In Design, Specification and Verification of Interactive Systems '99, Eds: D. Duke and A. Puerta, Springer-Verlag Wien, 1999

Haptic Interactions in the Real and Virtual Worlds

M A Srinivasan, C Basdogan, and C-H Ho
Laboratory for Human and Machine Haptics
Massachusetts Institute of Technology
77, Massachusetts Avenue, Cambridge, MA 02139 USA
e mail : srini, basdogan, chihhao@mit.edu

Abstract

In humans or machines, haptics refers to the use of hands for manual sensing and manipulation. Recently, haptic machines that enable the user to touch, feel, and manipulate virtual environments have generated considerable excitement. Synthesizing virtual haptic objects requires an optimal balance between the human haptic ability to sense object properties, computational complexity to render them in real time, and fidelity of the device in delivering the computed mechanical signals. In this paper, we primarily describe the progress made in our "MIT Touch Lab" over the past few years concerning the development of haptic machines, the paradigms and algorithms used in the emerging field of "Computer Haptics" (analogous to Computer Graphics), and experimental results on human perception and performance in multimodal virtual environments. Several ongoing applications such as the development of a surgical simulator and virtual environments shared by multiple users are also described.

1. Introduction

Virtual environments (VEs), generally referred to as virtual reality in the popular press, have caught the imagination of lay public as well as researchers working in a wide variety of disciplines. VEs are computer-generated synthetic environments with which a human user can interact to perform perceptual and motor tasks. A typical VE system consists of a helmet that can project computer-generated visual images and sounds appropriate to the gaze direction, and special gloves with which one can command a computer through hand gestures. The possibility that by wearing such devices, one could be mentally transported to and immersed in virtual worlds built solely through software is both fascinating and powerful. Applications of this technology include a large variety of human activities such as training, education, entertainment, health care, scientific visualization, telecommunication, design, manufacturing and marketing.

Virtual environment systems that engage only the visual and auditory senses of the user are limited in their capability to interact with the user. As in our interactions with the real world, it is desirable to engage the haptic sensorimotor system that not only conveys the sense of touch and feel of objects, but also allows us to manipulate them. In particular, the human hand is a versatile organ that is able to press, grasp, squeeze or stroke objects; it can explore object properties such as surface texture, shape and softness; it can manipulate tools such as a pen or a jack-hammer. Being able to touch, feel, and manipulate objects in an environment, in addition to seeing (and/or hearing) them, gives a sense of compelling immersion in the environment that is otherwise not possible. Real or virtual environments where one is deprived of the touch and feel of objects, seem impoverished and seriously handicap human interaction capabilities.

Haptic interfaces are devices that enable manual interactions with virtual environments or teleoperated remote systems. They are employed for tasks that are usually performed using hands in the real world, such as manual exploration and manipulation of objects. In general, they receive motor action commands from the human user and display appropriate tactual images to the user. Such haptic interactions may or may not be accompanied by the stimulation of other sensory modalities such as vision and audition. It is quite likely that much greater immersion in a virtual environment can be achieved by the synchronous operation of even a simple haptic interface with a visual and/or auditory display, than by large improvements in, say, the fidelity of the visual display alone.

Although computer keyboards, mice, trackballs, and even instrumented gloves available in the market can be thought of as relatively simple haptic interfaces, they can only convey the user's commands to the computer, and are unable to give a natural sense of touch and feel to the user. Recent advances in the development of force-reflecting haptic interface hardware as well as haptic rendering software have caused considerable excitement. The underlying technology is becoming mature and has opened up novel and interesting research areas. In this paper, we primarily describe the progress made in our "MIT Touch Lab" over the past few years on various aspects of haptics relevant to the development of multimodal VEs. In the next section, a basic introduction to *human haptics*, the study of the human hand-brain system relevant to manual exploration and manipulation, is provided. The subsequent section is on *machine haptics*, concerned with the electromechanical devices used as haptic interfaces. Next, the paradigms and algorithms used in the emerging field of *Computer Haptics* that deals with the software for haptic interactions are described. Subsequently, some of our recent experimental results on human perception and performance in multimodal virtual environments are summarized. Finally, several ongoing applications such as the development of a surgical simulator and virtual environments shared by multiple users are described.

2. Human Haptics

In order to develop haptic interfaces that are designed for optimal interactions with the human user, it is necessary to understand the roles played by the mechanical, sensory, motor and cognitive subsystems of the human haptic system. The mechanical structure of the human hand consists of an intricate arrangement of 19 bones connected by almost as many frictionless joints, and covered by soft tissues and skin.

The bones are attached to about 20 each of intrinsic and extrinsic muscles through numerous tendons which serve to activate 22 degrees of freedom of the hand. The sensory system includes large numbers of various classes of receptors and nerve endings in the skin, joints, tendons, and muscles. Appropriate mechanical, thermal or chemical stimuli activate these receptors, causing them to transmit electrical impulses via the afferent neural network to the central nervous system (of which the brain forms a part), which in turn sends commands through the efferent neurons to the muscles for desired motor action.

