HAPTICS IN VIRTUAL ENVIRONMENTS: TAXONOMY, RESEARCH STATUS, AND CHALLENGES

MANDAYAM A. SRINIVASAN[1,2†] and CAGATAY BASDOGAN[1,2]

[1]Laboratory for Human and Machine Haptics, Department of Mechanical Engineering, 36-769,
Massachusetts Institute of Technology, Cambridge, MA 02139, U.S.A.
e-mail: srini@mit.edu
[2]Laboratory for Human and Machine Haptics, Research Laboratory of Electronics,
36-769, Massachusetts Institute of Technology, Cambridge, MA 02139, U.S.A.

Abstract—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 has 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. In this paper, we clarify the terminology of human and machine haptics and provide a brief overview of the progress recently achieved in these fields, based on our investigations as well as other studies. We describe the major advances in 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. We also summarize the issues and some of our results in integrating haptics into multimodal and distributed virtual environments, and speculate on the challenges for the future. © 1997 Elsevier Science Ltd

1. INTRODUCTION

Haptics refers to manual interactions with environments, such as exploration for extraction of information about the environment or manipulation for modifying the environment. These interactions may be accomplished by human or machine hands and the environments can be real or virtual. Also, the interactions may or may not be accompanied by other sensory modalities such as vision and audition. Most of the virtual environments (VEs) built to date contain visual displays, primitive haptic devices such as trackers or gloves to monitor hand position, and spatialized sound displays. To realize the full promise of VEs, haptic displays with force and/or tactile feedback are essential. Being able to touch, feel, and manipulate objects in an environment, in addition to seeing (and hearing) them, provides a sense of immersion in the environment that is otherwise not possible. It is quite likely that much greater immersion in a VE can be achieved by the synchronous operation of even a simple haptic interface with a visual and auditory display, than by large improvements in, say, the fidelity of the visual display alone.

Exciting possibilities open up with the addition of haptics to various applications of virtual reality and teleoperation. Given below are some of the examples:

- *Medicine:* surgical simulators for medical training; manipulating micro and macro robots for minimally invasive surgery; remote diagnosis for telemedicine; aids for the disabled such as haptic interfaces for the blind.

- *Entertainment:* video games and simulators that enable the user to feel and manipulate virtual solids, fluids, tools, and avatars.

- *Education:* giving students the feel of phenomena at nano, macro, or astronomical scales; 'what if' scenarios for non-terrestrial physics; experiencing complex data sets.

- *Industry:* integration of haptics into CAD systems such that a designer can freely manipulate the mechanical components of an assembly in an immersive environment.

- *Graphic arts:* virtual art exhibits, concert rooms, and museums in which the user can log in remotely to play the musical instruments, and to touch and feel the haptic attributes of the displays; individual or co-operative virtual sculpturing across the Internet.

The subsystems and information flow underlying interactions between human users and force-reflecting haptic interfaces are shown in Fig. 1.

- *Human sensorimotor loop:* when a human user touches a real or virtual object, forces are imposed on the skin. The associated sensory information is conveyed to the brain and leads to perception. The motor commands issued by the brain activate the muscles and result in hand and arm motion. At our 'MIT Touch Lab', we have investigated various aspects of this sensorimotor process, such as the biomechanics of human finger pads, tactile

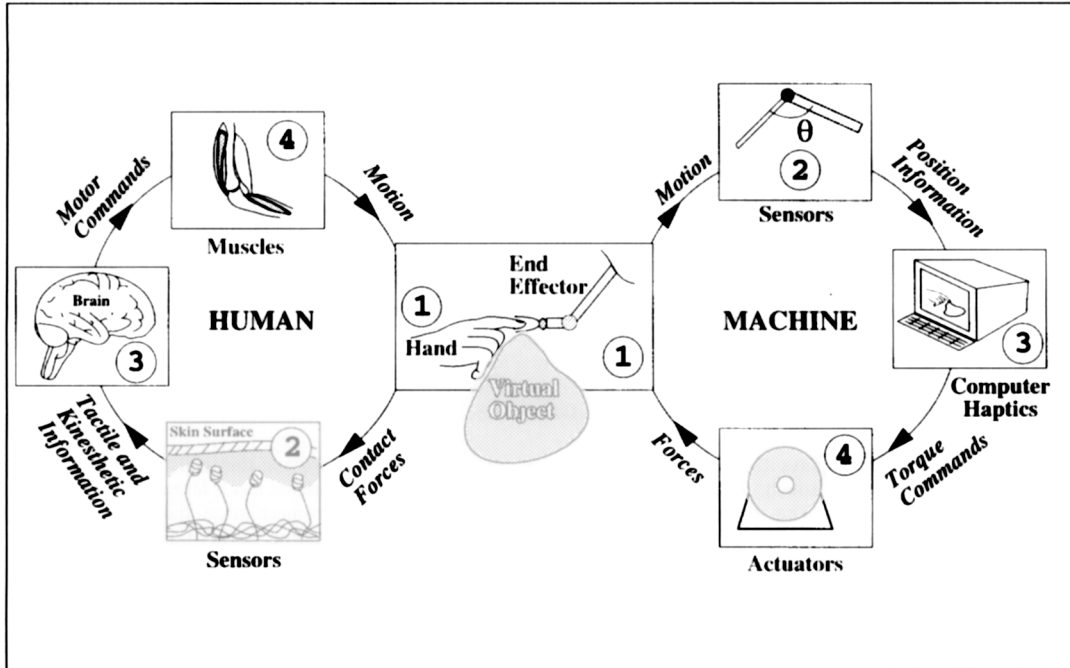---

† Author for correspondence.

Fig. 1. Haptic interaction between humans and machines.

neurophysiology, and human perceptual as well as motor capabilities (http://touchlab.mit.edu).

• *Machine sensorimotor loop:* when the human user manipulates the end-effector of the haptic interface device, the position sensors on the device convey its tip position to the computer. The models of objects in the computer calculate in real-time the torque commands to the actuators on the haptic interface, so that appropriate reaction forces are applied on the user, leading to tactual perception of virtual objects. In our laboratory, and in collaboration with Dr Salisbury's group in the MIT Lab, we have developed computer controlled electromechanical devices and the associated software to simulate the 'feel' of different objects. Studies are underway to investigate how controlled alterations in visual, auditory, and haptic displays affect human perception (refer to Section 5).

