

# Design of the *ImmersiveTouch*<sup>™</sup>: a High-Performance Haptic Augmented Virtual Reality System

Cristian Luciano, Pat Banerjee, Lucian Florea, Greg Dawe

Electronic Visualization Laboratory  
Industrial Virtual Reality Institute  
University of Illinois at Chicago  
842 West Taylor St., Chicago, IL 60607  
{clucia1, banerjee, lflore11}@uic.edu; dawe@evl.uic.edu

## Abstract

*ImmersiveTouch*<sup>™</sup> is the next generation of augmented virtual reality technology, being the first system that integrates a haptic device, with a head and hand tracking system, and a high-resolution and high-pixel-density stereoscopic display. Its ergonomic design provides a comfortable working volume in the space of a standard desktop. The haptic device is collocated with the 3D graphics, giving the user a more realistic and natural means to manipulate and modify 3D data in real time. The high-performance, multi-sensorial computer interface allows easy development of medical, dental, engineering or scientific virtual reality simulation and training applications that appeal to many stimuli: audio, visual, tactile and kinesthetic.

## 1 Introduction

*ImmersiveTouch*<sup>™ 1,2</sup> is a new haptics-based high-resolution augmented virtual reality system that provides an efficient way to display and manipulate three-dimensional data for training and simulation purposes. It is a complete hardware and software solution (Figure 1). The hardware integrates 3D stereo visualization, force feedback, head and hand tracking, and 3D audio. The software provides a unified API (Applications Programming Interface) to handle volume processing, graphics rendering, haptics rendering, 3D audio feedback, interactive menus and buttons.

This paper describes the design process of the hardware as well as the software of the *ImmersiveTouch*<sup>™</sup> prototype. The problems of current virtual reality systems and how they motivated the design of this system will be explained in the following section. The hardware constraints considered to achieve the optimal placement of its components will be described in section 3. How the *ImmersiveTouch*<sup>™</sup> API provides an easy workbench to develop haptics-based virtual reality applications integrating a set of C++ libraries will be clarified in section 4. The calibration procedure needed for a correct graphics/haptics collocation will be described in section 5. Finally, the system performance and possible future improvements will be seen in section 6.




Figure 1: The *ImmersiveTouch*<sup>™</sup> prototype

<sup>1</sup> <sup>™</sup> Board of Trustees of the University of Illinois

<sup>2</sup> Patent pending

## 2 Background and previous research

Rear-projection-based virtual reality (VR) devices, including the CAVE<sup>®</sup> [4] and the ImmersaDesk<sup>®</sup> [5], create a virtual environment projecting stereoscopic images on screens located between the users and the projectors. These displays suffer from occlusion of the image by the user's hand or any interaction device located between the user's eyes and the screens. When a virtual object is located close to the user, the user can place his/her hand "behind" the virtual object. However, the hand will always look "in front" of the virtual object because the image of the virtual object is projected on the screen. This visual paradox confuses the brain and breaks the stereoscopic illusion.


Augmented reality displays are more suitable for haptics-based applications because, instead of projecting the images onto physical screens, they use half-silvered mirrors to create virtual projection planes that are collocated with the haptic device workspaces. The user's hands, located behind the mirror, are integrated with the virtual space and provide a natural means of interaction. The user can still see his/her hands without occluding the virtual objects.

Another problem of regular VR devices displaying stereo images is known as the "accommodation/convergence conflict" [1] (Figure 2). The accommodation is the muscle tension needed to change the focal length of the eye lens in order to focus at a particular depth. The convergence is the muscle tension to rotate both eyes so that they are facing the focal point. In the real world, when looking at distant objects the convergence angle between both eyes approaches zero and the accommodation is minimum (the cornea compression muscles are relaxed). When looking at close objects, the convergence angle increases and the accommodation approaches its maximum. The brain coordinates the convergence and the accommodation. However, when looking at stereo computer-generated images, the convergence angle between eyes still varies as the 3D object moves back and forward, but the accommodation always remains the same because the distance from the eyes to the screen is fixed. When the accommodation conflicts with the convergence, the brain gets confused and causes headaches.

