An Intestinal Surgery Simulator: Real-Time Collision Processing and Visualization

Laks Raghupathi, Laurent Grisoni, François Faure, Damien Marchall,
Marie-Paule Cani, Christophe Chaillou

► To cite this version:

Laks Raghupathi, Laurent Grisoni, François Faure, Damien Marchall, Marie-Paule Cani, et al.. An Intestinal Surgery Simulator: Real-Time Collision Processing and Visualization. IEEE Transactions on Visualization and Computer Graphics, 2004, 10 (6), pp.708–718. 10.1109/TVCG.2004.36 . inria-00537466

HAL Id: inria-00537466

<https://hal.inria.fr/inria-00537466>

Submitted on 18 Nov 2010

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

An Intestinal Surgery Simulator: Real-Time Collision Processing and Visualization

Laks Raghupathi, Laurent Grisoni, François Faure,
Damien Marchal, Marie-Paule Cani, and Christophe Chaillou

Abstract—This research work is aimed toward the development of a VR-based trainer for colon cancer removal. It enables the surgeons to interactively view and manipulate the concerned virtual organs as during a real surgery. First, we present a method for animating the small intestine and the mesentery (the tissue that connects it to the main vessels) in real-time, thus enabling user interaction through virtual surgical tools during the simulation. We present a stochastic approach for fast collision detection in highly deformable, self-colliding objects. A simple and efficient response to collisions is also introduced in order to reduce the overall animation complexity. Second, we describe a new method based on generalized cylinders for fast rendering of the intestine. An efficient curvature detection method, along with an adaptive sampling algorithm, is presented. This approach, while providing improved tessellation without the classical self-intersection problem, also allows for high-performance rendering thanks to the new 3D skinning feature available in recent GPUs. The rendering algorithm is also designed to ensure a guaranteed frame rate. Finally, we present the quantitative results of the simulations and describe the qualitative feedback obtained from the surgeons.

Index Terms—Virtual reality, physically-based modeling, animation, curve and surface representation.

1 INTRODUCTION

MINIMALLY invasive surgical (MIS) procedures are gaining popularity over open procedures among surgeons and patients. This is mainly due to less post-operative pain, fewer infections, and an overall faster recovery. The tremendous success of laparoscopic cholecystectomy (gall bladder removal) has prompted surgical practitioners and educators to apply such techniques to other gastrointestinal procedures. In this paper, we focus on *laparoscopic colectomy* (colon cancer removal). Studies show that many patients undergoing this procedure benefit from the advantages of the MIS procedures listed above while sharing the same risks of the corresponding open procedure [1]. Yet, as with most laparoscopic procedures, it is difficult to master, with a very flat learning curve [2]. As part of the current training procedure, surgeons practice on pigs to get a feel for the organ's behavior. However, this technique is prohibitively expensive and also raises numerous ethical issues. We believe that a VR-based simulator platform can significantly help nonspecialist surgeons and medical residents to acquire the necessary surgical skills in a cost-effective way. This may well result in popularizing the use of the laparoscopic technique for this procedure, thus benefiting more patients. Thus, our aim is to simulate the behavior of the intestine when the surgeon is practicing in the virtual environment. Note that the current scope of this research work does not include the simulation of the cancer

removal itself. We focus on two important pedagogical problems: 1) *camera positioning* by allowing the trainee to visualize the relevant organs in 3D, 2) *manual dexterity* by letting them interactively manipulate these organs. For many surgeons who are trained primarily in open techniques, this may help to overcome the perceptual and motor challenges associated with MIS procedures.

We will first review the background of the problem and highlight the challenges posed. During this surgical procedure, the patient is lying on his back. As a result, the small intestine (henceforth simply referred to as *intestine*) is positioned just above the colon region, hiding the colon beneath (Fig. 1). The intestinal region of a human body is characterized by a very complex anatomy. The intestine is a tubular structure, about 4 m long, constrained within a small space of the abdominal cavity, resulting in the creation of numerous *folds*. This is further complicated by a tissue known as the *mesentery* that connects the intestine to the blood vessels [3] (Fig. 2). An important surgical task of this procedure is to displace the tissues and organs by pulling and folding them from the site of the operation [4]. As the surgeon manipulates these organs, they deform and collide with each other. Thus, the broad challenges here are the real-time animation and visualization of these organs at an acceptable frame rate (a minimum of 25 frames a second). Our overall approach to solving this problem consists of a layered model: a skeletal axis deformed using physically-based animation, rendered with a generalized cylinder-based skinning. Thus, in order to animate these organs in real-time, we propose to:

- efficiently model the intestine and the mesentery, taking into account its geometry,
- detect the collisions and self-collisions occurring in the intestinal region during animation,
- provide a fast and stable collision response.

• L. Raghupathi, F. Faure, and M.-P. Cani are with the GRAVIR/IMAG lab, 655 ave. de l'Europe, 38334 Montbonnet, France.

E-mail: {laks, francois.faure, marie-paule.cani}@imag.fr.

• L. Grisoni, D. Marchal, and C. Chaillou are with the LIFL lab, University of Lille 1, 59655 Villeneuve d'Ascq, France.

E-mail: {laurent.grisoni, marchal, christophe.chaillou}@lifl.fr.

Manuscript received 14 Nov. 2003; revised 19 Feb. 2004; accepted 26 Feb. 2004.

For information on obtaining reprints of this article, please send e-mail to: tvcg@computer.org, and reference IEEECS Log Number TVCG-0108-1103.

Fig. 1. Position of the intestine during the surgery.

The skeletal model used for animation should be covered with a triangulated mesh which can be realistically shaded or textured. However, a naive skinning approach will create tessellation problems in high-curvature regions. Hence, we present:

- an efficient method to detect high-curvature regions and to adapt the sampling of the axis accordingly,
- an algorithm that dynamically changes the skinning parameters to maintain a preimposed frame rate,
- a hardware-based skinning feature for fast and smooth rendering available in recent GPUs.

All these contributions have been implemented within a surgical simulator platform [5] that can be easily evaluated by the surgeons. Note that the first versions of our contribution on intestine animation and rendering appeared as separate conference proceedings [6], [7]. This paper gives an extended presentation of both, describes their integration into a simulator platform, provides the results of the integration, and summarizes the evaluation of the results, both qualitative and quantitative.

The remainder of this paper is organized as follows: In Section 2, we briefly describe the previous work related to our problem, from both the animation and rendering perspectives. We then elaborate on creating a virtual intestine and mesentery, an initial geometric model defining their shape, and a mechanical model for animation in Section 3. Section 4 presents our collision detection and response algorithms in detail. We then propose a method for fast rendering, along with a self-tuning algorithm for guaranteed frame rate in Section 5. Section 6 briefly describes the integration into a surgical simulator platform. The results and validation are discussed in Section 7, followed by the conclusions in Section 8.

