


Calibrating See-Through Head-Mounted Displays

TR95-034
January 1996



Anantha Kancherla
Mark Singer
Jannick Rolland

Department of Computer Science
University of North Carolina at Chapel Hill
Chapel Hill, NC 27599-3175



UNC is an Equal Opportunity/Affirmative Action Institution.

Abstract

A basic issue in the development of see-through display systems is obtaining precise calibration for the system in relation to the user. Unlike immersive *head-mounted displays* (HMDs) where only virtual objects are displayed, even small discrepancies in the location of the virtual images can be dis-orienting. Moreover, any discrepancies in global scaling or distortions of virtual objects impairs the registration of these objects with their real world counterparts. To properly calibrate displays it is necessary to take a number of precise measurements for each individual user. This process is tedious and time consuming and must be repeated not only for every individual user, but each time the HMD is donned by the same user. Taking into account other research into this topic we developed a simple series of steps for calibration that can be performed by end users in a relatively short period of time.

Contents

1	Introduction	1
2	What is Calibration ?	5
2.1	The Computational Model	5
2.2	Errors involved	7
2.3	Previous Work	10
3	Calibrating Off-the-Shelf Head-Mounted Displays	11
3.1	The I-glasses	11
3.1.1	A note on IPD adjustment	11
3.1.2	Tracking	12
3.2	Calibration Procedure	13
3.2.1	Experiment Setup	13
3.2.2	Wearing the HMD	17
3.2.3	Collecting Data	17
4	Results and Discussion	21
4.1	Evaluating the Calibration	21
4.2	Discussion	21
4.2.1	Errors	21
4.2.2	Possible improvements	22
A	Tracking Technologies	23
A.1	Mechanical Tracking	23
A.2	Magnetic Tracking	23

List of Figures

1.1	A video-see-through HMD.	2
1.2	Artist's rendering of what the teaching tool will eventually look like.	2
1.3	A bench optical see-through HMD built with custom components.	3
1.4	An off-the-shelf HMD sold by Virtual-IO, the I-glasses.	3
2.1	The set of transforms comprising the computational model.	6
2.2	The overscan.	8
2.3	The viewing frustum.	9
3.1	The i_glasses attached with a plexi-glass plate.	12
3.2	The setup for the calibration procedure.	14
3.3	The crate and the Reference CS.	15
3.4	The arrangement of nails on top of the crate.	16
3.5	The nail-sighting experiment, front view.	18
3.6	The step to determine frustum parameters.	19

1. Introduction

See-through HMDs (STHMD) are used in Augmented Reality (AR) applications to *augment* the real-world view with three-dimensional graphical objects. Using the principles of Virtual Reality (VR) two views of the graphical objects, one for each eye, are generated. These objects are also referred to as synthetic objects. The STHMD superimposes the synthetic objects over the real-world view to present the STHMD wearer with a combined view. The wearer thus perceives these synthetic objects to be located at specified locations in the real-world. A position-orientation sensor is mounted on the STHMD to report the current position and orientation of the STHMD in the real-world to the computer. The computer updates the two views of the synthetic scene using the position and orientation information.

The principle criterion regarding the effectiveness of an AR system is whether the synthetic objects appear where they are supposed to. Most often, the synthetic objects are used to represent some unknown/unseen aspect of some real object in the scene, eg. the teaching tool for radiologic positioning described in [KRWB95] and [WRK95], where the synthetic objects are the elbow-joint bones that are to be shown superimposed over the real joint to simulate x-ray vision (see Figure 1.2. These kinds of applications require that the synthetic objects be superimposed or *registered* very precisely over the real objects at all time.

Currently there exist two kinds of STHMDs: 1) Optical see-through HMD (OSTHMD), where the wearer obtains the real-world view through optics, and semi-transparent mirrors are used to combine the synthetic scene with the real scene. Figures 1.3 and 1.4 show two such OSTHMDs. 2) Video see-through HMD (VSTHMD), where the wearer obtains the real-world view from two video cameras mounted on the VSTHMD as shown in Figure 1.1, and the two scenes are computationally combined. Each method has its own merits/demerits.

One aspect that is common to both systems is that, both require the geometry of the STHMD to be communicated to the computer via an appropriate computational model. In the case of the OSTHMD this model includes the optics, the position of the eyes with respect to the virtual images of the screen and the position of the sensor on the STHMD. The VSTHMD requires the video cameras' placement and optics parameters in addition to the knowledge of the sensor position and eye position with respect to the HMD. In our work we have adopted the OSTHMD for the visualization purposes, the reasons for the choice are enumerated in [RHF94] and [KRWB95].

Two approaches could be taken towards building the computational model:

1. One could build an OSTHMD with specific components that are either known or measured, and build this knowledge in the computational model. Such a system built at University of North Carolina at Chapel Hill is shown in Figure 1.3.
2. Obtain an off-the-shelf OSTHMD and measure the required parameters to build the model (Figure 1.4 shows a typical off-the-shelf HMD built by Virtual-IO [vir] that has been used in this work).

Aligning the system components is termed as "calibration" for the first approach. Accurate measurement of the desired parameters is termed as calibration for the second approach. In this paper, we report the calibration techniques devised to cope with the latter approach.




Figure 1.1: A video-see-through HMD.




Figure 1.2: Artist's rendering of what the teaching tool will eventually look like.




Figure 1.3: A bench optical see-through HMD built with custom components.



