Six-DoF Haptic Interaction with Fluids, Solids, and their Transitions

Gabriel Cirio
INRIA Rennes, France
URJC, Madrid, Spain

Maud Marchal
INSA/INRIA Rennes, France

Miguel Otaduy
URJC, Madrid, Spain

Anatole Lécuyer
INRIA Rennes, France

ABSTRACT

Haptic interaction with different types of materials in the same scene is a challenging task, mainly due to the specific coupling mechanisms that are usually required for either fluid, deformable or rigid media. Dynamically-changing materials, such as melting or freezing objects, present additional challenges by adding another layer of complexity in the interaction between the scene and the haptic proxy. In this paper, we address these issues through a common simulation framework, based on Smoothed-Particle Hydrodynamics, and enable haptic interaction simultaneously with fluid, elastic and rigid bodies, as well as their melting or freezing. We introduce a mechanism to deal with state changes, allowing the perception of haptic feedback during the process, and a set of dynamic mechanisms to enrich the interaction through the proxy. We decouple the haptic and visual loops through a dual GPU implementation. An initial evaluation of the approach is performed through performance and feedback measurements, as well as a small user study assessing the capability of users to recognize the different states of matter they interact with.

Index Terms: I.3.5 [Computer Graphics]: Computational Geometry and Object Modeling—Physically based modeling; H.5.1 [Information Interfaces and Presentation]: User Interfaces—Haptics I/O; H.5.1 [Information Interfaces and Presentation]: Multimedia Information Systems—Artificial, Augmented, and Virtual Realities

1 INTRODUCTION

Our world is made of complex and rich materials which appear in different states of matter, and even change their state or mechanical properties when mixed, heated, or cooled. The state and mechanical properties of materials largely affect the way we interact with them; therefore, effective virtual environments require high-fidelity simulation of diverse materials and states of matter, their transitions, and their interactions. In particular, high-fidelity haptic interaction of fluids, solids, and their state transitions allows us to perceive intuitively their varying stiffness and/or viscosity, making our interaction with virtual environments more natural.

The computer graphics community has developed impressive solutions to simulate fluids, solids, their interactions and their changes of state [22]. However, many of these solutions have not made their way through computer haptics due mostly to two major challenges. First, haptic rendering requires specific coupling mechanisms between the simulation of the virtual environment and force feedback display to maximize rendering fidelity while guaranteeing stability [17]. Second, haptic rendering requires update rates at least one order of magnitude higher than visual rendering. These two challenges are exacerbated when simulating different and dynamically-changing materials, each with its own feedback coupling mechanism, and with complex interactions and state transitions.

We enable haptic interaction simultaneously with fluids, deformable bodies and rigid bodies, as well as their melting and freezing, by providing effective solutions to the two challenges described above. First, we adopt a representation of dynamic materials based

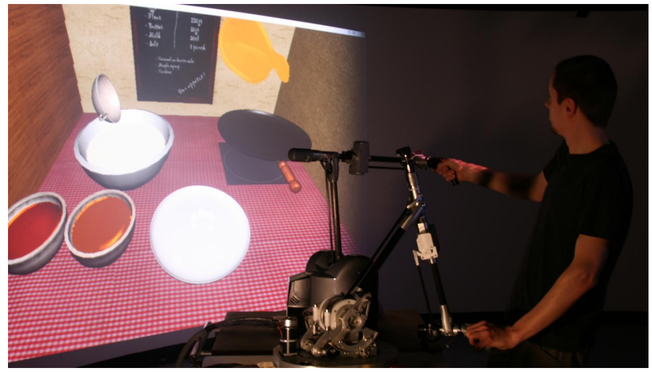

Figure 1: A user interacting with fluids, deformable bodies and rigid bodies through a virtual crepe making simulator, illustrating our multistate haptic interaction technique.

on Smoothed-Particle Hydrodynamics (SPH) [19, 27] that allows us to seamlessly integrate solids and fluids. Using this representation as a basis, we have developed a haptic coupling mechanism that handles in a unified manner the interaction with rigid, elastic, and fluid media.

Second, we have designed algorithms for very efficient simulation of combined fluid and solid media integrated with haptic interaction. One of these algorithms leverages dual GPU setups by processing the haptic and visual rendering components on separate GPUs, enabling efficient parallel computations. In addition, we present an efficient method to simulate changes in the state of matter, such as melting and freezing, without altering the haptic coupling mechanism, and enabling the perception of haptic properties during state changes.

Our simulation methods enable haptic interaction with a wide range of rigid, elastic, and fluid materials, empowering diverse applications in the industrial, medical and entertainment fields. In particular, we have successfully tested our methods on a pedagogical entertainment application of virtual pancake cooking, shown in Figure 1. We have evaluated our approach through performance and feedback measurements, as well as a small user study assessing the capability of users to recognize the different states of matter they interact with.