In computer graphics the stereo effect is achieved by defining a positive, negative, or zero parallax according to the position of the virtual object with respect to the projection plane. Only when the virtual object is located on the screen (zero parallax) the accommodation/convergence conflict is eliminated. In most augmented reality systems, since the projection plane is not physical, this conflict is minimized because the user can grab virtual objects with his/her hands nearby, or even exactly at, the virtual projection plane.

Current examples of these kinds of augmented reality devices are:

- PARIS<sup>™</sup> (Personal Augmented Reality Immersive System) [8]
- Reachin display [11]
- SenseGraphics 3D-MIW [13]



**Figure 2:** Accommodation/convergence conflict

## 2.1 PARIS™

PARIS™ is a projection-based augmented reality system that uses two mirrors to fold the optics and a translucent black rear-projection screen illuminated by a Christie Mirage 2000 stereo DLP projector (Figure 3). The user looks through the half-silvered mirror that reflects the image projected onto the horizontal screen located above the user's head. The screen is positioned outside the user's field of view, such that only the reflected image is viewable by the user looking at the virtual projection plane. This is important because since the mirror is translucent, the brightness of the image projected on the screen is higher than the brightness of the image reflected by the mirror. Otherwise, the screen would easily distract the user.

The essential idea behind haptic augmented reality systems is to keep the collocation of the graphical representation and the haptic feedback of the virtual object. To maintain certain realistic eye-hand coordination, the user has to see and touch the same 3D point in the virtual environment. In PARIS™ a head tracking system handled by a dedicated networked "tracking" PC enhances this collocation. The head position and orientation is continuously sent to the "rendering" PC over the network to display a viewer-centered perspective. This configuration is similar in the CAVE® and the ImmersaDesk®. In PARIS™, the tracking PC uses a pcBIRD, from Ascension Technologies Corp. for head and hand tracking.



**Figure 3:** PARIS™

Due to its large screen (58" x 47"), PARIS™ provides 120° of horizontal field of view (FOV) and, therefore, a high degree of immersion. The maximum projector resolution is 1280 x 1024 @ 108 Hz, which is adequate for a typical desk-top sized screen. However, since the screen used in PARIS™ is considerably larger, the pixel density (defined as the ratio resolution/size) is 22 pixels per inch (ppi), which is too low to distinguish small details.

Visual acuity is a measurement of a person's vision. Perfect visual acuity is 20/20. The limit for legal blindness in the US is 20/200, which means that a perfect eye can see an object at 200 feet that a legal blind can only see at 20 feet. According to [19], visual acuity for displays can be calculated as  $20 / (\text{FOV} * 1200 / \text{resolution})$ . In PARIS™, this is  $20 / (120 * 1200 / 1280 \text{ pixels}) = 20 / 112.5$ , which is close to the limit of legal blindness. Even though we can read the text shown by the image reflected on the half-silvered mirror (since the image is flipped by the projector), its poor visual acuity makes

reading very uncomfortable. This makes PARIS™ an inadequate choice for application development.

The workspace of the Sensable Technologies' PHANTOM® Desktop™ is approximately a six-inch cube. Therefore, the graphics volume exceeds the haptics volume considerably, causing not only a small portion of the virtual space to be touched with the haptic device, but also just a few pixels are essentially used to display the collocated objects.

Finally, due to the expensive stereo projector and cumbersome assembly, the cost to build a PARIS™ is high for a large-scale deployment.

## 2.2 Reachin display

The Reachin display is a low-cost CRT-based augmented reality system (Figure 4). One advantage of Reachin display with respect to PARIS™ is the fact that graphic and haptic workspaces match, so the user can touch all the virtual objects in the virtual environment.

