

VIRTUAL ENVIRONMENT TECHNOLOGY FOR TRAINING (VETT)

Sponsor

U.S. Navy Office of Naval Research

Grant Nos. N61339-96-K-0002, N61339-96-K-0003, N00014-97-1-0635, N00014-97-1-0655

Project Staff

Nathaniel I. Durlach, Mandayam A. Srinivasan, Thomas E. v. Wiegand, Loraine Delhorne, W. L. Sachtler, Cagatay Basdogan, David W. Schloerb, Adrienne Slaughter, Chih-Hao Ho, Steve Wang, Hanfeng Yuan, Wan-Chen Wu

1 Introduction

This work is being conducted within Virtual Environment Technology for Training (VETT), a large interdisciplinary, inter-institutional program which is studying the use of virtual environment (VE) technology to improve Navy training. At RLE, two mutually supporting components of this program are being pursued: (1) Enabling Research on the Human Operator (ERHO) and (2) Development of haptic interfaces and multimodal virtual environments. The ERHO component is concerned with how human perception and performance in virtual environments (VEs) depend upon (1) the physical characteristics of the VE system, (2) the task being performed, and (3) the user's experience with the system and the task. To the extent that the ERHO research is successful, the results will not only provide important information for the design and evaluation of VE training systems, but also for VE systems in general. The second component is focussed on the development of haptic interfaces that enable the user to touch, feel and manipulate objects in VEs. Software is being developed to generate haptic stimuli and to integrate visual, auditory, and haptic displays. Experiments on multimodal illusions due to interactions between haptic and visual or auditory displays have also been conducted. The progress in ERHO, haptic interface development, and multimodal VEs is described in the following subsections.

2 Visual Depth Perception in VEs

Background in this area is available in RLE Annual Report 141, pp 337-338. During the past year, attention in this area has been focused on further data analysis and the preparation of results for publication (v. Wiegand et al, 1999; Yuan et al, 1999; Schloerb, 1999).

3 Part-Tasks Trainer for Position-Velocity Transformations

Background in this area is available in RLE Annual Report 141, p 338. During the past year, extensive data have been collected on the ability of subjects to quickly and accurately estimate the effects of coordinate transforms on velocity vectors. Data analysis and preparation of a report on these results will be completed during the coming year.

4 The Role of Haptics in Learning to Perform Cognitive Tasks

Background in this area is available in RLE Annual Report 141, pp 340-345. The results of the most recent experiments conducted in this area are now being prepared for publication (Delhorne et al, 1999).

5 Conveying the Touch and Feel of Virtual Objects

Haptic displays are emerging as effective interaction aids for improving the realism of virtual worlds. Being able to touch, feel, and manipulate objects in virtual environments have a large number of exciting applications. The underlying technology, both in terms of electromechanical hardware and computer software, is becoming mature and has opened up novel and interesting research areas. The following sections summarize the progress over the past few years in our "Touch Lab" at RLE. A major advance has been the birth of a new discipline, *Computer Haptics* (analogous to computer graphics), that is concerned with the techniques and processes associated with generating and displaying haptic stimuli to the human user.

Over the past few years, we have developed device hardware, interaction software and psychophysical experiments pertaining to haptic interactions with virtual environments (recent reviews on haptic interfaces can be found in Srinivasan, 1995 and Srinivasan and Basdogan, 1997). Two major 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 along a linear track. 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 previously by Dr. Salisbury's group at the MIT Artificial Intelligence Laboratory, 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. 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.

Haptic Rendering Techniques: Point and Ray-Based Interactions

Haptic rendering, a relatively new area of research, is concerned with real-time display of the touch and feel of virtual objects to a human operator through a force reflecting device. It can be considered as a sub-discipline of Computer Haptics. A major component of the rendering methods developed in our laboratory is a set of rule-based algorithms for detecting collisions between the generic probe (end-effector) of a force-reflecting robotic device and objects in VEs. 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. Our haptic texturing techniques enable us to map surface properties onto the surface of polyhedral objects. 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. Since the virtual surfaces have finite stiffnesses, the end point of the haptic device penetrates into the object after collision. 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 (Ho, et al., 1997 and 1999). In ray-based haptic interactions, the generic probe of the haptic device is modeled as a finite ray whose orientation is taken into account, and the collisions are checked between the ray and the objects (Basdogan, et al., 1997 and Ho, et al., 2000). 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 servo 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.