2 PREVIOUS WORK

In this section, we survey the existing techniques allowing haptic interaction with either rigid bodies, deformable bodies or fluids, and efforts made towards multistate online or offline interactions.

2.1 Haptic Rendering of Different States of Matter

Rigid Bodies The challenge in 6DoF haptic rendering lies in the cost of collision detection computations, and the necessity of updating at over 1kHz [17]. Several techniques have tackled the problem using different approaches, such as a voxel-based haptic rendering [18] computing penalty forces between a point-sampled dynamic tool and a voxelized static virtual environment, a sensation-preserving simplification algorithm [24] trading accuracy for speed through multiple levels of detail while preserving the interaction forces, or a 6DoF God-object technique [23] that uses continuous collision detection and a constraint-based approach. For a detailed description of these techniques, we refer the reader to [17].

Deformable Bodies The haptic interaction with deformable bodies exacerbates the aforementioned collision and high update rate challenges due to the high complexity of physically based deformations. The majority of recent techniques leverage continuum mechanics through the widely established Finite Element Method (FEM). They use either geometrical deformation constraints [13], Linear Complementary Problem (LCP) formulations [12], precomputation of forces and displacements [10, 25], multi-rate approaches through different resolution meshes [1] or linearized approximations of the interaction forces [14], and voxel-based collision detection with a multiresolution pointshell [2].

Fluids Recent work has explored the possibilities of haptic interaction with fluids, although to a lesser extent than for the rigid and deformable states. Specific haptic interaction scenarios are possible in [11] by pre-computing the non-linear interaction forces between fluid and rigid bodies. Some physically based techniques using an Eulerian approach [3, 28] allow the haptic interaction with limited amounts of fluid. Recent work [7] describes a 6DoF approach for haptic interaction with viscous fluids, for a Lagrangian simulation based on Smoothed-Particle Hydrodynamics (SPH) [19] [20], allowing the use of arbitrary-shaped objects and large amounts of fluid with 6DoF haptic feedback.

2.2 Unified Models for Multi-state Media

Previous work on haptic interaction with deformable bodies or fluids only allowed the interaction with rigid bodies and one other state, with the exception of the pioneering CORDIS-ANIMA simulation framework [6], which was the first to simulate different states with haptic feedback.

Recent work on physically based simulations has used the Smoothed-Particle Hydrodynamics model for the simulation of multiple states of matter in a unified framework. However, update rates are usually not real-time, or fall short for haptic interaction. In [21], deformable bodies are modeled in real-time with SPH and the Moving Least Squares (MLS) algorithm to compute the elastic forces. The approach is improved in [15], allowing solid, deformable and fluid animation and interaction in a unified approach. A fully SPH-based approach is presented in [27], later corrected in [4] for a rotationally invariant formulation.

Our approach is the first to provide physically based haptic feedback for fluid, deformable and rigid states of matter in the same simulation. Previous haptic rendering techniques were focused on a single state, e.g. fluid, deformable or solid, making it difficult to have different types of media coexist in the same simulation while providing convincing haptic feedback. However, the SPH model provides a generic simulation framework for the simulation of fluid, elastic and rigid mechanics [27, 4]. This work, initiated in [8], leverages the SPH model to provide 6DoF haptic interaction with different states of matter through a unified haptic rendering mechanism. We further leverage this unified framework by proposing simple state change mechanisms and the use of proxies of different states. Combined to an efficient dual GPU implementation, we enable complex haptic interaction scenarios through complex tools, allowing the interaction with richer virtual environments.

3 SPH Physically Based Simulation

In this section, we briefly present the physically based simulation equations and algorithms that constitute the foundation of our approach, and detail our mechanism for changes of state.

3.1 Unified simulation for multi-state media

The SPH model [19] is based on a set of particles discretizing the simulated media and carrying different physical properties, such as mass, viscosity or elasticity parameters. The smoothed quantity Q_i of a particle i at any position $\mathbf{x}_i$ in space is computed through the general formula:

$$Q_i = \sum_j Q_j V_j W(\mathbf{x}_i - \mathbf{x}_j, h) \tag{1}$$

where Q_j is the discrete quantity Q sampled for neighboring particle j at position $\mathbf{x}_j$, V_j is the volume of j, and W is the smoothing kernel of support h, where particles farther than the distance h are not taken into account.