Its monitor resolution is 1280x720 @ 120 Hz. Since the CRT screen is 17 inches diagonal, the pixel density is higher than that of PARIS™: approximately 75



**Figure 4:** Reachin display

ppi. With a horizontal FOV of 35°, the visual acuity is  $20/(35^\circ * 1200/1280) = 20/32.81$ , resulting in a better perception of small details. However, the image reflected on the mirror is horizontally inverted; therefore, the Reachin display cannot be used for application development. In fact, to overcome this drawback, one has to use the proprietary Reachin API to display properly inverted text on virtual buttons and menus along with the virtual scene.

One of the main problems of the Reachin display is the lack of head tracking. It assumes the user's head is fixed all the time, so the graphics/haptics collocation is only achieved at a particular sweet spot, and totally broken as soon as the user moves his head to the left or right looking at the virtual scene from a different angle. In addition, the image reflected on the mirror gets out of the frame because the mirror is too small. In addition to that, unlike PARIS™, the position of the screen is inside the user's field of view, so it is very distracting.

### 2.3 SenseGraphics 3D-MIW

SenseGraphics is a portable auto-stereoscopic augmented reality display ideal for on-the-road demonstrations (Figure 5). It uses the Sharp Actius RD3D laptop to display 3D images without requiring wearing stereo goggles. It is relatively inexpensive and very compact.

However, it presents the following drawbacks. Like most auto-stereoscopic displays, the resolution in 3D mode is too low for detailed imagery: each eye sees 512x768 pixels. The pixel density is less than 58 ppi. With a FOV of 35°, the visual acuity is  $20/(35^\circ * 1200/512 \text{ pixels}) = 20/82.03$ . Like the Reachin display, the haptics/graphics collocation is poor because it lacks a head tracking system. Due to the orientation of the screen, only the reflected image is viewable. Even though, because of the short distance from the screen to the mirror, and their small sizes, the user's vertical FOV is too narrow to be comfortable. Once again, the image is inverted, so it is not suitable for application development.



Figure 5: SenseGraphics

## 3 Hardware of the *ImmersiveTouch*™

The drawbacks of current augmented reality systems and why they motivated the design of a new system were described in the previous section. This section covers the constraints taken into consideration for the design of the *ImmersiveTouch*™ hardware in such a way that the problems detected in current displays were solved, or at least, minimized. Parametric CAD software was used to set the constraints and analyze the design.

### 3.1 Haptic device and virtual projection plane

In order to design a haptic augmented reality system, we must first determine the haptic device position with respect to the user. Ergonomic analyses were performed to identify the ideal position of the haptic device considering its working volume and a comfortable user's posture having elbow and wrist support. Figure 7 shows the desired position of the haptic device and its workspace. Having the user looking directly at the stylus at the origin of the haptic coordinate system, we define a line between the position of the eyes and the center of the haptic workspace. The virtual projection plane needs to be located exactly at the center of the haptic workspace and oriented perpendicular to that line. The angle of the virtual projection plane with respect to the table resulted to be 45° (Figure 9). This fundamental constraint was maintained through the design process. We manipulated the positions and orientations of the monitor and the half-silvered mirror while maintaining the virtual projection plane in its optimal location. A height-adjustable chair is used to accommodate different users.





Figure 7: The haptic workspace




### 3.2 High-resolution monitor and half-silvered mirror

We decided to incorporate a 22" monitor with a maximum resolution of 1600 x 1200 @ 100 Hz. Since the monitor screen is 16" x 12", the pixel density is 100 ppi, which is higher than that of the Reachin display. The horizontal FOV is 33°. Therefore, the visual acuity is  $20/(33^\circ * 1200/1600 \text{ pixels}) = 20/24.75$ , which is close to the perfect vision. The refresh rate of 100 Hz diminishes the annoying flicker caused by the active stereo goggles, minimizing the strain on the eyes.