2 PREVIOUS WORK

2.1 Real-Time Animation of Deformable Models

Recently, many researchers have focused on the efficient simulation of deformable models, with some of them using adaptive, multiresolution techniques [8], [9], [10] and some applied to surgery simulators [11], [12], [13], [14], [15], [16]. In all these works, volumetric deformable bodies were either simulated in isolation or were interacting with a single rigid tool, enabling the use of very specific techniques for collision detection and response, such as methods based on graphics hardware [17]. The problem we have to solve here is different. As will be shown in the next section, no volumetric deformable model will be needed since the

Fig. 2. Anatomy of the intestinal region.

intestine and the mesentery can be represented as a 1D and a 2D structure, respectively. Moreover, the main issue is detecting and processing the collisions and self-collisions of the intestinal system in real-time. Accordingly, a simple spline model animated by mass-spring dynamics was used by France et al. [18], [19] for simulating the intestine. France et al. used a grid-based approach for detecting self-collisions of the intestine and collisions with its environment. All objects were first approximated by bounding spheres, whose positions were stored, at each time step, in the 3D grid. Each time a sphere was inserted into a nonempty voxel, new collisions were checked within this voxel. Though this method achieved real-time performances when the intestine alone was used, it did not handle the simulation of the mesentery.

A well-known technique for accelerating collision detection is to approximate the objects by a bounding volume hierarchy (BVH) [20], [21], [22], [23], [24], [25]. Such approaches provide simple tests for fast detection of noncolliding regions. The BVH can be recursively updated when the objects undergo small deformations. However, this is not suitable for intestine-mesentery interaction, where even a small local deformation can potentially cause a large movement of the folds. This creates a large scale global deformation, which prevents the BVH from being efficiently updated. An alternate multiresolution method, based on layered shells, was recently presented by Guy and Debunne [26]. It is well-suited for collision detection between deformable objects since the shells themselves are deformable structures extracted from a multiresolution representation of these objects. Though suitable for volumetric deformable bodies, this method is not appropriate for intestine and mesentery modeling since the time-varying folds cannot easily be approximated at a coarse scale. They also exploited temporal coherence following Lin and Canny's [27] idea of detecting collisions between convex polyhedra by tracking pairs of closest vertices. These pairs were efficiently updated at each time step by propagating closest distance tests from a vertex to its neighbors. Guy and Debunne adapted this technique for detecting collisions between their volumetric layered shells very efficiently. Since these shells were neither convex nor rigid, a stochastic approach was used at each time step to

generate new pairs of points anywhere on the two approaching objects. These pairs were made to converge to the local minima of the distance, disappearing when they reached an already detected minimum. Our work is inspired by this idea of *stochastic* collision detection exploiting temporal coherence. It has been adapted here to the specific processing of multiple collisions and contacts between the intestine and the mesentery folds.

2.2 Intestine and Mesentery Rendering

In addition to handling the collisions, we also need to provide a real-time, smooth rendering of the organs. The intestine and the mesentery are geometrically complex organs and also topologically distinct from each other: a deformable cylinder-like object and a deformable, nondevelopable tissue. Thus, it is natural to study their rendering separately. Many efficient techniques do exist for surfaces (a simple tessellation actually gives good results). Hence, we shall devote more preference to dealing with the much harder problem of intestine rendering.

A possible approach to intestine rendering is to use an implicit representation [18]. This is indeed a powerful representation, especially for handling topologically complex or deformable objects. Yet, it is not well-suited for our purpose since none of its classical features are useful here. Moreover, the blending property inherent to implicit surfaces implies that some control operations would have to be performed on the implicit intestine model. This additional processing is necessary in order to avoid unwanted blending between the intestinal folds.

A more natural approach to intestine modeling is to use generalized cylinders, also sometimes called *sweep*. To our knowledge, only a very few previous techniques dealt with the efficient rendering of such objects. Classical generalized cylinders are defined using a parameterized axis $C(u)$, along with a set of planar cross-section that can be represented as a continuous set $S(u)$ depending on the same u scalar value that parameterizes the axis [28]. The classical tessellation approach is as follows: The axis C is sampled either uniformly or adaptively, depending on some curvature isolation (see [7] for an extensive overview of the possible techniques). For each axis sample, section S is positioned on a local frame calculated along the axis [29] and then approximated by a polyline. Two consecutive polyline sections along the axis are then connected by a triangle stripset in order to form a local approximation of the sweeps. High curvature point detection becomes a critical step for such a process and any sharp angle is likely to be missed in the sampling. It is also likely that the tessellated sections can overlap, hence providing the user with a self-intersecting mesh (see Fig. 3).

Our solution, rather, comes from skinning and thus takes advantage of the OpenGL skinning feature provided by recent GPUs [30]. Skinning (also called *vertex blending*) is most commonly used in character deformation applications that allow continuous deformation of a skin mesh over an animated skeleton. The principle of this extension is as follows: It is a standard practice to define a simple geometric transformation in homogeneous coordinates by a 4×4 matrix. This composition of rotation, translation, and scaling allows an initial mesh to be positioned anywhere within a given scene, along with simple deformations. For an object, skinning uses two matrices (i.e., two different

Fig. 3. Classical tessellation problems on generalized cylinders. Top: Inadequate sweep tessellation (a) and correct one (b). Bottom: Mesh self-intersection when deforming the initial sweep.

positions of the object) and, for all the vertices of the object, a scalar value (called *weight*) is used to interpolate the vertex position. Accordingly, given an initial vertex ν^0 , with weight $\omega(\nu^0)$ and two transformations M_1 and M_2 , the resulting position ν of the vertex is given by:

$$\nu = \omega(\nu^0)M_1\nu^0 + [1 - \omega(\nu^0)]M_2\nu^0. \quad (1)$$

The above transformation can be done entirely by the GPU with no CPU computation involved. Setting optimal weights for a given deformation is, in general, a difficult problem. Recently, Bloomenthal presented a technique for automatically setting weights when skinning is used in character animation [31].

The above transformation process allows for simple deformation of the initial mesh. Fig. 4 shows a simple 2D example of what is obtained using an initial rectangle deformation with weights that only depend on the z position along its axis. In this figure, the weight function is similar to the one described for other purposes in [32]. Controlling the weight distribution curve determines the way the GPU will interpolate from one transformation to the other. Indeed, this technique cannot be directly applied to the rendering of a folded generalized cylinder controlled by a skeleton, the intestine in our case. Extending the skinning method to this case will be the main contribution of Section 5.