The motion of fluids is driven by the Navier-Stokes continuum equations, while the behavior of deformable bodies is driven by elasticity equations. Using the implementation of these equations in the SPH model [20] [27] [5], pressure, viscosity and elasticity forces are computed at each time step:

$$\mathbf{f}_i^{pressure} = -V_i \sum_j V_j \frac{P_i + P_j}{2} \nabla W(\mathbf{x}_i - \mathbf{x}_j, h) \tag{2}$$

$$\mathbf{f}_i^{viscosity} = \mu V_i \sum_j V_j (\mathbf{v}_j - \mathbf{v}_i) \nabla^2 W(\mathbf{x}_i - \mathbf{x}_j, h) \tag{3}$$

$$\mathbf{f}_{ji} = -2V_i^0 (\mathbf{I} + \nabla \bar{\mathbf{u}}_i^T) \sigma_i V_j^0 \nabla W(\mathbf{x}_i^0 - \mathbf{x}_j^0, h) \tag{4}$$

$$\mathbf{f}_i^{elasticity} = \sum_j \frac{-\mathbf{R}_i \mathbf{f}_{ji} + \mathbf{R}_j \mathbf{f}_{ij}}{2} \tag{5}$$

where ∇ and ∇^2 are respectively the gradient and Laplacian of the physical quantities. In the fluid equations, P is the pressure, $\mathbf{v}$ the velocity, and μ is the viscosity coefficient. In the elasticity equations, $\mathbf{x}_i^0$ and V_i^0 are respectively the initial position and volume of particle i, $\nabla \bar{\mathbf{u}}_i$ is the gradient of the locally rotated displacement field, σ is the stress and $\mathbf{R}$ is the rotation matrix computed through a corotational approach. These are <u>internal</u> forces, since they govern the internal behavior of each material.

Parameters such as the viscosity can be changed to obtain different fluid behaviors, while the Young modulus controls the stiffness of an elastic material.

Rigid bodies are also modeled with the SPH particles, with the main difference that there are no internal forces, and that bodies are subject to rigid dynamics. Using a particle based approach for rigid bodies allowing the seamless use of arbitrary-shaped rigid bodies such as concave objects, and provides a gain in efficiency by using a single model (the SPH model), making the approach unified with fluid and deformable bodies. In order to further embrace a unified model approach, interaction forces between fluid, deformable and solid media are computed using the same SPH force, namely fluid (pressure and viscosity) forces, providing a reasonable amount of control through density and viscosity values. Interaction forces are <u>external</u> forces, since they govern how different bodies (of equal or different material) interact with each other.

For further details about the foundations of our approach, we refer the reader to its initial implementation [8].

3.2 Changes of State

Changes of state can be easily added to a multistate simulation, as shown in [27]. In our approach, each particle carries a fluid state coefficient K^f and a deformable state coefficient K^d, both varying between 0 and 1 and with their sum equal to 1. These coefficients can be modified, producing changes of state. If $K^f = 1$ and $K^d = 0$, the particle behaves like a fluid. If K^f is lowered and K^d is raised, the particle starts moving to a deformable state. These changes are possible by multiplying fluid forces (pressure, viscosity) by the fluid state coefficient, and deformable body forces (elasticity) by the deformable state coefficient. Hence, all three forces are computed for each particle, and their magnitudes are scaled according to the aforementioned coefficients. However, in order to obtain symmetric forces, they are also scaled by the neighbor coefficients. Hence, the total force applied by particle j on particle i is:

$$\mathbf{f}_{ij}^{total} = K_i^f K_j^f (\mathbf{f}_{ij}^{pressure} + \mathbf{f}_{ij}^{viscosity}) + K_i^d K_j^d (\mathbf{f}_{ij}^{elasticity}) \qquad (6)$$

A body *id* is carried by each particle, so that only particles with the same *id* can trigger the computation of internal forces. At any time, a particle belonging to a body can be set loose by changing its body *id*. For example, when the K^f of a deformable body particle reaches 1, its body *id* can be changed so that external forces are computed instead of internal ones, thus make the particle leave the deformable body. Figure 2 illustrates a solidification mechanism, where some fluid batter is poured in a pan, then cooked into a deformable pancake and deposited on a plate, as part of a cooking simulation demonstration.

Rigidification is simply achieved in a single timestep, when a given variable (either a state coefficient or other physical attributes such as temperature) reaches a threshold. Particles belonging to the new rigid body get the same *id*, a rigid state, and are used to compute the inertia matrix of the body.

4 6DoF Haptic Rendering

In the previous section we showed how to simulate different states (fluid, deformable and rigid bodies), their interactions and their changes of state with the same SPH model. In this section, we detail our haptic coupling scheme allowing to convey force feedback to the user through a multiple state proxy and a 6DoF haptic device.

4.1 Multistate proxy

Everyday interactions with our surrounding environment are often achieved through solid tools, in either rigid or deformable states. While most haptic rendering mechanisms focus on the interaction through a proxy allowing only a single state, usually rigid, our framework allows the manipulation of a multistate proxy to interact with multistate environments.