The goals of this paper are (1) to clarify the terminology concerning both the human and the machine aspects of this rapidly developing field, (2) to provide pointers to the relevant literature, (3) to summarize the results of research in various multidisciplinary areas relevant to haptics in VEs, along with a guided review of our own research at the MIT Touch Lab, and (4) to discuss the challenges for the future. The next three sections describe the status of the three major components of haptics in VEs, namely, human haptics, haptic interfaces, and computer haptics. Although a large number of references are given to aid the reader, our goal is to summarize our research. We do provide references to

the related work by others, but do not claim to be exhaustive in covering the literature. In Section 2, we describe the salient terminology and quantitative results in human haptics. In Section 3, we give primary classifications of haptic interfaces and discuss the relevant issues briefly. Section 4 focuses on the recent advances in the software aspects of haptic displays. In the next two sections, we describe briefly the issues and our experiences in two areas: Section 5 is on multimodal VEs composed of visual, auditory, and haptic displays; Section 6 is on haptics across the Internet. Finally, Section 7 discusses the various challenges facing haptics in VE today.

2. HUMAN HAPTICS

The human haptic system consists of the mechanical, sensory, motor and cognitive components of the hand–brain system. Here, we give a brief summary to clarify the terminology and to provide quantitative performance specifications pertinent to haptic interfaces. More details and references can be found in Srinivasan [34].

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 tissue and skin. Altogether, 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, and chemical stimuli activate these receptors, causing them to transmit electrical impulses via the afferent (*i.e.* sensory) neural network to the central nervous system (of which the brain forms a part), which in turn sends commands through the efferent (*i.e.* motor) neurons to the muscles for desired motor action.

Tactual sensory information from the hand in contact with an object can be divided into two classes: (1) *tactile information*, referring to the sense arising from the skin in contact with the object; (2) *kinesthetic* (equivalently, *proprioceptive*) information, referring to the sense of position and motion of limbs along with the associated forces. In general, net forces of contact are sensed by both the systems, but the spatiotemporal force variations within the contact region are conveyed by the tactile system alone. Consequently, the fine shape, texture, and rubber-like compliance of the object within the contact region, as well as whether the object is slipping relative to the skin, are sensed by the tactile sensors in the skin. The coarser properties of objects such as large shapes (*e.g.* radius of about a meter) or spring-like compliances that require hand or arm motion for exploration are conveyed by the kinesthetic system.

Tactile sensory capabilities are most acute on the fingerpad. Spatial location of a point is detectable to within 0.15 mm [19] and the spatial resolution of two points is about a millimeter [14]. On a smooth surface, even a 0.06 micron high texture composed of a grating or a 2 micron high single dot is detectable [17]. Vibrations of up to 1 kHz are resolved by the tactile system, with highest sensitivity around 250 Hz where amplitudes less than a micron are detectable [7]. Kinesthetic resolution is about 2 degrees for the fingers and wrist, and about 1 degree for the shoulder [34]. Fingertip positional resolution is in the range of 0.5 to 2.5 mm during grasping objects of 1 to 80 mm length [33, 44]. The resolution for velocity and acceleration of the fingertip, measured in terms of the Just Noticeable Differences (JNDs), are about 11% and 17% of reference values, respectively [4].

The control of contact conditions is often as important as sensing those conditions for successful task performance. In humans, such control action can range from a fast muscle stretch reflex that acts in about 30 ms, a spinal reflex (about 70 ms), and a relatively slow conscious deliberate action. Human bandwidth for limb motions is found to be a function of the task being performed: 1–2 Hz for unexpected signals, 2–5 Hz for periodic signals, up to 5 Hz for internally generated or learned trajectories, and about 10 Hz for reflex actions [45]. It has also been observed that humans can produce actions such as drum rolls at over 40 Hz by allowing the drumstick to bounce through suitable control of the passive impedance of the hand joints [46]. The maximum controllable force that can be exerted by a finger is about 50 to 100 N, depending on whether only the finger muscles are allowed to be activated or if shoulder muscles can be used [36]. However, the

typical forces used in exploration and manipulation are in the range of 5 to 15 N. The force control resolution during visual tracking of constant forces is about 0.04 N or 1%, whichever is higher [31, 36].

When subjects are actively squeezing virtual objects under purely haptic conditions, the perceptual resolution in terms of JNDs has been found to be about 7% for force and elastic stiffness [37, 48], 12% for viscosity and 20% for mass [4]. The stiffness required to simulate a rigid wall has been estimated to be about 25 N/mm [36], but even about 5 N/mm can possibly be adequate for 'suspension of disbelief'. Such perceptual effects, however, can be considerably altered by additional visual and/or auditory stimuli, as described in Section 5.

3. HAPTIC INTERFACES

In interacting with VEs using 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. Computer keyboards, mice, and trackballs constitute relatively simple haptic interfaces. Gloves, body suits, and exoskeletons that only track hand postures are more recent examples of interfaces available in the market. These interfaces, however, do not convey the touch and feel of objects, which can be achieved only through tactile and/or force feedback.

An important distinction among haptic interfaces is whether they are tactile displays or net force displays. The corresponding difference in interactions with VEs is whether the direct touch and feel of objects contacting the skin is simulated or the interactions are felt through a tool. Simulation of interactions through a tool, such as feeling the virtual world through a rigid stick, requires only net force (and torque) display. Simulation of direct contact with objects is much more difficult since it requires a tactile display capable of distributing the net forces and torques appropriately over the region of contact between the object and the skin. At present, force display devices that can match at least some of the capabilities of the human haptic system are available. But the performance of the currently available tactile displays which use single or multiple stimulators composed of shape-memory alloys, pneumatic actuators, vibrotactile or electrotactile elements is inadequate in comparison to the human sensory capabilities.