In any task involving physical contact with an object, be it for exploration or manipulation, the surface and volumetric physical properties of the skin and subcutaneous tissues play important roles in its successful performance. For example, the fingerpad, which is used by primates in almost all precision tasks, consists of ridged skin (about 1 mm thick) that encloses soft tissues composed of mostly fat in a semi-liquid state. As a block of material, fingerpad exhibits complex mechanical behavior -- inhomogeneity, anisotropy, rate and time-dependence. The compliance and frictional properties of the skin together with the sensory and motor capabilities of the hand enable both gliding over a surface to be explored without losing contact, as well as stably grasping smooth objects to be manipulated. The mechanical loading on the skin, the transmission of the mechanical signals through the skin, and their transduction by the cutaneous mechanoreceptors are all strongly dependent on the mechanical properties of the skin and subcutaneous tissues.

Tactual sensory information conveyed to the brain from the hand in contact with an object can be divided into two classes: (i) *tactile information*, referring to the sense of the nature of contact with the object, mediated by the responses of low threshold mechanoreceptors innervating the skin (say, the fingerpad) within and around the contact region; (ii) *kinesthetic information*, referring to the sense of position and motion of limbs along with the associated forces, conveyed by the sensory receptors in the skin around the joints, joint capsules, tendons, and muscles, together with neural signals derived from motor commands. Only tactile information is conveyed when objects are made to contact passive, stationary hand, except for the ever-present kinesthetic information about the limb posture. Only kinesthetic information is conveyed during active, free (i.e., no contact with any object or other regions of skin) motion of the hand, although the absence of tactile information by itself conveys that the motion is free. Even when the two extreme cases mentioned above are included, it is clear that *all* sensory and manipulatory tasks performed actively with the normal hand involve both classes of information. In addition, free nerve endings and specialized receptors which signal skin temperature, mechanical and thermal pain, as well as chemogenic pain and itch are also present.

The control of contact conditions is as important as sensing those conditions for successful performance of any task. In humans, such control action can range from a fast muscle or spinal reflex to a relatively slow conscious deliberate action. In experiments involving lifting of objects held in a pinch grasp, it has been shown that motor actions such as increasing grip force are initiated as rapidly as within 70 msec.

after an object begins to slip relative to the fingerpad, and that the sensory signals from the cutaneous afferents are critical for task performance. Clearly, the mechanical properties of skin and subcutaneous tissues, the rich sensory information provided by a wide variety of sensors that monitor the tasks continuously, and the coupling of this information with the actions of the motor system are responsible for the human abilities of grasping and manipulation. A brief summary of the psychophysical and neurophysiological results available on the human haptic abilities in real environments and the references to the corresponding literature is given in [1].

3. Machine Haptics

Machine haptics refers to the design, construction, and use of machines to replace or augment human hands. Although such machines include autonomous or teleoperated robots, here we focus on haptic interfaces to VEs. Haptic interfaces are devices composed of mechanical components in physical contact with the human body for exchange of information with the human nervous system. In performing tasks with a haptic interface, the human user conveys desired motor actions by physically manipulating the interface, which, in turn, displays tactual sensory information to the user by appropriately stimulating his or her tactile and kinesthetic sensory systems. Thus, in general, haptic interfaces can be viewed as having two basic functions: (1) to measure the positions and contact forces (and time derivatives) of the user's hand (and/or other body parts) and (2) to display contact forces and positions (and/or their spatial and temporal distributions) to the user. Among these position (or, kinematic) and contact force variables, the choice of which ones are the motor action variables (i.e., inputs to the computer) and which are the sensory display variables (i.e., inputs to the human) depends on the hardware and software design, as well as the tasks the interface is employed for. At present, most of the force-reflecting haptic interfaces sense position of their end-effector and display forces to the human user.

A primary classification of our haptic interactions with real or virtual environments that affects interface design consists of the following three basic elements: (i) free motion, where no physical contact is made with objects in the environment; (ii) Contact involving unbalanced resultant forces, such as pressing an object with a fingerpad; (iii) Contact involving self-equilibrating forces, such as squeezing an object in a pinch grasp. Depending on the tasks for which a haptic interface is designed, some or all of these elements will have to be adequately simulated by the interface. For example, grasping and moving an object from one location to another involves all the three elements. The design constraints of a haptic interface are strongly dependent on which of these elements it needs to simulate. Consequently, the interfaces can be classified based on whether they are force-reflecting or not, as well as what types of motions (e.g., how many degrees of freedom) and contact forces they are capable of simulating.