A multistate proxy is achieved by defining a rigid handle, which is surrounded by matter in other states. The rigid handle is coupled to the haptic device through a virtual coupling mechanism, as described in [7]. Having a rigid handle simplifies haptic interaction through a haptic device with the degrees of freedom of a rigid body. In the case of a rigid proxy, the rigid handle is the entire proxy. In the case of a deformable proxy, the handle is a rigid core surrounded by matter in the deformable state, and the external (interaction) SPH forces between rigid and deformable parts are replaced by SPH elastic forces (Eq. 4 and 5). This simple approach has the advantage of requiring few implementation changes.

The proxy is therefore a rigid body as any other rigid body of the virtual world, and interacts with all media around it through the interaction forces described in Section 3. In this way, the user can feel the virtual world through an arbitrary-shaped rigid body in either rigid or deformable state. The resulting 6DoF haptic coupling scheme is unified, in the sense that it allows the interaction with different media (fluid, deformable and solid) without distinction, since the haptic interaction forces are computed in a unified way.

4.2 Dynamic effects

In addition to having a multistate proxy, and by leveraging the use of a rigid core and the state change mechanism, we can easily simulate proxy state changes. Dynamic effects such as tool melting (solid to fluid) can be simulated with 6DoF haptic feedback without any change in the implementation, only by allowing the media

Figure 2: Scenario illustrating the changes of state. The user can interact with fluid, deformable and rigid states through a cooking simulator. Some fluid batter is poured from a rigid cup into a pan (upper left), then progressively cooked into a deformable body (upper right), flipped (lower left), and then deposited on a solid plate (lower right).

surrounding the rigid core to melt. Future extensions of the framework could consider plasticity and tearing capabilities [27], so that deformable proxies could benefit from a wider range of dynamic effects, as well as state changes induced by the proxy.

5 GPU Implementation and Performance

In this work, and as many recent high-performance haptic rendering algorithms [14] [9], we decoupled the haptic and graphic rendering loops to achieve higher haptic update rates. The novelty of our implementation lies in the use of two GPUs for the decoupling.

The first GPU is used to run the simulation, while the second is used for visual rendering purposes. A CPU simulation thread running at a haptic rate controls the computations on the simulation GPU, and a CPU graphic thread running at a visual rate controls the OpenGL rendering operations of the second GPU. The simulation GPU reads and writes data to the graphic memory. After each iteration, only the data required for rendering is copied to central memory, namely positions and haptic coupling data. The CPU simulation thread then sends the coupling data from central memory to the corresponding haptic devices, while the CPU graphic thread copies the position data from central memory to the graphic memory of the second GPU.

This decoupled implementation results in a gain of up to one order of magnitude for the haptic loop, as seen in Section 5.1. Although our scheme requires writing the position data to central memory at each time step of the haptic loop, this operation accounts for less than 1% of the haptic loop computation time. Only 12 bytes have to be copied per particle (3 position floats), which corresponds to as little as 352KB for 30,000 particles. Copying the data at each time step allows complete independence between the haptic and graphic loops. The only inter-process communication mechanism is set up during memory copies to avoid concurrency issues. This performance boost is an essential improvement to obtain a realistic haptic feedback with rich and complex scenes.

The simulation is implemented on the GPU using the CUDA framework, making extensive use of texture lookups, and benefiting from texture data caching when retrieving the different data arrays.

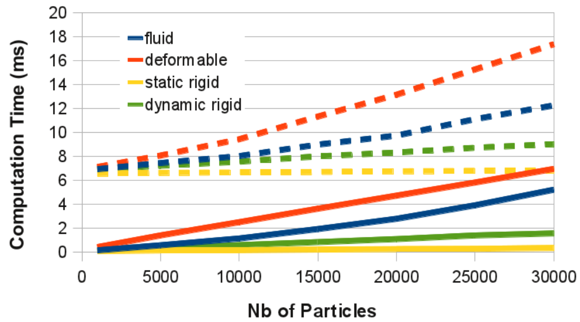

Figure 3: Haptic timestep in milliseconds for scenes made of either a rigid body (static and dynamic), a deformable body or a fluid volume, according to the number of particles, with coupled (dashed curves) and decoupled (continuous curves) GPU implementations.

5.1 Haptic Frequency

We measured the haptic time step of the simulation in order to evaluate the performance of our dual GPU implementation. We compare these results to a single GPU implementation. Four scenes were used, each with either a fluid volume, a deformable body, a rigid body, and a static rigid body that can exert, but not receive, interaction forces. The scenes have the same size and shape (a cuboid with a square base of 0.5m side length and a variable height according to the number of particles). We measured the time between two haptic updates (when data is sent to the haptic device), shown in Figure 3 in milliseconds, for a simulation ranging from 1,000 to 30,000 particles.

The performance improvement due to the decoupling of the haptic and visual threads is clearly visible for all the states of matter. The haptic timestep in the decoupled case corresponds to the haptic loop. For the coupled case, the haptic timestep corresponds to the program loop, hence inevitably including the visual rendering. This overhead in the coupled scheme is visible in the plot as a shift of the curves, with the visual rendering taking around 6ms to 10ms for the range of particles used in the measurements.

