Editors: Lawrence J. Rosenblum
and Michael R. Macedonia

Phantom-Based Haptic Interaction with Virtual Objects

J. Kenneth Salisbury and Mandayam A. Srinivasan

Massachusetts Institute of Technology

The Phantom haptic interface is a device that lets you literally feel virtual objects with your hands. It sits conveniently on your desktop next to your computer, looking a bit like a miniature desk lamp (see Figure 1). The Phantom haptic interface has a stylus grip or a fingertip thimble with which users can reach into virtual worlds to touch and interact with 3D objects. With hundreds of Phantom systems now in use throughout the world, physical interaction with virtual environments (VEs) is rapidly becoming a reality.

Touch, or *haptic* interaction is one of the most fundamental ways in which people perceive and effect changes in the world around them. Our very understanding of the geometry and physics of the world begins by touching and physically interacting with objects in our environment. Touch interaction differs fundamentally from all our other sensory modalities in that it is intrinsically bilateral. We exchange energy between ourselves and the physical world as we push on it and it pushes back on us. In this exchange, information and intent are conveyed in a physically direct and cognitively primal way. Our ability to express ourselves in musical performance, painting, sculpting, and gymnastics depends on physically performing the task and learning from the interactions.

Yet, in the world of computers, our primary mode of receiving information is still visual and depends on information abstractions such as words, diagrams, and pictures. Worse, to enter information into a computer, we are restricted to typing on a keyboard and clicking and dragging a mouse. Given the importance of touch in our lives, it is ironic that a convenient and truly bilateral means for haptic interaction with information has been, until recently, unavailable to computer users.

Mechanical devices that allow haptic interaction with remote and virtual objects have been around for decades, though unsatisfactory for widespread use. Early remote manipulation systems were used for handling hazardous substances as far back as the 1940s, when the danger of working with nuclear materials necessitated developing remote manipulation devices.

Today's surge of haptic research and commercial activity grew from the early efforts of designers who built the "master" input devices needed to control remote manipulators. In the 1960s, Knoll at Bell Labs was perhaps the first to demonstrate touching virtual shapes with a computer-controlled haptic interface. Since then, numerous devices have been built for haptic interaction, based on the recognition that adding haptics to graphic images significantly improves human-computer interactions.

In 1993, haptic interaction with computers took a significant step forward with the development of the Phantom haptic interface. This simple device has spawned a new field analogous to computer graphics—*computer haptics*—defined as the discipline concerned with the techniques and processes associated with generating and displaying synthesized haptic stimuli to the human user.

Inspired by our previous work in interpreting robot touch sensor information and study of human touch perception, the Phantom interface permits users to feel the forces of interaction they would encounter while touching objects with the end of a stylus or the tip of their finger. The resulting sensations prove startling, and many first-time users are quite surprised at the compelling sense of physical presence they encounter when touching virtual objects. To appreciate why the Phantom system succeeded where others failed, you need to understand the nature and functioning of the human haptic system.

How and what do humans feel?

We use our hands to explore and manipulate objects in our environment. Unique among our sensory modalities, the haptic system relies on action to stimulate perception and vice versa. While exploring an object, we actively scan our fingers across its surface and squeeze or palpate it to sense its physical properties. To manipulate an object, we apply forces to move it, while simultaneously sensing the results of these actions. Thus, in almost all of the hand's activities, either to extract information from or to alter the environment, we use both the sensory and motor parts of our haptic system.

Correspondingly, a haptic interface needs to sense our motor actions and display appropriate haptic "images." Keyboards and mice convey very restricted motor actions to the computer. Instrumented gloves sense more degrees of freedom, but with less precision. These passive interfaces, however, cannot display any synthesized haptic images to the user.

What should this haptic image consist of? In the real world, whenever we touch an object, it imposes forces on our skin. These net forces, plus the posture and motion of hands and arms, are transmitted to the brain as *kinesthetic* information. This is how we sense coarse properties of objects—such as large shapes or spring-like compliances—that require hand or arm motion in probing them.

In contrast, receptors embedded in the skin convey *tactile* information, such as spatial and temporal variations of force distributions on the skin, within the contact region with an object. Slipping of surfaces, fine textures, small shapes, and softness are felt through our tactile sensors. In addition, the skin's temperature, which relates to the temperature and thermal properties of the object, is also sensed through specialized tactile sensors.