An alternative distinction among haptic interfaces is whether the interface is ground-based or body-based. Force reflecting joysticks are examples of ground-based devices and exoskeletons represent body-based devices. Hybrid devices which combine both of these characteristics have also been built. Well-designed exoskeletal devices have the advantage that their kinematics and workspace coincide with those of the human. However, the design becomes complex because the unbalanced forces applied on,

say the user's fingerpad, must eventually be grounded somewhere. In addition, the need for the user to carry the mass of the device can interfere with feeling objects in the VE and can cause fatigue. At present, the best performance is achieved by ground-based, force-reflecting devices such as the PHAN-ToM (SensAble Technologies, Inc.) or the Impulse Engine (Immersion Corp.).

Ultimately, the performance specifications of haptic devices are set by the human abilities and limitations. Simulation of haptic interactions with VEs that are designed to mimic real environments will always be approximate, and which approximations are sufficient will be determined by the limits of human performance. The desirable features of force-reflecting haptic interfaces are as follows [20, 36, 47]:

(1) Low back-drive inertia and friction, and no constraints on motion imposed by the device kinematics, so that free motion feels free.
(2) The range, resolution, and bandwidth, both in terms of position sensing and force reflection, should match those of the human for the tasks for which the haptic interface is employed. The user (a) should not be able to go through rigid objects by exceeding the force range of the interface, (b) should not be able to feel unintended vibrations such as due to quantization of position or low servo rate, and (c) should not feel stiff objects as soft due to low structural and servo stiffness. These conditions are difficult to satisfy because of the fine sensitivity and ~ 1 kHz bandwidth of the human tactile system, but the fact that the human control bandwidth in haptic interactions is only of the order of 10 Hz is helpful.
(3) Ergonomics and comfort: making the human user comfortable when wearing or manipulating a haptic interface is of paramount importance, since pain, or even discomfort, supersedes all other sensations.

A more detailed review of various types of haptic interfaces and associated issues can be found in [34], [47] and [8].

4. COMPUTER HAPTICS

We define *Computer Haptics* as the discipline concerned with generating and rendering haptic stimuli to the human user, just as computer graphics deals with generating and rendering visual images. During the last two decades, the techniques of computer graphics have become a dominant part of human–computer interfaces in displaying and manipulating data and objects. Rapid improvements in computer graphics hardware have enabled the generation of photorealistic effects, animations, and real-time interactive simulations. We expect similar rapid progress to occur in computer haptics, as indicated by recent papers on PHANTOM-based applications [52]. As computers become more power-

ful and affordable, and sophisticated software tools and techniques are increasingly available to the human user, the need for more effective interactions between humans and computers becomes urgent. This demand, we believe, will lead to further development of immersive and interactive multi-modal virtual worlds that are focused on displaying visual, haptic, and auditory stimuli to the human operator.

4.1. Haptic rendering and interaction techniques

The basic process of haptically rendering objects in virtual environments with a force-feedback device is shown in Fig. 2. As the user manipulates the generic probe of the haptic device, the new position and orientation of the probe is sensed by the encoders. Collisions between the simulated stylus and virtual objects are detected. If the probe collides with an object, the mechanistic model calculates the reaction force based on the penetration depth of the probe into the virtual object. The calculated force vectors may then be modified by appropriately mapping them over the object surface to take into account the surface details. The modified force vectors are fed back to the user through the haptic device.

Several haptic rendering techniques have been developed recently to render 3-D objects. Just as in computer graphics, the representation of 3-D objects can be either surface-based or volume-based for the purposes of computer haptics. While the surface models are based on parametric or polygonal representations, volumetric models are made of voxels. An alternative way of distinguishing the existing haptic rendering techniques is based on the type of haptic interaction: *point-based* or *ray-based* (Fig. 3).

In point-based haptic interactions, only the end point of the 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. If so, the depth of indentation is calculated as the distance between the current HIP and a surface point [21, 29, 41], such as the nearest surface point. 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 [2]. The collision detection algorithms return the collision point, which is the intersection point between the ray and the surface of the object. For simulating rigid objects the component of the distance between the collision point and the HIP along the surface normal at the collision point is taken to be the depth of indentation (see Fig. 3).

In both point- and ray-based force reflection, the reaction force (F) is usually calculated using the linear spring law, $F = kx$, where k is the stiffness of
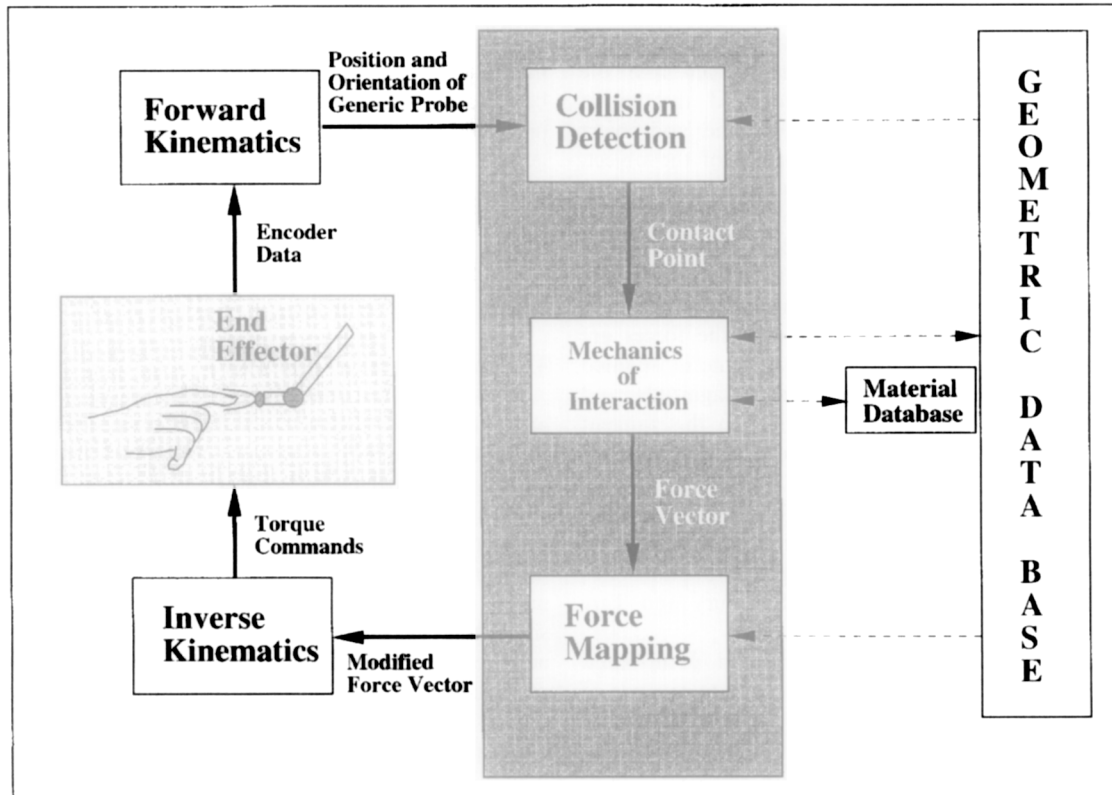
Fig. 2. The processes associated in haptic rendering with a force display. The solid and dashed lines represent the process flow and information exchange respectively.