In order to use *ImmersiveTouch*<sup>™</sup> as a regular workstation for application development, we must be able to read the text shown



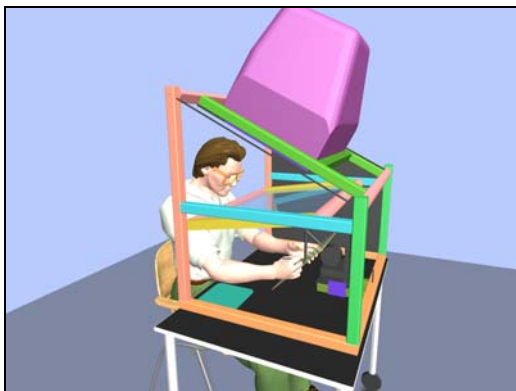
**Figure 9:** The virtual projection plane



**Figure 10:** What if the mirror is horizontal?

orientation of the virtual projection plane, the following step is to study all the possible configurations for the monitor and the mirror to maintain that fundamental constraint. The mirror corresponds to the bisector of the angle between the monitor screen and the virtual projection plane. Thus, both the mirror and monitor need to be coordinated in order to maintain the virtual projection plane at 45°.


Positioning the mirror horizontally, as in the Reachin display, moves the monitor inside the user's field of view when the user



**Figure 12:** The final design

by the image reflected on the mirror. In the case of *PARIS*<sup>™</sup>, the projector itself does the image inversion. In our case, regular CRT monitors do not provide that option. There are some hardware video converters that are connected between the graphic card and the monitor to mirror the image, but they are very expensive and most of them do not support a resolution of 1600x1200 @ 100 Hz. Therefore, we decided to modify the electronics of the CRT monitor. Replicating an old trick done in the consoles of arcade games in the '80s, we flipped the image simply reversing the wires of the horizontal deflector yoke of the monitor. That is a very simple, but effective, solution.

Having the desired position and




**Figure 11:** What if the screen is horizontal?

is looking at the top of the virtual projection plane (Figure 10). On the other hand, locating the monitor screen horizontally, as in *PARIS*<sup>™</sup>, the user's head occludes the monitor image reflected on the mirror (Figure 11). After analyzing both *PARIS*<sup>™</sup> and Reachin display, we arrived at a feasible solution in which the monitor is both outside the user's field of view and sufficiently separated from the user's head (Figure 12).


It is worth mentioning that the mirror is sufficiently wide (29"x21") to allow the user to view virtual objects from different viewpoints (displaying the correct viewer-centered perspective) moving his/her head up to one foot to the left and right without breaking the visual illusion (Figure 13).

### 3.3 Head and hand tracking system

To obtain a correct graphics/haptics collocation, the use of the head tracking system is fundamental. In addition to that, head tracking allows us to render a correct viewer-centered perspective, in which both left and right views are perfectly aligned with the user's eyes, even when the user tilts his/her head. The pcBIRD from Ascension Technologies, Corp., used by PARIS™, presents the drawback that it requires a legacy computer with an ISA slot. Instead, we use the pciBIRD, which is powered by the PCI bus, currently available in most of the new computers. Eliminating the latency caused by the network communication from a "tracking" PC to a "rendering" PC improves the real-time performance, whilst it decreases the cost of purchasing and maintaining two networked computers. In *ImmersiveTouch*™, a single dual-processor computer handles the graphics and haptics rendering as well as the head and hand tracking.



**Figure 13:** The viewer-centered perspective



**Figure 14:** The tracking system transmitter

pciBIRD receiver and provide access to 4 buttons through the serial port. Figure 14 shows the optimal location for the transmitter (at one side of the device) which affords sufficient tracking range for the hand and head while maintaining adequate distance from the monitor.