Figure 1.4: An off-the-shelf HMD sold by Virtual-IO, the I-glasses.

In the rest of this report: the theoretical basis of calibration is presented first. The various errors involved in making measurements and the computational model of the displays are described. This is followed by a description of the calibration procedure developed for the off-the-shelf-HMDs. Results of the calibration procedure are presented. Remaining problems are discussed and ways to further improve the calibration are suggested. The appendix contains technical details of the procedure such as the various tracking technologies used (or considered) in course of the research. Throughout this report, henceforth, HMD, STHMD and OSTHMD are used interchangeably.

2. What is Calibration ?

In this section the theoretical basis of calibration is discussed. As mentioned in the introduction, calibration is defined as a series of steps leading to an accurate determination of the various parameters required by the software to display the synthetic objects with minimum registration error. These steps must take into account the various errors that are made in measuring these parameters. The first section discusses what parameters need to be measured accurately. These parameters are essentially those needed in the computational model of the STHMD that has been adopted in our software. This is followed by a section addressing the various errors that are involved in the measuring process. Finally, a brief overview of previous work in this area is given.

2.1 The Computational Model

A computational model which matches the physical geometry of the HMD is used to generate the stereographics. This model includes the optics used to image the displays, the placement of the virtual-images with respect to the eyes and the placement of the tracker-receiver with respect to the HMD. To accurately project the virtual objects at specified depths, the model needs to match the physical geometry of the HMD as much as possible. Not meeting this criterion results in the virtual objects being displayed at locations other than those specified. Hence registration of the virtual objects with the real objects fails.

The model adopted by our software is adapted from the one discussed in [RR92]. A major difference being that, owing to the nature of the chosen STHMD, the optical distortion [RH93] has been omitted in our model. Figure 2.1 shows some of the various transforms that constitute this model. The epsilons in the figure indicate the error in the measurement of the corresponding transform. The objective of the calibration procedure is to minimize these errors so that the model mimics the physical HMD faithfully. Given below is, first, a list of all the coordinate-systems (CS) that exist in our computational model. This is followed by the actual parameters or transforms that need to be measured. Together they define the computational model. A brief description of each transform is then provided.

Notation

The notation followed in defining the transforms is that introduced in the text [FvDFH91]. A transform $T_{A \leftarrow B}$ means: a transformation from space B to space A.

Sensor CS It is the native CS of the receiver mounted on the HMD. It is also chosen as the representative CS of the HMD.

Eye CS This is the CS attached to the eye. The origin is chosen to be the center of rotation for the eye. This choice was made primarily for the ease of locating the point. A better choice is discussed in section 4.2. The orientation is: positive X to the right, positive Y upwards, positive Z into the eye. It should be noted that the notion of *upwards* is perceptual, and the calibration method should take into account this fact. There is one Eye CS for each eye.




Figure 2.1: The set of transforms comprising the computational model.

Head CS The head CS is chosen as the representative CS of the wearer's head. It is conveniently chosen to be at the point which lies exactly between the two eye points, with the same orientation as the eye CSs.

Image CS It is chosen as the representative CS of the virtual images of the display formed by the optics. The center of the virtual-image is chosen to be the origin. The positive X and Y axes correspond to the right and down directions when seen through the optics. The positive Z axis is away from the wearer.

The transforms that constitute the computational model of the HMD:

$T_{H \leftarrow R}$ The transform from the Receiver CS to the Head CS. Physically, this transform measures *how* the HMD is worn by the user.

IPD The interpupillary distance, i.e., the distance between the STHMD wearer's eyes. It is the distance between the origins of the two Eye CSs.

$T_{E \leftarrow I}$ This is equivalent to determining the viewing frustum. Includes values such as, the extents of the display in the eye-space, the near (same as the projection plane) and far clipping planes (see Figure 2.3).

$T_{I \leftarrow I'}$ (**Overscan**) This transform is used to determine the parameters of the view-port call made in the display program. Figure 2.2 illustrates the meaning of this transform.

2.2 Errors involved

A typical AR system involves the use of a STHMD to which a position/orientation sensor is rigidly attached. This sensor is used to report the position and orientation of the STHMD with respect to a fixed coordinate system – usually the transmitter coordinate system. The position of the sensor on the STHMD is also captured by the computational model. The sensor reports are used by the model to compute the correct view of the synthetic scene for each eye.

Since tracker reports are used by the model, the errors in the tracker report can cause misregistration. More specifically, there are three kinds of errors in the tracker report that can cause misregistration, namely:

systematic error This error can be measured and corrected for [LS95] [MDD⁺95]. This error is a static registration error since it is visible even when there is no motion involved.

random error The sensor reports varying values when placed at a single location. It can be considered to be “noise” in the tracker report. This error is usually corrected by taking a large number of sensor readings at every point and then averaging them. This error is also known as the jitter. Holloway quotes figures for some prominent tracking systems in his thesis [Hol95a]. This error can be considered to be a static error too since it appears even when there is no motion involved.

lag This error occurs because there is a finite amount of time involved in computing a tracker report. This error is not noticed unless a movement occurs and hence is termed a dynamic error in registration. It is the most destructive error in the AR environments. Sophisticated techniques like prediction are used to address this problem [Azu95].

The total registration error is affected not only by the kind of tracking used, but also by the accuracy of the computational model.