3 MODELING AND ANIMATING THE SYSTEM

Our overall aim is not to develop an accurate, patient-specific trainer but rather a generic simulator which can help the surgeons to practice the gestures used to manipulate the organs. The first problem we have to solve is to create a virtual model of the intestine and the mesentery which will serve as the basis for the animation and rendering stages. In this section, we describe the actual anatomy in brief, provide the basis of our approach, and present the details of our model.

Fig. 4. Simple skinning application to a polygonal primitive. Combination of M_1 and M_2 transformations on the object (top-left), using the weights (bottom-left), gives the results (right).

We further present the mechanical model needed for animating the intestinal system.

3.1 Anatomical Model

In order to extract the anatomy of the intestine and the mesentery, we sought the help of our medical collaborator, IRCAD, in Strasbourg (a digestive cancer research institute [33]). With the current imagery techniques, they found it impossible to extract the complex anatomy of organs such as the mesentery. Hence, we decided to come up with something simpler than the actual geometry, but which can still capture the overall behavior.

The mesentery is a folded membrane-like surface, approximately 15 cm wide, which links the intestine to the main vessels of 10 cm length (Fig. 2). Since the mesentery cannot be unfolded onto a plane, setting up its initial geometry free of self-intersections is quite difficult. We solved the problem by modeling a possible rest position for the intestine as a folded sine curve lying on the inner surface of a cylinder of radius 15 cm. With the axis of the cylinder representing the main vessel, the folds are placed on the cylinder such that their total length is 4 m (Fig. 5). Then, the mesentery can be defined as the surface generated by the set of nonintersecting line segments linking the cylinder axis to the curve. Though this initial geometry is too symmetric to be realistic, this model gives adequate local geometric properties to the mesentery membrane (Fig. 6). When animated under the effect of gravity, these collision-free initial positions will automatically move to their correct positions. The geometry of the intestine is defined by creating a piecewise tubular surface of radius 2 cm along its skeleton curve. The thickness of the mesentery surface, parameterized based on patient-specific data, was set to 1 cm.

3.2 Mechanical Considerations for Animation

For animation, our motivation to use a mass-spring approach was derived from the following arguments:

- The mass-spring method is more efficient and suitable over FEM for deforming bodies with large displacements and local elastic deformations.
- In addition, an organ such as the intestine can have an infinite number of rest states, whereas FEM is

Fig. 5. Initialization of the geometry of the intestine and mesentery. The intestine is drawn on the inner surface of a small cylinder. The figure is greatly simplified for clarity.

based on the notion of displacements w.r.t. a unique rest state thus making it unsuitable for our case.

- Mass-spring can provide stable simulation of deformable objects at moderate time steps.
- We can adjust the behavior of the system intuitively by adjusting the damping, stiffness, etc.
- Collision detection here is a much more complex task requiring more CPU-time over animation. So, a very complex mechanical model might slow down the simulation.
- Finally, the *perceived quality* of most interactive 3D applications does not depend so much on exact simulation, but rather on real-time response to collisions [34].

Accordingly, we designed a model consisting of mass points connected by damped springs. Since the mesentery has a much larger length (4 m near the intestine) than width (15 cm near the vessel), we sampled our model by four sets of 100 masses each (Fig. 6). The last set of masses requires no computation since they are attached to the main vessels, requiring only 300 masses to be integrated at each time step. In addition, no specific model is needed for the intestine since it can be simulated by adjusting the mass and stiffness values along the first bordering curve of the mesentery surface.

Fig. 6. Mechanical model of the organs (shown unfolded).

Fig. 7. Tracking the local minima of distance between nonneighboring segments.

4 REAL-TIME COLLISION PROCESSING

A major computational bottleneck in many animation/simulation systems is the handling of collisions between the objects under the influence of external forces (gravity, user-input, etc.). In our case of real-time simulation, we need to perform this as quickly as possible. Failure to detect the collisions will result in interpenetration—an unrealistic behavior. In this section, we describe our approach, which efficiently detects the colliding regions, and a response algorithm, which prevents/corrects the interpenetrations.

4.1 Collision Detection

Our method for real-time collision detection exploits *temporal coherence* [26], [27] (i.e., pairs of closest points between colliding bodies are tracked over time). The main differences here are: 1) the interacting objects have a tubular (intestine) or a membrane-like structure (mesentery) and 2) most collisions will be self-collisions between different folds of the same body. We first explain the collision detection method for the intestine alone and then explain the mesentery case.

Collision detection between cylinders can be processed by computing the closest distance between their axes [35], and comparing them to the sum of their radii. For the intestine, computing the distance between two segments is done by considering the distance between their principal axes (a segment here refers to a simple line segment and its end-points are parameterized by $(s, t) \in [0, 1]$). We then store the (s, t) value corresponding to the closest point within the segments and the actual minimum distance $d_{\min}$.

Adapting the notion of “closest element pairs” to this skeleton curve means that we track the local minima of the distance between nonneighboring segments along the curve (Fig. 7). Of course, only the local minima satisfying a given distance threshold are considered relevant. We refer to these pairs of segments as *active pairs*. Each active pair is locally updated at each time step in order to track the local minima when the intestine folds move. This is done by checking whether it is the current segment pair or a pair formed using one of their neighbors now corresponds to the smallest distance. This update requires *nine* distance tests (Fig. 8) and the pair of segments associated with the closest distance becomes the new active pair. When two initially distant active pairs converge to the same local minimum, one of them is suppressed. A pair is also suppressed if the associated distance is greater than a given threshold.

The above process tracks the existing regions of interest, but does not detect new ones. Since the animation of the intestine may create new collisions, a method for creating

Fig. 8. Distance computation between two intestine segments.

new active pairs of segments is needed. Our approach is inspired by the stochastic approach of [26]. At each time step, in addition to the update of the currently active pairs, n additional random pairs of segments, uniformly distributed between the end-points, but under the distance threshold, are generated. The update of these extra active pairs is similar to the update of the existing local minima. The complexity of the detection process thus linearly varies with the user-defined parameter n . At each time step, collision detection consists of selecting, from among the currently active pairs, the pairs of segments which are closer than the sum of their radii. Collision response (cf. Section 4.2), will then be applied between these segments.

For the mesentery, the total number of segments to be considered during each time step is too large for real-time computation. We use the following approximations to reduce the complexity of the problem: First, since the mesentery is very thin and soft compared to the intestine, its self-collisions will have almost no effect on the overall behavior of the system. Hence, we ignore the testing of these and only consider the nontrivial cases of intestine-intestine and intestine-mesentery interactions.