Overall, with a dual GPU implementation our multi-state haptic rendering approach can reach high frequencies, even with a time consuming visual rendering. For a body made of 10,000 particles, update frequencies are close to or over 1,000Hz for rigid bodies and fluids, and 400Hz for deformable bodies. As the frequency increases, the performance gain compared to the coupled scheme increases at a similar rate. This improved efficiency allows the design of complex scenes while preserving the update rates required for quality haptic feedback.

6 Evaluation

In this section we present an initial evaluation of the approach, through feedback measurements as well as a small user study assessing the capability of users to recognize the different states of matter they interact with.

The evaluations were carried out using a Virtuose 6DoF force-feedback device from Haption, and a computer with a Core 2 Extreme X7900 processor at 2.8GHz, 4GB of RAM memory, and two Nvidia GeForce 460 GT GPU with 1GB of graphic memory.

6.1 Haptic Feedback

In order to illustrate the typical haptic feedback pattern generated from the exploration of a multistate simulation, Figure 4 provides plots of force and torque feedback for each state of matter during a common interaction movement. The measurements were taken using a rigid body spoon as proxy coupled to the 6DoF haptic device, and a tray containing the media (Figure 6). A pre-recorded

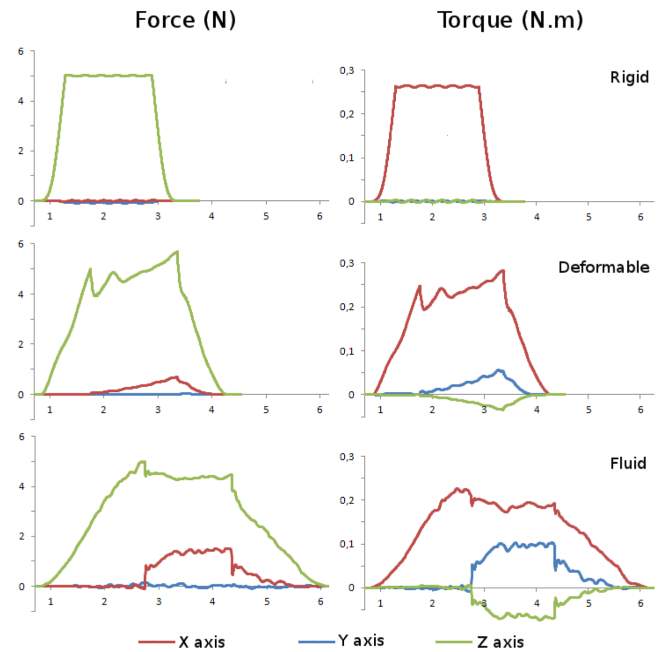

Figure 4: Plot of forces (in N) and torques (in N.m) sent to the haptic device during a pre-recorded motion of the spoon for each state of matter. X axis: horizontal. Y axis: vertical. Z axis: depth.

movement of the spoon was played for each state, consisting of a downwards motion until a force threshold of 5N was reached, followed by a lateral motion of 1.6 seconds, and by an upward motion until the spoon was out of contact with the media. We recorded the forces and torques sent to the haptic device for each state.

A first noticeable difference between states is the length in the downwards and upwards movements during contact with the media. For rigid bodies this period is very short (around 0.2s), being larger for deformable bodies (around 0.8s), and with the largest period for fluids (around 1.8s). Since the downwards movement stopped when the 5N force threshold was reached, the stiffer the media the faster the threshold is reached. We can thus observe how along the Y (vertical) axis this stiffness increases linearly with the spoon penetration depth, but at a different rate for each medium.

Another visible feedback difference between states is the force and torque generated during the lateral motion of the spoon. The force feedback along the X axis (horizontal) and the torque feedback along the Z axis (depth) are insignificant until the lateral motion is started. For the rigid state, the lateral motion does not produce any feedback along these axes since the contact is frictionless. For the deformable state, feedback increases linearly until the lateral motion stops, due to the decreasing distance to the fixed boundaries of the media, inducing a higher stress and the corresponding increase in elastic force. For the fluid state, there is a sudden increase in feedback due to the fluid surrounding the tip of the spoon. The feedback perceived by the user for a given motion is thus clearly different according to the type of media.