the object and x is the depth of indentation. For frictionless interactions, the reaction force (F) is normal to the polygonal face that the generic probe collides with. For rigid objects, the value of k is set as high as possible, limited by the contact instabilities of the haptic device. Studies show that addition of a damping term into the interaction dynamics improves the stability of the system and the haptic perception of 'rigidity' [9, 20, 23, 43]. Other forces such as those that arise from surface friction and roughness also need to be displayed to improve the realism of haptic interactions (see Section 4.4).


Fig. 3. Point- and ray-based haptic interactions.

4.2. Haptic display of shapes

Initially, haptic rendering methods focused on displaying simple, rigid, and frictionless objects. Massie and Salisbury [21] developed the PHANToM haptic interface device and proposed a point-based method for rendering primitive objects such as cube, cylinder, and sphere. The depth of indentation is taken to be the distance between the HIP and the nearest surface point (see Fig. 3). Computations for determining the force vector involve dividing the object into sub-spaces associated with particular portions of object surface. If the HIP penetrates into a region which is shared by multiple sub-spaces, then superposition of surface normals is used to calculate the direction of the resultant force vector. A virtual scene can be constructed from multiple primitives whose sub-regions can be defined in advance. However, there are many problems with this haptic rendering technique [42]: first, it is not easy either to divide an object into sub-spaces intuitively or to construct virtual environments from primitive objects, and second, the superposition of force vectors breaks down for thin or complex shaped objects.

Zilles and Salisbury developed a more sophisticated, point-based method to haptically render polygonal surfaces [41]. They defined a new point, namely the *god-object*, to represent the location of the surface point. The new location of the god-object

point is calculated each time the user manipulates the generic probe of the haptic interface by using a constrained optimization technique such that the distance between the god-object and the haptic interface point is minimized, and the god-object point always remains on the surface of the object, even though the HIP can penetrate it.

Another approach for point-based haptic rendering of surfaces is to use intermediate representations [18]. A local planar approximation to the surface (tangent plane) is computed at the collision point for each cycle of force loop. Although the forces are updated frequently ($\sim$1 kHz), the plane is updated more slowly. The reaction force magnitude and direction is calculated with respect to this tangent plane. However, undesirable force discontinuities may appear if the generic probe of the haptic device is moved over large distances before the new tangent plane is updated. An improved tangent plane method for surfaces represented by implicit functions has been recently implemented by Salisbury and Tarr [50].

Avila and Sobierajski implemented point-based interactions with volumetric objects that are made of voxels [1]. A total of 8 bytes of information is assigned to each voxel of the volumetric object that includes material density, density gradient, color, and haptic interaction properties such as stiffness and viscosity. During real-time computations, they first check to see if the HIP is inside a volumetric isosurface of the object. If so, the scalar density field at the haptic interface point is used to compute the stiffness and the reaction forces through a set of linear transformations. The gradient of the density field at haptic interface point is used to define the surface normal, which is computed using central difference approximation.

Basdogan *et al.* developed a ray-based haptic rendering technique for displaying 3-D objects in virtual environments [2]. The generic probe of the haptic device (stylus) is modeled as a line segment whose position and orientation are provided by the encoder signals in the haptic interface. We display the graphical model of the simulated stylus and update its tip and tail coordinates as the user manipulates the actual one, detect any collisions between the simulated stylus and the virtual object, estimate the reaction force typically by using the linear spring law, and finally reflect this force to the user via the haptic device. The detection of collisions occur in three consecutive steps: first, the collisions are detected between the simulated stylus and the bounding box of the virtual objects in the scene, then we check the collision between the simulated stylus and the bounding box of each triangular element. Finally, collisions between the simulated stylus and the triangular element itself are detected using computational geometry techniques. This multi-step collision detection technique significantly improves the update rate. Computations can be made even faster by utilizing space partitioning and advanced

search techniques used in computer graphics. The generic stylus or the other attached mechanical instruments such as gripping or cutting tools can be modeled using simple 2-D geometric primitives (*e.g.* line segments or triangles depending on the geometric shape of the attached mechanical instrument). The proposed method is more suitable for simulating *tool–object* interactions than earlier techniques, because the reflected forces are dependent on both the position and orientation of the generic probe.

4.3. *Haptic display of compliance*

Physically-based modeling of deformable objects and their real-time simulation have always been of interest to computer graphics researchers. A network of masses connected by springs and dampers, finite element techniques, parametric representations such as B-splines and Bernstein polynomials (free form deformation techniques) have been considered in modeling soft objects. The major challenge in simulating force-reflecting deformable models is to achieve the optimal balance between the complexity of models and the realism of the visual and haptic displays in real-time.

Swarup developed a model to represent the dynamic behavior of visco-elastic objects in VEs [35]. Virtual objects are modeled as a discrete network of masses connected by mechanical elements such as springs and dampers. Haptic interactions such as stroking, palpation, and plucking were demonstrated using this model.

Burdea and his colleagues have developed a model to simulate a virtual hand squeezing a virtual rubber ball [8]. Rutgers Master was used to reflect the reaction force to each finger that is proportional to the local deformation of the ball.