Second, we use a two-step convergence algorithm to reduce the number of distance computations required for the mesentery. Accordingly, given a segment pair (S_1, S_2) , we first find if there exists another segment S'_1 that is the closest to S_2 (S'_1 is S_1 if all other neighbors are further from S_2). Then, we find a segment S'_2 that is the closest to S'_1 . This update requires 13 distance computations at most (i.e., between an intestine segment and a nonneighboring mesentery segment). When a collision is detected, we apply a response force not only to the deepest penetrating segment-pair but also to the entire collision area (both for intestine and mesentery). A recursive algorithm searches the neighbors to find all the colliding pairs in the area.

4.2 Collision Response

We initiate the response whenever the distance between the two segments is less than the sum of their radii. The earlier approaches, such as the penalty method [36] and the reaction constraint method [37], implemented collision response by altering the force matrix in the mass-spring method. This force has to be of large magnitude in order to be effective in large time step scenarios. However, this may cause segment displacements several times larger than their thickness, thus creating new collisions and instabilities. Instead, our new method alters the displacements and velocities such that it instantaneously cancels the interpenetration while keeping a resting contact between the two colliding bodies with no bouncing effects.

Fig. 9. Collision response by displacement-velocity correction.

Let the end-point velocities of segment S_1 be $\mathbf{v}_1$ and $\mathbf{v}'_1$ and that of segment S_2 be $\mathbf{v}_2$ and $\mathbf{v}'_2$, respectively. Let $\mathbf{x}_1$, $\mathbf{x}'_1$, $\mathbf{x}_2$, and $\mathbf{x}'_2$ be the corresponding end-point positions. Let $\mathbf{v}_{c1}$ and $\mathbf{v}_{c2}$ be the velocities of the closest approaching points within each segment and $\mathbf{x}_{c1}$ and $\mathbf{x}_{c2}$ be the positions of the closest points (Fig. 9). Let $\bar{s} = 1 - s$ and $\bar{t} = 1 - t$. We have:

$$\mathbf{v}_{c1} = \bar{s}\mathbf{v}_1 + s\mathbf{v}'_1 \quad \mathbf{v}_{c2} = \bar{t}\mathbf{v}_2 + t\mathbf{v}'_2. \quad (2)$$

Let two impulses per time step, f and $f' (= -f)$ (one for each segment), be applied along the direction of collision $\mathbf{u}$ to cause a velocity change such that the relative velocities along $\mathbf{u}$ is zero. These impulses should set the new velocities $\mathbf{v}_{new1}$ and $\mathbf{v}_{new2}$ such that:¹

$$(\mathbf{v}_{new1} - \mathbf{v}_{new2}) \cdot \mathbf{u} = 0. \quad (3)$$

This cancels the penetration velocity and avoids any bouncing. The impulse f acting on the point of collision can be split between the end-points according to their barycentric coordinates. Then, we have:

$$\begin{aligned} \mathbf{v}_{new1} &= \mathbf{v}_1 + \frac{\bar{s}f}{m_1} \mathbf{u} & \mathbf{v}'_{new1} &= \mathbf{v}'_1 + \frac{sf}{m'_1} \mathbf{u} \\ \mathbf{v}_{new2} &= \mathbf{v}_2 + \frac{\bar{t}f}{m_2} \mathbf{u} & \mathbf{v}'_{new2} &= \mathbf{v}'_2 + \frac{tf}{m'_2} \mathbf{u}, \end{aligned} \quad (4)$$

where m_1 , m'_1 , m_2 , and m'_2 are the masses of the end-points. Again, expressing the new velocity of the colliding point $\mathbf{v}_{new1}$ in terms of $\mathbf{v}_{new1}$ and $\mathbf{v}'_{new1}$:

$$\begin{aligned} \mathbf{v}_{new1} &= \bar{s}\mathbf{v}_{new1} + s\mathbf{v}'_{new1} \\ &= \mathbf{v}_{c1} + \left(\frac{\bar{s}^2}{m_1} + \frac{s^2}{m'_1} \right) f \mathbf{u}. \end{aligned} \quad (5)$$

Similarly, for segment S_2 :

$$\mathbf{v}_{new2} = \mathbf{v}_{c2} - \left(\frac{\bar{t}^2}{m_2} + \frac{t^2}{m'_2} \right) f \mathbf{u}. \quad (6)$$

Substituting (5) and (6) into (3), we have:

$$f = \frac{(\mathbf{v}_{c2} - \mathbf{v}_{c1}) \cdot \mathbf{u}}{\frac{\bar{s}^2}{m_1} + \frac{s^2}{m'_1} + \frac{\bar{t}^2}{m_2} + \frac{t^2}{m'_2}}. \quad (7)$$

Using this value of f , we compute the new velocities of the end-points from (4). We use a similar formulation for

correcting the positions of the colliding segments. The only difference is in the condition for avoiding interpenetration, which considers the segment radii r_1 and r_2 :

$$(\mathbf{x}_{new1} - \mathbf{x}_{new2}) \cdot \mathbf{u} = r_1 + r_2. \quad (8)$$

The integrated impulse value g that changes the positions in accordance with the above condition is then:

$$g = \frac{(\mathbf{x}_{c1} - \mathbf{x}_{c2}) \cdot \mathbf{u} + r_1 + r_2}{\frac{\bar{s}^2}{m_1} + \frac{s^2}{m'_1} + \frac{\bar{t}^2}{m_2} + \frac{t^2}{m'_2}}. \quad (9)$$

g is used to modify the end-point positions as in (4).

We note that updating the position and the velocity of one edge may still create or cancel collisions among other segments. A possible solution is to repeat the collision check for all the pairs in our list in order to identify such pairs. But, this may result in an endless loop. So, in our implementation, we prefer using a fixed number of iterations, even if we miss some collisions during the current time step. Handling multiple collisions in general is a difficult problem even for rigid bodies, with no straightforward solution. Some approaches based on a global solution (simultaneous handling of all the collisions) have been proposed. Giang et al. [38] constructed this as a constrained optimization problem and used an LCP-based (Linear Complementarity Problem) solution. Needless to say, it is much more difficult for flexible bodies with real-time constraints and, to our knowledge, no such techniques exist for our case.