These different plots also allow to highlight some limitations of our approach. In the rigid state, when focusing on the force and torque feedback (Y and X axis respectively) during the lateral motion, we can observe a “bumpy” signal where a flat plateau would be expected. This noise is produced by the interpolation of values inherent to the SPH model over a discrete and sometimes coarse set of particles, as the probe follows a surface. Nevertheless, this noise remains small for reasonably strong forces, and subjects of the experiment never raised this issue during or after the trials. Another limitation is the slow propagation and bouncing of elastic waves

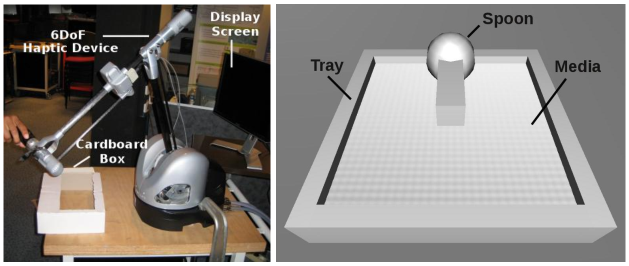

Figure 5: Experimental apparatus: display devices (left) and virtual scene (right) used in the experiment.

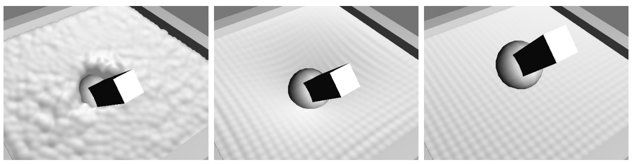

Figure 6: Left to right: fluid, deformable, and rigid states of matter.

within the material. The effect can be seen in the deformable state, and for the same movement and axes, as an oscillation in the signal once the downward motion has stopped. This behavior is expected, although not in the present form. The use of an explicit integration makes the elastic wave propagation much slower, and thus much more perceivable than what is should be. Smaller timesteps could reduce the impact of explicit integration.

6.2 User study

In order to qualitatively evaluate our haptic multi-state approach, we conducted a short user study where subjects interacted with different types of media. Our goal was to assess the capacity of subjects to recognize the state of the matter present in the virtual scene (fluid, deformable, rigid), in three different conditions: haptic feedback alone, visual feedback alone, and both at the same time. We recorded their answers on the state of matter they recognized.

Experimental Apparatus. The experiment was conducted with a Virtuose 6DoF force-feedback device for haptic interaction, and a 17' flat screen in front of the subject for visual display. A carved cardboard box was used to situate a virtual tray containing the media in the scene. The simulation was computed using the aforementioned configuration. Figure 5 shows the experimental conditions. The virtual scene consisted of a rigid body spoon coupled to the 6DoF haptic device positioned on top of a 60x60cm tray containing the media in either fluid, deformable or rigid state, as show in Figure 6. The media was ten-particle-layers thick. The entire scene was rendered with a gray texture.

Procedure. Before the beginning of the experiment, the subject was given the instructions in written form. Before each trial, the subject had to position the handle of the haptic device at the same starting position, instructed beforehand. The participant could explore the scene freely up to the cardboard limits. The subject had 10 seconds to detect the state of the media he was interacting with before the feedback (visual and haptic) was turned off. He would then select the answer from the 3 possible choices on the screen. The experiment lasted about 30 minutes.

Experimental Plan. Twelve participants (2 females and 10 males) aged from 23 to 26 (mean = 24.8, sd = 0.7), took part in this experiment. They were all naïve to its purpose. Participants completed all three feedback conditions (haptic, visual, and haptic+visual) and the order of the conditions was counterbalanced across participants. In each condition, the participants were exposed to 8 successive blocks of 3 trials corresponding to the three states, in random order. Participants completed a total of 72 trials (3 feedback conditions $\times$ 3 states of matter $\times$ 8 trials per state). For each trial and each subject, we recorded the answer given by the subject on the state of matter he recognized.

Table 1: Probabilities of correct answers for each state of matter according to the condition: visual (V), haptic (H) and haptic+visual (HV).

	Deformable	Fluid	Rigid
V	1	1	0.99
H	0.87	0.86	0.89
HV	1	1	1

Results and Discussion. We first evaluated the probability of correct answers for the different states of matter with one of the 3 conditions: haptic (H), visual (V) and haptic+visual (HV). It reveals that the probabilities of correct answers were the following: $p_H = 0.875$ for the haptic condition, $p_V = 0.99$ for the visual condition and $p_{HV} = 1$ for the haptic+visual condition. If we differentiate the different states of matter, we obtain the probabilities of correct answers shown in Table 1.

Overall, results suggest high recognition rates for all three conditions (H, V, HV), and for all three states of matter. In the context of the experiment, our haptic multi-state approach succeeds in simulating the three different states of matter to some extent. Results for V and HV indicate that recognizing the state of matter with visual feedback was a rather straightforward task, with almost perfect scores for both conditions, due to the choice of very representative conditions for each state of matter. These results suggest a realistic simulation in terms of physical model. However, the most interesting result of this experiment is the recognition rate for the Haptic only condition (H), which is around 87% in average. As this is the first approach allowing the haptic interaction with the three different states, no comparison can be made with previous studies. To the best of our knowledge, studies on the identification of the state of matter using the haptic modality alone have never been conducted, even for real materials. Nevertheless, one could suppose that, in real life, recognition would be very close to 100% for the three states when exploring with our bare hands. In our virtual setup, due to the use of a handle-based haptic device, we are exploring the material through a probe. Using a probe for exploration has been shown to significantly degrade material and shape cues gathered in the context of real object recognition tasks [16], with recognition of common objects dropping from 96% to 39%.