We have also developed a simple force reflecting deformable model for simulating compliant behavior of objects that are subject to external forces [3]. Using this model, the vertices of the object within a radius of the collision point can be deformed locally in real-time along the direction of the generic probe of the haptic device by using a second order polynomial function. In this model, the visual and haptic displays are decoupled: in the visual display, the vertices of the surface within a certain radius from the collision point are translated according to a second order polynomial such that the maximum depth of indentation is at the HIP; in the haptic display, the stiffness is set to a low value so that the object feels soft.

4.4. *Haptic display of surface details*

Our research on haptic display of surface details can be summarized into three main groups: (a) haptic smoothing of object surfaces, (b) rendering of haptic textures, and (c) haptic rendering of surfaces with friction. To convey to the user the tactual feeling of smooth surfaces, we compute the force vector for each vertex and smoothly interpolate its direction

over the polygonal surfaces. In order to display rough surfaces haptically, we either perturb the force vector using the local gradient of texture field (force perturbation) or divide the surface of the object into smaller polygons to change its topology (displacement mapping).

4.4.1. *Haptic smoothing of object surfaces.* 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, Morgenbesser and Srinivasan proposed *force shading* [25, 42]. 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 [Fig. 4(a)]. Consequently, the surfaces of virtual objects feel smoother than their original polyhedral representations. This technique is analogous to Phong shading in computer graphics. Force shading is quite useful in rendering multiple objects in virtual environments. One can develop geometric models of objects in various levels of detail and display the high fidelity model or a force-shaded low fidelity copy of the model, only as it is required by the application. Perceptually, the force shading technique may not be very effective for polygons of a surface where the angle between them is less than a certain threshold value [26]. However, once the critical separation angle is known, the polyhedral object can be tessellated into smaller polygons until the separation angles between the polygons are higher than the critical value.

4.4.2. *Rendering of haptic textures.* Visual rendering of surface details such as textures and bumps, to add realism to the appearance of the 3-D objects, has been a challenging research topic because of its variety and complexity. In multimodal VEs, computational burden is further increased because of the need to represent and synchronously display the macro and micro surface details through touch and/ or sound. Minsky implemented and performed perceptual experiments on a variety of textures with

a 2-D haptic display [24]. We developed two haptic rendering techniques that are similar to those used in computer graphics, and are suitable for 3-D haptic displays:

(a) force perturbation;
(b) displacement mapping.

(a) *Force perturbation*: Force perturbation is a technique of modifying the direction and magnitude of the force vector to generate surface effects such as roughness [Fig. 4(b)]. Analogous to the 'bump mapping' technique of computer graphics [6], force perturbation provides the user with a sense of surface details such as textures and bumps. To generate visual bumps on the surface of the object, Max and Becker developed a formulation which is based on the original surface normal and the local gradient of the height field, as follows [22]:

$$M = N - \nabla h + (\nabla h \cdot N)N$$

where M is the perturbed surface normal, N is the original surface normal, and ∇h is the gradient of the texture height field. We used the same approach and perturbed the direction and magnitude of the force vector to generate bumpy or textured surfaces that can be sensed tactually by the user [2].

(b) *Displacement mapping*: In the displacement mapping technique of computer graphics, the actual geometry of the object is modified to visually display the surface details [6, 13]. In order to generate microtextures, the geometry of the surface has to be composed of tiny polygons so that its surface can be modified point by point. During the rendering of new surface, normals need to be updated in the geometrical database. However, graphical rendering of objects with many polygons adds considerable computational load on the graphics engine. Haptic rendering of such surfaces is not easy as well [Fig. 4(c)]. Virtual objects and their surfaces cannot be infinitely stiff, hence HIP penetrates into the surfaces. Detecting collisions or locating the new position of the surface point each time the generic probe is moved creates ambiguities due to the existence of fine

(a) force shading   (b) force perturbation   (c) displacement mapping

Fig. 4. Haptic rendering techniques for simulating surface details in virtual environments. The arrows represent the reflected force vectors. In each case, the shaded area represents the displayed geometry of the object and the dark boundary line represents the geometry of the surface intended to be haptically perceived by the user.

details on the surface geometry. Force discontinuities occur as the user strokes the surface of the object with the probe to feel the textures. Yet, this technique needs to be explored more and can be valuable in haptic rendering of macro textures.

Modeling of haptic textures requires an understanding of textures that exist in nature and their interactions with the human tactile sensory system. We know that textures in nature come in many varieties, hence no single model can represent all the textures. Various graphical techniques have been developed in the past to generate photorealistic textured objects [11, 13, 39]. Studies also exist on the range and limits of the human tactile sensory system [17]. In our studies, we first focused on *frequency* and *height* as the two major indicators of the haptic textures. We then classified the haptic texturing techniques into two groups, as in computer graphics:

(a) image-based haptic texturing;
(b) procedural haptic texturing.

(a) *Image-based haptic texturing:* in computer graphics, the images of textures are wrapped around 3-D objects to make them look more realistic. In haptic texturing, the aim is to make the user feel the textures of the visual image that has already been mapped onto the 3-D surface. To simulate image-based textures, we construct a texture field from the 2-D image data. Assuming that gray scale intensities of the 2-D image can be used directly as height indicators to generate a haptic texture field, we can associate each texture coordinate (u, v) with the coordinate of each vertex (x, y, z). We then map these heights onto the object surface using the two-stage texture mapping techniques used in computer graphics [5, 39]. For haptic rendering, we compute the local gradient of the height field at the collision point and perturb the force vector direction accordingly [2]. Obviously, this method of associating heights to images is neither unique nor fool-proof.

(b) *Procedural haptic texturing:* the aim of procedural haptic texturing is to generate synthetic texture fields using mathematical functions for the height field. Once the function that describes the texture field is known, the gradient vector (∇h) at the point of contact can be easily calculated either to perturb the reaction force vector or to modify the surface geometry of the object as in displacement mapping. For example, Fritz and Barner and Siira and Pai perturbed the force vectors to simulate stochastic haptic textures in virtual environments [12, 30]. We have generated a variety of synthetic haptic textures by using various procedural texturing techniques of computer graphics in conjunction with force perturbation [2]. Examples of formulations we have used are based on Fourier series [15], Stochastic functions such as Perlin's [27] 'Noise' and Wijk's [40] 'Spot Noise' functions, Reaction-Diffusion textures [38], and Fractals (Musgrave in [11]).

