riences the feel of virtual information. Extant systems with both visual and haptic output usually employ an *spatially offset* workspace configuration in which the hand's operation of the haptic device is in a space separate from where the eyes are looking. In contrast, our spatially coincident workspace permits the hand, feeling due forces, to be part of the visual display.

Visual sensing, haptic sensing, and motor planning and actions are indeed intimately linked, and in a unique way. The primary human *haptic sensing* hardware (the hands) are also the primary *motor effectors*. Motor actions are usually necessary to bring about manual sensing of specific object properties, and visual assessment of a scene is usually necessary to plan most of the hand/arm movements we make when interacting with objects in a workspace.

As part of most interactions with objects in manipulatory space, one must determine the distance from oneself to the object of interest, so that a reach may be planned and executed without constant visual monitoring of the hand and the object. Many visual cues to depth, both monocular (motion parallax, occlusion, perspective, texture, size, vertical position, color, shadows, aerial perspective and accommodation) and binocular (stereopsis, vergence movements and vertical disparity) can help us to determine the necessary depth/distance information. Binocular visual cues have specifically been shown to improve on-line control of prehensile movements[6]. Thus, providing binocular visual feedback for any tasks which require spatial hand/arm movements may improve performance, or at least may provide a more comfortable and natural feel to the tasks a person must accomplish. The question remains as to whether actually *feeling* a simulated thing exactly where it is *seen* to be in physical space will even more strongly verify the simulation, heighten physical believability or improve performance.

7.0 CONCLUSIONS AND FUTURE WORK

We have described the construction of a coincident visuo-haptic workspace that employs edge-illuminated holographic stereograms for visual display and high quality haptic simulation with the PHANToM force feedback interface. The architectures of the software systems that produce the holographic stereograms and the haptic simulations were discussed, and two example tangible holograms were presented. Problems revealed during experimentation with these holograms were described as well as suggestions for minimizing their impact.

All work described in this paper deals with static scenes, though we are currently developing interactive applications using stereo computer graphics and holographic video and dynamic simulations.

ACKNOWLEDGEMENTS

This work was funded by the Honda R&D Company, IBM Corporation, NEC Corporation and Interval Research Corporation. Equipment was provided by Intel Corporation.

REFERENCES

1. Halle, M., Kropp, A., (1997). "Fast Computer Graphics Rendering for Full-Parallax Spatial Displays". in: S.A. Benton, ed., *SPIE Proc. Practical Holography XI*. (Feb. 1997, current issue).

2. Halle, M.H., (1994). "Holographic stereograms as discrete imaging systems," in: S.A. Benton, ed., *SPIE Proc.* Vol. #2176: *Practical Holography VIII*, pp. 73-84.

3. Katz, D., (1989). **The World of Touch**. Transl. by Kreuger, L. E., Lawrence Erlbaum Associates: Hillsdale.

4. Klug, M.A., Klein, A., Plesniak, W., Kropp, A., Chen, B., (1997). "Optics for full-parallax holographic stereograms". in: S.A. Benton, ed., *SPIE Proc. Practical Holography XI*. (Feb. 1997, current issue).

5. Plesniak, W., Underkoffler, J., (1996). "SPI Haptics Library". in: Salisbury, J.K., Srinivasan, M.A., eds., The Proceedings of the First PHANToM Users Group Workshop, published as M.I.T. Artificial Intelligence Laboratory Technical Report AITR-1596 and M.I.T. Research Laboratory for Electronics Technical Report No. 612.

6. Servos, P., Goodale, M.A., Jakobson, L.S., (1991). The Role of Binocular Vision in Prehension: A Kinematic Analysis. *Vision Research*. Vol 32, No. 8, pp. 1513-1521.

7. Benton S.A., Birner, S.M., Shirakura, A., (1990). "Edge-Lit Rainbow Holograms," in: S.A. Benton, ed., SPIE Proc. Vol. #1212, Practical Holography IV. pp. 149-157.

report, the surface qualities change substantially with depth though their haptic quality remains the same.