An alternative but important distinction in our haptic interactions with real or virtual environments is whether we touch, feel and manipulate the objects directly or with a tool. The complexity in the design of a haptic interface is seriously affected by which of these two types of interactions it is supposed to simulate. Note that an 'ideal' interface designed to provide realistic simulation of direct haptic exploration and manipulation of objects, would be able to simulate handling with a tool as well. Such an interface would measure positions of, say, the user's hand and display forces, and would have a single hardware configuration (e.g., an exoskeleton) that could be adapted to different tasks by changes in software alone. For example, the act of grasping a hammer would be simulated by monitoring the position and posture of the hand and exerting the appropriate forces on the fingers and palm when the fingers and palm are in the appropriate positions. However, the large number of degrees of freedom of the hand, extreme sensitivities of cutaneous receptors, together with the presence of mass, friction and limitations of sensors and actuators in the interface make such an ideal impossible to achieve with current technology. In contrast, an interface in the form of a tool handle, where reconfigurability within a limited task domain is achieved through both hardware and software changes are quite feasible. Thus, one of the basic distinctions among haptic interfaces is whether they attempt to approximate the ideal exoskeleton or employ the tool handle approach.

Another set of important distinctions concerning haptic interfaces results from a consideration of the force display sub-systems in an interface. Broadly speaking, force display systems can be classified as either (1) ground-based or (2) body-based. Frequently, the distinction between grounding sites is overlooked in the literature. For example, exploration or manipulation of a virtual object requires that force vectors be imposed on the user at multiple regions of contact with the object. Consequently, equal and opposite reaction forces are imposed on the interface. If these forces are self-equilibrating, as in simulating the contact forces that occur when we squeeze an object, then the interface need not be mechanically grounded. However, if the forces are unbalanced, as in pressing a virtual object with a single fingerpad, the equilibrium of the interface requires that it be attached somewhere. A force-reflecting joystick attached to the floor would be a ground-based display, whereas a force reflecting exoskeletal device attached to the user's forearm would be a body-based display. If such an exoskeleton is used to simulate the act of pressing an object with a fingerpad while standing, the stesses that would have been normally experienced by the entire musculoskeletal system would be absent beyond the forearm. The perceptual consequences of such an alteration are not known and warrant investigation.

A survey of the haptic interface devices developed so far is beyond the scope of this paper, but a relatively recent one can be found in [1]. In our MIT Touch Lab, we have developed device hardware, interaction software and psychophysical experiments pertaining to haptic interactions with virtual environments (recent reviews can be found in [1] and [2]). Two specialized devices for performing psychophysical experiments, the linear and planar graspers, have been developed. The linear grasper is capable of simulating fundamental mechanical properties of objects such as compliance, viscosity and mass during haptic interactions. Virtual walls and corners were simulated using the planar grasper, in addition to the simulation of two springs within its workspace. The PHANToM, another haptic display device developed at the MIT Artificial Intelligence Laboratory [3], has been used to prototype a wide range of force-based haptic display primitives. A variety of haptic rendering algorithms for displaying the shape, compliance, texture, and friction of solid surfaces have been implemented on the PHANToM [2, 4]. All the three devices have been used to perform psychophysical experiments aimed at characterizing the sensorimotor abilities of the human user and the effectiveness of computationally efficient rendering algorithms in conveying the desired object properties to the human user.

4. Computer Haptics

Computer Haptics is a rapidly emerging area of research that is concerned with the techniques and processes associated with generating and displaying the touch and feel of virtual objects to a human operator through a force reflecting device. Analogous to computer graphics, it deals with models and behavior of virtual objects together with rendering algorithms for real-time display. It includes the software architecture needed not only for haptic interactions but also their synchronization with visual and other display modalities.

In order to develop effective software architectures for multimodal VEs, we have experimented with multi-threading (on Windows NT platform) and multi-processing (on UNIX platform) techniques and have successfully separated the visual and haptic servo loops. Our experience is that both techniques enable the system to update graphics process at almost constant rates, while running the haptic process in the background. We are able to achieve good visual rendering rates (30 to 60 Hz), high haptic rendering rates (more than 1 kHz), and stable haptic interactions. Although creating a separate process for each modality requires more programming effort, it enables the user to display the graphics and/or haptics on any desired machine(s), even those in different locations, as long as the physical communication between them is provided through a cable. Programming with threads takes less effort, but they are not as flexible as processes.

We have also developed a graphical interface that enables a user to construct virtual environments by means of user-defined text file, toggle stereo visualization, save the virtual environment and quit from the application. This application program was written in C/C++ and utilizes the libraries of (1) Open Inventor (from Silicon Graphics Inc.) for graphical display of virtual objects, (2) ViewKit (from Silicon Graphics Inc.) for constructing the graphical user interface (e.g. menu items, dialog boxes, etc.), and (3) Parallel Virtual Machine (PVM), a well-known public domain package, for establishing the digital communication between the haptic and visual processes. The user can load objects into the scene, and assign simple visual and haptic properties to the objects using this text file. Following the construction of the scene using the text file, the user can interactively translate, rotate, and scale objects, and the interface will automatically update both the visual and haptic models.