In our work, we concentrated on errors caused by the inaccurate modeling of the system, i.e., the errors caused by the computational model. In this work, by “calibration” we mean determining the computational model accurately. The tracker errors have been enumerated here because of their importance in designing a proper calibration technique.

From this, a few observations about the desired properties of the calibration technique could be made. Namely:




Figure 2.2: The overscan.

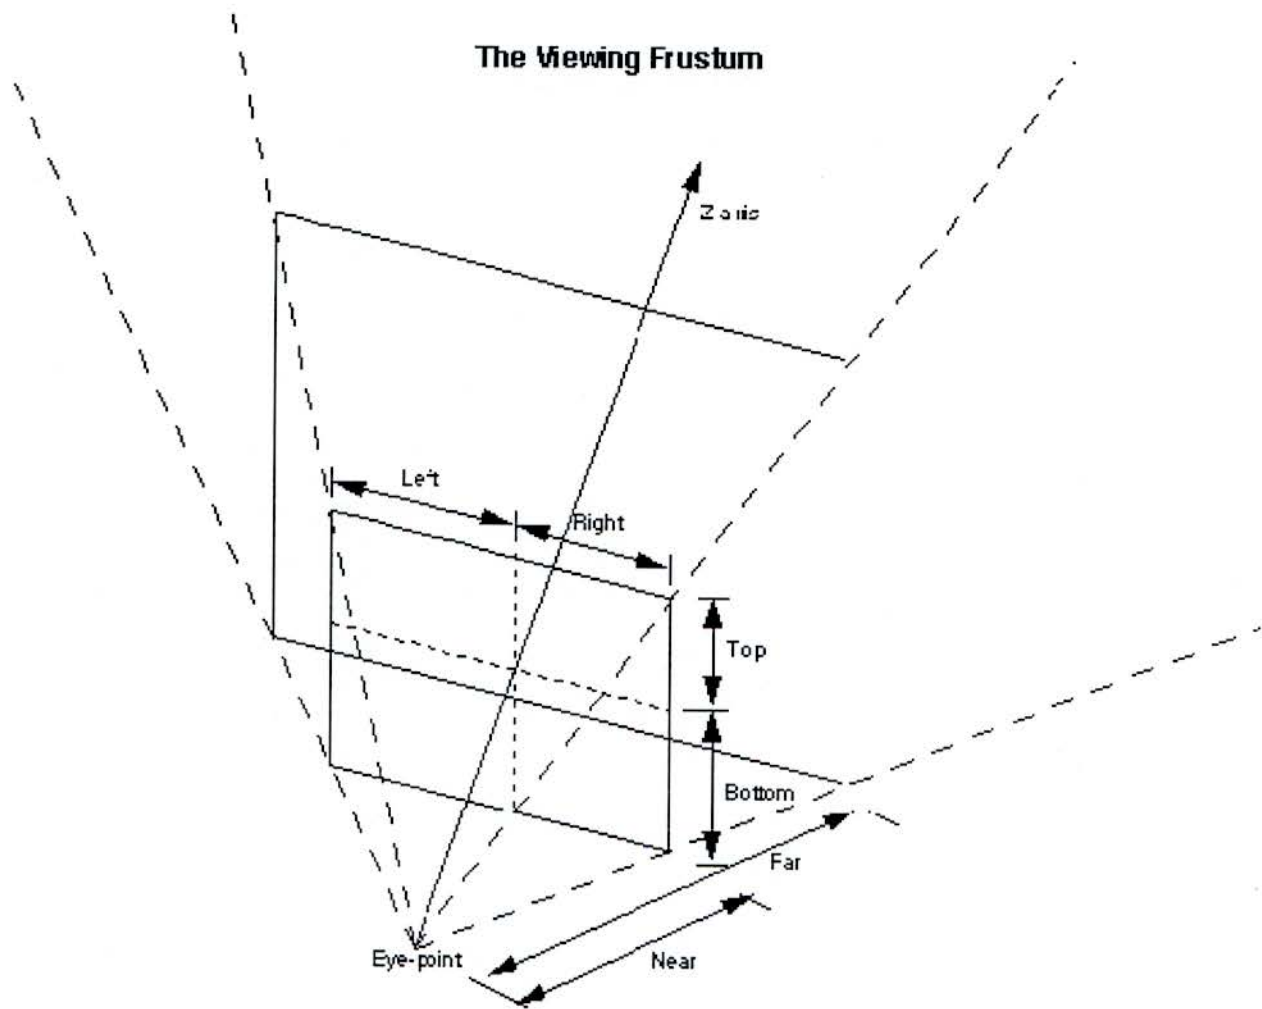


Figure 2.3: The viewing frustum.

- One should always average out the readings to minimize noise.
- A more philosophical note about the systematic error in the tracker: Systematic error of the tracker is a kind of nebulous term since one has to define “error with respect to some gold standard”. What is more important is that the software be aware how numerical values (in some units, say meters) translate into physical distances when displayed (or read). One measuring¹ system will never be the same as another and hence if we use more than one measuring method in the course of calibration, then it will cause more errors. Ideally, during calibration one should use the same measuring mechanism as used during tracking (which is measuring at run-time). If that is not possible, whatever tracking device is going to be used in the application needs to be calibrated with respect to the one used during the calibration of the computational model.
- Involve no motion, or else the dynamic errors will creep in.

