Imaging Group and the community at the MIT Media Laboratory, especially Professor Stephen Benton, Professor Hiroshi Ishii, Carlton Sparrell, Michael Halle, and John Underkoffler. We also thank Bill Verplank for inspiration and valuable discussion and Yael Maguire for his 3D printing effort.

REFERENCES

1. Joel S. Kollin, Stephen A Benton, and Mary Lou Jepsen, Real-Time Display of 3-D Computed Holograms by Scanning the Image of an Acousto- Optic Modulator, in: G.M. Morris, ed., *SPIE Proc. Vol. #1136, Holographic Optics II: Principles and Applications* (Soc. Photo-Opt. Instr. Engrs., Bellingham, WA, 1989), paper #1136-60.

2. St. Hilaire, P., "Scalable Optical Architectures for Electronic Holography", Ph.D. Thesis, MIT Program in Media Arts and Sciences, Massachusetts Institute of Technology, 1994.

3. Pappu, R., *et al.*, A generalized pipeline for preview and rendering of synthetic holograms, in S.A. Benton, ed., *Proceedings of the IS&T/SPIE's Symposium on Electronic Imaging, Practical Holography XI*, 1997.

4. Yokokohji, Y., Hollis, R.L., Kanade, T., Vision-based Visual/Haptic Registration for WYSIWYF Display. International *Conference on Intelligent Robots and Systems,* pp. 1386-1393, 1996.

5. Wiegand, T.E.v., The Virtual Workbench & the Electronics Training Task. Internal memo, 1994. VETT project at MIT/RLE, 77 Massachusetts Ave. MIT, Cambridge MA 02139-4307, http://mimsy.mit.edu/.

6. Jones, M.R.E., The Haptic Hologram, *Proceedings of SPIE, Fifth International Symposium on Display Holography*, Vol. 2333, pp. 444-447, 1994.

7. Plesniak, W., Klug, M., Tangible holography: adding synthetic touch to 3D display, in S.A. Benton, ed., *Proceedings of the IS&T/SPIE's Symposium on Electronic Imaging, Practical Holography XI*, 1997.

8. Watlington, J., *et al.*, A hardware architecture for rapid generation of electro-holographic fringe patterns, in S.A. Benton, ed., *Proceedings of the IS&T/SPIE's Symposium on Electronic Imaging, Practical Holography IX*, 1995.

9. Deering, M., High Resolution Virtual Reality. Proceedings SIGGRAPH'92, *Computer Graphics*, Vol. 26, No.2, pp. 195-202. July 1992.

an operator can *see* the stylus tip penetrating into the holographic surface before the surface is apparently subtracted away. Higher bandwidth spatial light modulators, efficient data compression techniques, improvements in computation speed, and higher bandwidth data pipelines will help diminish this problem in future generations of holovideo.

Because the visual display is holographic, the full range of horizontal parallax is always available in the viewzone; no visual lag is encountered with motion of the operator's head. Additionally, no special eyewear is necessary to perceive the stereo information.

Additionally, there is also some latency in the system. Both latency and lag are principally due to the fact that the Cheops display routine has to receive and act on a message from the SGI host machine. Some improvement in performance may be achieved by sending the updated hololines *directly* to the Cheops memory module via a high bandwidth link and directly displaying them. This does not require maintaining an expensive open connection between the host and Cheops.

5.2 Sensory conflicts

In this overlapping workspace format, some more subtle conflicts between what is seen and felt do arise. The things which most degrade the desired impression of the hologram/haptic simulation as a single multimodal event are spatial misregistrations and occlusion violations. At the moment when an operator *feels* that the stylus tip is in contact with the surface, if the tip is *seen* either penetrating the surface or not making contact at all due to misregistration of the visual and haptic output, the visual discrepancy is striking. Due to the lag present in the holovideo pipeline, our simulation is vulnerable to this problem when the operator is actively carving the surface.

Allowing interaction between the output of optical projection systems, like holograms, and an operator's hands (plus physical objects or instruments), permits object depth relationships to occur which violate occlusion rules obeyed by the physical world. Normally, when we see the image of one object blocking the image of another from view, we understand that the occluded object is farther from our eyes. In our system, it is possible to interpose part of the stylus between the holographic image and physical output plane of the holovideo optical system, thus blocking the image from the line of sight between the viewer and stylus. In this event, it appears that the farther object (stylus) occludes the nearer (holographic image). This anomalous cue is strong enough to confuse perception, even when correct depth reporting from stereopsis and motion parallax is available.