5 SKINNING AND RENDERING

In this section, we present our approach for intestine and mesentery rendering. We shall deal with them separately as explained before. Medical experience showed that it is very hard to perceive how the mesentery actually deforms. However, we observed that this surface most often lies behind the intestine during surgical manipulations. Thus, the main purpose of simulating the mesentery is to ensure proper mechanical behavior of the intestine. As a result, we chose a very coarse rendering representation for the mesentery. The rendering primitive is simply a triangle stripset whose points are the intestine mechanical sampling points and the mesentery root. Using two separate models for mechanics and rendering did not attract any negative comments from the surgeons during demonstration, which, in a way, vindicates our presumption. We now describe the approach we used for efficiently rendering the intestine.

5.1 High-Curvature Detection and Adaptive Sampling

We first convert the mechanical sampling points of the intestine axis into a cardinal spline approximation (an *interpolation* spline), which is a straightforward process [39]. This is done in order to counterbalance the piecewise linearity of the mechanical model. The underlying idea behind the sampling process is fairly simple. Many interesting spline models have a regularity property [40]. This analytic property has among its consequences that no undesired oscillation appears in a given spline segment. In other words, a curve's overall shape can be somewhat predetermined by studying its control point configuration. Although we used cardinal

1. Where $\cdot$ denotes a vector dot product.

Fig. 10. Spline segment coloration to illustrate our curvature isolation process, during interactive sweep deformation.

splines in our experiments, most of these techniques can be easily generalized to other cases [7].

The sampling process uses the following input variables: the control point sequence $(P_0, P_1, \dots, P_n)$ of C , the corresponding key parameter values $(u_0, u_1, \dots, u_n)$ (uniformly sampled with $u_k = k\delta$, where δ is a constant value), and some threshold value ε . Recall that these control points are generated from the mechanical sampling points using spline approximation. The sampling algorithm is split into two steps:

Step 1: A first pass calculates the indicators of the local spline curvature. The spline regularity ensures that such a perturbation can be studied by considering only the control point configuration. A sequence of scalar values $\{c_i\}_{i=0\dots n}$ is defined using:

$$c_i = \frac{1}{2} \left(\frac{\mathbf{P}_{i-1,i} \cdot \mathbf{P}_{i,i+1}}{\|\mathbf{P}_{i-1,i}\| \|\mathbf{P}_{i,i+1}\|} + 1 \right), \quad (10)$$

where $\mathbf{P}_{i,j} = P_j - P_i$. Extremal values c_0 and c_n are treated by considering two additional ghost sampling points, generated using the Bessel technique [39], i.e., the ghost point at u_{-1} (similarly u_{n+1}) is calculated using the unique parabola that interpolates P_0, P_1, P_2 (similarly P_n, P_{n-1}, P_{n-2}). Note that c_i , which measures the local control point curvature, is proportional to the dot product (cosine of the angle). Hence, a higher curvature will result in a smaller angle between the segments and, consequently, a higher c_i value (10). Fig. 10 shows a visual evaluation of this step, associating different colors to each segment, depending on the numerical values of c_i . Here, the darker regions correspond to those segments that the algorithm selects as high curvature ones.

Step 2: A second pass further processes the segments that potentially contain high-curvature points (only one such point is possible per spline segment, once again thanks to regularity). For each parameter interval $[u_i, u_{i+1}]$, if $|c_i| < \varepsilon$ and $|c_{i+1}| < \varepsilon$, then the spline segment $[u_i, u_{i+1}]$ is treated as a whole by the display process. Otherwise, the corresponding spline segment potentially contains a high-curvature point which can be detected by symbolic root extraction. In the case of uniform cardinal splines, the maximal curvature parameter point m_i on a segment satisfies a linear equation. Using this, the parameter segments that are actually drawn are $[u_i, m_i - \delta_i]$, $[m_i - \delta_i, m_i + \delta'_i]$, and $[m_i + \delta'_i, u_{i+1}]$, with

$$\delta_i = \frac{m_i - u_i}{3} \text{ and } \delta'_i = \frac{u_{i+1} - m_i}{3}. \quad (11)$$

Such a separation ensures that the frame calculation (cf. Section 5.2) is always done at low curvature regions of the spline, ensuring numerical stability. Thus, the regularity property ensures that the high curvature segments are detected during the first pass of the process itself. It only involves a simple and fast computation, thus fitting the requirement for real-time rendering. The second pass positions the additional sampling points in the high-curvature segments in order to get a more stable tessellation. This sampling process is somewhat related to the surface analysis technique described in [41], where both symbolic and numerical tools are combined for solving the addressed problem. Here, the high curvature segments are isolated using numerical evaluators c_i and the maximal curvature points are calculated using symbolic extraction on the considered spline segment. The main difference is the simplification done in order to achieve faster performance. However, with this modification (especially regarding c_i), no theoretical validation is available. In addition, our scheme does not guarantee the minimality of the number of samples with respect to any error criteria. It simply ensures that no high curvature point will be missed. Our method makes sure that a spline curve is turned into a set of interest points along the curve (i.e., those corresponding to segment extremities and high-curvature regions). The resulting curve segments are displayed using a hardware-based technique described in the following section.

5.2 Hardware-Based Rendering Using Skinning

The skinning definition (cf. Section 2) provides us with a quite powerful tessellation primitive, more complex than an ordinary polygon. In this section, we will define a *tessellation elementary cell* from a regular cylinder with a standard weight function and two matrices M_1 and M_2 . We shall also see how combining this simple, precomputed cylindrical shape with the skinning deformation process can significantly reduce the software side of the tessellation process. We now describe the overall tessellation algorithm in two steps:

Step 1: The sampling process gives a set of parameter intervals $[a_j, a_{j+1}]$ ($\{a_j\}$ is a reordered set of $\{u_i\}$ with the inserted parameters of the high-curvature region).

Step 2: For each interval generated by the sampling algorithm, two frames (one for each segment extremity) are calculated along the axis of C as follows [29]: Each frame $(\mathbf{t}_j, \mathbf{k}_j, \mathbf{b}_j)$ corresponding to the parameter value a_j is calculated from the previous frame using:²

$$\mathbf{b}_j = \mathbf{t}_j \times \mathbf{k}_{j-1} \text{ and } \mathbf{k}_j = \mathbf{b}_j \times \mathbf{t}_j. \quad (12)$$

The $\mathbf{t}_j$ vectors are always assigned the first derivative of the local curve. This reduces the frame twist along the curve, which is not the case when using a simple Frenet frame. However, for the first frame alone (corresponding to a_0), we use the Frenet frame. Each frame gives the local orientation of the sweep cross-section. These frames generate the homogeneous transformation matrices M_1 and M_2 for the considered parameter interval. They are then used by the blending technique (cf. Section 2) on an elementary