Courtesy of SensAble Technologies

1 **The Phantom haptic interface.**

Because haptic images are composed of both tactile and kinesthetic information, each arising from multiple sensory channels, they can seem quite complex compared to visual and auditory information. To be successful, however, a VE does not have to perfectly replicate reality; it only needs to match the abilities and limitations of the human sensory, motor, and cognitive systems. For example, because of the limitations of human vision, graphic images displayed at about 30 frames per second seem continuous and can even convey a sense of telepresence.

It's difficult to quantify the human haptic abilities that prescribe the design specifications of haptic interfaces, because of the multiplicity of the subsystems and the sensorimotor nature of the tasks. We know that the tactile system can resolve vibrations of up to 1 kHz, with submicron amplitudes detectable around 250 Hz. The kinesthetic resolution in sensing the position of our fingertips is about 1 mm, with an ability to discriminate differences of about 10 percent for velocity and 20 percent for acceleration. The motor system's bandwidth for controlled motions is less than 10 Hz, and the maximum controllable force exerted through the fingers is 50 to 100 newtons.

Creating the feel of objects

In matching human capabilities, the Phantom haptic interface succeeds by trading off complexity to ensure high fidelity. One key observation in simplifying haptic interface design is that people perform quite well in exploring and manipulating the world through a stick or a rigid thimble. The enabling insight in the device is that forces generated through point contacts, especially during active exploration, contain significant spatial and temporal information that humans can easily understand. You need only close your eyes and probe objects with the tip of a pencil to understand the basic mode of touch interaction that the Phantom system uses.

To mimic these point contact interactions, the Phantom interface allows and measures motion along six degrees of freedom and can exert controllable forces on the user along three of those freedoms. Because the device does not constrain motions within its workspace, and because its inertia and friction are low, free motion feels free and comfortable. These device characteristics also let users feel the objects without being distracted by the device. The relatively large dynamic range in force output—the ratio of the largest to the smallest displayable force—plus a good match with human resolution and bandwidth provides enough contrast in force sensations to convincingly display impact, rigidity, texture, complex shapes, and a range of compliances.

To evoke the sensation of touching objects, we must first model the geometric, material, kinematic, and dynamic properties of the world we wish to represent. We then must devise computational methods ("haptic rendering") to determine the forces that result when we interact with objects. The algorithms rely on the model used and must be considered carefully to meet the real-time needs. Due to the high servo rates required to generate smooth transitions and distinct sensations (Phantom rendering programs typically run at 500 to 2,000 Hz), it is important that the rendering algorithms be efficient.

One of the first rendering methods we implemented used potential fields defined by planar and spherical surfaces. Evoking the sensation of touching these surfaces simply required exerting on the user a force proportional to the penetration depth of the Phantom endpoint beneath the surface. Unfortunately, this method is limited in expressiveness and does not prevent users from pushing entirely through small spherical and thin composite shapes.

We developed two new rendering methods based on polyhedral models to take advantage of their rich ability to encode 3D shapes. The first method ensures the sensation of solidity while touching these polyhedra. It defines a local tangent plane constraint surface that is impenetrable and coplanar with the currently touched facet of the object. This constraint method permits rendering solid objects with a great variety of shapes.