Another issue to be considered is the location of the transmitter of the electromagnetic tracking system. Since the pciBIRD lacks a mechanism to synchronize the I/O reading with the monitor refresh rate (unlike pcBIRD, miniBIRD, nest of Bird, and Flock of Birds), if the transmitter is located close to the monitor, it incorporates magnetic noise to the monitor. On the other hand, if the transmitter is located far away from the receivers the accuracy of the tracking system decreases while its jitter increases.

Hand tracking is very useful because it allows users to use both hands to interact with the virtual scene. While they can feel tactile sensations with the hand holding the haptic stylus, they can use the tracked hand to move the 3D objects, manipulate lights, or define clipping planes in the same 3D working volume.


For hand tracking, we use the SpaceGrips® [10] that holds a

### 3.4 Summary of features of *ImmersiveTouch*™ and alternative systems

Feature	PARIS™	Reachin display	SenseGraphics	<i>ImmersiveTouch</i> ™
Display resolution	1280x1024	1280x720	512x768	1600x1200
Display refresh rate	108Hz	120Hz	60Hz	100Hz
Pixel density	22 ppi	75 ppi	58 ppi	100 ppi
Visual acuity (20/20 = perfect)	20/112.5	20/32.81	20/82.03	20/24.75
Haptic and graphic volumes match	No	Yes	Yes	Yes
Head and hand tracking	Yes	No	No	Yes
Number of computers required	Two (one legacy PC for tracking)	One	One	One
Comfortable wide mirror	Yes	No	No	Yes
Suitable for application development	No	No	No	Yes
Only reflected image is viewable	Yes	No	Yes	Yes

## 4 Software of the *ImmersiveTouch*<sup>TM</sup>

Since the PARIS<sup>TM</sup> evolved from the CAVE<sup>®</sup> and the ImmersaDesk<sup>®</sup>, both invented at the Electronic Visualization Laboratory, the applications developed for PARIS<sup>TM</sup> use VRCO's CAVELib<sup>TM</sup> for the graphics rendering, and Trackd<sup>®</sup> [18] for the head and hand tracking system. Even though they are excellent libraries, they require users to purchase royalties and pay maintenance fees. In the case of the Reachin display and the SenseGraphics 3D-MIW, users are encouraged to purchase the Reachin API and the SenseGraphics H3D-API respectively. Since these libraries are not open source, the users are limited to use only the functions provided by those APIs and rely on their customer support to fix bugs or implement improvements.



**Figure 15:** Haptic Visible Human

Instead, we use freely-available and/or open source libraries, and combine them so users do not have to worry about performing cumbersome integrations and they can focus on the development of haptics-based applications for *ImmersiveTouch*<sup>TM</sup>. In this way, we offer not only an open architecture but also a way to implement enhancements towards a bug-free library.

Two applications being currently under development at the University of Illinois at Chicago (UIC) using the *ImmersiveTouch*<sup>TM</sup> API are the “Haptic Visible Human”, which helps Medicine students to learn human anatomy “touching” the Human Visible Project<sup>®</sup> dataset [16] (Figure 15), and the “Periodontal Training Simulator”, which is a joint project with the Department of Periodontics at UIC to teach Dentistry students to detect calculus and cavities, and to measure depths of

dental pockets based on their sense of touch (Figure 16).

*ImmersiveTouch*<sup>TM</sup> API integrates the following libraries:


- VTK 4.5 for volume processing and surface extraction [9]
- Coin 2.0 (Open Inventor) for graphics rendering [15]
- GHOST 4.0 SDK for haptics rendering [12]
- pciBIRD API for head and hand tracking [2]
- FLTK for the GUI and the OpenGL interface [7]
- OpenAL for the 3D audio [3]

### 4.1 VTK

The Visualization ToolKit (VTK) is an open source, freely-available, cross-platform C++ library that supports a wide variety of advanced visualization and volume processing algorithms. We use VTK to read and process volumetric data obtained by Magnetic Resonance Imaging (MRI) or Computer Tomography (CT) scanners, applying a marching cube algorithm to generate isosurfaces from certain sections of the volume with homogeneous density. For example in the “Haptic Visible Human” application, we extract the skin and the bone surface from MRI data. The isosurfaces generated with VTK are polygonal meshes that can be quickly rendered and manipulated in real time.