Two types of haptic rendering techniques have been developed: point-based and ray-based. In point-based haptic interactions, only the end point of haptic device, also known as the end effector point or haptic interface point (HIP), interacts with objects

[5, 6]. Since the virtual surfaces have finite stiffnesses, the end point of the haptic device penetrates into the object after collision. In our laboratory, a set of rule-based algorithms has been developed for fast detection of collisions. We use a hierarchical database, multi-threading techniques, and efficient search procedures to reduce the computational time and make the computations almost independent of the number of polygons of the polyhedron representing the object. Each time the user moves the generic probe of the haptic device, the collision detection algorithms check to see if the end point is inside the virtual object. In ray-based haptic interactions developed in our laboratory, the generic probe of the haptic device is modelled as a finite ray whose orientation is taken into account, and the collisions are checked between the ray and the objects [7]. Both techniques have advantages and disadvantages. For example, it is computationally less expensive to render 3D objects using point-based technique. Hence, we achieve higher haptic update rates. On the other hand, the ray-based haptic interaction technique handles side collisions and can provide additional haptic cues for conveying to the user the shape of objects.

Once the software and hardware components were put together for integrating multiple modalities, we focussed on developing techniques for generating multimodal stimuli. Our interest in generating multimodal stimuli is three-fold: (a) to develop new haptic rendering techniques to display shape, texture, and compliance characteristics of virtual objects, (b) to utilize these techniques in our experiments on human perception and performance to study multimodal interactions, and (c) to explore a variety of applications of multimodal virtual environments where haptics adds value. Our progress in the area of haptic rendering is summarized under three headings: shape, texture, and compliance.

4.1 Shape

When smooth and continuous object shapes are approximated by polyhedra for haptic rendering, the user does not perceive the intended shape. Instead, the discrete edges between polygons as well as the planar faces of the polygons are felt. To minimize such undesirable effects, we have proposed *force shading* [8]. In this method, which falls within the general class of force mapping techniques, the force vector is interpolated over the polygonal surfaces such that its direction varies continuously. Consequently, the surfaces of virtual objects feel smoother than their original polyhedral representations. This technique is analogous to Phong shading in computer graphics.

4.2 Texture

Since a wide variety of physical and chemical properties give rise to real-world textures, a variety of techniques are needed to simulate them visually and haptically in VEs. Haptic texturing is a method of simulating surface properties of objects in virtual environments in order to provide the user with the feel of macro and micro surface textures. We have developed two basic approaches: *force perturbation*, where the direction of the displayed force vector is perturbed, and *displacement mapping*, where the microgeometry of the surface is perturbed [7]. Using these methods, we have successfully displayed textures based on Fourier series, filtered white noise, and fractals. But the display of haptic textures using the force perturbation technique was effective only in a certain range (0.5 mm to 5.0 mm in height). To extend the range of haptic textures that can be displayed, we have modified the algorithm to include the calculation of the location of the point closest to the object surface prior to collision detection. Using this additional information, we are now able to render macro textures (> 5.0 mm height) as well. We have also experimented with 2D reaction-diffusion texture models used in computer graphics and successfully implemented them for haptics to generate new types of haptic textures. The reaction-diffusion model consists of a set of differential equations that can be integrated in time to generate texture fields. Moreover, we have developed techniques to extend our work on 2D reaction-diffusion textures to three dimensional space. We have also studied some of the image and signal processing techniques frequently used in computer graphics to convolve 2D images of spots (i.e. simple 2D geometric primitives such as circles, squares, and triangles) with noise functions in order to generate a new class of haptic textures.

In summary, the following texture rendering techniques have been developed: a) force perturbation, b) displacement mapping. Using these rendering techniques, we can display the following types of synthetic haptic textures: a) periodic and aperiodic haptic textures based on Fourier series approach, b) noise textures (based on the filtered white noise function), c) fractal textures, d) reaction-diffusion textures (a set of differential equations are solved in advance to generate a texture field that can be mapped onto the 3D surface of the object), and e) spot-noise textures (the noise function is convolved with 2D images of spots to generate distorted spots that can be displayed haptically). In addition, we have developed image-based haptic textures (the grey scale values of an image are used to generate texture fields that can be mapped onto the surface of 3D objects) as well as methods to display static and dynamic friction.

4.3 Compliance

We have developed procedures for simulating compliant objects in virtual environments. The developed algorithms deal directly with geometry of 3D surfaces and their compliance characteristics, as well as the display of appropriate reaction forces, to convey to the user a feeling of touch and force sensations for soft objects. The compliant rendering technique has two components: (1) the deformation model to display the surface deformation profile graphically; and (2) the force model to display the interaction forces via the haptic interface. The deformation model estimates the direction and the amount of deformation (displacement vector) of each node (i.e. a vertex) of the surface when it is manipulated with the generic probe of the haptic interface device. We utilize a polynomial model or a spline-based model to compute the displacement vector of each node and to visually display deformations. In the force model, a network of springs is utilized to compute the direction and magnitude of the force vector at the node that is closest to the contact point. The techniques described here enable the user to interactively deform compliant surfaces in real-time and feel the reaction forces.

