coarsening: attributes elasticity to the coarse mesh. Cube 1 (C1) represents the initial distribution of the elasticity. Red and yellow elements represent respectively stiff and soft elements, and the different tints for each color represent close elasticity values. Cube 2 (C2) represents the cube after the initialization phase, with information added to the tetrahedra: all tetrahedra are labelled with their elasticity bins (1 or 2); and barycentric clusters have been computed for all elasticity bins except one (here two clusters for the yellow elements). After the homogenization phase, in Cube 3 (C3), all tints for each color have been modified to the homogenized elasticity value for each elasticity bin. Tetrahedra are then regrouped for each cluster around the cluster barycenter in Cube 4 (C4), creating two large homogeneous volumes. A coarser mesh is finally used for Cube 5, with fewer elements, and the elasticities are chosen to match at best the elasticity from the fine resolution.

proper approximation. Figure 3, Cube 3 shows the results of this phase.

Aggregation

In this phase, the number of different elasticities taken into account is already reduced, from $numberTetrahedra$ to $numberElasticity$. The next step of simplification for the object is then to regroup tetrahedra with same elasticity. The tetrahedra have already been regrouped inside clusters during the Initialization phase. The only remaining step is thus to aggregate them together, in order to create important homogeneous volumes inside the object. To do so, the tetrahedra inside the object are not actually displaced, but the elasticity of each tetrahedron is modified to build the clusters. This process is only a displacement of elasticity values and does not modify the geometry of the object, which would be too costly and would not be meaningful in our context. It is also important to note that the overall volume for each elasticity remains the same, in order to preserve an overall equal elasticity.

The process is described in Algorithm 1: each elasticity bin is aggregated in turn, meaning that clusters are aggregated inside each elasticity bin. In order to keep the overall elasticity similar, the volume of modified tetrahedra for each elasticity E_{bin} must be equal to the volume of tetrahedra with elasticity E_{bin} before the aggregation process, as depicted in Figure 4. The volume of modified tetrahedra is recorded in order to enforce this condition.

In order to preserve notable features of the original distribution of elasticity, two properties need to be considered. First, the barycenter of the elasticity cluster in original and aggregated states should be at the same place, in order to preserve the locality of

Algorithm 1: Aggregation of elasticity bins

```

forall ElasticityBin do
    forall TetrahedronCluster do
        ModifiedVolume = 0;
        Tetrahedron center =
            getTetrahedron(ClusterBarycenter);
        TetrahedronList InterestList = List(center);
        while ModifiedVolume < ClusterVolume do
            Tetrahedron interest = InterestList.pop();
            for Tetrahedron neighbour  $\in$ 
                interest.NeighbourTetrahedra do
                ComputeInterest(neighbour);
                InterestList.add(neighbour); (Sorted by
                    measure of interest)
            if NotAlreadyModified(interest) then
                interest.ModifyElasticity(ElasticityBin.Elasticity);

                ModifiedVolume += interest.Volume;
                NotAlreadyModified(interest) = false;
    
```

elements with this elasticity. Second, the distribution of the elasticity around this barycenter in the aggregated version should also correspond to the distribution in the original version of the elasticity. To do so, the elasticity is incrementally modified, starting at the barycenter of the cluster. For this purpose, a measure of interest is created for a tetrahedron, based on the distance to the cluster

barycenter and on the relative orientation between the tetrahedron to be modified and the original tetrahedron w.r.t the cluster barycenter. The final measure of interest is a weighted sum of the distance factor and the orientation factor. The tetrahedra are then incrementally modified, starting at the center of the barycenter according to this measure of interest. A tetrahedron which elasticity has been modified can no longer be modified.




Figure 4: Example of clustering: our algorithm transforms the elasticity of mesh elements: from 3 sub parts of same elasticity (left image), there is only one homogeneous sub part after the end of the algorithm (right image). The aggregated part has similar volume compared to previous sub parts, and keeps distribution features from the previous position of the sub parts.

Since each tetrahedron can only be modified once, the last elasticity bin to be considered can only modify the tetrahedra that have not been modified by any previous elasticity bin. The location of the tetrahedra might therefore not depend on the distribution of the tetrahedra from this cluster in the object. This raises the fact that the order of aggregated elasticity bins has a non-negligible impact on the final configuration.