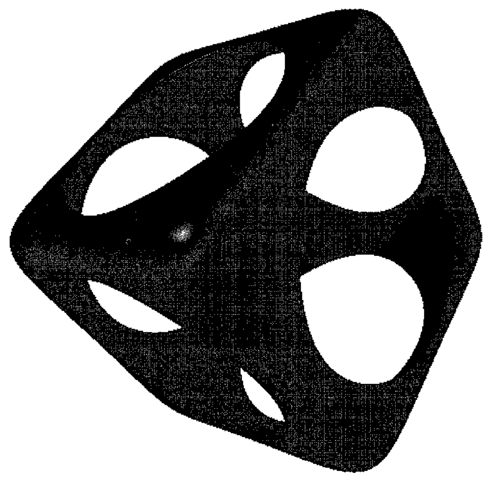

2 An object modeled and rendered using an implicit function.

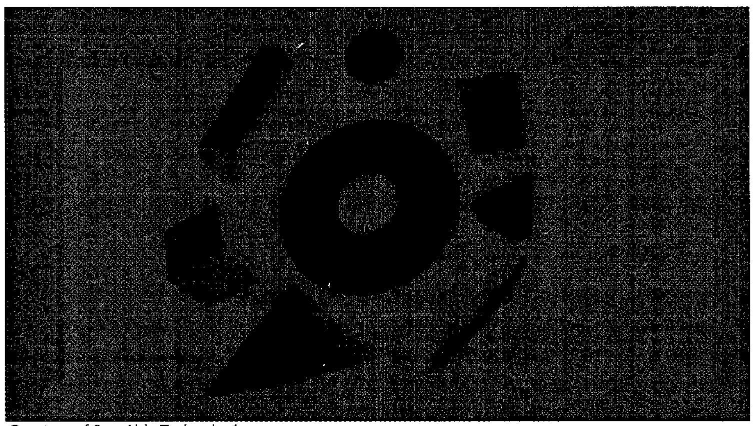

Courtesy of SensAble Technologies

3 Examples rendered using the Ghost software toolkit.

4 A compliant heart model using a surface mesh representation.

The second method, a ray-based technique, models the Phantom stylus as a line segment to take into account its orientation in reflecting the forces. Haptic interactions with polyhedral surfaces are simulated using ray-tracing techniques that can find surface intersections.

In either of these methods the faceted shapes can be made to feel smooth by perturbing the effective surface normal by a technique we call force shading, similar to Phong shading in graphics. You can add friction and simple textures by perturbing the normal and tangential contact forces.

A third class of shapes for which we have developed a rendering method is defined by implicit functions (such as $f(x, y, z) = 0$). This compact, yet expressive, representation has the advantage of continuity and a well-defined surface normal (see Figure 2). As with the first polyhedral rendering method, it employs a local tangent plane defined at the point of contact that prevents penetration and tracks contact across the surface. Amenable to surface texture and friction force perturbations, this method also lets you feel the exquisitely fine details, including cusps and singularities. Extensions of the method enable it to render piecewise continuous surfaces such as NURBS.

Versions of some of the methods above have been incorporated in SensAble Technologies' Ghost (General Haptics Open Software Toolkit), which lets users instantiate and interact with a variety of rigid objects (see Figure 3). The toolkit also gives objects mobility, enabling creation of switches and other dynamic objects.

Rendering compliant objects is another growing interest in the computer haptics community. Of particular value to medical simulation (see Figure 4), methods for rendering compliant objects place even more stringent demands on efficient representation and rendering. Both surface and volumetric representations are employed among the many models used in this context. In general, the more precise the model, the more complex and time-consuming the computation. Thus, the choice of modeling and rendering methods forces a trade-off between fidelity and complexity. In rendering compliant objects for surgical training, it is unclear how much fidelity is needed for effective training.

The current Phantom interface design precisely displays the forces resulting from point interactions. If the simulated tool has a nonpoint, extended object geometry, then torques that the Phantom system cannot display will arise as the tool contacts objects in the environment. Surprisingly, even in this situation users can deduce the state and contact geometry of these objects by performing small exploratory motions.

By itself, a haptic display can give users a realistic feel of objects. Combining it with visual and auditory feedback lets us take advantage of the human ability to synergistically integrate sensory information into a compelling cognitive experience.

The ultimate goal of VR researchers and designers is to make immersive environments seem as natural as the real one. Appropriately synchronized multimodal displays—which would include visual, auditory, and haptic modalities—seem essential in achieving this dream. We have shown that if you reduce or eliminate object deformation in the visual display when haptically deforming the virtual object, or if impact sounds appropriate to stiff objects are displayed when tapping virtual objects, users perceive the objects to be stiffer than the haptic feedback alone would indicate. Such multimodal illusions can be taken advantage of to overcome some of the technological limitations, enhancing the perceived fidelity and richness of users' experiences in VEs.