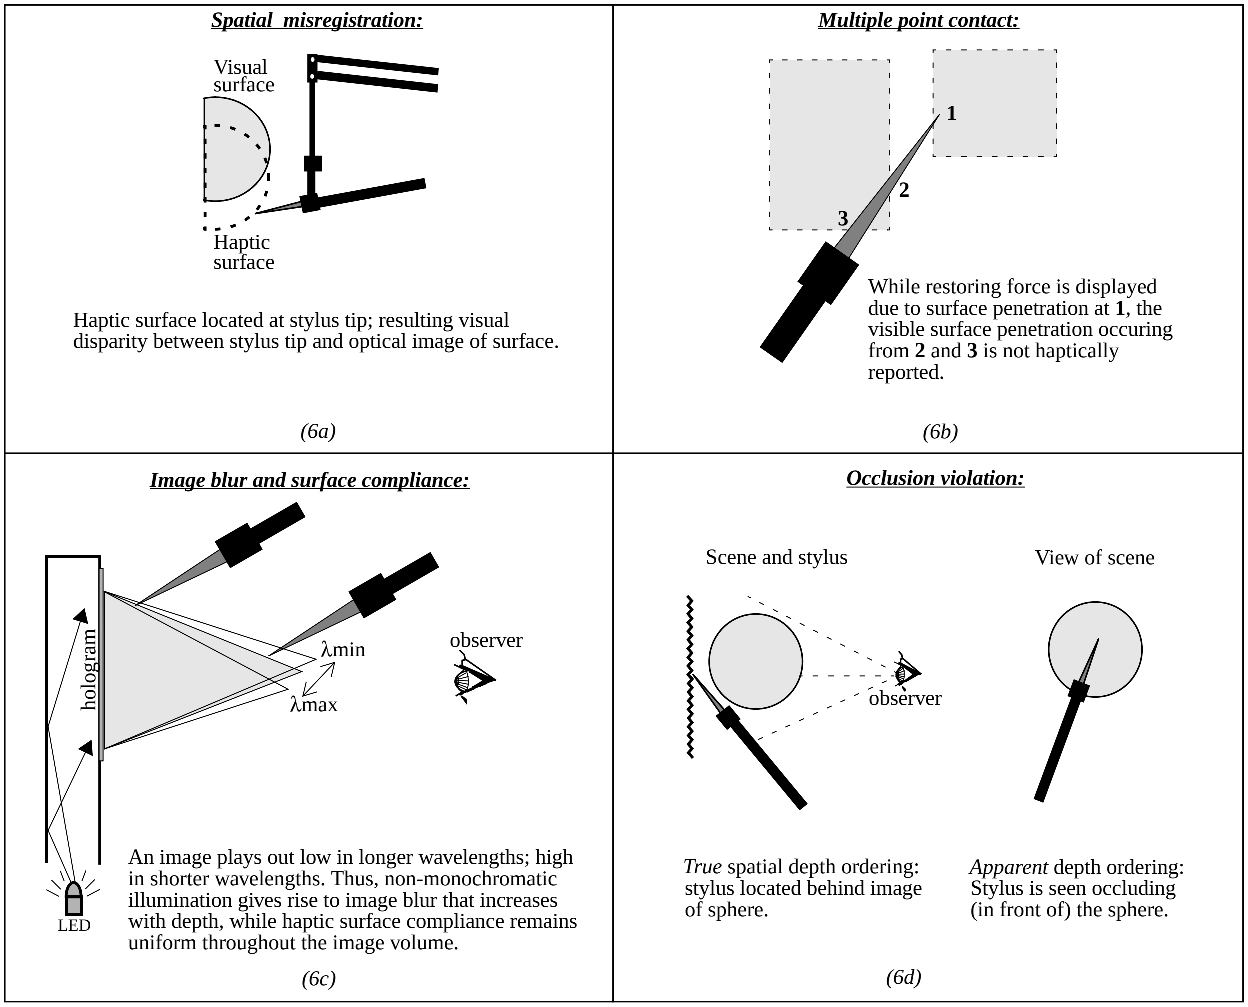

Figure 6. Conflicting sensory cues

These four conflicting cue situations are specific to coincident visuo-haptic display; none of these problems occur when using an offset configuration. Most of the problems can be addressed to diminish their deleterious effects. Spatial registration problems are perhaps the easiest to attend to; careful design of the hologram and the haptic model usually solves them entirely. A haptic model can be scaled to account for the spectral image blur, so that the boundaries of the blurred image volume match the boundaries of the haptic geometry. Surface compliance could be modeled to vary with image blur, so that surface stiffness might decrease in blurry image areas. Most simply, a monochromatic illumination source could be used. Additionally, using a transparent stylus tip, a thin glass rod for instance, appears to slightly minimize the disturbance of the occlusion violations, though these remain problematic.

6.0 DISCUSSION

Aside from the conflicting cue situations which arise during interaction, the effect of haptically interacting with a simple static aerial image is compelling. Visual lag, such as we see with a slowly updated computer-rendered position-marker, cannot impede one's knowledge of hand position in this setup; movement throughout the scene can be as rapid as haptic simulation will permit and as natural as reaching out and apprehending an object in the near environment.

The psychologist David Katz maintained that no sense better convinces us of the reality of objects in our environment than does touch[3]. Indeed, the merits of adding high-fidelity haptic display become instantly obvious when one expeperceived to be co-located with the haptic surface. During contact, if the holographic surface and the haptic surface are not precisely aligned, the misregistration is strikingly obvious to vision. These conflicting visual cues erode the impression of sensing a single object. Instead, the impression of two separate representations is evident. This condition is shown in Figure 6a.

5.2 Occlusion violations

A powerful visual cue to depth in a scene is occlusion. When we see the image of an object being blocked by the image of another, we understand the occluded object to be farther from our eye than the occluding one. Holograms and other optical projection systems permit object relationships to occur which violate the occlusion rules normally obeyed by the physical world.