This entire process can be performed multiple times, with decreasing number of elasticities and of clusters, for further simplification. Figure 3, Cube 4 shows the aggregated cube, with only 2 large homogeneous volumes.

Geometric coarsening

The last aspect of this method is to attribute elasticities to the elements of coarser resolution meshes. The object is already decomposed into large homogeneous volumes, so the elasticity of the coarse elements inside these volumes is straightforwardly set to the volume elasticity. At the limit of these volumes however, the same procedure cannot be used. Existing methods compute the new mesh based on the elasticity in order to keep homogeneous regions in the coarser mesh, requiring either a computation on-the-fly of the new

geometry or a precomputation based on a previous known elasticity pattern. These methods are not adapted to an arbitrary elasticity distribution inside the mesh.

A first precomputation phase computes the correspondence between the tetrahedra from each mesh resolution. In order to determine the elasticity of coarser elements, the elasticity of the finer elements located at the same place of the object must be known. This precomputation computes the correspondence C_{ij} for each pair of tetrahedra T_i and T_j respectively taken from the fine and the coarse tetrahedra, $C_{ij} = \frac{\text{volume}(T_i \cap T_j)}{\text{volume}(T_i)}$. This step provides a comparison matrix between the two meshes, that can be arbitrary meshes sharing the same surface.




Figure 5: Computation of the intersection between triangles from the finer resolution (dotted lines) and from the coarser resolution (full lines). The area between the intersection with stiff elements (in red/stripes) and soft elements (in yellow/squares) is compared in order to determine the elasticity of the coarse element.

Using this comparison matrix, the elasticity of the coarse elements is determined as the elasticity of the elements occupying the maximum space in the element. This process is illustrated in Figure 5. This approach has the advantage of preserving the number of elasticities from the fine aggregated object. An alternate approach computing the average elasticity value from the fine elements inside the coarse element was also considered, yet preliminary tests showed no significant improvement of this approach. Figure 3, Cube 5 shows the final result, with a coarser mesh. While the result is different from the original cube, an evaluation is performed using a perception-based quality criterion, as explained in Section 4. It shows in most cases a non-significant perceptual difference between the two objects.

4. Evaluation

Objective

We assess the fact that an object coarsened by our method is sufficiently close from the original one, from a perceptual point-of-view. In order to evaluate this, the initial object is compared to the simplified one, for example both cubes from Figure 7, using forces obtained through the interaction with both objects. A perception-based criterion is used, which means that in order for the simplification to be valid, the difference between the two forces should

not be noticeable. A usual metric in order to enforce this is the Just Noticeable Difference (JND), which is the maximum relative difference between two stimuli – here forces – under which there is not perceived difference. For forces, the JND is known to be close to 10% [HHC*08]. We later refer to *acceptable difference*, a difference to the reference inferior to the JND, and *noticeable difference*, a difference to the reference superior to the JND. If the force from the simplified object has a difference to the force from the initial object inferior to the JND, we consider that our method has been successful.

In order to evaluate it, our method is used on several scenarios, representing general distribution patterns, and from which most distributions can be deduced. All scenarios involve the interaction with a cube in which tetrahedra are associated with various Young Modulus. The cube is a $10\text{cm} \times 10\text{cm} \times 10\text{cm}$ cube, Poisson's ratio is 0.3, and the total mass is 1kg . Each scenario has a specific elasticity distribution pattern. For each scenario, the reaction force is recorded on the interaction with a reference fine-resolution cube, without the method being applied. The method is then applied, on various resolution levels of the mesh. Each recorded force is then compared to the reference one. The reference is also compared to the force corresponding with the interaction with the fully homogenized cube, which represents the simplest pattern from the initial configuration.

The scenarios involve a distribution of soft tetrahedra, with elasticity $2\text{kPa} \leq E_{soft} \leq 6\text{kPa}$, and stiffer tetrahedra, with elasticity $38\text{kPa} \leq E_{stiff} \leq 42\text{kPa}$, except for the last scenario, involving a random distribution of tetrahedra with elasticities $4\text{kPa} \leq E \leq 200\text{kPa}$. All the distributions are displayed in Figure 6. A stiff zone contains a majority (2/3) of stiff elements, and a soft zone contains a majority (2/3) of soft elements.