2.3 Previous Work

The calibration technique reported in this paper is mainly based on the works reported by Holloway and Azuma. A brief discussion of these works is presented here. For more details, the reader should refer to [Hol95a] and [AB94].

Both Azuma and Holloway used a custom STHMD that was built in the University of North Carolina. The HMD was quite heavy, hence the problems associated with a light STHMD, described in section 4.2.1 do not impede calibration. Another difference is that the virtual image location could be controlled by a knob provided unlike the STHMD used by us. Finally, the size and weight contributed to significant sag in the STHMD structure when it is worn on the head, leading to some change in the parameters measured by the calibration.

The computational model assumed by Holloway is the same as the one assumed in this calibration. However, the computational model assumed by Azuma assumes an on-center projection, and hence he measures only the FOV and aspect ratio instead of measuring the complete viewing frustum. The approaches to computing the other parameters differed only in small implementation details. For instance the boresight procedure is used to compute the Eye-Tracker transform in all three cases. However, in our work and Azuma’s, this procedure is used to directly provide the transform, where as Holloway uses this step to correct the transform that he measures using another technique.

Our calibration technique is more similar to Holloway’s than to Azuma’s. The major difference being:

- Since the virtual images formed by STHMD used by Holloway are located roughly an arm’s length away, he uses a parallax technique to determine the Image-Tracker transform. The STHMD used in our experiment formed the images eleven feet away and hence such a procedure is non-applicable.
- Holloway’s calibration accounts for the sag in the STHMD structure. The compact nature of our HMD spared us of this requirement.
- Holloway’s technique addresses the problem of optical distortion, whereas our’s didn’t require it.

By far the most significant difference was that we designed a technique that could be executed by a naive user quickly as compared to ones mentioned in this section.

¹Ultimately, tracking is one form of making measurements

3. Calibrating Off-the-Shelf Head-Mounted Displays

In this section, the calibration of a typical off-the-shelf HMD is described. The *i_glasses* developed by Virtual i-O [vir] were chosen for use in our application. In this section, calibration refers to measuring various parameters of the STHMD, and adjusting the computational model accordingly to match these values in order to obtain registration of the real and synthetic images. Calibration must take into account measurements made by the tracker and the position of the tracker-receiver on the STHMD, because of the reasons iterated in section 2. Finally, of importance is the fact that this calibration needs to be performed every time the STHMD is donned, and hence it should be a simple series of steps that can be performed unassisted by a reasonably naive user.

This section begins with a brief discussion about the *i_glasses*. An overview of the entire calibration is given next. This is followed by an in-depth description of each step in the calibration.

3.1 The I-glasses

In the experiment conducted we have implemented the calibration procedure using the *i_glasses* manufactured by Virtual i-O. All of the applications and tests referred in this report were performed utilizing this headset. Since we are developing software that is intended to be used by a broad audience, we opted for a consumer headset which provided optical see-through. Virtual i-O's product, *i_glasses*, is a small, lightweight headset that allows some ambient light to filter through its lenses, letting the user perceive his actual surroundings. The device also features a fixed IPD. This particular HMD has a number of worthwhile features such as low cost and light weight. There are also a number of possible problems associated with this device. Though it does have optical see-through, it is designed as a safety feature to prevent simulation sickness. Therefore the system favors the synthetic images. Real world images are somewhat muddy with a purple hue, and require additional room lighting. The commercial tracker that can be purchased for the device provides only head orientation (yaw, pitch and roll). But since applications for see-through devices require six *degrees-of-freedom* (DOF) tracking, it was modified to have a 6-DOF receiver mounted on it. When making the mount we planned for a simple design that could mount several different tracker types.

3.1.1 A note on IPD adjustment

In both see-through and non see-through devices, determination of IPD for the user and the viewing lenses is critical to producing realistic 3-D images. Some commercial devices offer adjustments for the viewing lens IPD, but many, including *i_glasses*, do not. The lack of IPD adjustment can be compensated for by appropriately adjusting the computational model – by having an off-center projection (refer to [NDW] or [Wat93] for a good treatment on projections) instead of an on-center projection (where the eyes are centered on the optics). However, this is not the best solution since vignetting¹, causes a reduction in the user's

¹some light rays are excluded at the entrance pupil




Figure 3.1: The i_glasses attached with a plexi-glass plate.

effective *field of view* (FOV). However those off the shelf systems which have an IPD adjustment often do not have a vernier scale to determine its exact value. With most of these devices, only a rough adjustment can be made and hence calibration is still required for optimized accuracy of displayed depth.

3.1.2 Tracking

We had available, two tracking technologies:

- Magnetic, more specifically, the Extended range Flock of Birds manufactured by Ascension Technologies [asc].
- Mechanical, more specifically, the Faro-arm [far].

Each tracker had its own advantages and disadvantages, specifically:

Magnetic • Multiple receivers can be tracked simultaneously.

- They are light-weight.
- However, they suffer from significant noise.
- Usually report systematic errors, because the magnetic field is distorted by metal in the environment.

Mechanical • Extremely accurate.

- Low noise.
- Very cumbersome, cannot be mounted easily, especially on a light STHMD like the i_glasses . They cause the STHMD to move about the head, causing more misregistration.

- Cannot track multiple objects simultaneously.