An example of the sensory conflict during an occlusion violation is shown in Figure 6d. The figure shows the true spatial depth ordering of a scene in which the stylus is farther from the viewer than the closer image of a sphere. Stereopsis and motion parallax would also confirm this depth ordering. Yet since the stylus is interposed between image and hologram, the holographic image that would otherwise occupy the line of sight between the stylus and the viewer's eye will be blocked. As it appears, the object closer to the viewer is visually occluded by the farther object. In most cases the erroneous occlusion cue is strong enough, even in the presence of correct depth accounting from stereopsis and motion parallax, to confound perception.

5.3 Contact with multiple surface points

Obviously, holograms present spatial images which can exhibit no restoring force when penetrated by an object. With no haptic simulation running to detect collisions with model surfaces and to display contact forces, the haptic apparatus is free to pass through the holographic image undeterred. While applying haptic simulation can prevent a single point on the stylus from penetrating the model, device limitations preclude emulation of the kind of multi-point contact that may occur in the physical world.

During each haptic control loop cycle, the simulation checks for a surface collision all along the stylus probe; even if it finds many, it can only compute and display forces for one. If a model surface has been penetrated by the stylus tip, it is assumed the viewer's primary attention is focused there, and forces due to this collision are computed and displayed. However, if not the tip, but other points along the probe have penetrated the model, then the collision closest to the tip is used for computation and display.

The situation permits another kind of occlusion violation to occur as shown in Figure 6b. When the stylus tip is seen and felt in contact with some geometry, the stylus may be rotated around its tip and swept through proximal holographic image volume. Seeing the stylus and holographic image coexist in the same physical volume presents a very confusing visual cue. It remains to be seen whether any modeling tricks can be applied to render a plausible multi-point contact force output, without causing device instability.

5.4 Surface compliance and blur

An artifact of a hologram's diffractive properties is the chromatic blurring that occurs with non-monochromatic illumination. In the transmission edge-illuminated holograms used in this work, the holographic image plays out high and farther from the hologram in wavelengths shorter than the recording wavelength, and lower and closer in longer ones. If the illumination source used in hologram reconstruction is not monochromatic, chromatic blurring will be evident in the final image. Image elements close to the hologram plane will be quite clear, but those farther from the hologram plane will exhibit blur in accordance with source bandwidth.