Using the user interface and haptic rendering techniques described in the previous sections, we have designed experiments to investigate human performance involving multimodal interactions in virtual environments. The user interface has enabled several experimenters to rapidly load virtual objects into desired experimental scenarios, interactively manipulate (translate, rotate, scale) them, and attach sophisticated material and visual properties to the virtual objects.

5. Experiments on Human Perception and Performance

Concurrent with the technology development that enables one to realize a wider variety of haptic interfaces, it is necessary to characterize, understand, and model the basic psychophysical behavior of the human haptic system. Without appropriate knowledge in this area, it is impossible to determine specifications for the design of effective haptic interfaces. In addition, because multimodal sensorimotor involvement constitutes a key feature of VE systems, it is obviously important to understand multimodal interactions. Furthermore, because the availability of force feedback in multimodal VE interfaces is relatively new, knowledge about interactions involving force feedback is relatively limited. In general, research in this area not only provides important background for VE design, but the availability of multimodal interfaces with force feedback provides a unique opportunity to study multimodal sensorimotor interactions.

5.1 Purely Haptic Interactions

Using the Linear Grasper, a haptic interface device, psychophysical experiments have been carried out to measure human haptic resolution in discriminating fundamental physical properties of objects through active touch. The subjects utilized their thumb and index fingers to grasp and squeeze two plates of the Linear Grasper, which was programmed to simulate various values of the stiffness, viscosity, or mass of virtual objects. During the experiments, haptic motor performance data in terms of applied forces, velocities, and accelerations were simultaneously recorded.

The Just Noticeable Difference (JND), a commonly accepted measure of human sensory resolution, was found to be about 7% for stiffness, 12% for viscosity, and 20% for mass. The motor data indicated that subjects used the same motor strategy when discriminating any of these material properties. Further analysis of the results has led to the postulation of a single sensorimotor strategy capable of explaining both the sensory resolution results and motor performance data obtained in the experiments. This hypothesis, called the "Temporal force control - spatial force discrimination (TFC-SFD) hypothesis," states that subjects apply the same temporal profile of forces to all stimuli and discriminate physical object properties on the basis of differences in the resulting spatial profiles of these forces. A special case of this hypothesis is that when humans discriminate stiffness, viscosity or mass, they do so by discriminating the mechanical work needed for actually deforming the objects. Implications of these results to the design of virtual environments include specifications on how accurately the dynamics of virtual objects need to be simulated and what parameter values will ensure discriminable objects.

To explore the possibility that multisensory information may be useful in expanding the range and quality of haptic experience in virtual environments, experiments have been conducted to assess the influence of auditory and visual information on the perception of object stiffness through a haptic interface, as described in the next two sections.

5.2 Haptic-Auditory Interactions

We have previously shown that when virtual objects are tapped through a haptic interface, contact sounds can influence the perception of object stiffness [9]. In another series of experiments, we investigated the effect of the timing of a contact sound on the perception of stiffness of a virtual surface. The PHANToM was used to display virtual haptic surfaces with constant stiffness. Subjects heard a contact sound lasting 130 ms through headphones every time they touched a surface. Based on our earlier work on stiffness discrimination, we initially hypothesized that presenting a contact sound prior to actual impact creates the perception of a less stiff surface, whereas presenting a contact sound after actual impact creates the perception of a stiffer surface. However, the findings indicate that both pre-contact and post-contact sounds result in the perceptual illusion that the surface is less stiff than when the sound is presented at contact.

5.3 Haptic-Visual Interactions

Previously we have shown how the perception of haptic stiffness is strongly influenced by the visual display of object deformation [10]. An important implication of these results for multimodal VEs is that by skewing the relationship between the haptic and visual displays, the range of object properties that can be effectively conveyed to the user can be significantly enhanced. For example, although the range of object stiffness that can be displayed by a haptic interface is limited by the force-bandwidth of the interface, the range perceived by the subject can be effectively increased by reducing the visual deformation of the object.

In continuing this line of investigation on how vision affects haptic perception, we have conducted two new sets of experiments to test the effect of perspective on the perception of geometric and material properties of 3D objects [11]. Virtual slots of varying length and buttons of varying stiffness were displayed to the subjects, who then were asked to discriminate their size and stiffness respectively using visual and/or haptic cues. The results of the size experiments show that under vision alone, farther objects are perceived to be smaller due to perspective cues and the addition of haptic feedback reduces this visual bias. Similarly, the results of the stiffness experiments show that compliant objects that are farther are perceived to be softer when there is only haptic feedback and the addition of visual feedback reduces this haptic bias. Hence, we conclude that our visual and haptic systems compensate for each other such that the sensory information that comes from visual and haptic channels is fused in an optimal manner. In particular, the result that the farther objects are perceived to be softer when only haptic cues are present is interesting and suggests a new concept of *haptic perspective*. To ensure that this result was not an artifact of the robot arm (i.e. position and force errors due to the kinematics of the haptic device) or our experimental design, we performed three different tests, but the result did not change.