We chose the mechanical tracker because for us the promise of accuracy far outweighed any other inconvenience presented by the use of such a heavy piece of equipment. To overcome the problem of the weight of the tracker yanking the HMD off the wearer's head, we adopted another approach. A small (65mm square) plexi-glass plate with minute holes on a regular grid pattern (14mm apart) was mounted on the HMD in place of a regular tracker (Figure 3.1 illustrates this clearly). This plate is considered to be the receiver, (the receiver coordinate-system is as shown in that figure), the $T_{T \leftarrow R}$ transform² could be obtained by digitizing the plate with the stylus tipped mechanical tracker (see Figure 3.2) in the following manner:

1. Digitize points C, X and Y on the plate (see Figure 3.1). Point C is the center of the plate (the origin of the receiver CS).
2. These points define the x-y plane of the receiver CS. The vectors $v_{X \leftarrow C}$ and $v_{Y \leftarrow C}$ are determined by simply subtracting the reading for point C from those of points X and Y respectively. The normal to this plane can then be determined by:

$$n = v_{X \leftarrow C} \times v_{Y \leftarrow C} \quad (3.1)$$

3. The x, y and z axes can be determined normalizing the vectors $v_{X \leftarrow C}$, $v_{Y \leftarrow C}$ and n respectively.
4. These vectors are initially expressed in the Tracker CS, hence the rotation matrix relating the receiver and tracker space is given by:

$$R_{T \leftarrow R} = [x \mid y \mid z] \quad (3.2)$$

The translation is simply given by the reading for point C. Putting these together we can obtain $T_{T \leftarrow R}$.

3.2 Calibration Procedure

The aim of this process is to obtain the following parameters (refer to section 2.1). Specifically:

- The position of the receiver with respect to the head. This is given by the transforms $T_{H \leftarrow R}$.
- The viewing frustum, i.e. the extents of the display in the eye-space and the near (also the projection plane) and far clipping planes. See Figure 2.3.
- The overscan of the displays in the STHMD (see Figure 2.2).

In this section, data collection is first described. How the desired parameters were computed from the collected data is described next.

3.2.1 Experiment Setup

We were able to use a fairly simple apparatus as shown in Figure 3.2. The pieces of test equipment are positioned on a table top to allow the test subject to be seated while having everything at eye level. A chin rest is used to stabilize the users head and minimize extraneous head movements. At an arms length away from the user, on the table, a wooden crate is placed³. The crate has a large grid on the front and two pairs of nails arranged as in Figure 3.4. The top-left corner of the grid on the crate facing the user is a reference coordinate system for the Image CSs (the x-axis going right, y-axis down and the z-axis away from the user). This coordinate system is henceforth referred as the Reference CS (see Figure 3.3). Note: The nails have to be driven into the wooden frame of the crate so that equal lengths of nail stick vertically out of the crate. Each nail-head had a small hole from a metal stylus to keep the Faro point-tip from slipping during digitization.

²The tracker reports the $T_{T \leftarrow R}$ transform which is the position and orientation of the receiver in the tracker CS.

³This distance is desired in the application being developed [KRWB95]




Figure 3.2: The setup for the calibration procedure.




Figure 3.3: The crate and the Reference CS.




Figure 3.4: The arrangement of nails on top of the crate.

3.2.2 Wearing the HMD

The user wears the HMD while in the chin-rest and sights different features on the crate. Though the chin rest has to move as the user is asked to locate different features, the rest of the apparatus is fastened to the table to prevent slippage. There is a small strategy involved in wearing the HMD. The goal is to wear the HMD in such a way that:

- there is no vertical disparity.
- the full extents of the displays are visible.
- the distortion at the fringes of the display is minimized.

With trial and error it has been discovered that the proper wearing of the HMD is greatly aided if a grid pattern is displayed on both the displays. Ideally, the wearer should fuse these images and perceive a flat grid in the plane in front of them – this plane is the one where the virtual image is formed by the optics. So the user adjusts the HMD on his/her head until all the above three conditions are satisfied. It has been observed that there are very few positions of the HMD when this happens and hence it takes a bit of an effort to get it right the first time. The grid lines also help in detecting distortion when the HMD is improperly mounted.

3.2.3 Collecting Data

The various data collection steps are described in the order they were performed. Some of the steps need to be repeated every time the user wears the HMD and the others need to be done once for a particular setup. Yet some others need to be performed once for the HMD. The calibration programs were run on an SGI Onyx.

Determining overscan

The user wears the HMD and a vertical and a horizontal line are displayed on each of the displays. The user first moves the lines with the help of mouse-button presses till they appear to intersect at the top-left hand corner, then with another button press, the values of the point of intersection of the two lines in the screen space are output. The process is repeated for the bottom-right hand corner, and then for the two corners in the other display. The values obtained are the area of the frame-buffer that are visible on the displays, i.e. the overscan parameters. This step is performed once for a given HMD and need not be repeated again.

Digitizing the crate

First, the transform between the Tracker CS and the Reference CS ($T_{T \leftarrow Ref}$) is computed. The points C, X, and Y as shown in Figure 3.2 are digitized. The transform $T_{T \leftarrow Ref}$ is computed in the way described in section 3.1.2.

Nail-sighting to compute $T_{R \leftarrow H}$

Since the transform $T_{R \leftarrow H}$ varies each time the HMD is worn, this procedure has to be performed every time the user dons it. There are two parts to this procedure. First, the four nails on the crate are digitized in the order shown in Figure 3.4. The point of intersection of the two lines(E) $Nail_1-Nail_2$ and $Nail_3-Nail_4$ computed⁴ determines where the eye is placed in the second part of this procedure described next.