To assess the significance of the recognition rates, we performed likelihood ratio tests (LRT) to compare three statistical models: (1) pV $\neq$ pH, pV $\neq$ pHV , pH $\neq$ pHV ; (2) pH $\neq$ pV , pH $\neq$ pHV ,pV = pHV ; (3) pV = pH = pHV. LRT revealed that the second model was significantly more likely than the two other models. Thus, without taking into account the different states of matter (fluid, deformable or rigid), the probability of correct answers for H condition was significantly lower than the ones for V and HV conditions (p-value= $2.21.10^{-13}$). The probability of correct answers for HV was not significantly different from the V condition (p-value= 0.24).

Results are not significant in terms of state recognition when comparing both conditions with visual feedback since, again, visual cues seem to be enough to convey the state of matter to the users. Hence, the addition of haptic feedback did not significantly contribute to the recognition of the type of matter users were interacting with, when the visual component was present. However, subjective feedback from the participants indicated that the addition of haptic feedback on top of visual rendering increased both the degree of realism and the general appreciation of users. Their feedback was very positive regarding the combination of both modalities, and is in accordance with previous studies evaluating the combination of haptics with other modalities [26].

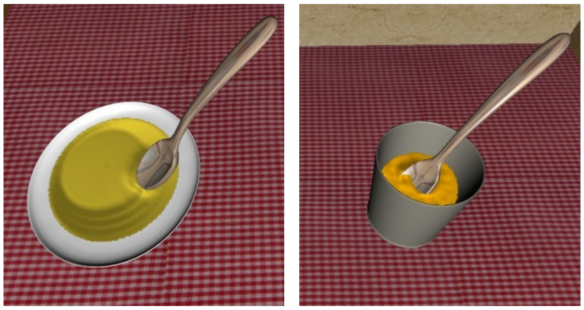

Figure 7: Scenario illustrating a cooking simulator. The user, through a spoon coupled to the 6DoF haptic device, can touch and play with a custard tart, a glass and its liquid content. Haptic feedback is seamlessly computed for the three states of matter simultaneously.

7 Applications and Conclusion

Potential applications of our technique span from medical training to industrial and entertainment simulations. In this scope, we designed a virtual kitchen scenario with familiar objects and materials. Users can freely explore the environment and better perceive each object through the sense of touch. Figure 7 shows how a user interacts with a solid glass full of thick orange juice and a wobbly custard dessert, using a spoon coupled to the 6DoF haptic device, while everything follows a physically based behavior. Figure 1 shows a virtual crepe making simulator, which illustrates state changes as the one in Figure 2. These scenarios can be seen in the accompanying video.

In this paper, we presented a simulation framework based on Smoothed-Particle Hydrodynamics that enables 6DoF haptic interaction simultaneously with fluid, elastic and rigid bodies, as well as their melting or freezing. Haptic rendering is achieved through a unified mechanism, allowing the seamless interaction with bodies of different state as well as the use of rigid, deformable and state changing proxies. Combined to an efficient dual GPU implementation, we enable a richer haptic interaction with complex environments and tools. An initial evaluation of the approach showed the importance of the GPU decoupling scheme, and suggested that haptic feedback computed within our framework can effectively convey material state cues to the user, even without visual rendering.

Future work will focus on extending the range of haptic sensations than can be simulated, such as friction, plasticity and tearing. We would also like to conduct further evaluations to fully assess the rendering quality perceived by users, particularly when using different viscosity and stiffness values for fluid and deformable states, and when undergoing changes of state.

References

[1] O. Astley and V. Hayward. Multirate haptic simulation achieved by coupling finite element meshes through norton equivalents. In Proceedings of IEEE Robotics and Automation, 1998.

[2] J. Barbič and D. L. James. Six-DoF haptic rendering of contact between geometrically complex reduced deformable models. IEEE Transactions on Haptics, 1(1):39–52, 2008.

[3] W. Baxter and M. C. Lin. Haptic interaction with fluid media. In Proceedings of Graphics Interface, 2004.

[4] M. Becker, M. Ihmsen, and M. Teschner. Corotated SPH for deformable solids. In Proceedings of the Eurographics Workshop on Natural Phenomena, 2009.

[5] M. Becker, H. Tessendorf, and M. Teschner. Direct forcing for lagrangian Rigid-Fluid coupling. IEEE Transactions on Visualization and Computer Graphics, 15(3):493–503, 2009.