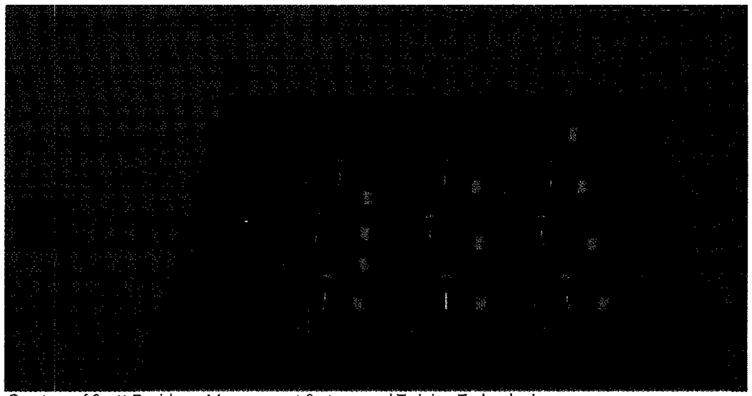

Courtesy of Scott Davidson, Management Systems and Training Technologies

5 A virtual circuit board for training electronic technicians.

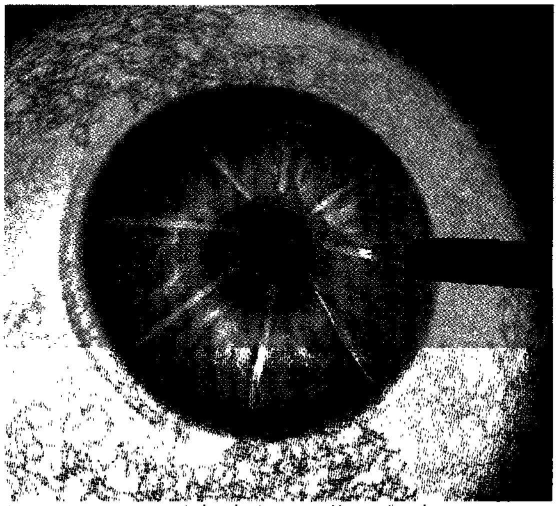

6 Simulation of radial keratotomy on a virtual eye.

Applications

Because haptics forms an essential part of most of our interactions with the real world, any VR application that involves simulating the real world benefits from haptic interactions—especially when combined with visual and auditory displays.

One of the first broad application areas of haptic technology is training people to perform real-world tasks. By providing a simple physical interface to computer-mediated training environments, the Phantom haptic interface enables the development of reconfigurable training systems that can easily be deployed and upgraded. The first application taking advantage of this capability is the virtual workbench for training electronic technicians, developed through a collaboration between researchers at the Massachusetts Institute of Technology and Naval Air Warfare Command/Training Systems Division (NAWC/TSD). This VE testbed uses a semisilvered mirror so that the virtual visual image of an electronic circuit board (see Figure 5) lies within the Phantom interface's workspace. Trainees can see the circuit board along with their own hand on it, feel the components of the circuit board with a probe, use a virtual multimeter at various contact points, and even haptically operate switches on it to observe changes in the circuit's electrical behavior.

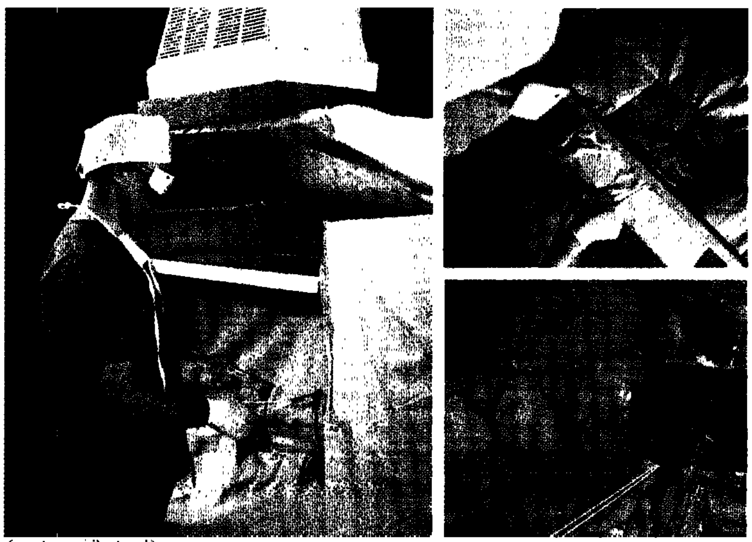

7 Simulation of blood vessels in an anastomosis procedure.

While this application takes advantage of the Phantom haptic interface in training cognitive skills, the force interaction capabilities of the device can also be used to train and evaluate sensorimotor skills. Researchers at the Center for Human Simulation at the University of Colorado Health Center have simulated both the look and feel of performing surgical procedures on the eye and the knee, based on highly detailed and realistic anatomical models (see Figure 6). Boston Dynamics has developed a surgical simulation of an anastomosis procedure with high-quality graphics and a Phantom-based haptic display (see Figure 7). The user can look at and feel a virtual blood vessel, use forceps to grasp a needle, and even suture the vessel.

Such training systems use the Phantom's force display capabilities to let medical trainees experience and learn the subtle and complex physical interactions needed to become skillful in their art. Because the Phantom can also act as a recorder and player of mechanical signals, it can be used to monitor a trainee's skill acquisition, customize the training program, or even let a trainee feel the prerecorded actions of an expert. Eventually, it may even help recertify medical practitioners by verifying their skill levels.