Since a viewer normally expects scene elements closer to the eye to be more keenly resolved, the blurry image elements near the viewer challenge the impression of image "solidity". This condition is recognized as problematic on its own, but adding coincident haptic display adds further difficulty. Usually an object's bulk material properties, compliance for instance, remain uniform throughout the display volume. If the haptic and visual output are precisely in register, then near the hologram plane the stylus will be exactly coincident with an imaged surface during contact. However, far from the image plane, the stylus will visually penetrate the blurry image of the surface by a substantial distance before contact is felt. Not co-locating image surface and stylus tip during contact, especially when close to the viewer's eye, can diminish the simulation quality. In addition, visual and haptic information conflicts; by visual and haptically texture mapped with a vertical grating with spatial frequency 2.52 cy/cm, as in the hemisphere hologram. On all object surfaces, we haptically model static and dynamic friction, damping and compliance.

The full-parallax master was 30cm x 20.1cm, providing a much narrower view zone for the final hologram, approximately 35 degrees horizontally. The master is comprised of 6700 exposures, 3mm x 3mm in size, of pseudoscopically-rendered frames. The final edge-illuminated hologram was produced in two optical transfer steps. The recording wavelength was 528 nm in the mastering and transfer steps.

The total depth of the final hologram is approximately 3.5 centimeters, and the entire model is imaged on and in front of the hologram plane. Image plane width and height are each 10 centimeters. Like the hemisphere hologram, the maze hologram is also illuminated with an LED centered at 520nm.

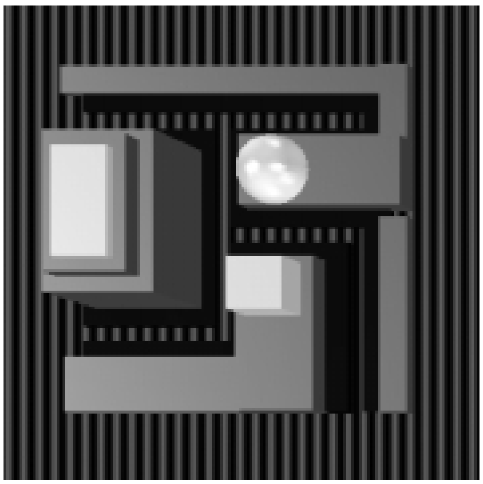

Figure 5a. Rendered frame (40, 50).

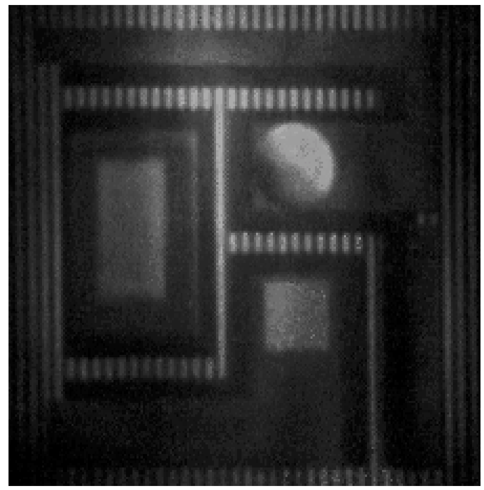

Figure 5b. Maze holographic stereogram

The model presented contains more image features (edges) spaced at varying distances. The maze hologram gives us a chance to compare maze-tracing performance in this coincident and an offset visuo-haptic workspace configuration. Additionally, within this slightly more complex volume of information, there is more opportunity for us to examine the kinds of sensory conflicts occurring in this workspace.