4.4.3. Haptic rendering of surfaces with friction. The addition of friction improves the haptic interactions and makes the simulations more realistic. For example, without friction, we will not be able to push virtual buttons even with a small tangential force, because the probe will slip off the surface. Many researchers have proposed Columb and viscous friction models for haptic interactions [12, 18, 20]. Columb friction has static and dynamic components and it is applied in a direction tangential to the normal force. Viscous friction depends on the velocity of the object which can either be directly measured or approximated numerically from the position information. If the measured velocity signal contains noise or the rapid changes in the position add noise to the approximated velocities, the haptic interactions can be unstable. Salisbury *et al.* have implemented a stable stick-slip model to simulate Columb friction that does not depend on velocity [29]. We have successfully generated spatially periodic waves of friction on the surface of 3-D objects by modifying the static and dynamic friction coefficients of this friction model.

5. HAPTICS IN MULTIMODAL VIRTUAL ENVIRONMENTS

Multimodal VEs that combine the visual, haptic, and auditory sensory information are essential for designing immersive virtual worlds. We believe that it is time to make a more concerted effort (1) to bring the three modalities together in VEs and (2) to study their interactions in affecting human perception and performance.

5.1. An interactive user interface for multimodal virtual worlds

We have constructed a PC-based software application and library to generate virtual scenes from object primitives such as cube sphere, cylinder, sphere, *etc.* (MAGIC Tool-kit [16]). Currently, the visuals are drawn as a two-dimensional projection on the monitor, while the haptics is displayed in a three-dimensional field. Using multi-threaded programming techniques, we have successfully separated the visual and haptic servo loops to control the flow of visual and haptic information. We have also integrated haptic and sound displays to perform perceptual experiments (see below). Prerecorded sounds of contact between several pairs of objects were played to the user through the headphones to stimulate the auditory senses.

A friendly and powerful UNIX-based user interface is under development to enable the user to rapidly create a 3-D virtual environment and to interact with objects in it visually and haptically. The graphics interface enables the user to load a virtual environment from a text file, toggle on/off stereo visualization, save the virtual environment and quit from the application. The user can intuitively add (subtract) 3-D polygonal objects into (from) the virtual scene, and assign visual and haptic properties to the objects using a simple text file. This interface is

flexible enough to be extended by the end-user. The graphical scene and the virtual objects were created using Open Inventor (Silicon Graphics Inc.). PHAN-ToM was used as a haptic device to convey to the user a tactual sense of object shapes and surface details.

5.2. Perceptual interactions among visual, auditory and haptic modalities

It is known that an individual's perceptual experience can be influenced by interactions among various sensory modalities. For example, in real environments, visual information has been shown to alter the haptic perception of object size, orientation, and shape [49]. Similarly, it is likely that appropriate scaling and timing of each display modality in virtual environments would influence human perception and performance.

We have conducted experiments to assess the influence of visual and auditory information on the perception of object stiffness through a haptic interface. We have shown that visual sensing of object deformation dominates over kinesthetic sense of hand position and results in a dramatic mis-perception of object stiffness when visual display is intentionally skewed [32]. Similarly, contact sounds influenced the perception of object stiffness during tapping of virtual objects through a haptic interface [10]. However, the effect of sound in this experiment was not as dramatic as the effect of visual display in the previous experiment. An important implication to virtual environments is that by skewing the relationship between the haptic and visual and/or auditory 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 stiffnesses 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 or eliminating visual deformation of the object.

Using the MAGIC Tool-kit we have also conducted human experiments to investigate the role of the visual–haptic size ratio, the visual and haptic sensory feedback in isolation or together, the effects of visual scaling on training, and various cursor control paradigms [16]. The task was to navigate through a maze with visual and/or haptic feedback. Results show that subjects preferred large visual–haptic ratios, small haptic workspaces, and a position controlled cursor. Subjects performed best with a large visual display and a small haptic workspace. Also, subjects performed best when both visual and haptic cues were given, with a slight decrease in performance when only haptic cues were given, and with a significant decrease in performance when only visual cues were given. The performance of the subjects improved linearly with increases in visual display size, when subjects were initially trained on the largest visual size in the series. Among various cursor control methods, subjects performed

best when there was a high correlation in cursor position and movement between the visual and haptic workspaces.

We believe that the results mentioned above are only the beginnings of a large body of results yet to be discovered. The effects of multimodal interactions on human perception need to be investigated in more detail through a study of normal and altered relationships among haptic, visual, and auditory displays. This will then lead to a rational basis upon which multimodal VEs can be designed and implemented.

6. HAPTICS IN DISTRIBUTED VIRTUAL ENVIRONMENTS

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 described above using the PHANToM haptic interface 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 has been 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 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. The 'Haptics Group' at the MIT AI Lab, in collaboration with our group, is also looking into the best ways of embedding the haptics into Virtual Reality Modeling Language (VRML) through its extension mechanism [28]. In order to support haptics in VRML, we have extended the VRML language specifications and also developed external Java scripts.

7. CHALLENGES FOR THE FUTURE

In spite of the recent progress, the incorporation of haptics into VEs is in its infancy. Given the continued rapid development of computers, the challenges for the future are likely to be in the areas of haptic interface devices, computer haptics, and human–machine interactions, as detailed below.

7.1. Haptic interfaces

As described in Sections 2 and 3, the force-reflecting interfaces still need to be improved to match the human range, resolution, and bandwidth, both in terms of forces and displacements. In addition to net forces, better torque displays, which at the same time

do not increase the friction and inertia of the devices, need to be developed. Interface devices suitable for multi-finger tasks as well as for large workspaces are also of interest. Body-based higher performance exoskeletons can be of help in large scale virtual worlds. More challenging is the design of sophisticated tactile displays (including temperature stimuli) that match the human perceptual capabilities. Once such tactile displays are developed and integrated with powerful, multiple degrees of freedom force-feedback devices, we can expect an explosion in the applications of haptics in VEs.

7.2. 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.