[6] C. Cadoz, A. Luciani, and J. L. Florens. CORDIS-ANIMA: modeling and simulation system for sound and image synthesis - the general formalism. Computer Music Journal, 17(1), 1993.

[7] G. Cirio, M. Marchal, S. Hillaire, and A. Lecuyer. Six Degrees-of-Freedom haptic interaction with fluids. IEEE Transactions on Visualization and Computer Graphics, 17(11):1714–1727, 2011.

[8] G. Cirio, M. Marchal, A. Le Gentil, and A. Lecuyer. Tap, squeeze and stir the virtual world: Touching the different states of matter through 6DoF haptic interaction. In Proceedings of IEEE Virtual Reality Short Papers, 2011.

[9] H. Courtecuisse, H. Jung, J. Allard, C. Duriez, D. Y. Lee, and S. Cotin. GPU-based real-time soft tissue deformation with cutting and haptic feedback. Progress in Biophysics and Molecular Biology, 103(2-3):159–168, 2010.

[10] H. Delingette, S. Cotin, and N. Ayache. A hybrid elastic model allowing real-time cutting, deformations and force-feedback for surgery training and simulation. In Proceedings of Computer Animation, 1999.

[11] Y. Dobashi, M. Sato, S. Hasegawa, T. Yamamoto, M. Kato, and T. Nishita. A fluid resistance map method for real-time haptic interaction with fluids. In Proceedings of ACM Virtual reality software and technology, 2006.

[12] C. Duriez, F. Dubois, A. Kheddar, and C. Andriot. Realistic haptic rendering of interacting deformable objects in virtual environments. IEEE Transactions on Visualization and Computer Graphics, 12(1):36–47, 2006.

[13] C. Forest, H. Delingette, and N. Ayache. Surface contact and reaction force models for laparoscopic simulation. Lecture Notes in Computer Science, 3078:168–176, 2004.

[14] C. Garre and M. A. Otaduy. Haptic rendering of complex deformations through handle-space force linearization. In Proceedings of IEEE World Haptics, 2009.

[15] R. Keiser, B. Adams, D. Gasser, P. Bazzi, P. Dutre, and M. Gross. A unified lagrangian approach to solid-fluid animation. In Proceedings of Eurographics/IEEE Point-Based Graphics, 2005.

[16] S. J. Lederman and R. L. Klatzky. Haptic identification of common objects: effects of constraining the manual exploration process. Perception & Psychophysics, 66(4):618–628, 2004.

[17] M. Lin and M. A. Otaduy. Haptic Rendering: Foundations, Algorithms and Applications. A K Peters, Illustrated Edition, 2008.

[18] W. A. McNeely, K. D. Puterbaugh, and J. J. Troy. Six degree-of-freedom haptic rendering using voxel sampling. In Proceedings of ACM SIGGRAPH, 1999.

[19] J. J. Monaghan. Smoothed particle hydrodynamics. Annual Review of Astronomy and Astrophysics, 30(1):543–574, 1992.

[20] M. Muller, D. Charypar, and M. Gross. Particle-based fluid simulation for interactive applications. In Proceedings of ACM SIGGRAPH/Eurographics Symposium on Computer Animation, 2003.

[21] M. Muller, S. Schirm, M. Teschner, B. Heidelberger, and M. Gross. Interaction of fluids with deformable solids. Computer Animation and Virtual Worlds, 15(3-4):159–171, 2004.

[22] M. Muller, J. Stam, D. James, and N. Thürey. Real time physics. In ACM SIGGRAPH Courses, 2008.

[23] M. Ortega, S. Redon, and S. Coquillart. A six Degree-of-Freedom God-Object method for haptic display of rigid bodies with surface properties. IEEE Transactions on Visualization and Computer Graphics, 13(3):458–469, 2007.

[24] M. A. Otaduy and M. C. Lin. Sensation preserving simplification for haptic rendering. In Proceedings of ACM SIGGRAPH, 2003.

[25] I. Peterlík, M. Sedef, C. Basdogan, and L. Matyska. Real-time visio-haptic interaction with static soft tissue models having geometric and material nonlinearity. Computers & Graphics, 34(1):43–54, 2010.

[26] P. Richard, G. Birebent, P. Coiffet, G. C. Burdea, D. Gomez, and N. A. Langrana. Effect of frame rate and force feedback on virtual object manipulation. Presence, 5(1):95–108, 1996.

[27] B. Solenthaler, J. Schläfli, and R. Pajarola. A unified particle model for fluid-solid interactions. Comput. Animat. Virtual Worlds, 18(1):69–82, 2007.

[28] M. Yang, J. Lu, A. Safonova, and K. J. Kuchenbecker. GPU methods for real-time haptic interaction with 3D fluids. In IEEE Haptic Audio visual Environments and Games, 2009.