5.0 SENSORY CONFLICTS

As we readily observe in our everyday interactions, harmonious multisensory stimulation usually gives rise to correct perception of objects and events. The broad body of work on sensory interaction indicates that some disparity between visual and haptic information can distort the overall percept while still being tolerated. The ability of sensorimotor systems to adapt to discordant sensory input permits us to perform well even in the presence of distortion, so long as sensory feedback is available. This fact is extremely useful in offset visuo-haptic workspace configurations, wherein the tracked hand or device position is represented as a graphical element on the visual display and the user never actually visually observes her hand. In such workspace configurations, slight spatial misregistrations, or changes in scale between the visual and haptic display can be virtually unnoticeable. Yet too much intermodality disparity can cause the visual and haptic cues to be perceived as entirely separate events, and may be quite confusing or annoying.

Tolerances are lower still when visual and haptic workspaces are superimposed. In our coincident workspace format, anomalous sensory cues and several conflicts between what is seen and what is felt do arise and appear quite curious to the user. These unnatural cues and intersensory conflicts are described in turn below.

5.1 Spatial registration

When exploring a surface with the PHANToM and visually monitoring the device, simultaneous visual and haptic cues to the surface location are available. The location of the stylus tip is seen when we feel contact, and the tip is

The haptic workspace must be precisely isometric along all axes. Whereas slightly imperfect x, y, z grids are often hard to notice in spatially offset workspace configurations, even slight curvatures of space are quite distracting when the visual and haptic displays are coincident. For this reason, the haptic device kinematics must be computed so that the device reports precise measurements along a rectilinear grid.

So, an object modeled to have a specific extent in x, y and z will also have a holographic image of the same extent, and an equal haptic extent.

4.0 RESULTS

4.1 Hemisphere hologram

The simplest model we have worked with is a hemisphere affixed to a vertically oriented plane. The plane is both visually and haptically textured with a vertical grating with spatial frequency 2.52 cy/cm. The hemisphere is visually texture mapped so that its three-dimensionality is more clearly apprehended, but it has no haptic surface relief. On all surfaces, we haptically model static and dynamic friction, damping and compliance. A rendered component frame and the final hologram are shown below in Figure 4a and b.

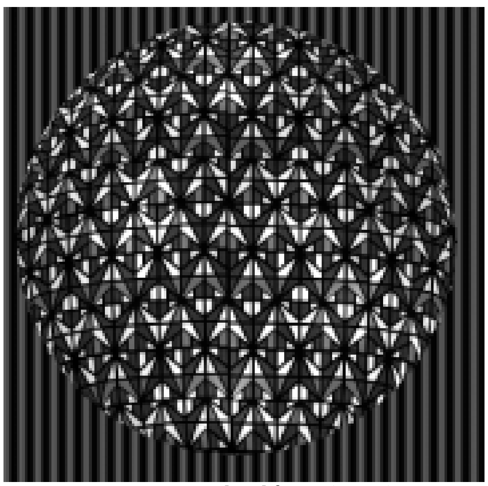

Figure 4a. Rendered frame (75, 50).

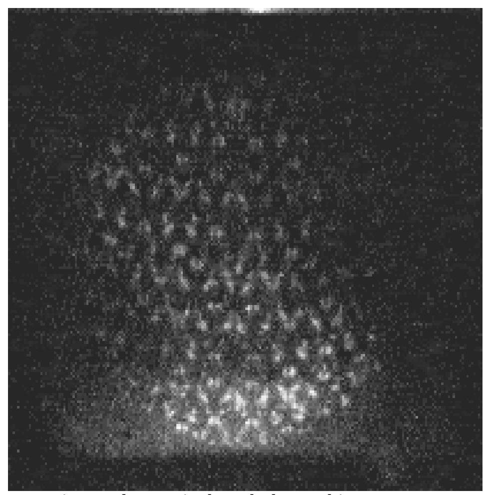

Figure 4b. Hemisphere holographic stereogram

The full-parallax master hologram was 30 x 45 cm, providing a broad view zone of approximately 50 degrees horizontally, 30 degrees vertically. The master contains 15,000 3mm x 3mm exposures of pseudoscopically-rendered frames[1] and required approximately 7 hours to print. A final edge-illuminated hologram was produced in an additional optical transfer step[4]. Both the mastering and transfer steps used 514 nm as the recording wavelength.