### 4.2 Coin

Coin is an open source high-level 3D graphics library that uses scene-graph data structures to render real-time graphics. It is an Open Inventor implementation, ideal to develop scientific and engineering visualization applications. Coin is free under the GPL for Free Software development, and requires an annual fee per developer for commercial use.



**Figure 16:** Periodontal Training Simulator

VTK also has graphics rendering capabilities. However, Coin is optimized for real-time polygonal rendering and provides more sophisticated interaction nodes. Therefore, we use Coin for rendering the isosurfaces generated with VTK.

The *ImmersiveTouch*<sup>TM</sup> API provides a camera node that computes the correct viewer-centered perspective projection on the virtual projection plane. This new camera is an extension of the native Open Inventor *SoPerspectiveCamera* node. It properly renders both left and right views according to the position and orientation of the user's head given by the tracking system. The specialized camera node is based on the work done by [17] for the *CAVELib*<sup>TM</sup>.

### 4.3 GHOST

The General Haptic Open Software Toolkit (GHOST) is a cross-platform library commercialized by SensAble Technologies. Even though we would rather use open source libraries for haptics rendering as well, there is none currently available. Recently, SensAble Tech. has released a new haptics library, which, although it is called "Open Haptics", is not open source. Unlike GHOST, Open Haptics does not provide VRML support. In our case, VRML is fundamental to transfer 3D models from VTK to the haptic library and Coin. So we rely on GHOST to interact with the PHANTOM<sup>®</sup> device and to compute the collision detection.

Using GHOST, we can define different haptic materials to each 3D object in the virtual scene specifying four coefficients: stiffness, viscosity, static and dynamic frictions (Figure 17). Once the collision between the tip of the probe held by the user and any virtual object is detected, GHOST computes the reaction forces the haptic device needs to apply to give the user the illusion of "touching" the object. Both Coin and GHOST must be synchronized with the head tracking system so the user can see and touch exactly at the same 3D point, no matter from which viewpoint he/she is looking.

### 4.4 pciBIRD API

Ascension Technologies Corp. provides the freely-available pciBIRD API to control the data acquisition from the tracking system. The pciBIRD is Windows and Plug & Play compatible. It gives us the positions and orientations of the user's head and hand. As we stated above, head tracking is fundamental to obtain perfect graphics/haptics collocation; hand tracking provides a more natural interaction with the 3D virtual models.

In order to minimize the noise caused by the CRT, we set the measurement rate to 85 Hz, which is different from the monitor horizontal refresh rate (100 Hz).

### 4.5 FLTK

Since the monitor image is horizontally flipped, the image reflected on the mirror can be read normally. Therefore, we can use any library we want to create the graphical user interface (GUI). We use the Fast Light ToolKit (FLTK) because it is a small and modular freely-available cross-platform C++ GUI that supports 3D graphics via OpenGL<sup>®</sup> and its built-in GLUT emulation. With FLTK we can incorporate all of the usual widgets to develop our applications (menus, buttons, sliders, etc.). It even has a Fast Light User-Interface Designer (FLUID), which is useful to easily draw the user-interface and to define functions, classes and variables as needed. FLUID creates C++ source and header files that can be included in our application. The control panel shown in Figure 17 is an example of the GUI implemented on FLTK.