Some of the research problems that, we believe, may be of interest to the researchers are as follows:

- Parametric surfaces need to be considered for displaying haptic representation of 3-D objects in virtual environments. For example, NURBS surfaces are commonly used in CAD systems for modeling 3-D objects. Therefore, the development of rendering techniques for parametric surfaces, as in [51] will enable us to integrate haptic interfaces into CAD systems in the near future.
- Physically-based deformable models that can reflect haptic interaction forces to the user need to be developed. Exciting applications of this technology, such as surgical simulation and medical training, makes this topic very appealing to work on. Special finite element modeling (FEM) techniques can be explored to simulate the dynamics of deformable objects in real-time. As the interest in medical simulation grows, we also believe that volumetric-based haptic rendering techniques and models will gain more acceptance, since volume data is currently available from medical image scanners.

- Rapid prototyping of multimodal virtual worlds that quickly integrates visual, haptic, and auditory modalities is of interest. For example, GHOST toolkit developed by Sensible Technologies is able to display the visual and haptic attributes of a library of objects. The improvements to the multimodal virtual environments can be to modify the existing shape, property, group, and sensor nodes of the visual scene and add new nodes for incorporating haptic and auditory displays.
- A more comprehensive study of haptic texturing is required. The flexibility in designing and generating textures with only a few modifications of the texture function is crucial. The generated haptic textures must also match the human haptic resolution. New haptic texturing techniques that can take advantage of the image-based and procedural techniques will be valuable.

7.3. Human–machine interaction

VE is fundamentally an immersive and interactive medium through which the user's sensorimotor and cognitive functions are affected. Therefore any rational basis for the design of both the hardware and the software has to depend on human perception and performance. Our current understanding of human abilities and limitations that is useful to VE design is quite limited. But to increase this knowledge base, human studies using VEs would be very useful. Therefore a bootstrap approach where the current VEs help perform human experiments, which, in turn, help design the next generation of VE systems seems to be necessary. Extracting the general principles of human perception and performance by designing and conducting such experiments over a wide enough population may very well be the ultimate challenge.

REFERENCES

1. Avila, R. S. and Sobierajski, L. M., A haptic interaction utilizing a volumetric representation. In Proceedings of the First PHANToM Users Group Workshop, eds. J. K. Salisbury and M. A. Srinivasan. Dedham, MA, September 1996.
2. Basdogan, C., Ho, C. and Srinivasan, M. A., A ray-based haptic rendering technique for displaying shape and texture of 3-D objects in virtual environments, Proceedings of the ASME Dynamic Systems and Control Division, Dallas, TX, November 1997 (in press).
3. Basdogan, C. and Srinivasan, M. A., Haptic rendering of deformable objects. Unpublished document (http://touchlab.mit.edu/ people/cagatay).
4. Beauregard, G. L. and Srinivasan, M. A., Sensorimotor interactions in the haptic perception of virtual objects. RLE TR-607, MIT, Cambridge, MA, 1996.

5. Bier, E. A. and Sloan, K. R., Two-part texture mapping. *IEEE Computer Graphics and Applications*, 1986, September, 40–53.

6. Blinn, J. F., Simulation of wrinkled surfaces. *ACM (Proceedings of SIGGRAPH)*, 1978, **12**(3), 286–292.

7. Bolanowski, S. J., Gescheider, G. A., Verillo, R. T. and Checkosky, C. M., Four channels mediate the mechanical aspects of touch. *Journal of the Acoustical Society of America*, 1988, **84**(5), 1680–1694.

8. Burdea, G. D., *Force and Touch Feedback for Virtual Reality*. Wiley, New York, 1996.

9. Colgate, J. E., Stanley, M. C. and Brown, J. M., Issues in the haptic display of tool use. In *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems*. Pittsburgh, PA, 1995, pp. 140–145.

10. DiFranco, D., 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*, Dallas, TX, November 1997 (in press).

11. Ebert, D. S., Musgrave, F. K., Peachey, D., Perlin, K. and Worley, S., *Texturing and Modeling*. AP Professional, Cambridge, MA, 1994.

12. Fritz, J. P. and Barner, K. E., Stochastic models for haptic textures. In *Proceedings of SPIE's International Symposium on Intelligent Systems and Advanced Manufacturing—Telemanipulator and Telepresence Technologies III*. Boston, MA, November 1996.

13. Foley, J. D., van Dam, A., Feiner, S. K. and Hughes, J. F., *Computer Graphics: Principles and Practice*. Addison-Wesley, Reading, MA, 1995.

14. Johnson, K. O. and Phillips, J. R. Tactile spatial resolution—I. Two point discrimination, gap detection, grating resolution and letter recognition. *Journal of Neurophysiology*, 1981, **46**(6), 1177–1191.

15. Gardner, G. Y., Visual simulation of clouds. In *ACM (Proceedings of SIGGRAPH)*, 1985, **19**(3), 297–303.

16. Hou, I. A. and Srinivasan, M. A., Multimodal virtual environments: the MAGIC toolkit and visual–haptic interaction paradigms. Unpublished paper.

17. LaMotte, R. H. and Srinivasan, M. A., Surface microgeometry: neural encoding and perception. In *Information Processing in the Somatosensory System*, eds. O. Franzen and J. Westman. Wenner–Gren International Symposium Series, Macmillan, New York, 1991.

18. Mark, W. R., Randolph, S. C., Finch, M., Van Verth, J. M. and Taylor, II, R. M., Adding force feedback to graphics systems: issues and solutions, In *ACM (Proceedings of SIGGRAPH)*. New Orleans, August 447–452, 1996.

19. Loomis, J. M., An investigation of tactile hyperacuity. *Sensory Processes*, 1979, **3**, 289–302.

20. Massie, T. H., Initial haptic explorations with the PHANToM: virtual touch through point interaction. MS Thesis, Department of Mechanical Engineering, MIT, Cambridge, MA, 1996.

21. Massie, T. H. and Salisbury, J. K., The PHANToM haptic interface: a device for probing virtual objects. In *Proceedings of the ASME Dynamic Systems and Control Division*, DSC-Vol. 55-1, Chicago, IL, 1994, pp. 295–301.

22. Max, N. L. and Becker, B. G., Bump shading for volume textures. *IEEE Computer Graphics and Applications*, 1994, **14**, 18–20.