6 Constructing Multimodal Virtual Environments

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.

Using the haptic rendering techniques and the user interface described above, we have designed experiments to investigate human performance involving multimodal interactions in virtual environments. The user interface 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.

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 two-fold: (a) we would like to develop new haptic rendering techniques to display shape, texture, and compliance characteristics of virtual objects, and (b) utilize these techniques in our experiments on human perception and performance to study multimodal interactions. Our progress in this area is summarized under two headings: texture and compliance.

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. Using these methods, we have successfully displayed textures based on Fourier series, filtered white noise, and fractals. 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.

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 gray scale values of an image are used to generate texture fields that can be mapped onto the surface of 3D objects).

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. These techniques enable the user to interactively deform compliant surfaces in real-time and feel the reaction forces.

This set of algorithms has been integrated with GHOST, a commercially available software package for haptic rendering. This configuration combines the higher rendering speed of our algorithms with the object-oriented flexibility of the commercial software, facilitating development of new applications, such as a simulation of a spinal injection.

7 Experimental Studies on Interactions Involving Force Feedback

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.

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 visual and auditory information on the perception of object stiffness through a haptic interface. We have previously shown that visual sensing of object deformation dominates kinesthetic sense of hand position and results in a dramatic misperception of object stiffness when the visual display is intentionally skewed (Srinivasan, et al, 1996). However, the influence of contact sounds on the perception of object stiffness is not as dramatic when tapping virtual objects through a haptic interface (DiFranco, et al, 1997). Over the past year, we have designed and conducted more experiments to explore the human haptic resolution as well as the effect of haptic-auditory and haptic-visual interactions on human perception and performance in virtual environments.

7.1 Haptic Psychophysics

Human abilities and limitations play an important role in determining the design specifications for the hardware and software that enable haptic interactions in VE. With this viewpoint, psychophysical experiments have been carried out over the past few years with a haptic interface to measure human haptic resolution in discriminating fundamental physical properties of objects through active touch. A computer controlled electromechanical apparatus, called The Linear Grasper, was developed and used in these experiments. 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.

7.2 Haptic-Auditory Interactions

In this 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, a six degree of freedom haptic interface with 3 degrees of active force feedback, 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.

7.3 Haptic-visual interactions

Previously we have shown how the perception of haptic stiffness is influenced by the visual display of object deformation (Srinivasan, et al, 1996). 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.

In continuing this line of investigation on how vision affects haptic perception, we designed two sets of experiments to test the effect of perspective on the perception of geometric and material properties of 3D objects (Wu, et al., 1999). 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 perhaps suggests a new concept of *haptic perspective*.

7.4 Learning ship dynamics through haptic sensing

We have designed experiments to understand whether the sensing of environmental forces through a haptic device improves our ability to manipulate objects and navigate in VEs. In experiments, the user controls the direction of a surface ship to navigate inside a complex virtual environment where there are other ships and bridges around. Preliminary results indicate that sensing and/or observing environmental forces that act on the simulated objects provides the user with additional information on how to navigate in unstructured environment. Although we have not observed a significant difference between visual only feedback and haptic only feedback of environmental forces, either one of them was better than no feedback.

The preliminary results suggest that the learning of ship dynamics can be improved even if the environmental forces acting on the ship are passively displayed to the trainee through a haptic device (i.e. the haptic interaction that does not involve the active manipulation of objects through the device).

7.5 Control of Object Dynamics via Visual and Haptic Feedback

We have designed motor-control experiments to understand the nature of force control strategies in manipulating objects through an active joystick in VEs. The experimental task involves the moving of a virtual object from one point to another by following a pre-defined path under six different sensory conditions. Preliminary results suggest that the best performances were obtained when the joystick center and the cross bar that provides visual cues about the direction and magnitude of the applied force were transferred with the object center. Moreover, as is to be expected, the performance was better with force feedback than no force feedback. The results of the study will enable us to find the best combination of haptic and visual feedback for controlling the dynamics of objects with a joystick that can function in various ways.

8 Investigations on Expanding the Perceived Workspace

A larger haptic workspace would be useful in several haptic applications. Unfortunately, the physical haptic workspace that is provided by the currently available haptic devices is limited due to several reasons (e.g. dimensions of the mechanical links of the device for effective force feedback and the reachable distances by the human user). We are currently working on the concept of expanding the perceived haptic workspace using our rendering algorithms and the user interface.