6. ONGOING / FUTURE WORK

The system stands implemented as described above. Now that we have a working interactive holo-haptic system, our research is focusing on improving modeling and rendering and increasing display update rates (for both visual and haptic channels). Specifically, we are devoting effort to improving the fidelity of materials simulation, and modeling a more realistic haptic representation of lathe carving. We are also working on computing smooth-shaded and visually textured holographic images; better visual and haptic rendering are necessary to improve overall display quality. We are modifying our pipeline to write hologram lines directly to the memory module as well, with the intent of increasing our visual display update rate. When the holovideo display can be updated more rapidly, simulating more complicated dynamics becomes a tractable pursuit. Additionally, it is necessary to make more precise measures of system lag and latency.

7. CONCLUSION

We have described an interactive multimodal spatial imaging system which combines holography and haptics in this paper. We have situated the system in a research context and outlined the requirements of such a coincident workspace system. The component haptic and holographic subsystems were described and the implementation of the whole system was detailed. This is the first time that a truly interactive holo-haptic system has been built; this offers a first glimpse of what full-fledged, interactive holographic systems might be like in the future. Our belief in the importance of high-quality visual and haptic cues as well as the attentive dialog between vision, hands, tools, and the material they manipulate is at the root of the work presented here. This prototype system presents, for further inquiry, a possible interaction technology for future digital design studios and visualization laboratories.

ACKNOWLEDGEMENTS

We would like to thank the Honda R&D Company, NEC, IBM, the Digital Life Consortium at the MIT Media Lab, the Office of Naval Research (Grant N0014-96-11200), Mitsubishi Electric Research Laboratories, and the Interval Research Corporation for their support of this research. We would also like to acknowledge valuable conversations with members of the Spatial

displayed. As the initial haptic cylinder is carved, a visual approximation to the resulting surface of revolution is assembled on the display by loading the appropriate lines from each of these five separate holograms.

First we must determine *how many* and *which* lines should change on the holovideo display. The number of display lines that require updating will vary, depending on exactly which model control points are displaced. In regions near the top or bottom of the carved shape, a smaller region of the curve contributes to the visible extent of the shape, so fewer display lines will require change. The new radius values in **R** corresponding to changed display lines are quantized to match one of the set of five holographic cylinder radii, and each is assigned a radius code based on its quantized value as shown below:

radius (mm)	25.0	22.5	20.0	17.5	15.0
code	5	4	4	2	1

A message, which contains the number of the hololine marking the start of the update region, the number of lines that need to be updated, and the radius codes of each new line, is sent to the holovideo output module on Cheops. In order to minimize the display update time, we are currently updating a maximum of 32 hololines per cycle, a compromise in accuracy which represents only the display lines between the original six control points sent by the haptics module.

4.3 Holovideo indexing

Upon receiving the update message, the holovideo output module must instruct Cheops to collect the appropriate hololines and dispatch them to the display. This is accomplished by indexing into the memory module with the radius code to determine the correct cylinder to display, and then writing the corresponding hololine to the output card (Figure 4). The final holographic image is assembled using hololines from the five individual holograms. It must be noted that this method of hologram assembly is valid only for HPO holograms; for full-parallax holograms, the entire hologram would have to be recomputed. In the absence of the computation and communication bandwidth necessary to update fully-computed holograms in real-time, pre-computed hologram indexing enables rapid, local updating.

5. RESULTS AND DISCUSSION

5.1 System lag and latency

A compelling multimodal representation depends heavily on minimizing, to imperceptible levels, the time lag between the operator effecting changes in the haptic model and the result of that change appearing on the visual display. A reasonable visual update rate (20+ frames per second) is not currently possible with the MIT holovideo system, principally due to the speed at which we can communicate with and update the display. The effect of the resulting system lag, on the order of 0.5 sec., is that