23. Minsky, M. M., Ouh-young, M., Steele, O., Brooks, F. and Behensky, M., Feeling and seeing: issues in force display. In *Proceedings of Symposium on Interactive 3-D Graphics, ACM SIGGRAPH*, Snowbird, UT, 1990, pp. 235–243.

24. Minsky, M. M., Computational haptics: the Sandpaper system for synthesizing texture with a force-feedback haptic display. PhD Thesis, MIT, Cambridge, MA, 1995.

25. Morgenbesser, H. B. and Srinivasan, M. A., Force shading for haptic shape perception. In *Proceedings of the ASME Dynamic Systems and Control Division*. DSC-Vol. 58, Atlanta, GA, 1996, pp. 407–412.

26. Morgenbesser, H. B. and Srinivasan, M. A., *Force shading for haptic shape perception in virtual environments*. RLE TR-606, MIT, Cambridge, MA, 1996.

27. Perlin, K., An image synthesizer. *ACM (Proceedings of SIGGRAPH)*, 1985, **19**(3), 287–296.

28. Salisbury, J. K. and Srinivasan, M. A., *Development of haptic interfaces*. Technical Report to NAWS/TSD, 1997.

29. Salisbury, J. K., Brock, D., Massie, T., Swarup, N. and Zilles, C., Haptic rendering: programming touch interaction with virtual objects. In *Proceedings of the ACM Symposium on Interactive 3-D Graphics*. Monterey, CA, 1995.

30. Siira, J. and Pai, D. K., Haptic texturing—a stochastic approach. In *Proceedings of the IEEE International Conference on Robotics and Automation*. Minneapolis, MN, 1996, pp. 557–562.

31. Srinivasan, M. A. and Chen, J. S., Human performance in controlling normal forces of contact with rigid objects. In *Advances in Robotics, Mechatronics, and Haptic Interfaces*. DSC-Vol. 49, ASME, New Orleans, CA, 1993.

32. Srinivasan, M. A., Beauregard, G. L. and Brock, D. L., The impact of visual information on haptic perception of stiffness in virtual environments. In *Proceedings of the ASME Dynamic Systems and Control Division*. DSC-Vol. 58, Atlanta, GA, 1996, pp. 555–559.

33. Ho, C. and Srinivasan, M. A., *Human Haptic Discrimination of Thickness*. RLE TR-608, MIT, Cambridge, MA, 1996.

34. Srinivasan, M. A., Haptic interfaces. In *Virtual Reality: Scientific and Technical Challenges*, eds. N. I. Durlach and A. S. Mavor. National Academy Press, Washington, DC, 1995, pp. 161–187.

35. Swarup, N., Haptic interaction with deformable objects using real-time dynamic simulation. MS Thesis, Mechanical Engineering Department, MIT, Cambridge, MA, 1995.

36. Tan, H. Z., Srinivasan, M. A., Eberman, B. and Cheng, B., Human factors for the design of force-reflecting haptic interfaces. In *Proceedings of the ASME Dynamic Systems and Control Division*, Vol. 1, DSC-Vol. 55-1, ed. C. J. Radcliffe. ASME, Chicago, IL, 1994, pp. 353–359.

37. Tan, H. Z., Durlach, N. I., Beauregard, G. L. and Srinivasan, M. A., Manual discrimination of compliance using active pinch grasp: the roles of force and work cues. *Perception and Psychophysics*, 1995, **57**(4), 495–510.

38. Turk, G., Generating textures on arbitrary surfaces using reaction-diffusion. *ACM (Proceedings of SIGGRAPH)*, 1991, **25**, 289–298.

39. Watt, A. and Watt, M., *Advanced Animation and Rendering Techniques*. Addison-Wesley, New York, 1992.

40. Wijk, J. J., Spot noise: texture synthesis for data visualization. *ACM (Proceedings of SIGGRAPH)*, 1991, **25**, 309–318.

41. Zilles, C. B. and Salisbury, J. K., A constraint-based god-object method for haptic display. IEEE International Conference on Intelligent Robots and Systems, 1995.

42. Zilles, C. B., Haptic rendering with the toolhandle haptic interface. MS Thesis, Mechanical Engineering Department, MIT, Cambridge, MA, 1995.

43. Massie, T., Taking the mush out of haptics with infinitely stiff walls. In *Proceedings of the First PHANToM Users Group Workshop*, eds. J. K. Salisbury and M. A. Srinivasan. Dedham, MA, September 1996.

44. Durlach, N. I., Delhorne, L. A., Wong, A., Ko, W. Y., Rabinowitz, W. M. and Hollerbach, J. Manual discrimination and identification of length by the finger-span method. *Perception and Psychophysics*, 1989, **46**(1), 29–38.

45. Brooks, T. L., Telerobotic response requirements. In *Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics*. Los Angeles, CA, 1990, pp. 113–120.

46. Hajian, A. Z., Sanchez, D. S. and Howe, R. D., Drum roll: increasing bandwidth through passive impedance modulation. Submitted to the 1997 IEEE Robotics and Automation Conference, Albuquerque, NM, 1997.

47. Hasser, C. J. and Massie, T. H., The haptic illusion. In *Digital Illusion*, ed. C. Dodsworth. Addison-Wesley, Reading, MA, 1996.

48. Pang, X. D., Tan, H. Z. and Durlach, N. I., Manual discrimination of force using active finger motion. *Perception and Psychophysics*, 1991, **49**(6), 531–540.

49. Welch, R. B. and Warren, D. H., Intersensory interactions. In *Handbook of Perception and Human Performance, Vol. 1, Sensory Processes and Perception,* Chapter 25. Wiley, New York, (1986).

50. Salisbury, J. K. and Tarr, C., Haptic rendering of surfaces defined by implicit functions. *Proceedings of the ASME Dynamic Systems and Control Division,* Callas, TX, November, 1997 (in press).

51. Thompson, T. V., Johnson, D. E. and Cohen, E., Direct haptic rendering of sculptured models. *Proceedings of Symposium on Interactive 3D Graphics,* Providence, RI, 1997, April, pp. 167–176.

52. Salisbury, J. K. and Srinivanam, M. A., Eds, *Proceedings of the First PHANToM Users Group Workshop,* Dedham, MA, AI TR No. 1596 and RLE TR No. 612, MIT, December 1996.