2. Where $\times$ denotes a vector cross product.

nondeformed cylinder aligned along the z-axis, described in our implementation as:

$$\begin{cases} x^2 + y^2 = 1 \\ z \in [0, 1]. \end{cases} \quad (13)$$

This nondeformed cylinder is approximated by polygons using a positive integer n as the *approximation resolution*. In our implementation, we used 2^{n+1} sections for the cylinder, uniformly positioned along the cylinder axis. All the sections are discretized at a given constant resolution. The higher the n value, the smoother the interpolation between the frames represented by M_1 and M_2 . Many weight functions are available for smooth deformation results (see [7] for possible choices). In our case, the weight function should be smooth, locally supported, and monotonically decreasing from 1 to 0. Accordingly, we used the following:

$$\omega(\nu) = \begin{cases} 1 - 2z_\nu^2, & z_\nu^2 < \frac{1}{4} \\ \frac{(1-z_\nu^2)^2}{0.75+1.5z_\nu^2}, & z_\nu^2 \geq \frac{1}{4}, \end{cases} \quad (14)$$

where z_ν is the z-coordinate of the vertex ν on the nondeformed cylinder defined in (13). This weight function is a particular case of the more general function proposed in [32] for soft objects blending.

Fig. 11 shows an example of a tessellation achieved using this technique (hence, the display primitive is a deformed cylinder). We can observe that the results of this new technique are visually more pleasing than the standard approach (brute-force tessellation with piecewise linear approximation) and the high-curvature points are handled in a much better way (see the box in Fig. 11). The software part of this tessellation only involves the calculation of the geometric transformations and no vertex position is explicitly evaluated and transmitted to the GPU during the display process. The tessellation primitive for a given resolution n is constant over time since only the matrix changes from one parameter interval to another. Thus, it can be stored in a *displaylist* using *vertexArray* technology [30] to take full advantage of the data transmission optimizations of the GPU. Note that there is no mathematical proof that such a deformation tool can properly handle all the classical tessellation problems. We also observe that increasing the number of vertices on the tessellation primitive results in smoother deformations (Fig. 12). Though this may seem quite obvious, the smoothing involves only a small CPU overhead since the deformation process is performed entirely by the GPU. The only overhead needed is the transmission of the tessellation primitive to the GPU.

5.3 Guaranteed Frame Rate

The accuracy of the entire visualization process described above depends on two parameters. The first parameter, ε , determines how many parameter intervals will be used in the tessellation of the skeleton curve (this number should at least be equal to the number of spline segments on C). The second parameter n defines the complexity of the cylinder primitive used for tessellation. In our implementation, we actually use several versions of the tessellation primitive. The number of polygons in each version is $O(2^i)$, $i = 0, \dots, n$ (cf. Section 5.2). For a given parameter interval,

Fig. 11. (a) Standard tessellation result and (b) high quality rendering using our new technique.

increasing the number of vertices on the tessellation primitive produces smoother interpolation with only a small measured computational overhead. Accordingly, we combine the tessellation scheme with an algorithm that compares the current frame rate with a given reference value and tunes the ε and n values automatically. In principle, this can be seen as a feedback control loop. More precisely, the iterations from one display configuration (ε_k, n_k) to the next one $(\varepsilon_{k+1}, n_{k+1})$ use the following relations:

$$\begin{aligned} \varepsilon_{k+1} &= \varepsilon_k + \alpha(f_0 - f_k) \\ \rho_{k+1} &= \rho_k + \beta(\varepsilon_{k+1} - \varepsilon_k) \\ n_{k+1} &= \lfloor \rho_{k+1} \rfloor, \end{aligned} \quad (15)$$

where ρ_k is a floating point version of the integer variable n_k , α and β are two arbitrary constants, f_k is the current measured frame rate, and f_0 is the desired one. In our implementation, we used $\alpha = 10^{-4}$ and $\beta = 5$. These values were experimentally determined in order to provide a satisfactory frame rate control (Fig. 12). The tessellation process stabilizes itself at the desired frame rate in less than a second.

6 SYSTEM INTEGRATION

The models presented so far have been integrated into a simulation platform for seamless interoperability with existing models and tools. We now briefly describe the SPORE library (see [5], [15] for more details). Developed in C++, this library provides a fully functional environment

Fig. 12. (a) Low quality tessellation and (b) high quality tessellation obtained using our automatic frame rate adaptation system at 900 and 460 frames/sec, respectively, on a GeForce 4-based system.

for the simulation of minimally invasive ovarian surgery. All the objects within this library are decomposed into three different models: a geometric model used for rendering, a mechanical model for animation, and a collision detection model for evaluating the influence of the environment. The three layers of an object share the necessary information in order to access them (e.g., external forces are transferred from the collision model to the mechanical model, which in turn transfers the state variables to the geometric model for rendering). The library is designed such that all the interactions between objects are sphere-based, which is a good trade off between speed and being generic. As a result, the system composed of the intestine and the mesentery has been included within the simulator as a whole and is treated by other simulated objects as a single entity. SPORE is organized as follows: A global animation process iteratively processes the user interactions and collisions between objects in order to calculate the forces that will be applied to each object. It then gives an object the possibility to evaluate its own behavior, by calling an ad hoc method of the object class, transferring the applied external forces to itself. This way, the system composed of intestine and mesentery can calculate the forces due to self-collision and the response to the global force system by itself using the techniques described earlier. With this system integration, we can evaluate the simulator for its performance and realism, the results of which are presented in the following section.

7 RESULTS AND VALIDATION

Simulator Results: Snapshots from the real-time animation of the simulator are shown in Fig. 13. In all cases, the organs were subject to forces due to gravity, user-input, and collisions. Fig. 13a shows the case of an isolated intestine and mesentery manipulated by a virtual probe (the tiny sphere). Note that our displacement-velocity method for collision response produces fairly stable simulations (Fig. 13b) with some undesired vibrations. Fig. 13c and Fig. 13d show the simulations inside a simulated abdominal cavity. All the snapshots were captured from our simulator in real-time on a standard PC with Bi-Athlon 1.2 GHz, 512 MB RAM and nVIDIA GeForce 3 graphics card. A dynamic real-time demonstration of our simulator is available as a companion to this paper at: <http://www-evasion.imag.fr/Membres/Laks.Raghupathi/intestine.html>.