We briefly summarize our approach here on how the perceived space can be extended. Currently, we are testing the feasibility of these approaches for integration into our existing multimodal software.

8.1 Scaling the cursor movement

If the cursor movements, determined by the end point of the haptic device, are scaled appropriately before they are displayed in the visual space, then the haptic space may be perceived as bigger than the actual one. We define a scaling ratio as the size of visual workspace to the size of haptic workspace (SVW/SHW) and we set its value higher than one to make sure that the haptic workspace perceived by the human user becomes bigger than its physical dimensions. Hence, the user will feel as if he/she is traveling faster in the visual workspace as a result of this scaling.

8.2 Scaling the visual size of stylus

Similarly, the visual size of the stylus can be scaled by the ratio of workspaces to extend the reachable visual workspace. In the past, we have developed a ray-based rendering technique to haptically interact with 3D objects in virtual environments. In this technique, the stylus of the haptic interface device is modeled as a line segment and the collisions are detected between the line segment and the 3D objects in the scene. Since our ray-based rendering technique relies on detecting collisions between a line segment model of the stylus and the 3D objects and it can handle side collisions, small movements of the hand in haptic space will result in the display of longer stylus movements in visual workspace. This can be imagined by assuming that a blind person is exploring the surface of a 3D object using a special stick that can be extended automatically as he/she presses a button on the stick to feel the objects that are further away.

8.3 Haptic Perception of Force and Torque on a Stylus

To reach objects outside the limited workspace of our haptic device, a virtual stylus is constructed by linearly extending the stylus. Some causes of vibration have been studied and solved to improve the quality of display. Human psychophysical experiments were designed and conducted to investigate: (1) the feasibility of using one Phantom (a haptic device with only force reflection) or combining two (specifying extra force to result in corresponding torque) in applying ray-based rendering (2) the effects of force/torque on exploring objects with a stylus held in hand. Two Phantoms were connected by their styluses (the pen-like end effectors) and ray-based rendering was applied. In this way, the force and torque can be split successfully and applied individually. Besides, force display with either one or two Phantoms was also applied to evaluate the importance of using one/two Phantoms in virtual object exploration. The displayed object in this experiment was a thin vertical plate switching to one of three possible positions (front, at, and behind the hand). The results have shown that using one Phantom can only be applied in exploring object close to the tip of the stylus and using two Phantoms combined can give the best cues due to its consistent force-torque display. Torque was found to provide vital cues for detecting object positions.

Two Phantoms were connected by their styluses and activated. By allocating different ratios of forces sent to each, four haptic feedback conditions were displayed: "tip force" is the case when the resultant force was sent through the Phantom connected to the tip of the stylus; "force+torque" is the case when the forces sent to each machine constitute the resultant force and torque; "force only" is the case when the force sent to each machine result in zero torque respective to the hand holding position; and "torque only" is the case when the forces sent to machines result in zero resultant force.

The experimental results showed that only when force and torque were both applied can the subjects correctly detect the location of the object relative to the hand.

9 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 in Srinivasan, et al., (1996) was developed to be used 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 web access. Remote users can 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 are 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. In a related project, the industry-standard Virtual Reality Modeling Language (VRML version 2.0) was extended to accommodate the haptic modality, allowing the rapid prototyping and development of multi-modal applications.

10 The Role of haptics in shared virtual environments

We are conducting a set of human experiments to investigate the role of haptics in shared virtual environments (SVEs). Our efforts are 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. If so, (2) 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 experiments concern a scenario where two people, at remote sites, co-operate to perform a joint task in a SVE. 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 modeled using a spring-damper system and a point-based haptic rendering technique (Ho et al., 1998 and Basdogan, et al., 2000). 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:

- (a) both visual and haptic feedback provided to the participants
- (b) only visual feedback was provided to the participants

Performance and subjective measures were developed to quantify the role of haptic feedback in SVEs and the results suggest that haptic feedback significantly improves the performance and contributes to the feeling of "sense of togetherness" in SVEs.