5.4 Haptics Across the World Wide Web

In order to make haptics and our research studies accessible and transferable to the others, we opted to integrate haptics into the Web. A demonstration version of the visual-haptic experiment as described above using the PHANToM haptic interface was developed for use across the World Wide Web. The program was written in Java, using multi-threading to create separate visual and haptic control loops, thereby increasing the speed of the haptics loop to keep the program stable despite its graphics overhead. The application program was placed on the Laboratory of Human and Machine Haptics web page (http://touchlab.mit.edu), to be executed by any remote user with a PHANToM and a Windows NT computer running Netscape for WWW access. Remote users could download a dynamic link library and some Java classes from the web page to their computer, and then run the program in their web browser. Users were asked to discriminate the stiffness of sets of two springs, displayed visually on the screen and haptically with the PHANToM, and to send in their responses via an e-mail window in the web page. Thus, we now have the ability to perform perceptual experiments with multimodal VEs across the internet.

6. Applications

We describe below two examples of how multimodal virtual environment systems are being used in our laboratory to explore novel application areas.

6.1 Simulation of Minimally Invasive Surgical Procedures

Research in the area of computer assisted surgery and surgical simulation has mainly focused on developing 3D geometrical models of the human body from 2D medical images, visualization of internal structures for educational and preoperative surgical planning purposes, and graphical display of soft tissue behavior in real time. Conveying to the surgeon the touch and force sensations with the use of haptic interfaces has not been investigated in detail. We have developed a set of haptic rendering algorithms for simulating "surgical instrument - soft tissue" interactions. Although the focus of the study is the development of algorithms for simulation of laparoscopic procedures, the developed techniques are also useful in simulating other medical procedures involving touch and feel of soft tissues. The proposed force-reflecting soft tissue models are in various fidelities and have been developed to simulate the behavior of elastically deformable objects in virtual environments. The developed algorithms deal directly with geometry of anatomical organs, surface and compliance characteristics of tissues, and the estimation of appropriate reaction forces to convey to the user a feeling of touch and force sensations [12].

The hardware components of the set-up include a personal computer (300 MHz, dual Pentium processor) with a high-end 3D graphics accelerator, a force-feedback device (PHANToM from SensAble Technologies Inc.) to simulate haptic sensations. During the simulations, the user manipulates the generic stylus of the force-feedback device to simulate the movements of a surgical instrument and to feel its interactions with the computer generated anatomical organs. The associated deformations of the organs are displayed on the computer monitor and reaction forces are fed back to the user through the haptic interface. The software was written in C/C++, using multi-threading techniques to create separate visual and haptic control loops, thereby increasing the haptics servo rate (varies from 500 Hz to 2 kHz) while simultaneously satisfying the requirements of graphics update rate of at least 30 Hz. Recently, we have made progress in two areas:

Development of "thin-walled" models to simulate tissue deformations and to compute reaction forces: In the language of mechanics, the "thin-walled" structures are broadly classified into membranes, structures with essentially no bending stiffness compared to the in-plane stiffness, and shells, structures in which bending behavior is also important. We have used such structures as organ models to simplify the mechanistic computations while at the same time retaining some of the physics of tool–tissue interactions. In our implementation, triangular elements are used to represent the organ geometry and the virtual work principle is used to derive the incremental equations of motion [13]. The initial results suggest that "thin-walled" models can predict nonlinear behavior of tissues.

Development of algorithms to simulate tissue cutting and bleeding: We have developed computationally fast algorithms to display (1) tissue cutting and (2) bleeding in virtual environments with applications to laparoscopic surgery. Cutting through soft tissue generates an infinitesimally thin slit until the sides of the surface are separated from each other. Simulation of an incision through tissue surface is modelled in three steps: first, the collisions between the instrument and the tissue surface are detected as the simulated cutting tool passes through. Then, the vertices along the cutting path are duplicated. Finally, a simple elastic tissue model is used to separate the vertices from each other to reveal the cut. Accurate simulation of bleeding is a challenging problem because of the complexities of the circulatory system and the physics of viscous fluid flow. There are several fluid flow models described in the literature, but most of them are computationally slow and do not specifically address the problem of blood flowing over soft tissues. We have reviewed the existing models, and have adapted them to our specific task. The key characteristics of our blood flow model are a visually realistic display and real-time computational performance. To display bleeding in virtual environments, we developed a surface flow algorithm. This method is based on a simplified form of the Navier-Stokes equations governing viscous fluid flow. The simplification of these partial differential equations results in a wave equation that can be solved efficiently, in real-time, with finite difference techniques. The solution describes the flow of blood over the polyhedral surfaces representing the anatomical structures and is displayed as a continuous polyhedral surface drawn over the anatomy [14].