The principal benefit of using edge-illuminated holograms in haptic applications lies in the steep-angle lighting that they incorporate; in the final viewing configuration, the hand and haptic apparatus do not block the illumination source as they interact with the image. The total depth of the final hologram is approximately 4 centimeters, all in front of the image plane. Image plane width and height are each 10 centimeters. The hologram is illuminated with an LED centered at 520nm, which yields a bright but slightly blurry image. Using a 532 nm illumination source instead of the LED improves image clarity substantially.

The scene presented was intended to have very few formal features; with the final tangible hologram we have a simple example with which to explore the spatial coincidence of the visual and haptic displays, examining tolerances for spatial misregistration and curvature mismatch (by replacing the hemisphere with a haptic model of a hemi-ovoid).

4.2 Maze hologram

A slightly more complex example is based on a model of a maze which is comprised of blocks and oriented against a vertical back plane. The blocks vary in size and spacing, and the channels are narrow. The back plane is also visually

2.2 Simulating and displaying forces

All haptics rendering is displayed using the PHANToM haptic interface which allows us to precisely track position and rotation at the center of the gimbal joint at the end of the linkage. We use the device to display a force vector that depends on the positional and rotational encoder output. For haptics applications in which the PHANToM is controlled by a PC, we have built a client-server system[5] so that high-quality stereo computer graphic rendering, in lieu of a hologram, can also accompany haptic simulation. The system currently consists of a 586 (Pentium) PC to drive the PHANToM and an SGI to render stereo computer graphics. The machines communicate over Ethernet using TCP/IP.

The haptics server recognizes commands from a client machine for object creation and deletion, activation and deactivation, setting global drift forces and global viscosity, specifying surface damping and friction (static and kinetic), creating textures on objects, and for position-query. In turn, the client application receives confirmations of successfully-completed commands, requested position information, and automatic notifications when objects tagged as "active" have been touched. The basic system is shown below [Figure 3].

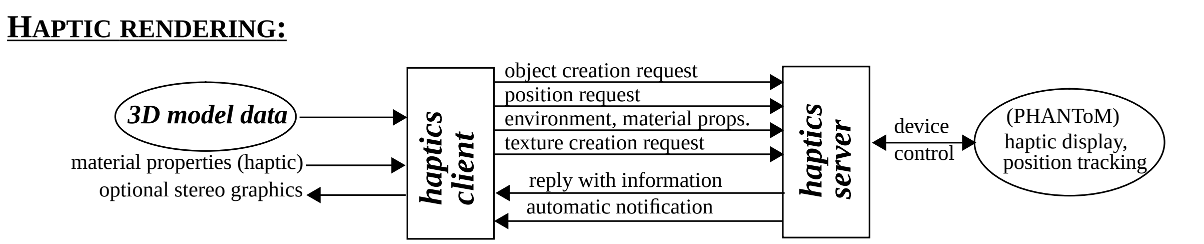

Figure 3. Haptic simulation pipeline.

Within each haptic control loop cycle, which accommodates a servo rate of about 4000 pts/sec, the location of the stylus is reported, and its collision with all objects is determined. If the stylus is in contact with any object, the force in the direction of the object surface normal is computed as well as forces from specified surface properties (friction, damping, texture). Force contributions from all objects are added to global drift or damping forces, and the result is displayed to the user. Object state is updated as necessary. Any requests from the client are then addressed before position polling repeats.

There are obviously many ways to model the surface and bulk properties of materials, and the computation of forces generated when we interact with them. Because the only directly measured state provided by the PHANToM is position, indirect methods must be used for computing physical quantities involving velocity and acceleration. Our haptics library uses a variety of simple and computationally fast models[5] to provide a flexible platform for mocking up the haptic simulations that accompany visual display of the same 3D model. The complexity of stereo computer graphic rendering, handled by the client, does not interfere with server control loop speed; thus visual simulations rendered by the client can be arbitrarily elaborate.