Quantitative Validation: The mechanical and collision detection model were implemented and tested independently prior to the integration. The results were fast enough to run at 30 Hz on a standard PC and the overall behavior was good. However, due to the stochastic nature of our algorithm, we do not have any theoretical proof that all the collisions are detected. However, in Fig. 14, we present one of the several evaluations which we performed on the collision processing algorithm (see [6] for detailed results). Here, as the object is deformed, we plotted the colliding regions as detected by our method and a naive $O(n^2)$ method. Results show that our method is able to identify all the *active colliding regions* by evaluating a far fewer number of points. Though not the same as finding the *exact collisions*, this is good enough for our case since we can quickly narrow down to these points using temporal coherence. Nevertheless, we miss collisions in some cases between intestine-mesentery due to an entirely different reason (Fig. 13d). The interpenetrations occurred because we performed only segment-segment collision detection. Note that this technique worked well for intestine-intestine collisions and we do not see any collisions missed here. But, the mesentery was modeled as a triangulated mesh and, when there are cases of segment-triangle or point-triangle collisions, the above approach in its present form was not effective. More recently, we extended this approach by adding point-triangle, edge-triangle cases and tested it successfully with a mesentery-like surface.

Qualitative Validation: The prototype simulator was demonstrated to the surgeons at IRCAD in July 2003. The surgical educators practiced on our simulator and explained to us the good points as well as the drawbacks of the current model. The overall intestine behavior and contact modeling was observed to be very good. Some of the instabilities observed in certain cases of the intestine's motion actually turned out to simulate a patient who is spasmodic (suffering from intestinal convulsions). Although this happens quite frequently in a real patient, this problem can be fixed by increasing the damping of the system. They also suggested that the mesentery should be less elastic. With our mass-spring model, the characteristics demanded by the surgeons can be obtained simply by tuning the simulation parameters. Though these parameters can be varied intuitively, it is difficult to incorporate those parameters obtained from biometric studies. Finally, a small error was detected in the geometric design—the mesentery should have a zero width at the two extremities, where the intestine is directly attached to the main vessels.

8 CONCLUSIONS

Our novel anatomical model greatly simplifies the problem of simulating a complex deformable organ in real time. The collision model efficiently determines all the active colliding regions and provides a realistic and stable response. Our rendering technique quickly generates high-quality image sequences with no tessellation artifacts. All of our three contributions were integrated into a surgical simulator platform with real-time performance. The feedback from the surgeons is positive in general, with suggestions for improvements. The corrections are not hard to incorporate into our system and we are currently working on them. We believe that our simulator can be used as a generic trainer that can mimic the behavior of the organs very well, thus eventually replacing the practice on animals. The results encouraged us to include triangle tests for more robust

Fig. 13. Snapshots from our intestinal surgery simulator. (a) Intestine and mesentery on a plane pulled by a probe. (b) Stable rest position. (c) Inside the virtual abdominal cavity. (d) Case when collision detection fails (see encircled region).

collision detection. We are also working on adapting continuous collision detection techniques for handling fast motion of thin objects such as membranes, vessels, etc.

Fig. 14. Quantitative evaluation of the collision detection method. Points in green (light circles) detected by the $O(n^2)$ method and blue (dark points) by our approach.

ACKNOWLEDGMENTS

This work is supported by INRIA (National Institute for Research in Computer Science and Control) under the ARC SCI grant (Intestine Surgery Simulator project). The authors would like to thank Luc Soler (IRCAD) for his help on anatomy and for facilitating the feedback from the surgeons.

REFERENCES

- [1] A.P. Fine, "Patient Information: Laparoscopic Colon Surgery" Soc. Laparoendoscopic Surgeons, Inc., 2004, <http://www.sls.org/patientinfo/colon.html>.
- [2] S. Sgambati and G. Ballantyne, "Minimally Invasive Surgery for Diseases of the Colon and Rectum: The Legacy of an Ancient Tradition," *Laparoscopic Colectomy*, R. Jager and S. Wexner, eds., pp. 13-23, New York: Churchill & Livingstone, 1995.
- [3] L. Augusten, R. Bowen, and M. Rouge, "Pathophysiology of the Digestive System," Colorado State Univ., 2004, <http://arbl.cvmbs.colostate.edu/hbooks/pathophys/digestion>.
- [4] "Laparoscopic Colectomy," Center for Videoendoscopic Surgery, Univ. of Washington, 2004, <http://depts.washington.edu/cves/lapcol.html>.
- [5] P. Meseure, J. Davanne, L. Hilde, J. Lenoir, L. France, F. Triquet, and C. Chaillou, "A Physically Based Virtual Environment Dedicated to Surgical Simulation," *Surgical Simulation & Soft Tissue Modeling*, pp. 38-46, Springer, 2003.
- [6] L. Raghupathi, V. Cantin, F. Faure, and M.-P. Cani, "Real-Time Simulation of Self-Collisions for Virtual Intestinal Surgery," *Surgical Simulation & Soft Tissue Modeling*, pp. 15-26, 2003.
- [7] L. Grisoni and D. Marchal, "High Performance Generalized Cylinder Visualization," *Proc. Shape Modeling Int'l '03*, pp. 257-263, 2003.
- [8] S. Capell, S. Green, B. Curless, T. Duchamp, and Z. Popovic, "Interactive Skeleton-Driven Dynamic Deformations," *Proc. SIGGRAPH '02*, pp. 586-593, 2002.
- [9] G. Debunne, M. Desbrun, M.-P. Cani, and A.H. Barr, "Dynamic Real-time Deformations Using Space and Time Adaptive Sampling," *Proc. SIGGRAPH '01*, pp. 31-36, 2001.
- [10] E. Grinspun, P. Krysl, and P. Schröder, "CHARMS: A Simple-Framework for Adaptive Simulation," *Proc. SIGGRAPH '02*, 2002.
- [11] M. Bro-Nielsen, "Finite Element Modeling in Surgery Simulation," *Proc. IEEE*, vol. 86, no. 3, pp. 490-503, 1998.
- [12] D.L. James and D.K. Pai, "ArtDeform—Accurate Real Time Deformable Objects," *Proc. SIGGRAPH '99*, pp. 65-72, 1999.
- [13] G. Székely, C. Brechbühler, R. Hutter, A. Rhomberg, N. Ironmonger, and P. Schmid, "Modelling of Soft Tissue Deformation for Laparoscopic Surgery Simulation," *Medical Image Analysis*, no. 4, pp. 57-66, 2000.
- [14] S. Cotin, H. Delingette, and N. Ayache, "A Hybrid Elastic Model for Real-Time Cutting, Deformations, and Force Feedback for Surgery Training and Simulation," *The Visual Computer*, vol. 16, no. 8, pp. 437-52, 2000.