6.2 The Role of Haptics in Shared Virtual Environments

We have conducted a set of human experiments to investigate the role of haptics in shared virtual environments (SVEs). Our efforts were aimed at exploring (1) whether haptic communication through force feedback can facilitate a sense of togetherness between two people at different locations while interacting with each other in SVEs, (2) if so, what types of haptic communication/negotiation strategies they follow, and (3) if gender, personality, or emotional experiences of users can affect the haptic communication in SVEs. The experiment concerns a scenario where two people, at remote sites, co-operate to perform a joint task in a SVE. The experiments are abstractions from real situations in order to create a more controlled environment suitable for explanatory studies in the laboratory. During the experiment, subjects were not allowed to meet their remote partner, and did not know where their partner was located. The participants were in different rooms but saw the same visual scene on their monitor and felt the objects in the scene via a force feedback device, the PHANToM.

The goal of the task was to move a ring with the help of another person without touching a wire. A ring, a wire, and two cursors attached to the ring were displayed to the subjects. Haptic interactions between cursors as well as between cursor and the ring were modelled using a spring-damper system and a point-based haptic rendering technique [15]. Subjects were asked to move the ring back and forth on the wire many times, in collaboration with each other such that contact between the wire and the ring was minimized or avoided. If the ring touched the wire, the colors of the ring and the surrounding walls were changed to red to warn the subject of an error. They were changed back to their original colors when the subjects corrected the position of the ring. To hold the ring, both subjects needed to press on the ring towards each other above a threshold force. If they did not press on the ring at the same time, the ring did not move and its color was changed to gray to warn them. To move the ring along the wire, they each needed to apply an additional lateral force.

Two sensory conditions have been explored to investigate the effect of haptic communication on the sense of togetherness: (1) both visual and haptic feedback provided to the participants; (2) only visual feedback was provided to the participants. Performance and subjective measures were developed to quantify the role of haptic feedback in SVEs. Performance measure was derived from the following measurements: (1) total amount time takes to complete task, (2) the ratio of erroneous-time to total time. Several *subjective* questions were asked, through a questionnaire, in four categories including their (1) performance, (2) their sense of 'being together', (3) emotional reactions, and (4) personality profile. Each of the questions in categories 1,2, and 3 were rated on a 1-7 scale. Subjective measures were correlated with the performance measures to deduce conclusions on the effect of haptic feedback to the task performance and the sense of being with someone in SVEs . The results suggest that haptic feedback significantly improves the performance and contributes to the feeling of "sense of togetherness" in SVEs.

7. Future Research

Many of the issues concerning the development of haptic interfaces and computer haptics are summarized in [2]. Although both ground-based and exoskeletal force-reflecting haptic interface devices are available in the market, further improvements in range, resolution, and frequency bandwidth of these devices are needed to match their performance with that of the human user. In moving towards realistic haptic displays that mimic direct natural touch, tactile displays are probably the most challenging among the technologies that need to be developed. The emerging field of micro-mechanical systems holds promise for providing very fine arrays of tactile stimulators. In collaboration with researchers at Carnegie-Mellon University, we are developing electrostatic-pneumatic actuators for use in sensors and as surface-normal actuators. Although capable of relatively small forces and deflections, arrays of such actuators integrated with addressing electronics would be inexpensive, light-weight, and compact enough to be worn without significantly impeding hand movement or function. In addition, the current technology makes feasible a 20 x 20 array of individually controlled stimulators on a 1 cm x 1cm chip.

In the area of computer haptics, the current models of virtual objects that can be displayed haptically in real-time are quite simplistic compared to the static and dynamic behavior of objects in the real world. Computationally efficient models and interaction techniques that result in real-time haptic displays that match the human perceptual capabilities in accuracy and resolution will continue to be a challenge, even with the current rate of increase in processing speeds. This is because the complexity of the models, such as in detecting collisions of moving multiple objects or in performing a mechanistic analysis of a deformable object in real-time, can be arbitrarily high. Synchronization of the visual, auditory and haptic displays can be problematic, because each modality requires different types of approximations to simulate the same physical phenomenon. Use of multiple processors with shared memory and/or multi-threading seems to be essential. To have haptics across the internet in a manner that is useful to a large number of users, standardized protocols for distributed VEs should include haptics explicitly.

