

The Hedgehog: A Novel Optical Tracking Method for Spatially Immersive Displays

A. Vorozcovs*

A. Hogue†

W. Stuerzlinger‡

Department of Computer Science and Engineering and the Centre for Vision Research
York University, Toronto, Ontario, Canada.

Abstract

Existing commercial technologies do not adequately meet the tracking requirements for fully-enclosed VR displays. We present the Hedgehog which overcomes several limitations imposed by existing sensors and tracking technology. The tracking system robustly and reliably estimates the 6DOF pose of the device with high accuracy and a reasonable update rate. The system is composed of several cameras viewing the display walls and an arrangement of laser diodes secured to the user. The light emitted from the lasers projects onto the display walls and the 2D centroids of the projections are tracked to estimate the 6DOF pose of the device. The system is able to handle ambiguous laser projection configurations, static and dynamic occlusions of the lasers, and incorporates an intelligent laser activation control mechanism that determines which lasers are most likely to improve the pose estimate. The Hedgehog is also capable of performing autocalibration of the necessary camera parameters through the use of the SCAAT algorithm. A preliminary evaluation of the accuracy reveals the system to have an angular resolution of 0.01 degrees RMS and a position resolution of 0.2 mm RMS.

1 Introduction

Spatially immersive displays (SIDs) have recently become popular for scientific visualization, training, entertainment and VR research. These types of displays provide a way to fully immerse the user into the virtual world allowing them to be more accurate and productive at the given task. In order to create a compelling visual world, the VR display must produce correct visual cues (perspective, parallax, stereo). If any of these cues are incorrect then the user may not feel ‘present’ [21] in the virtual environment. This can have catastrophic effects if the user is being trained for a real-world task such as helicopter landing.

These visual cues are highly dependent upon the position and orientation (pose) of the user’s head, hence it is absolutely necessary to have a system that can accurately and robustly track head pose. If the head pose is estimated poorly, the user is more likely to experience discomfort (headaches, nausea, disorientation; symptoms collectively known as cybersickness [20]).

In non-fully-enclosed displays, it is trivial to use existing commercial head tracking systems since the tracking equipment can be positioned in such a way that it does not interfere with the user’s view of the scene (i.e. behind the user). Such an approach is not possible in a fully-enclosed spatially immersive display.

Tracking a user within fully-enclosed SID’s such as COSMOS [29, 8], HyPi-6 [13, 22], PDC VR-CUBE [3], C6 [15], ALICE [25], and IVY [24] is a more complex task than typical single wall displays. The user is confined in a fully-enclosed volume and there is no acceptable location for visible tracking equipment as it interferes with the desired immersive experience, effectively *removing* the user from the virtual world.

1.1 Commercial Tracking Systems

The most popular tracking technology used in existing fully-enclosed displays is electromagnetic technology, such as Ascension Technologies’ MotionStar Wireless® tracking system. It is known that electromagnetic tracking systems behave poorly in the presence of metallic objects [18, 19]; accuracy and signal strength degrades with distance from the base emitter. This precludes the use of this type of technology in IVY (the display we intend to use, Figure 1) which contains enough metallic material in its frame to render electromagnetic trackers useless.

Commercial optical tracking systems, such as the HiBall®[28] from UNC and available through 3rdTech™ work on principles that cannot be easily adapted for use in the aforementioned displays. The HiBall is an inside-out optical tracking system; photodiodes are placed on the user and infrared LEDs are positioned on the ceiling of the tracking volume. The photodiodes track the LEDs and a single LED measurement is used to update the estimated pose of the device. At the heart of the HiBall is the SCAAT [26] tracking method implemented in the extended Kalman filter [27] framework. The HiBall requires that tiles with LEDs be placed in the tracking volume, thus it is not possible to employ the HiBall in fully-enclosed SIDs.

Acoustic tracking systems, such as the IS-900 [14] motion tracking system from Intersense, rely on time-of-flight of ultrasonic chirps to determine the pose of the user. Unfortunately the ultrasonic chirps are absorbed by diffuse fabric projection screen material precluding this technology from working within a fully-enclosed display with fabric screens. Both the sensors and emitters would be required to be placed within the display volume.

Although various commercial tracking technologies have been tried in spatially immersive displays, no single tracking device exists that constitutes a completely effective tracking technology for fully-enclosed CAVE™-like [2] displays.

1.2 Research Laboratory Tracking Systems

Very few non-commercial tracking systems, typically developed in university research labs, exist. Of these, the most notable were designed with Augmented Reality applications in mind [30, 5, 23, 7]. These technologies are not directly applicable to fully-enclosed spatially immersive displays due to the inherent sensor limitations.