In every case, the cube is compressed from the top surface by 10% of its height over 15 seconds, followed by a static phase for 5 seconds, in order to reach equilibrium state, with corresponding force. The tested resolutions are cubes with 6,000, 3,072 and 1,296 tetrahedra. These cubes were obtained with a decomposition of the initial cube into $n \times n \times n$ smaller cubes, each divided into 6 tetrahedra. The process is illustrated in Figure 7, with the scenario involving a majority of stiff tetrahedra on the top part of the cube. An example of the recorded forces can be seen in Figure 8. The obtained behavior is really stable and consistent over the simulation, i.e. the ratio between all forces remains similar over the simulation.

Apparatus

Our method was implemented using SOFA framework [ACF*07], using corotational FEM for the deformation [MG04]. It was run on a PC (CPU : Intel Core i7 - 4800MQ 2.7 GHz, GPU : NVIDIA Quadro K3100M, Memory : 16GB).

Results

Performance of the simulation with the computation time required for the deformation is presented in Table 1. Computation time is compared between 3 mesh resolutions, the reference object has 6,000 tetrahedra, and meshes at a coarser resolution contain respectively 3,072 and 1,296 tetrahedra. The recorded interaction is a compression of the object, depicted in Figure 7. As expected, the simulation gets more efficient with a lower number of tetrahedra in

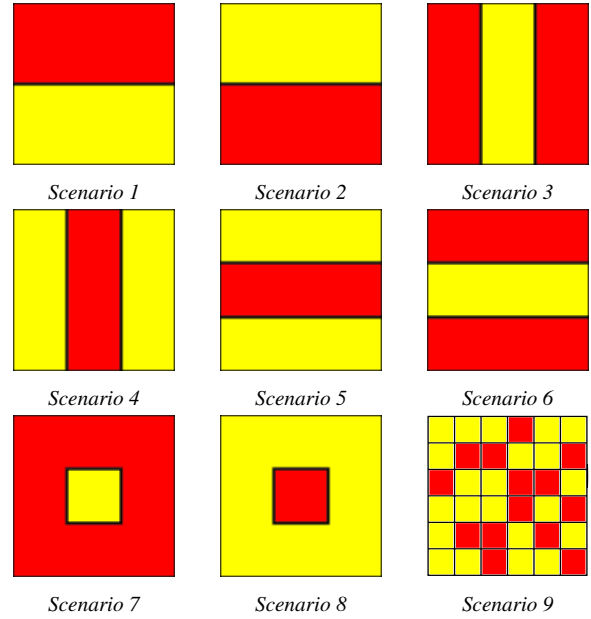


Figure 6: 2D Illustration of the 9 scenarios with different elasticity distributions chosen to evaluate the method. The red part represents a stiff zone, and the yellow part represents a soft zone. Scenario 9 is a more heterogeneous distribution using a randomized filling of the cube. The presented squares are a cross-section from the actual cubes, for a force applied on the top.




Figure 7: Scenario 1 close-up: The cube is compressed from a pressure on the top surface. The left cube (the input cube) has a majority of stiff (red) tetrahedra on the top part, and on the right the corresponding aggregated version, generated by phase P3.

the mesh. The coarsest resolution shows a speedup factor of more than 6 compared to the reference. The computation time is independent from the elasticity distribution inside the object.

The full result of the forces computed with the previously described scenarios can be found in Table 2. In addition to the reference force and the force recorded from the fully homogenized cube, several levels of precision of the method are presented, each in all 3 resolutions. Since the different scenarios involve two different ranges of elasticity (stiff elements and soft elements), the results




Figure 8: Example of recorded forces for Scenario 2: reference force and forces with a fully homogenized cube, and using 1 and 3 clusters. The JND in force is also displayed. Only the force with 3 clusters is found here in acceptable perceptual range.

Table 1: Average computation time required for the deformation of cube with various number of tetrahedra. The mesh for the reference object has 6,000 tetrahedra, compared to 3,072 and 1,296 for the coarser object simplified by our method.

	reference	our method	
Number of Tetrahedra	6,000	3,072	1,296
Computation time (ms)	13.0	5.2	2.2

using the method are all computed using 2 elasticity bins. The results are compared at each resolution using 1, 3 and 6 clusters for each elasticity. The results are displayed as a percentage difference to the reference force. As a reference, forces recorded on fully homogeneous cubes for these resolutions displayed force difference inferior to 2%.

