


It's a Hit

Real-time collision detection brings virtual surgery closer to reality

Diana Phillips Mahoney

There's no question that virtual reality has come a long way since its inception more than a decade ago. The digital worlds through which users move are more sophisticated, thanks to advanced modeling, animation, and rendering tools, and innovations in positional tracking and haptic interfaces have resulted in more natural interaction with virtual objects. Additionally, the compute engines driving these developments are getting more powerful by the day.



A virtual laparoscope collides with a deformable liver model. The interaction and subsequent deformation is processed and represented in real time using a hardware-based collision-detection method.

VR's journey from nascent concept to practical technology is far from complete, however. There are many roadblocks yet to overcome before the technology makes good on many of the promises set forth by early visionaries. This is particularly evident in the area of surgical simulation. While virtual surgery has the potential to fundamentally change the teaching and practice of medicine, doing so requires more from the technology than current tools are able to provide.

To be useful, a surgical simulation has to meet an almost impossibly high standard--the one set by Mother Nature, who did not design human anatomy so that it could be easily mimicked via geometric primitives and standard animation and interaction techniques. To be effective, a surgical simulation must not only realistically reproduce the complexity of human organs, but it must also represent the visual and haptic impact of each and every interaction with the digital objects. And it must do so in real time. Because of the computational intensity of this task, one or more of the critical design factors--real-time operation, photorealistic rendering, or accurate haptic simulation--must often be sacrificed for the sake of another, making the overall experience less-than-real.

A new technique developed by a group of French researchers may bring the technology one step closer to fulfilling its potential. Marie-Paule Cani of the National Polytechnic Institute of Grenoble (INPG), along with Christophe Lombardo and Fabrice Neyret, researchers with iMAGIS (a joint project between the French National Scientific Research Center, INPG, and the Gravir-Imag Lab of the University Joseph Fourier), have developed a hardware-based solution for performing real-time collision detection, one of the most imposing obstacles in surgical simulation.

Accurate, real-time collision detection is an important element of all VR applications, and a critical one to virtual surgery. If a system can't detect and process collisions, synthetic objects will move through each other when in motion rather than react to interference with each other, precluding both useful interaction and a sense of realism. Simulating collision detection is computationally expensive, however. It requires testing for intersections between each pair of virtual objects in a scene, as well as in self-collisions between different parts of the same object. Typically, objects are described using polygonal meshes, and algorithms detect the points at which the faces of the meshes collide. The more objects, or faces of objects, that collide, the more demanding the algorithm.



To test for collisions between a surgical tool and the virtual liver, the graphics hardware provides a list of faces in the organ intersected by the tool.

In addition to calculating the visual impact of such collisions, surgical simulations require computing the haptic response as well. "In surgery simulators, the user trains for surgical gestures on virtual models of the human body by interacting with a moving tool. He or she needs to be able to feel as well as see the objects the tool is colliding with," says Cani. This increases the processing load of the simulation by an order of magnitude, because the human haptic response is far more sensitive than the visual one, requiring 500 calculations per second to simulate reality, as opposed to 25 or 30 for a realistic visual response.

Another, more subtle problem developers must contend with is that time in the world of computers is discrete. An object can be in front of an obstacle at one time step and behind it at the next, while never actually "colliding" with it. This occurs when the virtual objects are thin or when the motions are quick. As a result, the system needs to be able to test for collisions "in the meantime" between two steps. "Strictly speaking, this means [the system] is no longer testing for collisions between pairs of faces, but between moving faces, or prisms in a 4D space," says Cani. Once again, this significantly adds to the computational burden.

Although there have been a number of developments in the area of collision detection in the past few years, most standard approaches are ill suited to virtual surgery applications, primarily because human organs are non-convex and they deform over time. For example, one "family" of solutions relies on specific data structures that store the closest points and faces for each pair of moving objects. These methods exploit temporal coherence, in that the stored features vary only slightly at each time step. This approach is only useful, however, when dealing with convex-shaped objects, however. For non-convex objects (those with boundaries not described by outward curves), there are locations where the closest points change suddenly, even for a slight location change.

Another approach involves precomputing hierarchies of oriented bounding volumes (spheres, boxes, or grids) that tightly fit the geometry of the virtual objects. The volumes perform efficient rejection tests to determine if a given series of computations will be useful before implementing the simulation. In applications in which the virtual objects deform over time, the hierarchy has to be recomputed at each time step, considerably impacting the performance of the simulation.

The approach the French researchers have taken avoids these pitfalls. Rather than trying to make some of the existing software-based geometric and algorithmic optimization methods "fit" the difficult surgical simulation problem, the team uses a hardware solution to simulate real-time interaction between virtual objects and rigid, geometrically simple tools. With respect to surgical simulation, the researchers' efforts focus on laparoscopic surgery--a type of minimally invasive procedure in which a camera-tipped probe and surgical instruments are inserted into a patient through small incisions. The surgeon manipulates the devices based on views captured by the probe and delivered to an adjacent video monitor. This type of procedure is well suited to virtual representation because of its monitor-based paradigm and the limited range of motion of the surgical tools. In this case, the laparoscope also fits the bill in terms of being a rigid tool that is defined by a simple geometry.

A New Approach

"The method the researchers have developed relies on an analogy between collision detection and the first part of a traditional projective rendering process," says Cani. The collision-detection process--testing for intersections by finding object faces that penetrate each other--is similar to the use of clipping planes in rendering, whereby faces of objects are clipped toward a viewing volume, characterized by the location, orientation, and projection of a camera. This is done before projection and rasterization as a way to render only the intersection between the objects in the scene and the viewing volume, and it is a function that is handled well by specialized graphics hardware.

The researchers borrowed this model to achieve accurate, real-time collision detection. The first step is to specify a viewing volume that corresponds to the shape of the tool or to the volume covered by the tool between consecutive time steps, and use hardware to "render" the main object relative to the "camera," says Cani. "We can get the hardware to provide a feedback buffer containing the list of faces intersected by the tool. If nothing is visible, then there is no collision."

The advantage of this approach is twofold. First, it detects collisions between the tool and arbitrarily complex objects without the need for precomputation, so objects can deform over time without any loss of performance. Second, because the tool trajectory between time steps can be modeled as a viewing volume, the resulting collision detection is highly accurate, again without increasing the computational load.

This collision-detection method is simple to program, notes Cani, requiring only a dozen lines of code in any OpenGL application with appropriate 3D graphics acceleration capabilities. Additionally, the method is portable, thanks to the availability of OpenGL on many architectures, and is "open" in that it can use any geometric primitive (such as NURBS) that is treated by the graphics library. In principle, the method will work with other graphics libraries, such as Direct3D, as long as the graphics card accelerates the geometric portion of the rendering process as well as the rasterization component.

To demonstrate the efficacy of their hardware-based collision-detection process, the researchers have developed a virtual liver model that deforms in real-time based on the action of the surgical tool. The user interacts with the organ using SensAble Technology's Phantom force-feedback device. The researchers are currently extending the simulation to handle incisions made into the liver so that, as with "the real thing," the simulated organ undergoes significant deformations and topological changes as parts of it are cut and removed. Cani expects to have a version of the system in the hands of surgeons within one year for evaluation.

Diana Phillips Mahoney is chief technology editor of Computer Graphics World.

Computer Graphics World November, 1999