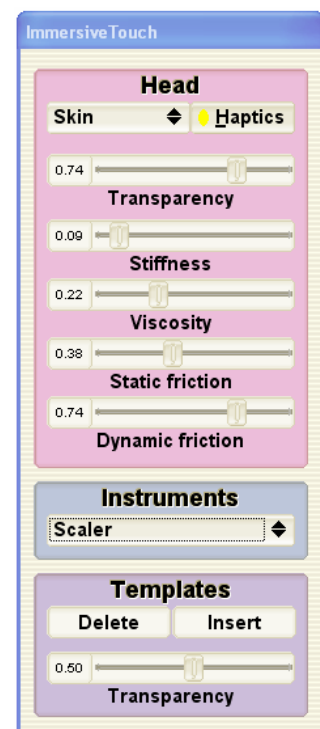


Figure 17: Control panel



## 4.6 OpenAL

Open Audio Library (OpenAL) is a freely-available cross-platform 3D audio API that serves as a software interface to audio hardware. OpenAL is a means to generate arrangements of sound sources around a listener in a virtual 3D environment. It handles sound-source directivity and distance-related attenuation and Doppler effects, as well as special effects such as reflection, obstruction, transmission, and reverberation.

In *ImmersiveTouch*<sup>™</sup>, even though the OpenAL works fine with a pair of regular loudspeakers, the half-silvered mirror presents certain barrier for high frequency sounds. Therefore, the most realistic results are obtained when wearing headphones. Since we track the user's head position and orientation, we can render listener-centered 3D audio in a similar way we render stereoscopic viewer-centered perspective projection. This allows us to achieve a more comprehensive graphics/haptics/audio collocation.

## 5 Calibration of *ImmersiveTouch*<sup>™</sup>

*ImmersiveTouch*<sup>™</sup> includes many elements that must be calibrated to provide a correct graphics/haptics collocation. The virtual projection plane and the haptic workspace need to be expressed in terms of the tracking coordinate system, whose origin is located in the transmitter. It is done as follows:

Since we can measure the size of the physical screen, we know the dimensions of the virtual projection plane. From the fundamental design constraint, we also know that the projection plane orientation is 45°. Then, we should measure the distance from the center of the projection plane to the transmitter. Since the projection plane is virtual, a physical measurement is very cumbersome to perform. Instead, we take advantage of the tracking system. We measure it simply holding a tracking sensor (receiver) at the projection plane until it is superimposed with a point displayed at the center of the projection plane. Then we read the position given by the tracking system.

The measurement of the offset from the center of the haptic workspace to the transmitter is done interactively moving the haptic stylus and leaving the graphics rendering fixed until the haptic stylus coincides with the virtual probe. This is done only at the center of the projection plane. However, for a better calibration, we should repeat this procedure at many points in the haptic workspace to create a correction table as done by [6]. This will be done in future research.

The interocular distance, and the offset from the head sensor to the center of the head, as well as the offset from the hand sensor to the center of the SpaceGrips<sup>®</sup> are specified manually, similar to the CAVE<sup>®</sup> and ImmersaDesk<sup>®</sup> applications.

## 6 Conclusions and future research

We have designed and built a high-performance haptic augmented reality system which compares favorably with currently available alternative systems, presenting significant advantages, including more accurate graphics/haptics/audio collocation, higher display resolution, higher pixel density, better visual acuity, and more comfortable workspace. Also we have developed an API, by integrating many open source and/or freely-available libraries, that efficiently performs volume processing, graphics rendering, haptics rendering, head and hand tracking, graphical user interface, and 3D audio.

Implementing a more sophisticated calibration procedure to improve graphics/haptic collocation thorough the haptic work volume remains at the core of future work. A virtual globe might also be incorporated to the system to provide more sophisticated 3D data manipulation by performing gesture recognition. The point-based collision detection provided by GHOST is extremely fast but not realistic enough for many haptic applications. We will evaluate newly available object-to-object collision detection libraries and their real-time performance in the context of haptic applications.

## Acknowledgements

This research was supported by NSF grant DMI 9988136, NIST ATP cooperative agreement 70NANB1H3014, the Link Foundation Fellowship of the Institute for Simulation and Training at the University of Central Florida, the Department of Mechanical and Industrial Engineering, and the Department of Periodontics at the University of Illinois at Chicago (UIC).