In the second part of this step, a modified a version of the bore-sight procedures of Holloway [Hol95a] and Azuma [AB94] was used to determine the eye's position with respect to the HMD. While wearing the

⁴Since these are lines in 3D it is not guaranteed that they intersect. Hence the nearest point of approach of the two lines is determined on the line $Nail_1-Nail_2$.

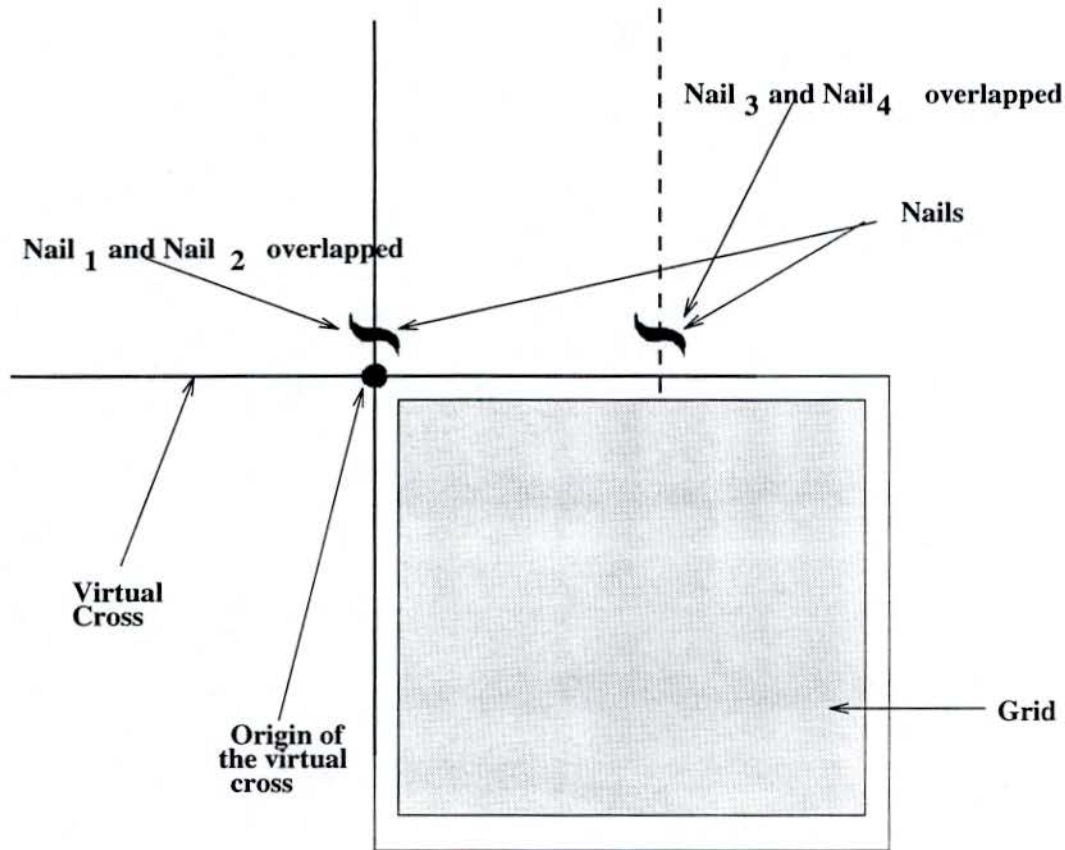


Figure 3.5: The nail-sighting experiment, front view.

HMD and using the chin-rest, the user is asked to sight down the first set of nails with one eye. After the user finds the point at which the front nail completely obscures the second, he is then asked to find the point where the second pair of nails line up. When both pairs are aligned simultaneously, the head tracker location (T_1) is recorded⁵. During this procedure a faint cross is displayed on the HMD to aid the user to keep his/her head vertical so that the orientation of the Eye CS matches that of the Reference CS on the crate⁶. The procedure is repeated for the second eye. Again the tracker location (T_2) is recorded when the desired sighting is achieved.

From these two tracker measurements the receiver-to-head transform ($T_{R \leftarrow H}$) is derived in the following way:

- The Eye position in Tracker CS is simply the point E . This is the origin of the Eye CS in Tracker CS.
- The orientation of the Eye CS in the Tracker CS is the same as the rotation part of $T_{T \leftarrow Ref}$.
- Putting these two together gives $T_{T \leftarrow E}$.
- $T_{R \leftarrow E_1}$ is obtained by composing T_1^{-1} by $T_{T \leftarrow E}$.
- $T_{R \leftarrow E_2}$ is obtained similarly.
- $T_{R \leftarrow H}$ is obtained by combining the rotation part of either $T_{R \leftarrow E_1}$ or $T_{R \leftarrow E_2}$, and the mean of the translation part of the same two transforms.

⁵As described in section 3.1.2

⁶Except that the Z axis is flipped.




Figure 3.6: The step to determine frustum parameters.

- As a sanity check, the transform $T_{E_1 \leftarrow E_2}$ can be computed, and verified that the rotation part of this transform tends to identity and that the magnitude of the translation part tends to the IPD of the user.