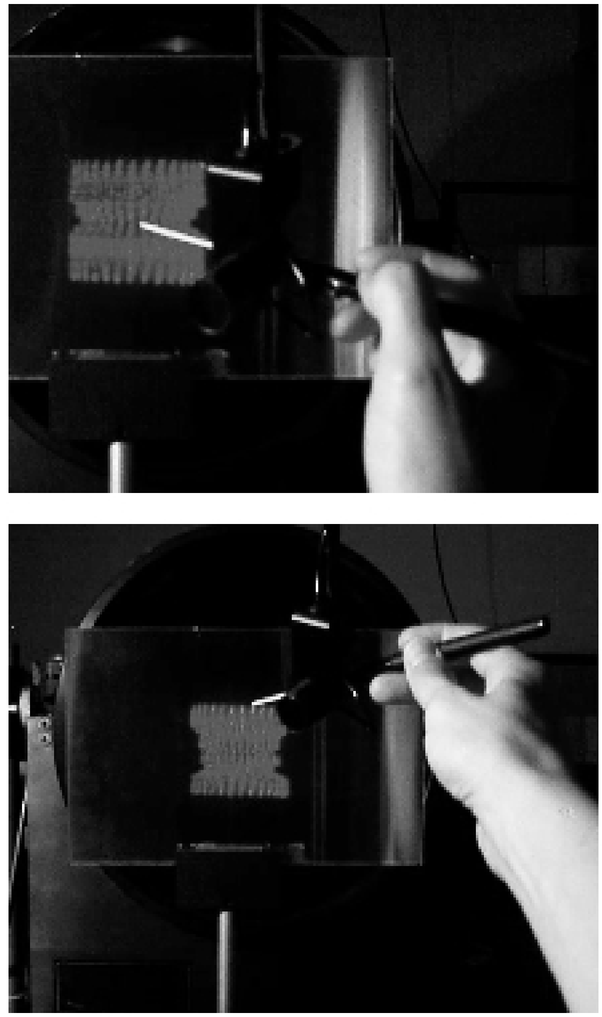

Figure 5a. Interactive carving in the coincident workspace

Figure 5b. Physical hardcopy from a Stratasys 3D printer

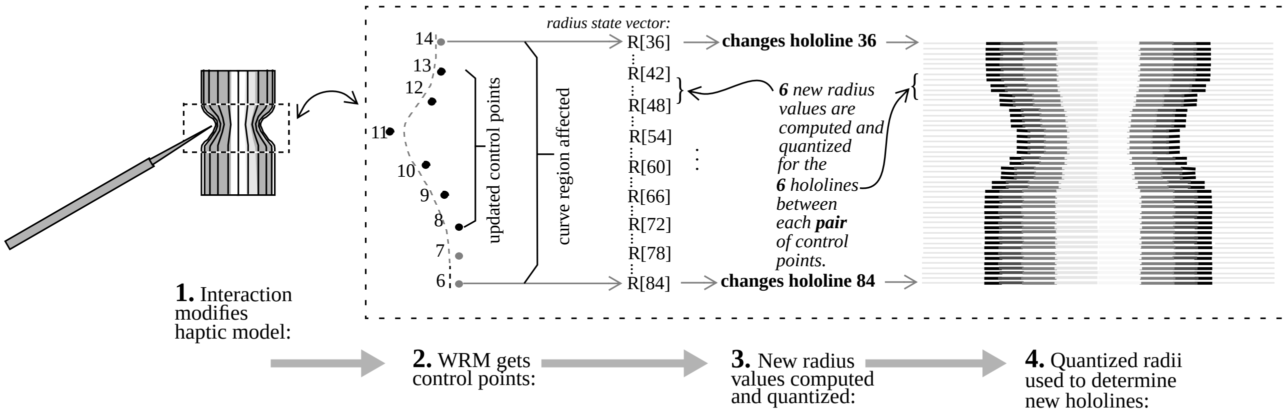

Figure 3. Determining which display lines are affected by changes in haptic model

within an update interval. Though this assumption usually puts us within the realm of normal interaction speed, eventually, communicating a variable number of control points to reflect the precise region of change would be more robust, and future work will implement this change.

4.2 Workspace Resource Manager

Once the WRM receives the message, the changed control points are used to update its own representation of the radius profile. The WRM determines which lines of the holovideo display will be affected by the updated region of the curve. Since the final holographic image will span 120 lines of the display, we maintain a state vector, **R**, with 120 elements whose values represent the exact radii of the surface of revolution at corresponding display lines. A set of six holovideo display lines correspond to the space between any two adjacent control points in the WRM's model. If as many as six control points have changed, it is necessary to recompute radii for the 48 display lines spanning *eight* control points, between which the curve will have been affected (Figure 3). These new radius values are reflected in the state vector **R**. In the current implementation, the WRM's model can also be rendered to a graphics display using SGI's Graphics Library for debugging purposes, and to provide a means for remotely monitoring a user's performance.