3.0 COINCIDENT SPATIAL DISPLAY

Presenting two spatially coincident simulations requires that each be constructed with the same final spatial metric in mind. Thus, three important conditions must be met:

> *The 3D model and rendering parameters must precisely match the hologram geometry.* Hologram parameters must be designed into the computer graphic rendering geometry[2] and any distortion, such as that produced by imperfect conjugation, must be also taken into account when rendering the stereogram's perspective views.
>
> *The holographic printing process must be distortion-free.* Any factor known to cause distortion during the printing of the master or transfer holograms, such as shrinkage, optical magnification of the master, or imperfect conjugation must be addressed[4].

tics simulation are briefly described below.

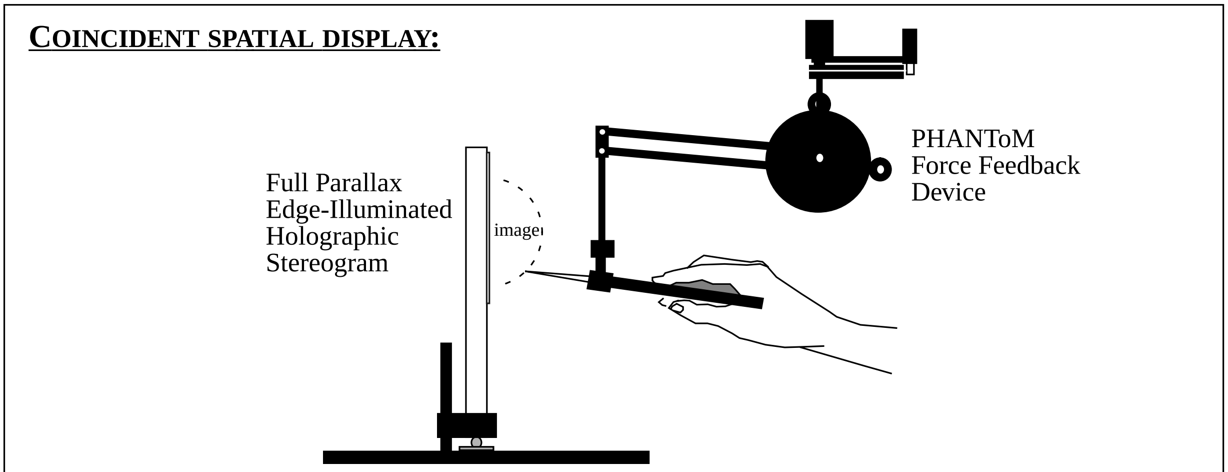

Figure 1. PHANToM stylus probing holographic image.

2.1 Designing and printing the hologram

We have developed a computer graphics rendering server and client applications for both scene design and holographic stereogram printing[4].

2.1.1 Scene design:

The client scene design application makes it possible to import 3D model data, interactively light the scene, tailor the computer graphics view and camera to match hologram geometry, change surface rendering parameters, and request that a perspective view be rendered by the render server. All final model and render parameters are saved in a configuration file which is read directly by the render server at printing time.

Both the scene design application and the render server have been designed to accommodate a wide variety of hologram viewing geometries, and both one-step and multiple-step printing processes[1]. Thus, one particular printer may produce a stereogram master which requires one transfer step; another printer setup may produce a one-step white-light viewable holographic stereogram. For any given master setup, the specific geometries of the printer can be specified during the design process, and then recorded in the configuration file.

2.1.2 Printing:

During the hologram printing process, the printing application requests that a perspective view be rendered by the server using the appropriate configuration file, and upon notification of the renderer's success, sends appropriate exposure and frame-advance control sequences to the hologram printer. In this manner, the newly rendered frame is recorded and the printer is readied for the next exposure. The render server, as well as the scene design client, are usually run on an SGI Onyx Reality Engine, and the printing client can run on any unix machine. The processes communicate over Ethernet using remote procedure calls (rpc). The process of printing is shown below [Figure 2].