Computation of the frustum parameters

The aim of this procedure is to obtain the frustum parameters, i.e. the left, right, top, bottom, near and far as illustrated in Figure 2.3, with respect to the Eye CS. The idea is to project the image of the screen on a plane parallel to the image-plane and digitize the top-left and bottom-right corners of the projection and express them in the Eye CS. An orthogonal grid is displayed on the HMD to assist the user to keep the image-plane parallel to the projection plane – the face of the crate painted by the grid. While this grid is displayed, the user sights the real grid painted on the crate with one eye. If the image-plane is not parallel to the projection-plane, the user perceives a keystone distortion of the graphical-grid with respect to the real-grid. A part of the sighting task for the user is to superimpose the top-left corner of the image with the top-left corner of the real-grid (the origin of the Reference CS). Once the desired sighting is achieved, the tracker measurement (T_1) is taken and the point where the bottom-right corner of the image projects on the crate is digitized (P_1). The procedure is repeated for the other eye (T_2 and P_2).

From these readings, the frustum parameters for Eye1 are computed in the following way:

- $T_{T \leftarrow E_1}$ is computed by composing T_1 and $T_{R \leftarrow E_1}$
- The top-left corner of the image projection is transformed to Eye CS by composing translation part of $T_{T \leftarrow Ref}$ with $T_{T \leftarrow E_1}^{-1}$. Thus the left, top and near are the x, y and z coordinates of this transformed point respectively.
- The bottom-right corner of the image projection is transformed to Eye CS by composing P_1 with $T_{T \leftarrow E_1}^{-1}$. Thus right and bottom are the x and y coordinates respectively of this transformed point.

4. Results and Discussion

The results of the calibration are first presented. The problems encountered in the procedure and the error that resulted are discussed next. Finally, a couple of suggestions for future improvement of the procedure are given.

4.1 Evaluating the Calibration

The calibration was evaluated for two virtues: time and accuracy. The initial run of the calibration (when all the steps have to be executed) took about ten minutes to execute. Subsequent calibrations took five minutes.

The accuracy of the calibration was evaluated by superimposing a virtual crate over the real crate using the parameters obtained in the calibration process plugged into the computational model. The user viewed the crate from various viewpoints. At each viewpoint, the head was stabilized by using a chin-rest. The best calibration achieved afforded an accuracy in the order of 5mm.

4.2 Discussion

4.2.1 Errors

The error in the registration is the cumulative error of all the steps in the calibration procedure. Sanity checks were introduced at each step to give us an idea of the amount of error introduced in that step. If the error tended to be large, e.g. In the nail-sighting step, the rotation part of $T_{E_1 \leftarrow E_2}$ is not identity, that step was repeated.

Even though the mechanical tracker being used had extremely low noise [LS95], the digitizing procedure used to obtain the tracker readings (refer to 3.1.2) introduced some noise. To alleviate this problem, multiple readings were taken and averaged. However, more readings meant more digitization. And because digitization was done by hand, the process introduced more error. Thus we settled for a trade-off of 4 readings.

The following were some other observations we made about the error in the calibration procedure:

1. The nail-sighting step seemed to introduce the maximum error. This was because of the extreme sensitivity to the orientation of the head when taking measurements.
2. The tracking used in our experiment was quite crude – digitizing with hand. This usually resulted in slight head movement. Also, it was nearly impossible to apply the stylus in a consistent manner while digitizing.
3. There was some amount of parallax error in the frustum calculation step.

A more detailed, theoretical and numerical treatment of the error can be found in [Hol95a] and [Hol95b].

Other problems

1. Clearly, the Faro-arm was a very heavy system with a limited range. Some positions were quite difficult to digitize without disturbing the user.
2. The digitizing process also tended to move the HMD with respect to the head. Since the $T_{R \leftarrow H}$ transform is assumed to be a constant, this movement is a serious problem.
3. In spite of using the chin-rest, the user's head tended to move slightly. The roll and pitch of the head were especially apparent.

4.2.2 Possible improvements

The calibration technique is very promising, however the factors described above are quite distracting. At this point we are able to offer the following suggestions for improvement:

- A better tracking system like the Optotrak [opt] optical tracker can be used to eliminate the various ills that plague the current system. Not only is it more accurate¹ the entire digitization process can be eliminated by attaching an LED at every point that needs to be digitized. This will also reduce the calibration time to a minimum.
- The right choice for the origin of Eye CS is the entrance-pupil of the eye-lens. However, there are a few practical problems with this choice given the current state of technology. Firstly, detecting that point is more difficult. Secondly, unlike the center of rotation of the eye, this point changes when the eyeball rotates for a different convergence. Thus imposing the necessity of eye-tracking to correctly obtain the current position of the Eye CS with respect to the Head CS. A simple improvement would be to compute the center of rotation of the eye as before. Then correct this value to obtain the entrance-pupil location when the eye is looking straight ahead. The correction value can be obtained from any optics literature. The IPD will then be the distance between the entrance-pupils of the two eyes.
- A more stable system like a bench-HMD or a boom system should be used to constrain the user while performing the calibration.
- Not all HMDs are devoid of optical distortion, hence the future techniques should incorporate a procedure to quickly detect and correct optical distortion.
- One practical improvement could be to pre-calibrate the HMD for different user parameters (the IPD and the depth of the eye from the user's forehead). At the time of use, the user's IPD and eye-depth could be determined and the system configured for the closest matching setting (very much like selecting the right shoe size in the shoe-store). The advantage is that the user is not involved in the calibration steps.