Because it is not yet possible to compute 36 Mbyte holograms in real time [3], we decided to pre-compute five cylinder holograms for use in updating the display, as explained shortly. Each hologram displays a cylinder with a different radius, the initial cylinder, and four progressively smaller ones, r_{cyl} (mm) = {25.0, 22.5, 20.0, 17.5, 15.0}, ending with the minimum-radius cylinder. All holographic cylinders are 47.9 mm high. These holograms, from largest to smallest radius, are loaded sequentially into the Cheops memory module. It would be possible to compute a fewer *total* number of lines if we omitted visual texture from the object or restricted texture to be periodic. At system start-up, the cylinder with the largest radius is

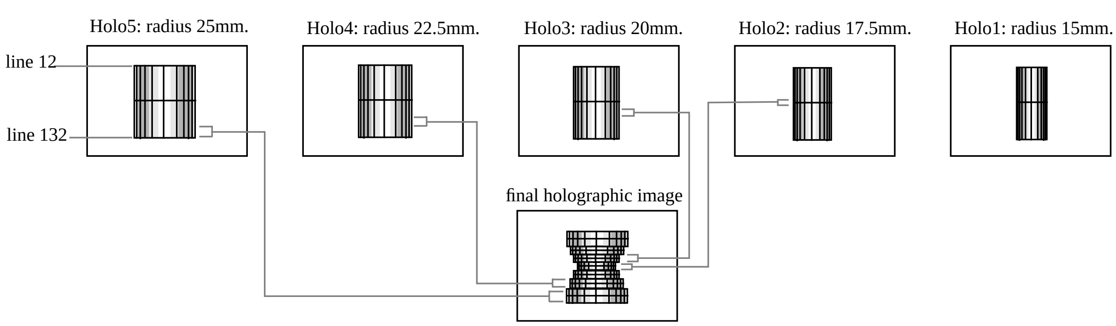

Figure 4. Method of assembling final holographic image from pre-computed holograms

modeled with the same bulk and frictional properties as the cylinder. Currently, the haptics simulation is run on a Pentium PC with an average servo rate of 1500Hz.

The radius profile of the surface of revolution is represented as a cubic B-spline curve with 28 control points, all of which are initially set to the same radius value (25mm) to let us begin with a cylinder. The curve evaluated between the middle 21 points defines the profile of the cylinder body; the remaining top three and bottom four points lie beyond the actual extent of the cylinder, and serve to "lock" the shape at its top and bottom, respectively. Control points are modified as force is exerted on the shape at height *h*, between control points P_i and P_{i+1}. A new radius for the surface of revolution at this height can be computed by evaluating the nonuniform, rational B-spline formulation.

The cylinder can be felt spinning beneath the user's touch, and when pressed with enough force (*i.e.*, when the surface has been penetrated by some threshold distance Δ) the surface deforms. A very simple method for surface deformation is used: the two control points straddling the penetration location are displaced toward the central cylinder axis by a fraction of the penetration distance. The upper point is displaced by $tk\Delta$, and the lower by $(1-t)k\Delta$, with *t* being the normalized distance between the contact point and the lower control point, used in the B-spline formulation. The closer control point is displaced by a greater distance. If contact occurs directly on a control point, then that point alone is displaced by $k\Delta$. Thus, control point displacement modifies the circumference of the cylinder at height *h*, as force is interactively applied.

The parameters *k* and Δ can be adjusted to make carving the rotating cylinder require more or less force. A minimum radius of 15mm is enforced, so that once the surface has deformed this much, the control points update no further. The control point density, 4.17 points/cm, was experimentally determined to be high enough to accommodate local model changes, yet sparse enough to avoid unstable deep notching of the haptic surface.

3.2 Holographic video modeling

We employ the second generation of holovideo in this work [2]. This system is capable of displaying monochromatic, horizontal-parallax-only (HPO) images in a volume of 150 mm x 57.5 mm x 150 mm, and the viewing angle is 30˚. The 3D image produced by holovideo supports the most important depth cues: stereopsis, motion parallax, occlusion, and many pictorial and physiological cues to depth.