From these data, it is possible to draw several conclusions. First of all, just homogenizing does not produce an acceptable force in most of the cases, as shown in Table 2 Line 3. In most of these cases, the homogenized cube is much stiffer than the reference cube, which justifies the use of our method. Second the force is very stable from one resolution to another, independently from the number of clusters used or the scenario involved.

More dependent on the initial distribution, a few observations can be made, representing the different general behavior of our method. In most cases, due to the simplicity of certain patterns, 3 clusters are enough to reproduce a force similar to the original one. For instance, scenarios 1, 2 and 5 have the same behavior. The force on the homogenized cube is above the reference force by an amount exceeding the JND, the force with only one cluster not important enough, and the force with 3 clusters is really close from the reference force. While the homogenized elasticity is too important, one cluster produces a large volume with a too low stiffness, creating a too low force. Those scenarios have in common to have a uniform elasticity distribution on the interaction surface, and to have a large enough soft volume in the center in order to compensate for the post-aggregation soft region.

Since stiff elements are aggregated first, leaving the remaining space for soft elements, scenarios with an important proportion of stiff elements require a more important precision for the algorithm. For instance, Scenario 3 and 6 both require 6 clusters in each

elasticity value to provide acceptable results. Conversely, scenarios with a great proportion of soft elements require a lesser precision of our method in order to achieve satisfactory results, such as Scenario 4 and 5.

Having a great majority of the object with a unique elasticity and a small portion of a different stiffness inside does not benefit much from our method. In the case of a mostly stiff object, such as Scenario 7, a small part of soft material inside goes completely undetected and any simplification returns an acceptable result. Conversely, a stiff element inside a mostly soft object such as Scenario 8 remains challenging for our method. While an existing stiffer core makes the homogenized cube too stiff, no reasonable (under 10) number of clusters can re-create a distribution close to the initial one, because the initial soft part included stiff elements that are later aggregated. This creates an external softer part that produces forces really inferior to the reference one. A completely random distribution (Scenario 9) does not benefit much from an important precision either, since every distribution provides a good result, because there was no clear pattern to recover, and this distribution does not create uniform soft parts that would make the homogenized cube too stiff.

As a summary, except for the case of small stiff elements inside the object, provided the number of clusters is chosen accordingly to the complexity of the initial distribution, our method provides a simplification of the elasticity that is not noticeable, yet allows an important coarsening of the object, leading to 6 times faster simulation.

5. Illustrative use case: cooking scenario

In order to illustrate our method, we implemented a use case, of a steak being cooked. The heterogeneous object, the steak, is attached to a pan, and interaction is possible through a cooking tool, a spatula, as depicted in Figure 9.




Figure 9: Use case: A steak on a pan being cooked. The user controls a spatula to interact with the steak.

Haptic interaction for this use case is achieved using a Geomagic Touch (3D Systems, Rock Hill, USA), and implemented using again the SOFA framework, as depicted in Figure 1. We implemented an impedance haptic coupling as explained in [LO08]. The position of the haptic device defines a constraint on the deformable

Table 2: Percentage force difference compared to the reference for each scenario, for the homogenized cube, and under the different conditions of clusters (1, 3 and 6) and resolutions (6,000, 3,072 and 1,296 tetrahedra). Green indicates an acceptable difference, and red a noticeable difference to the reference force. The threshold used to discriminate acceptable difference and noticeable difference is the JND in force, i.e. 10%.