Additional support was obtained by the virtual reality and advanced networking research, collaborations, and outreach programs at the Electronic Visualization Laboratory (EVL) at the UIC, which were made possible by major funding from NSF awards EIA-9802090, EIA-0115809, ANI-9980480, ANI-0229642, ANI-9730202, ANI-0123399, ANI-0129527 and EAR-0218918, as well as the NSF Information Technology Research (ITR) cooperative agreement (ANI-0225642) to the University of California San Diego (UCSD) for "The OptIPuter" and the NSF Partnerships for Advanced Computational Infrastructure (PACI) cooperative agreement (ACI-9619019) to the National Computational Science Alliance. EVL also receives funding from the US Department of Energy (DOE) ASCI VIEWS program. In addition, EVL receives funding from the State of Illinois, Microsoft Research, General Motors Research, and Pacific Interface on behalf of NTT Optical Network Systems Laboratory in Japan.

## References

1. Accommodation/Convergence conflict, [http://vresources.jump-gate.com/articles/vre\\_articles/stereo/sterean2.html](http://vresources.jump-gate.com/articles/vre_articles/stereo/sterean2.html)
2. Ascension Technologies Corp., pciBIRD API, <http://www.ascension-tech.com/products/pcbirebird.php>
3. Creative, OpenAL, <http://www.openal.org/>
4. Cruz-Neira, C., Sandin, D., DeFanti, T., Kenyon, R., and Hart, J.C., The CAVE: Audio Visual Experience Automatic Virtual Environment, *Communications of the ACM*, Vol. 35, No. 6, 1992, pp. 65-72.
5. Czernuszenko, M., Pape, D., Sandin, D., DeFanti, T., Dawe, G., Brown, M., The ImmersaDesk and Infinity Wall Projection-Based Virtual Reality Displays. *Computer Graphics*, 1997.
6. Czernuszenko, M., Sandin D., DeFanti, T., Line of Sight Method for Tracker Calibration in Projection-Based VR Systems, *Proceedings of 2nd International Immersive Projection Technology Workshop*, Ames, Iowa, 1998.
7. Fast Light ToolKit, <http://www.fltk.org/>
8. Johnson, A., Sandin, D., Dawe, G., DeFanti, T., Pape, D., Qiu, Z., Thongrong, S., Plepys, D., Developing the PARIS: Using the CAVE to Prototype a New VR Display, *Proceedings of IPT 2000: Immersive Projection Technology Workshop*, Ames, IA., 2000.
9. Kitware Inc., Visualization ToolKit 4.5, <http://www.vtk.org/>
10. LaserAid, SpaceGrips, <http://www.spacegrips.com/spacegrips.htm>
11. Reachin Display, <http://www.reachin.se/>
12. SensAble Technologies, GHOST 4.0, <http://www.sensable.com/>
13. SenseGraphics 3D-MIW, <http://www.sensegraphics.se/3DMIW.pdf>
14. Stereographics theory, <http://astronomy.swin.edu.au/~pbourke/stereographics/vpac/theory.html>
15. Systems in Motion, Coin 2.3, <http://www.coin3d.org/>
16. The Visible Human Project, [http://www.nlm.nih.gov/research/visible/visible\\_human.html](http://www.nlm.nih.gov/research/visible/visible_human.html)
17. Pape D., Sandin, D., Transparently supporting a wide range of VR and stereoscopic display devices, *Proceedings of SPIE, Stereoscopic Displays and Virtual Reality Systems VI (The Engineering Reality of Virtual Reality 1999)*, vol 3639, San Jose, CA
18. VRCO, CAVELib™ and Trackd®, <http://www.vrco.com/>
19. Zwern, A., How to select the right head-mounted display, *Meckler's VR World*, 1995, <http://www.genreality.com/howtochoose.html>