For the present purpose, we may consider holovideo to be a black box which accepts two inputs: a *computer-generated hologram* (CGH) and light. The output of the black box is a 3D *holographic image* whose visual and geometrical characteristics depend on how the CGH was computed. Each CGH contains an enormous amount of data—36 megasamples (at 1 byte per sample) apportioned into 144 hololines of 256 kilosamples each. The CGH is made available to the display via a framebuffer. Because holovideo has a non-standard display format, an image-processing system developed at the MIT Media Lab, dubbed *Cheops*, was extended to support it. Cheops has three different module types: processor, input/memory, and output, and an optional memory module provides up to 0.5 Gbytes local to the system. These modules are interconnected by two linear buses. One of these buses, the Nile bus, is capable of sustained high bandwidth (>100 Mbyte/sec.) transfer of samples and the second, the Global bus, is capable of 32 Mbyte/sec. transfer [8].

4. SYSTEM IMPLEMENTATION

4.1 Haptics module

The Workspace Resource Manager (WRM) running on the SGI/Onyx initializes its own model of the surface of revolution, which starts as a cylinder of desired height and radius. It then initiates the haptic simulation by making client calls to the haptics server on the Pentium PC. These calls request creation of a haptic cylinder of the same height and radius at a desired location. The haptics module commences physical simulation of this spinning cylinder, and computes collisions of the Phantom tip with the computational model. Based on these collisions, forces are computed and displayed to the operator's hand, and any resulting shape modifications are reflected in the model update.

Any changes in the cylinder's underlying B-spline representation are automatically communicated from the haptics module to the WRM approximately 30 times per second. The information sent contains the location where change begins on the curve (the number of the bottom-most control point), and values of the six affected control points, ordered from bottom to top. It is assumed that model changes occur reasonably slowly, so that no more than six control points are updated within 0.033 second. Since computing a deformation means updating at most two control points surrounding the point of contact, our communication rate means that we can only guarantee reporting accurate model changes from contact in a region 6.9 mm high space is to visually monitor the hand (or hand-held tool) and its interaction with the object or material. Consequently, the ability to keep both the displayed object and the hand in visual focus is essential, and careful design must be employed to render it so.

Holographic displays eliminate this particular design problem by permitting a viewer to freely converge and accommodate to any point in the display volume. The combination of haptics and holography was first investigated by researchers at De Montfort University for an object inspection task [6]. Visual display was provided by a reflection transfer hologram which presented an aerial image of a control valve. A Computer Controlled Tactile Glove (CCTG) provided coincident haptic display of the same data. Similar informal experiments in combining reflection transfer holograms with force-feedback were also performed by researchers at the MIT Media Laboratory's Spatial Imaging Group. In all of these efforts the interacting hand could literally block the reflection hologram's illumination and prevent image reconstruction.

This problem was addressed by employing full-parallax edge-illuminated holograms in combination with the Phantom for the inspection of static 3D models [7]. The edge-illuminated hologram format allowed hand movements in any part of the visual workspace. Thus a viewer could haptically explore the spatially-registered force model while visually inspecting the holographic image details over a wide field of view. All of these displays were static, however; no dynamic modification could be made to the data presented.

3. SYSTEM DESIGN

Two separate modules comprise the computation which feeds the displays; a *haptics module* that performs force modeling, and the *holovideo module* which pre-computes holograms and drives rapid local holographic display updates based on changes to the model. The haptics and hologram modules are organized by the *Workspace Resource Manager* (WRM) which is notified of geometry changes imparted to the spinning cylinder by the user's hand, and requests hologram updates to local regions of the visual display where changes have occurred. The haptics and hologram modules rely upon separate and characteristically different representations of the cylinder, which are carefully spatially and metrically registered. From the point of view of the user, who is holding the stylus and pressing it into the holographic image, a single multimodal representation of the simulation can be seen and felt changing in response to the applied force. The system overview is shown below in Figure 2.