We have also observed strong indications of the influence of haptics on gender, personality, and emotional experience of the participants. For example, force feedback seems to be associated with male gender and expert behavior can be 'felt' through a haptic device. Specifically, this study provides a necessary foundation for the incorporation of 'personal' touch and force feedback devices into multimodal networked VR systems and the internet. The outcomes of this research can also have an impact on the enhancement of virtual environments for performing collaborative tasks in shared virtual worlds on a daily basis such as co-operative teaching, planning, and training.

Publications

1. v. Wiegand TE, Schloerb DW, and Sachtler WL (1999). "Virtual Workbench Near-Field Virtual Environment System With Applications," Presence: Teleoperators and Virtual Environments, 8(5), 492-519.
2. Yuan HF, Sachtler WL, and Durlach NI. (1999). "Effects of Time Delay on Depth Perception via Motion Parallax in Virtual Environment Systems, " Presence: Teleoperators and Virtual Environments, under revision.
3. Schloerb DW (1999). "Adaptation to Perceived Depth Related to Changes of the Effective Interpupillary Distance in Computer-Generated Stereoscopic Displays," Perception & Psychophysics, to be submitted.
4. Delhorne L, Aviles W, and Durlach NI (1999). "Note on the Role of Haptics in Learning to Perform Cognitive Tasks," Presence: Teleoperators and Virtual Environments, under revision.
5. Ho C-H, Basdogan C, and Srinivasan MA, Efficient point-based rendering techniques for haptic display of virtual objects, Presence, Vol. 8, No. 5, pp. 477-491, 1999.
6. Srinivasan MA, Basdogan C, and Ho C-H, Haptic Interactions in the Real and Virtual Worlds, Design, Specification and Verification of Interactive Systems '99, Eds: D. Duke and A. Puerta, Springer-Verlag Wien, 1999.
7. Wu W-C, Basdogan C, and Srinivasan MA, Visual, Haptic, and Bimodal Perception of Size and Stiffness In Virtual Environments, DSC-Vol. 67 Proceedings of the ASME Dynamic Systems and Control Division, Ed. N. Olgac, pp. 19-26, ASME, 1999.
8. Srinivasan MA, Basdogan C, and Ho C-H, Haptic Interactions in Virtual Worlds: Progress and Prospects, Proceedings of the International Conference on Smart Materials, Structures, and Systems, Indian Institute of Science, Bangalore, India, July, 1999.
9. Salisbury JK and Srinivasan MA (Eds.), Proceedings of the Fourth Phantom Users Group Workshop, AI Technical Report no. 1675 and RLE Technical Report no. 633, MIT, November, 1999.
10. Basdogan, C, Ho C-H, Srinivasan, MA, and Slater M, An experimental study on the role of touch in shared virtual environments, 2000 (accepted for publication in the ACM Transactions on Human Computer Interactions).
11. Ho C-H, Basdogan C, Srinivasan MA, Ray-based haptic rendering: force and torque interactions between a line probe and 3D objects in virtual environments, International Journal of Robotics Research, 2000 (in press).

References

1. Srinivasan MA, 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 MA and Basdogan C, Haptics in virtual environments: taxonomy, research status, and challenges, Computers and Graphics, Vol. 21, No. 4, 1997 (won the 1997 Computers & Graphics Best Paper Award).
3. Basdogan C, Ho C-H, and Srinivasan MA, A ray-based haptic rendering technique for displaying shape and texture of 3D objects in virtual environments, DSC-Vol. 61, Proceedings of the ASME Dynamic Systems and Control Division, Ed. G. Rizzoni, pp. 77-84, ASME, 1997.

4. Srinivasan MA, Beauregard GL, and Brock, DL, 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.
5. DiFranco DE, Beauregard GL, and Srinivasan MA, The effect of auditory cues on the haptic perception of stiffness in virtual environments, DSC-Vol. 61, Proceedings of the ASME Dynamic Systems and Control Division, Ed. G. Rizzoni, pp. 17-22, ASME, 1997.
6. Ho C-H, Basdogan C, and Srinivasan MA, Haptic rendering: Point- and ray-based interactions, Proceedings of the Second PHANToM User's Group Workshop, October, 1997.
7. Ho C-H, Basdogan C, Slater M, Durlach M, and Srinivasan MA, The Influence of Haptic Communication on the Sense of Being Together, Workshop on Presence in Shared Virtual Environments, BT Labs, Ipswich, UK, June, 1998.