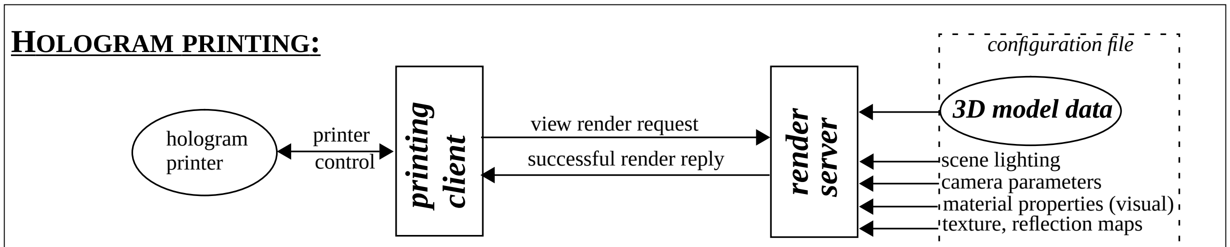

Figure 2. Holographic stereogram printing pipeline

Published in *Proceedings of the IS&T/SPIE's Symposium on Electronic Imaging, Practical Holography XI,* 1997

Tangible holography: adding synthetic touch to 3D display

Wendy Plesniak[1], Michael Klug[2]

MIT Media Laboratory, Cambridge, MA, USA

ABSTRACT

Just as we expect holographic technology to become a more pervasive and affordable instrument of information display, so too will high fidelity force-feedback devices. We describe a testbed system which uses both of these technologies to provide simultaneous, coincident visuo-haptic spatial display of a 3D scene.

The system provides the user with a stylus to probe a geometric model that is also presented visually in full parallax. The haptics apparatus is a six degree-of-freedom mechanical device with servomotors providing active force display. This device is controlled by a free-running server that simulates static geometric models with tactile and bulk material properties, all under ongoing specification by a client program. The visual display is a full parallax edge-illuminated holographic stereogram with a wide angle of view. Both simulations, haptic and visual, represent the same scene. The haptic and visual displays are carefully scaled and aligned to provide coincident display, and together they permit the user to explore the model's 3D shape, texture and compliance.

1.0 INTRODUCTION

Given the dexterous and precise way people naturally interact with hand-held objects or tools, we see that interaction mediated by current interface technology in many computer-interactive systems is still a very poor approximation to the real thing. In most computer-interactive systems, often a representation of one's positional input, like a cursor or a computer-rendered stylus or hand, visually reports the hand's movement and interaction with the image. The visuo-haptic workspace described here represents an effort toward a more natural theater for interaction; participants are able to see their true hand interacting with the image, while also feeling the appropriate forces.

We have configured a prototype visuo-haptic workspace that exhibits many of the same properties as a true physical workspace; for manipulatory tasks, we have the ability to move into this space with our effectors, and see our effectors affecting objects or agents represented within. Precisely registered force feedback is essential, particularly when effectors are *seen* making contact with things.

The prototyped workspace has several interesting characteristics to offer future interactive applications: first, the simulation is imaged into physical space with a hologram to diminish the boundary between physical space and ‘‘display space”. Next, using a 6 degree-of-freedom force-feedback display to simulate physical contact with static and active objects in a workspace permits interaction that is more Newtonian than symbolic. And finally, by rendering spatially coincident visual and haptic simulations, we begin to enable the looking, reaching and touching strategies that people use with physical objects.

Our testbed system currently uses a PHANToM haptic interface, a mechanical linkage and a three degree-of-freedom passive gimbal that supports a thimble or stylus used by the hand. Six encoders on the device provide positional information resolved to approximately 0.1 mm, and three servo motors provide force display up to roughly 8 Newtons. We combine force feedback with either stereo computer graphics or three dimensional holographic edge-illuminated display[7], to present simulated static scenes [Figure 1]. We describe in this paper two *tangible holograms* and the systems built to accomplish the holography and haptic simulation.

2.0 HOLOGRAM AND HAPTIC DISPLAY DESCRIPTION

The systems for producing the holograms and the haptics simulations are separate, but both take the same 3D geometry description as input. The hologram-production pipeline also requires specification of camera geometry, scene lighting information and all visual material properties for rendering, whereas the haptics pipeline requires tactual and bulk material property specification. The general system for hologram rendering and printing, and the system for hap-

1. wjp@media.mit.edu
2. klug@media.mit.edu