Due to inherent hardware limitations, haptic interfaces can only deliver stimuli that approximate our interactions with the real environment. It does not, however, follow that synthesized haptic experiences created through the haptic interfaces necessarily feel unreal to the user. Consider an analogy with the synthesized visual experiences obtained while watching television or playing a video game. While visual stimuli in the real world are continuous in space and time, these visual interfaces project images at the rate of about 30 frames/sec. Yet, we experience a sense of realism and even a sense of telepresence because we are able to exploit the limitations of the human visual apparatus. The hope that the necessary approximations in generating synthesized haptic experiences will be adequate for a particular task is based on the fact that the human haptic system has limitations that can be similarly exploited. To determine the nature of these approximations, or, in other words, to find out what we can get away with in creating synthetic haptic experiences, quantitative human studies are essential to assess which types of stimulation provide the most useful and profound haptic cues for the task at hand.

Acknowledgements

This work was carried out under NAWC/TSD and ONR contracts. Authors would like to acknowledge Lee Beauregard, David Brock, Suvranu De, David DiFranco, Alexandra Hou, Hugh Morgenbesser, and Wan-Chen Wu for various parts of the work described in here. We also want to thank Ken Salisbury and his group at the MIT Artificial Intelligence Laboratory for their collaboration in various projects.

References

[1] Srinivasan, M A, Haptic Interfaces, In Virtual Reality: Scientific and Technical Challenges, Eds: N. I. Durlach and A. S. Mavor, Report of the Committee on Virtual Reality Research and Development, National Research Council, National Academy Press, 1995.

[2] Srinivasan, M A and Basdogan, C, Haptics in Virtual Environments: Taxonomy, Research Status, and Challenges, Computers and Graphics, Vol. 21, No. 4, 1997.

[3] Massie, T H and Salisbury, J K, The PHANToM Haptic Interface: A Device for Probing Virtual Objects. Proceedings of the ASME Dynamic Systems and Control Division, Vol. 55(1), pp. 295-301, 1997.

[4] Salisbury, J K and Srinivasan, M A, Phantom-Based Haptic Interaction with Virtual Objects, IEEE Computer Graphics and Applications, Vol. 17, No. 5, 1997.

[5] Ho, C-H, Basdogan, C and Srinivasan, M A, Haptic Rendering: Point- and Ray-based Interactions, Proceedings of the Second PHANToM User's Group Workshop, October, 1997.

[6] Ho, C-H, Basdogan, C and Srinivasan, M A, Efficient Point-based Rendering Techniques for Haptic Display of Virtual Objects, Presence, 1999 (in press).

[7] Basdogan, C, Ho, C-H and Srinivasan M A, A Ray-based Haptic Rendering Technique for Displaying Shape and Texture of 3D Objects in Virtual Environments, Proceedings of the ASME Dynamic Systems and Control Division, Ed. G. Rizzoni, DSC-Vol. 61, pp. 77-84, ASME, 1997.

[8] Morgenbesser, H B and Srinivasan, M A, Force Shading for Haptic Shape Perception, Proceedings of the ASME Dynamic Systems and Control Division, DSC-Vol. 58, pp. 407-412, ASME, 1996.

[9] DiFranco, D E, Beauregard, G L and Srinivasan, M A, The Effect of Auditory Cues on the Haptic Perception of Stiffness in Virtual Environments, Proceedings of the ASME Dynamic Systems and Control Division, Ed. G. Rizzoni, DSC-Vol. 61, pp. 17-22, ASME, 1997.

[10] Srinivasan, M A, Beauregard, G L and Brock, D L, The Impact of Visual Information on Haptic Perception of Stiffness in Virtual Environments, Proceedings of the ASME Dynamic Systems and Control Division, DSC-Vol. 58, pp. 555-559, ASME, 1996.

[11] Wu, W-C, Basdogan, C and Srinivasan, M A, Effect of Visual Perspective on the Visual and Haptic Perception of Size and Stiffness in Virtual Environments, Proceedings of the ASME Dynamic Systems and Control Division, 1999 (in press).

[12] Basdogan, C, Ho, C-H, Srinivasan, M A, Small, S D and Dawson, S L, Force Interactions in Laparoscopic Simulations: Haptic Rendering of Soft Tissues, Proceedings of the Medicine Meets Virtual Reality (MMVR'98) VI Conference, San Diego, CA, pp. 385-391, January, 1998.

[13] De, S and Srinivasan, M A, Thin Walled Models for Haptic and Graphical Rendering of Soft Tissues in Surgical Simulations, Proceedings of Medicine Meets Virtual Reality Conference 7, San Francisco, CA., January, 1999.

[14] Basdogan, C, Ho, C-H and Srinivasan, M A, Simulation of Tissue Cutting and Bleeding for Laparoscopic Surgery Using Auxiliary Surfaces, Proceedings of Medicine Meets Virtual Reality Conference 7, San Francisco, CA., January, 1999.

[15] Ho, C-H, Basdogan, C, Slater, M, Durlach, M, and Srinivasan, M A, The Influence of Haptic Communication on the Sense of Being Together, Workshop on Presence in Shared Virtual Environments, BT Labs, Ipswich, UK, June, 1998.