Haptic interfaces have the potential to radically alter and augment human-computer interactions. They can be used, for example, to enhance the naturalness of conveying a user's commands to the computer. They can also reduce the informational load on the visual system of a sighted user or provide an alternative display to a vision-impaired user.

Phantom technology can be effective in applications that generate, modify, and interact with shapes, textures, and material properties. For example, by adding feel to the acts of placing, arranging, cutting, joining, shaping, and sculpting, we should be able to greatly enhance the expressiveness and facility with which we create 3D objects and scenes in industrial and artistic contexts.

Opportunities exist for even more ambitious applications of haptic interfaces in general and the Phantom interface in particular such as

- science and business—enabling users to experience and manipulate complex multidimensional data sets;
- commerce—permitting customers to feel and interact with products;
- entertainment—allowing users to feel and manipulate different environmental behaviors, interaction tools, and avatars;
- education—giving students a feel of realistic and nonrealistic phenomena at a variety of spatial and temporal scales; and
- the arts—creating individual or collaborative (across the Web) virtual works of art through manual interactions.

Conclusion

The success and growing use of the Phantom haptic interface can be attributed to its simplicity and high performance. It is safe, convenient to use, and easy to program. Yet, there's certainly room to grow. We have experimented with adding powered two-finger grasp interfaces to the Phantom to explore force-reflecting grasp interactions. High-frequency vibration stimulators show promise for increasing the realness of impact, texture, and hardness sensations. Temperature displays can convey a dimension of information about thermal state and material properties that force alone cannot reveal. True tactile array stimulators are currently under consideration as a means for conveying intricate details of objects. We don't yet know which of these submodalities will provide the most useful next step in the Phantom technology. It will depend as much on technical advances as on understanding how to concisely convey haptic information.

The ability to record, display, transmit and edit visual and auditory information has had a profound effect on our society from sociological, technological, artistic, and economic points of view. Virtual environments go one step further in allowing real-time human interactions involving these modalities. The addition of haptics to VEs enhances both the quality and the quantity of information transmitted from and to the user.

Computer haptics—still in its infancy— is perhaps akin to the first telephone by Bell, the early sound recordings of Edison, or the early vector graphics images of simple geometric objects. The compelling, interactive nature of emerging haptic technologies suggests that they may well enjoy a similar evolution into becoming a part of our everyday lives. ■

Acknowledgments

The work we have reported here is a product of the efforts of many people. In particular, we wish to thank Walter Aviles, Cagatay Basdogan, Lee Beauregard, David Brock, David DiFranco, Nat Durlach, Donald Green, Chih-Hao Ho, Alexandra Hou, Craig Latimer, Trudy Liu, Thomas Massie, Hugh Morgenbesser, Mark Ottensmeyer, Nick Swarup, Chris Tarr, Tom Wiegand, Evan Wies, Craig Zilles, and numerous other students and colleagues who have been partners in helping us explore the world of haptics. Support for our research has come from a variety of sources including the Office of Naval Research, NAWC/TSD, DARPA, NSF, NIH, NASA, and a variety of fellowship programs.

Further Reading

J.K. Salisbury and M.A. Srinivasan, eds., *Proc. First Phantom Users Group Workshop*, AI Lab Tech. Report 1596 and RLE Tech. Report 612, Massachussetts Institute of Technology, Dec. 1996.

J.K. Salisbury and C. Tarr, "Haptic Rendering of Surfaces Defined by Implicit Functions," *Proc. Sixth Annual Symp. on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, ASME, Fairfield, N.J., 1997 (in press).

M.A. Srinivasan and C. Basdogan, "Haptics in Virtual Environments: Taxonomy, Research Status, and Challenges," *Computers and Graphics*, Special Issue on Haptic Displays in Virtual Environments, Vol. 21, No. 4, 1997 (in press).

Contact Salisbury at jks@ai.mit.edu (http://www.ai.mit.edu/people/jks/haptics.html) and Srinivasan at srini@mit.edu (http://touchlab.mit.edu).

To submit a column to Projects in VR, contact editors Lawrence J. Rosenblum (rosenblum@ait.nrl.navy.mil) or Michael R. Macedonia (mmacedon@crcg.edu). To obtain an author guide go to http://computer.org/cga/edguide.htm#vr