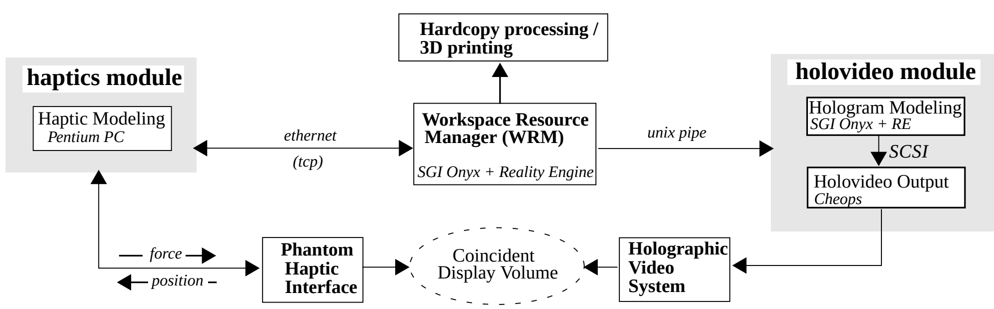

Figure 2. System overview

3.1 Haptic modeling

Research in haptics modeling is yielding methods for modeling the surface and bulk properties of materials, and the computational forces generated as we mechanically interact with them. The fidelity with which computational haptics is currently able to render both the pleasing *feel* of a material interacting with our tools, and the mechanical *cues* that relay information about object and material integrity is rapidly improving. Our haptic device can display force to the user's hand according to its position-based interaction with the computational models describing the object's geometry, bulk and tactual properties. Six encoders on the device provide positional information resolved to approximately 0.1 mm, and three servo motors provide force display up to roughly eight Newtons, within a workspace of about 290 mm x 400 mm x 560 mm.

The haptic cylinder, initially and in subsequent stages of "carving," is represented as a surface of revolution with two caps. It has a mass of 1 gram, an algorithmically defined vertical grating as surface texture, static and dynamic frictional properties, stiff spring bulk resistance, and rotates about its vertical axis at one revolution per second. The cylinder model straddles a static haptic plane (which spatially corresponds with the physical output plane of the holovideo optical system); the haptic plane is

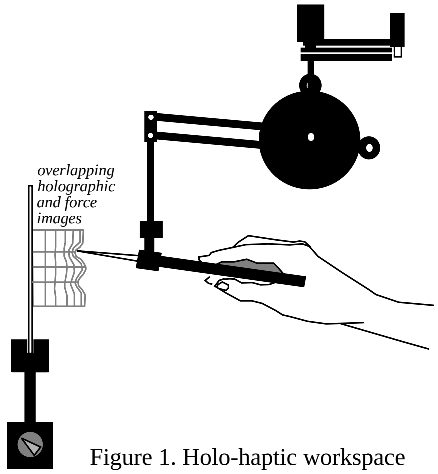

Figure 1. Holo-haptic workspace

The paper is structured as follows. In section 2, we discuss some previous efforts to combine *computational haptics* and *visual feedback* in a coincident workspace format. An overview of the HHS is provided in section 3 and its implementation is detailed in section 4. In section 5, we provide system results and discussion. We describe ongoing and future work in section 6 and conclude in the final section.

2. BACKGROUND

A wide variety of virtual reality (VR) application areas such as telesurgery, training, computer modeling and entertainment employ computational haptics and high-quality computer graphics to study, interact with or modify data. Most of these applications offset the visual and manual workspaces, but several efforts to conjoin eye and hand in interactive applications exist.

An example thematically related to our work is the compelling "Virtual Lathe" described and presented at the SIGGRAPH'92 conference by Michael Deering [9]. In this demonstration, a head-tracked stereo display showed a virtual stock, spinning about its long axis, which a person could interactively lathe using a 3D mouse in the shape of a rod. The demonstration used LCD shutter goggles for stereo viewing, but had no provision for force feedback.

Another interesting example has been demonstrated by researchers at Carnegie Mellon University [4] called the WYSIWYF (What You See Is What You Feel) display. The visual display behaves like a moveable "magic window," interposed between the viewer's eyes and hand, and through which the hand can be seen interacting with a virtual, tangible scene. The work employs a six degree-of-freedom haptic manipulator and monographic visual rendering to combine three pieces of information in this final coincident display: a video image of the operator's hand/arm, the computer graphically rendered scene, and the accompanying force model. The visual display is a color LCD panel with a CCD camera attached to its backplane. This display/camera unit can be moved with respect to the physical scene, while vision-based pose estimation is employed to determine its new orientation. The display shows a computer graphic view of the synthetic scene generated from the newly-computed viewpoint, composited with a live chroma keyed image of the operator's hand/arm moving behind the display and interacting with the haptic device. This display cannot currently reproduce correct occlusion relationships between the hand/arm and virtual objects and provides only monocular cues to scene depth (no stereo viewing or head-tracked motion parallax is available).