¹Positional accuracy and resolution are 0.1mm and 0.01mm respectively. Angular(Orientation) accuracy and resolution are 0.07 degrees and 0.007 degrees respectively.

A. Tracking Technologies

A few salient points about the tracking technologies we encountered are mentioned in these sections. For a more detailed discussion on the pro and cons of various tracking technologies see [HL94] and [MAB92].

A.1 Mechanical Tracking

Mechanical tracking is by far the most accurate method commonly in use. High precision trackers can report minute changes in position and orientation. The drawback is with greater accuracy often comes greater size and weight. Metal arms and gears making up the tracker usually make for a bulky machine. Another drawback is that the tracker must remain attached to some stationary object. The user is limited in movement by the trackers's reach.

In trying to use the Faro arm mechanical tracker with the *i-glasses* we found both the weight and limited range to be nearly insurmountable. In order for a see-through device to function properly, it must remain in place on the head rigidly and not stray from its position even slightly. We found that the weight of the arm would have a tendency to push the HMD down over the eyes of the user. Additionally, the arm has stop points for each of its segments to keep them from going to far in any one direction. If, while moving the head with the arm attached to the HMD, the arm hit a stop point, the HMD would be pulled in the opposite direction. In either of these cases, the head mount would have to be adjusted back to the normal position and then re-calibrated to compensate for the new location. Though the accuracy is valuable, there seems to be no practical way to implement mechanical head tracking for such a lightweight device.

A.2 Magnetic Tracking

Magnetic trackers completely do away with the problems associated with mechanical tracking. They are lightweight, and though they must be cabled to a source, they are not inseparably tied to the recording unit. The only major question is that of accuracy. Tests which were run on a number of *Flock of Birds* units consistently showed errors in position and orientation. These errors were usually on the order of several millimeters (and centimeters in some cases), well outside of the acceptable range for realistic super-imposition of real and graphic objects.

Bibliography

- [AB94] Ronald Azuma and Gary Bishop. Improving static and dynamic registration in an optical see-through hmd. In *Computer Graphics (SIGGRAPH '94 Proceedings)*, pages 197–204, July 1994.
- [asc] Ascension technology co. P.O.Box 527 Burlington, VT 05402, USA.
- [Azu95] Ronald Azuma. *Predictive Tracking for Augmented Reality*. PhD thesis, University of North Carolina at Chapel Hill, 1995.
- [far] Faro technologies inc. 125 Technology Park, Lake Mary, Florida 32746, USA.
- [FvDFH91] James D. Foley, Andries van Dam, Steven K. Feiner, and John F. Hughes. *Computer Graphics Principles and Practice*. Addison-Wesley, 1991.
- [HL94] Richard L. Holloway and Anselmo Lastra. Virtual environments: A survey of the technology. Technical Report TR93-039, University of North Carolina at Chapel Hill, 1994.
- [Hol95a] Richard L. Holloway. *An analysis of Registration Errors in See-Through Head-Mounted Display System for Craniofacial Surgery Planning*. PhD thesis, University of North Carolina at Chapel Hill, 1995.
- [Hol95b] Richard L. Holloway. Registration error analysis for augmented reality. Paper submitted for SIGGRAPH '95, 1995.
- [KRWB95] Anantha Kancherla, Jannick P. Rolland, Donna L. Wright, and Grigore Burdea. A novel virtual reality tool for teaching dynamic 3d anatomy. In *Proc. CVRMed'95*, pages 163–169, 1995.
- [LS95] Mark Livingston and Andrei State. Improved registration for augmented reality systems via magnetic tracker calibration. Technical Report TR95-037, University of North Carolina at Chapel Hill, 1995.
- [MAB92] Kenneth Meyer, Hugh L. Applewhite, and Frank A. Biocca. A survey of position trackers. *Presence*, 1:173–189, spring 1992.
- [MDD⁺95] M.Ghazisaedy, D.Adamczyk, D.J.Sandin, R.V.Kenyon, and T.A.DeFanti. Ultrasonic calibration of a magnetic tracker in a virtual reality space. In *Proc. of VRAIS'95*, pages 179–188, March 1995.
- [NDW] Jackie Neider, Tom Davis, and Mason Woo. *OpenGL Programming Guide*.
- [opt] Optotrak animation engineering. 870 W Center Street, Orem, Utah 84057, USA.
- [RH93] Jannick Rolland and Terry Hopkins. A method of computational correction for optical distortion in head-mounted displays. Technical Report TR93-045, University of North Carolina at Chapel Hill, 1993.
- [RHF94] Jannick Rolland, Richard L. Holloway, and Henry Fuchs. A comparison of optical and video see-thru head-mounted displays. In *Proc. SPIE 2351*, pages 293–307, 1994.

- [RR92] Warren Robinett and Jannick Rolland. A computational model for the stereoscopic optics of a head-mounted display. *Presence: Teleoperators and Virtual Environments*, 1:45-62, winter 1992.
- [vir] Virtual i-o. 1000 Lenora St., Suit 600, Seattle, WA, 98121, USA.
- [Wat93] Alan Watt. *3D Computer Graphics*. Addison-Wesley, 1993.
- [WRK95] Donna L. Wright, Jannick P. Rolland, and Anantha Kancherla. Using virtual reality to teach radiographic positioning. *Radiologic Technology*, 66(4):167-172, 1995.