Tetrahedra	Clusters	Sc. 1	Sc. 2	Sc. 3	Sc. 4	Sc. 5	Sc. 6	Sc. 7	Sc. 8	Sc. 9
6,000	reference	20.7116	20.8187	25.0896	19.8527	18.9669	22.6011	27.9015	16.5248	107.695
6,000	homogenized	+19.6%	+17.6%	+9.6%	+6.5%	+16.2%	+16.7%	+6.3%	+13.0%	+3.4%
6,000	1	-22.7%	-22.7%	-35.1%	-8.5%	-30.5%	-12.1%	+2.2%	-27.9%	-5.1%
3,072	1	-22.6%	-22.3%	-35.6%	-9.4%	-32.8%	-10.0%	+2.5%	-31.4%	-5.8%
1,296	1	-15.9%	-16.6%	-35.1%	-8.4%	-30.4%	-8.3%	+3.0%	-29.4%	-5.7%
6,000	3	-2.1%	+5.1%	-16.4%	-4.6%	-4.6%	-16.9%	+0.1%	-26.7%	-3.6%
3,072	3	-4.4%	+5.0%	-19.0%	-5.6%	-8.5%	-19.0%	-1.7%	-29.8%	-4.6%
1,296	3	-0.8%	+6.7%	-14.9%	-4.4%	-4.3%	-16.7%	+0.3%	-26.2%	-3.9%
6,000	6	+6.6%	-0.0%	-5.8%	-5.8%	-5.7%	+9.5%	+3.6%	-21.2%	-3.7%
3,072	6	+6.3%	-0.6%	-7.3%	-7.3%	-6.4%	+7.6%	+2.4%	-24.0%	-5.4%
1,296	6	+6.8%	+0.5%	-4.4%	-4.4%	-3.3%	+8.5%	+2.1%	-22.7%	-4.0%

object. The force is then computed based on the deformation of the object, and returned to the user by the haptic device.

Two different resolutions are used for the steak, 2,894 and 1,207 tetrahedra, as depicted in Figure 10. Two levels of precision of our methods are compared, using 1 and 3 clusters for each elasticity. A cooking steak is expected to have a greater concentration of stiff elements on the bottom part, corresponding to Scenario 2. The interaction consists in a progressive constraint on the top of the steak by the spatula, similar to the interaction used in the previous section. Results show an improvement in performances, from 70Hz on the fine mesh to 120Hz on the coarse mesh. Interaction with the object with 3 clusters creates an acceptable difference (5% difference), while there was a noticeable difference in the interaction with the object with only 1 cluster (50% difference).

6. Discussion and Conclusion

In this paper, we presented a novel multiresolution approach for haptic interaction with heterogeneous deformable objects. Our method is based on the design of elasticity clusters that could be used for simulating the same object with a coarser mesh. While our method enables to improve computation time performances, it still keeps important features of the elasticity distribution of the original object. We evaluated our method on diverse scenarios using a perception-based quality criterion. We found that the recorded force is independent from the multiresolution level, thus allowing the use of a much coarser geometry, and subsequent important gain in performances, at a non-noticeable perceptual cost.

Future work on our approach could first include the possibility to provide haptic rendering with different tools. As of today, the elasticity is modified over the entire object, and local interactions on a small part of the object are greatly impacted by our method. In order to provide a similar interaction over the different mesh resolutions, our method requires a tool that interacts with a significant part of the object at once, for example a large spatula pressing a great part of the top surface. In order to overcome this limitation, a domain decomposition on the object could be performed,

and our method and the corresponding geometric coarsening could be applied on all domains except for the one in direct interaction. It would provide the accurate deformation on the surface with the precise mesh, and the coarse interaction between domains represents a good approximation of the force from the remaining of the object. With such a decomposition, topology modifications such as cutting gestures could thus broaden the spectrum of possible applications, with domains reverting to non-modified state as soon as they are in the cutting path. Further validation could also include a thorough comparison with the different methods from the literature, both for computation time performances and for the perceived interaction force. A comparison with different spatial clustering methods could also prove to be interesting. Moreover, as of today, the optimal simplification level relies on a prior knowledge of the elasticity distribution inside the object, and an analytical tool computing the optimal clustering level for a given object based on its distribution could be of interest. Furthermore, while depending on well-studied results on perception, we did not formally validate our method with a user study, which could evaluate the haptic perception of our approach.

Thanks to the performance gain obtained by geometric coarsening at no perceptual cost, our method allows the interaction with a great number of heterogeneous objects. This could lead to potential applications in the entertainment field, with haptic interaction with a large environment, or to virtual prototyping applications.