In other systems which employ a coincident workspace, the use of a half-silvered mirror to combine an image of the CRT's pixel plane and the haptic workspace is a historically popular and frequently used technique. One such project is the "Virtual Workbench" project, developed by the The Virtual Environment Technologies for Training (VETT) Group at MIT's Research Lab for Electronics [5]. This system, used to study human sensorimotor capabilities and to develop training applications, employs a Phantom haptic interface and the half-silvered mirror technique for coincident stereo display. Representing correct occlusion relationships between the hand and simulated objects is a problem in this display configuration too. Additionally, the workspace that can actually be shared by the visual display and the hand is depth-limited in stereoscopic systems; inherent in these displays is an accommodation-convergence mismatch—a functional disengagement of several systems of the eye which normally function in cooperation. If scene depth is not designed well for the display's particular viewing geometry, eye strain, headaches and unfuseable stereo images can result. Of course, the very purpose of combining the manual and visual work-

Published in *Proceedings of the IS&T/SPIE's Symposium on Electronic Imaging, Practical Holography XII*, 1998

Haptic interaction with holographic video images

Ravikanth Pappu and Wendy Plesniak
Spatial Imaging Group,
MIT Media Laboratory, Cambridge, MA.
{pappu, wjp}@media.mit.edu

ABSTRACT

We describe the implementation of a system which enables a user to interact with and modify an electronic holographic image using a force-feedback device. The force-feedback (or *haptic*) device is capable of sensing and reporting the 3D position of its hand-held stylus and displaying appropriate forces to the user. Thus, a user can feel and modify algorithmically specified shapes in the haptic workspace. We precisely register the haptic workspace with the free-standing, spatial image displayed by the MIT second-generation holographic video system (*holovideo*). In the coincident visuo-haptic workspace, a user can see, feel, and interact with synthetic objects that exhibit many of the properties one expects of real ones, and the spatial display enables synthetic objects to become a part of the user's manipulatory space. To the best of the authors' knowledge, this is the first time that such an interactive holographic system has been built.

Keywords: holography, computer-generated holograms, haptics, interactivity, force-feedback, coincident workspace

1. INTRODUCTION

Throughout the history of holography, there has been considerable interest in building real-time, interactive holographic displays. The problem was recognized to be a difficult one and it is only very recently that quasi real-time holographic displays have made their appearance [1, 2]. Making these displays interactive has proved to be a challenging engineering problem owing to the large computation, communication, and modulation bandwidths involved. The most recent incarnation of holovideo is capable of displaying up to 3 pre-computed 36 megabyte holograms per second. Computing a single hologram still requires about 5 seconds on our fastest computing hardware [3]. These computational and display update rates still short of those required for real-time interactivity. However, by making use of the horizontal parallax only (HPO) nature of our computed holograms and by making only local changes to the hologram, we have assembled a system that enables a user to rapidly interact with and modify an electronic holographic image using a force-feedback device.

The haptic interface used in the system, the Phantom ™, is a three degree-of-freedom (d.o.f) mechanical linkage with a three d.o.f passive gimbal that supports a simple thimble or stylus used by the hand. It is capable of reporting the 3D position of the stylus tip and displaying a force to the user's hand. This allows a user to feel and modify algorithmically specified shapes in the haptic workspace — the space which is capable of being addressed by the force-feedback device. The haptic workspace is precisely registered with the free-standing, spatial image displayed by holovideo. The combined visuo-haptic workspace is referred to as the *coincident workspace* wherein a user can see, feel, and interact with synthetic objects that exhibit many of the properties one expects of real objects.

In the workspace, a single multimodal image of a cylinder to be interactively "carved" is presented. The user sees the hand-held stylus interacting with the holographic image while feeling forces that result from contact with the haptic model (Figure 1). As the user pushes the stylus into the simulated cylinder, its haptic model deforms in a non-volume-conserving way and a simple surface of revolution can be interactively shaped. This change is also reflected in the HPO holographic image. Ultimately, the finished computer model can be dispatched to a 3D printer to produce physical hardcopy of the design (Figure 5b).

Although the holo-haptic system (HHS) reported in this paper is preliminary, it assumes importance for several reasons. First, to the best of the authors' knowledge, this is the first time that a holographic system which permits *model geometry to be interactively modified* has been built. This system offers a first glimpse of what a full-fledged HHS might be like. Second, this paper is not about holovideo *per se*; the HHS marks the first time that holovideo has been used as a *display device* in an interactive system and hints at a time when holographic displays might routinely be used as design and visualization tools. Finally, future engineering challenges which must be overcome to make the HHS a useful tool emerge from the process of its implementation in this form.