- [15] P. Meseure and C. Chaillou, "A Deformable Body Model for Surgical Simulation," *J. Visualization and Computer Animation*, vol. 11, no. 4, pp. 197-208, 2000.
- [16] M. Kauer, V. Vuskovic, J. Dual, G. Székely, and M. Bajka, "Inverse Finite Element Characterization of Soft Tissues," *Proc. Int'l Conf. Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 2001.
- [17] J.C. Lombardo, M.-P. Cani, and F. Neyret, "Real-Time Collision Detection for Virtual Surgery," *Proc. Computer Animation '99*, pp. 82-91, 1999.
- [18] L. France, A. Angelidis, P. Meseure, M.-P. Cani, J. Lenoir, F. Faure, and C. Chaillou, "Implicit Representations of the Human Intestines for Surgery Simulation," *Proc. Modeling & Simulation for Computer Aided Medicine & Surgery (ESAIM)*, M. Thiriet, ed., pp. 42-47, 2002.
- [19] L. France, J. Lenoir, P. Meseure, and C. Chaillou, "Simulation of a Minimally Invasive Surgery of Intestines," *Proc. Fourth Virtual Reality Int'l Conf.*, June 2002.
- [20] G. Bradshaw and C. O'Sullivan, "Sphere-Tree Construction Using Dynamic Medial Axis Approximation," *Proc. Symp. Computer Animation*, pp. 33-40, 2002.
- [21] J.D. Cohen, M.C. Lin, D. Manocha, and M.K. Ponamgi, "I-COLLIDE: An Interactive and Exact Collision Detection System for Large-Scale Environments," *Proc. ACM Symp. Int'l 3D Graphics*, pp. 189-196, 1995.
- [22] S. Gottschalk, M. Lin, and D. Manocha, "OBBTree: A Hierarchical Structure for Rapid Interference Detection," *Proc. SIGGRAPH '96*, 1996.
- [23] I.J. Palmer and R.L. Grimsdale, "Collision Detection for Animation Using Sphere Trees," *CG Forum*, vol. 14, no. 4, pp. 105-116, 1995.
- [24] G. van den Bergen, "Efficient Collision Detection of Complex-Deformable Models Using AABB Trees," *J. Graphics Tools*, vol. 2, no. 4, pp. 1-14, 1997.
- [25] J. Klosowski, M. Held, J. Mitchell, H. Sowizral, and K. Zikan, "Efficient Collision Detection Using Bounding Volume Hierarchies of k-DOPs," *IEEE Trans. Visualization and Computr Graphics*, vol. 4, no. 1, pp. 21-36, Jan.-Mar. 1998.
- [26] S. Guy and G. Debunne, "Monte-Carlo Collision Detection," INRIA, Technical Report RR-5136, 2004, <http://artis.imag.fr/Publications/2004/GD04/>.
- [27] M. Lin and J. Canny, "Efficient Collision Detection for Animation," *Proc. Third EG Animation & Simulation Workshop*, 1992.
- [28] G. Agin and T. Binford, "Representation and Description of Curved Objects," *IEEE Trans. Computers*, vol. 25, no. 4, pp. 439-449, Apr. 1976.
- [29] J. Bloomenthal, "Calculation of Reference Frames along a Space Curve," *Graphics Gems*, A. Glassner, ed., pp. 567-71, New York: Academic Press, 1990.
- [30] OpenGL Vertex Weighting, NVIDIA Corp., 2004, http://developer.nvidia.com/object/GL_EXT_vertex_demo.html.
- [31] J. Bloomenthal, "Medial-Based Vertex Deformation," *Proc. Symp. Computer Animation*, pp. 147-151, 2002.
- [32] C. Blanc and C. Schlick, "Extended Field Functions for Soft Objects," *Proc. Implicit Surface '95*, pp. 21-32, 1995.
- [33] L'Institut de Recherche contre les Cancers de l'Appareil Digestif, Strasbourg, 2004, <http://www.ircad.org>.
- [34] S. Uno and M. Slater, "The Sensitivity of Presence to Collision Response," *Proc. IEEE Symp. Virtual Reality and Its Application in Industry (VRAI)*, pp. 95-103, 1997.
- [35] D.H. Eberly, *3D Game Engine Design: A Practical Approach to Real-Time Computer Graphics*. Morgan Kaufmann, 2000.
- [36] D. Baraff and A. Witkin, "Large Steps in Cloth Simulation," *Proc. SIGGRAPH '98*, pp. 43-54, 1998.
- [37] J. Platt and A. Barr, "Constraints Methods for Flexible Models," *Proc. SIGGRAPH '88*, pp. 279-288, 1988.
- [38] T. Giang, G. Bradshaw, and C. O'Sullivan, "Complementarity Based Multiple Point Collision Resolution," *Proc. Fourth Irish Workshop Computer Graphics*, pp. 1-8, 2003.
- [39] G. Farin, *Curves and Surfaces for Computer Aided Geometric Design*. New York: Academic Press, 1990.
- [40] L. Piegl and W. Tiller, *The NURBS Book*. New York: Springer, 1995.
- [41] G. Elber and E. Cohen, "Second Order Surface Analysis Using Hybrid Symbolic and Numeric Operators," *ACM Trans. Graphics*, vol. 12, no. 2, pp. 160-178, 1993.

graphics. He is a member of Tau Beta Pi and Eta Kappa Nu.

Laurent Grisoni received the diploma in computer science from the ENSEIRB national engineering school in 1996 and the PhD degree from the University of Bordeaux I in 1999. He is currently an assistant professor at the Polytech'Lille engineering school, University of Lille 1, and is a member of the INRIA ALCOVE project. His research interests include spline and implicit modeling, animation, and multiresolution aspects.

François Faure received the PhD degree in computer science from the University of Joseph Fourier, Grenoble, France, in 1997, where he is currently an assistant professor. He previously studied mechanical engineering at the École Normale Supérieure de Cachan. His research interests are about physically-based animation of complex scenes.

Damien Marchal received the Master's degree from the University of Lille 1 in 2003. He is currently working on the PhD degree at the LIFL lab and is a member of the INRIA ALCOVE project. His research interests include collision detection, deformable models, and hardware-accelerated computation.

Marie-Paule Cani is a graduate of the École Normale Supérieure and received the PhD degree from the University of Paris Sud in 1990. She is a professor of computer science at the Institut National Polytechnique de Grenoble, France. Her main research interests cover physically-based simulation, implicit surfaces applied to interactive modeling and animation, and the design of layered models incorporating alternative representations and LODs.

Christophe Chaillou has been a professor at the Polytech'Lille engineering school, University of Lille 1 since 1995 and is the research director of the INRIA ALCOVE project. His research interests include haptic peripheral design and collaborative virtual environments.