References

- [ACF*07] ALLARD J., COTIN S., FAURE F., BENSOUSSAN P.-J., POYER F., DURIEZ C., DELINGETTE H., GRISONI L.: Sofa-an open source framework for medical simulation. In *Proceedings of Medicine Meets Virtual Reality* (2007), vol. 125, IOP Press, pp. 13–18. 6
- [CLSM15] CHEN D., LEVIN D. I., SUEDA S., MATUSIK W.: Data-driven finite elements for geometry and material design. *ACM Transactions on Graphics (TOG)* 34, 4 (2015), 74. 2
- [GFBP13] GILLES B., FAURE F., BOUSQUET G., PAI D. K.: Frame-based interactive simulation of complex deformable objects. In *Deformation Models*. Springer, 2013, pp. 145–166. 1, 2




Figure 10: Different resolutions and number of clusters achieved by our method: The initial textured object is shown in a), with the corresponding elasticity distribution in b). Green represents stiffer tetrahedra, and blue softer tetrahedra. The coarsened mesh using 1 cluster is shown in c), and using 3 clusters in d). The interaction force with the steak with 1 cluster has a noticeable difference (50%) to the reference, whereas the interaction force with the steak with 3 clusters had an acceptable difference (5%).

- [HHC*08] HINTERSEER P., HIRCHE S., CHAUDHURI S., STEINBACH E., BUSS M.: Perception-based data reduction and transmission of haptic data in telepresence and teleaction systems. *IEEE Transactions on Signal Processing* 56, 2 (2008), 588–597. 6
- [IYCN13] IMAI S., YUE Y., CHEN B.-Y., NISHITA T.: Adaptively simulating inhomogeneous elastic deformation. In *Proceedings of GRAPP/IVAPP* (2013), Citeseer, pp. 237–244. 2, 3
- [KMOD09] KHAREVYCH L., MULLEN P., OWHADI H., DESBRUN M.: Numerical coarsening of inhomogeneous elastic materials. *ACM Transactions on Graphics* 28, 3 (2009), 51. 2, 3
- [KS04] KOLCÁREK P., SOCHOR J.: Haptic rendering using velocity driven level of detail. In *Proceedings of Virtual reality interaction and physical simulation* (2004), pp. 149–158. 2
- [LO08] LIN M. C., OTADUY M. A.: *Haptic Rendering : Foundations, Algorithms, and Applications*. A K Peters, Ltd., 2008. 7
- [MG04] MÜLLER M., GROSS M.: Interactive virtual materials. In *Proceedings of Graphics Interface* (2004), pp. 239–246. 3, 6
- [MWN*16] MANTEAUX P.-L., WOJTAN C., NARAIN R., REDON S., FAURE F., CANI M.-P.: Adaptive physically based models in computer graphics. *Computer Graphics Forum* (2016). 2
- [NKJF09] NESME M., KRY P. G., JEŘÁBKOVÁ L., FAURE F.: Preserving topology and elasticity for embedded deformable models. *ACM Transactions on Graphics* 28, 3 (2009), 52. 1, 2
- [NMK*06] NEALEN A., MÜLLER M., KEISER R., BOXERMAN E., CARLSON M.: Physically based deformable models in computer graphics. *Computer Graphics Forum* 25, 4 (2006), 809–836. 2
- [NPF06] NESME M., PAYAN Y., FAURE F.: Animating shapes at arbitrary resolution with non-uniform stiffness. In *Proceedings of Virtual reality interaction and physical simulation* (2006). 3
- [OL05] OTADUY M. A., LIN M. C.: Sensation preserving simplification for haptic rendering. In *Proceedings of ACM SIGGRAPH Courses* (2005), p. 72. 2
- [PDC11] PETERLIK I., DURIEZ C., COTIN S.: Asynchronous haptic simulation of contacting deformable objects with variable stiffness. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems* (2011), pp. 2608–2613. 2
- [PDZ05] PAYANDEH S., DILL J., ZHANG J.: A study of level-of-detail in haptic rendering. *ACM Transactions on Applied Perception* 2, 1 (2005), 15–34. 2
- [TMD*15] TALVAS A., MARCHAL M., DURIEZ C., OTADUY M., ET AL.: Aggregate constraints for virtual manipulation with soft fingers. *IEEE Transactions on Visualization and Computer Graphics* 21, 4 (2015), 452–461. 2
- [TOT13] TAGAWA K., OISHI T., TANAKA H. T.: Adaptive and embedded deformation model: An approach to haptic interaction with complex inhomogeneous elastic objects. In *Proceedings of World Haptics Conference* (2013), IEEE, pp. 169–174. 2
- [TYT12] TAGAWA K., YASUYUKI S., TANAKA H. T.: Online re-mesh

and multi-rate deformation simulation by gpu for haptic interaction with large scale elastic objects. In *Proceedings of Haptics Symposium* (2012), IEEE, pp. 533–540. 2