*e-mail:av@cs.yorku.ca

†e-mail:hogue@cs.yorku.ca

‡e-mail:wolfgang@cs.yorku.ca

Tracking systems developed for AR are mainly categorized into being either marker-based or LED-based. Marker-based systems require that visual markers, or fiducials, be placed within the tracking volume. Cameras attached to the user visually track the markers and use either the known geometry of multiple markers to estimate the pose, or the pose is inferred by the marker itself. These types of systems are known to decrease their accuracy when viewing the markers at oblique angles. A similar problem is present with LED-based systems. The LEDs emit light in a non-uniform manner, thus when viewed from oblique angles they become more difficult to detect. These tracking methods are not applicable for fully-enclosed displays since the markers or LEDs would need to be placed between the user and the display (you cannot place markers within the display).

The most effective method to date is a hybrid optical-inertial tracking system previously developed by one of the co-authors in this paper and is discussed in [11, 12]. The inertial system is comprised of a commercially available InertiaCube2TM from Intersense which provides the system with fast relative motion information. The secondary “outside-in” optical system utilizes a set of cameras outside of the display viewing the screens. An arrangement of laser diodes in a known geometry is attached to the user and the projections of the lasers are tracked within each image. The absolute pose is directly computed from the laser projections using geometric constraints and the system is able to provide very accurate pose at a low update rate of 15fps due to limitations of the imaging system.

This optical approach requires that all laser spots be visible by the cameras in each frame. Since there are places where the lasers cannot be seen by the cameras (i.e. corners of the display), this is the major limitation of the system. When all lasers are not available the pose cannot be computed and the head-tracker must be re-initialized.

1.3 Motivation

The lack of commercial tracking technology that exists for fully-enclosed displays stimulates the development of a novel tracking system. Several issues must be addressed when developing a new tracking technology [4]. The system must

- work in a fully-enclosed display environment
- be easy to construct from low cost components
- track reliably in the presence of noise
- be easy to calibrate
- track robustly in the presence of occlusions
- have a high accuracy, low latency, and reliable update rate

1.4 Contributions

We introduce the **Hedgehog** which overcomes several limitations imposed by the previously discussed optical approach. A novel tracking hardware device has been developed comprised of multiple (more than the minimum of 3) laser diodes. As before, the lasers are tracked visually on each of the display walls. In order to disambiguate and reliably label the laser projections, the diodes’ activation state is changed periodically. This is synchronized with the image capture enabling the algorithm to determine exactly which laser produced the currently tracked projection in the image.




Figure 1: IVY: The Immersive Visual environment at York. IVY is shown here with the rear (entry) wall removed in order to show the structure of the device more clearly.

By examining the maximum angular velocity and acceleration that a user can actually induce by moving their head, we are able to maximize the number of lasers that can be reliably and robustly identified. The lasers are intelligently controlled on a per frame basis increasing the update rate of the tracker to an appropriate level for a VR system. By fully utilizing the SCAAT framework, the camera parameters required for the laser projection can also be optimized online increasing the accuracy of the tracking system.

Our system successfully addresses the above concerns for tracking systems since the Hedgehog

- works in a fully-enclosed display environment
- is relatively easy to construct and contains inexpensive materials
- performs automatic online calibration of projection parameters
- is able to track reliably in the presence of noise
- has an intelligent laser activation control that is robust in the presence of static and dynamic occlusions
- provides a high accuracy, low latency, and reasonable update rate for a VR tracking system

2 Tracking Approach

The basic approach for tracking within a fully-enclosed display (see Figure 2 for an illustration and [11] for more information) is to use the projective surfaces outside of the view of the user to estimate and track their head pose within the environment. A fixed arrangement of low power laser diodes is attached to a helmet worn by the user. Using lasers allows us to extend the line-of-sight to the user from outside of the display making it possible to *indirectly* observe the user’s motion. Cameras are positioned behind each screen of the display viewing the projection surface. The projections of the laser beams are visually tracked as they strike the projective surfaces. We exploit the known geometry of the lasers, and using the 2D measurements of each laser projection we apply each measurement as a constraint that unique determines the pose of the device. By securing the device to




Figure 2: Optical Tracking Approach. The user wears many low power laser diodes whose projections on each screen surface are tracked via cameras outside of the display. The head pose is determined from these projections alone.

a helmet worn by the user, the device becomes a practical head-tracker. Intelligently controlling the activation state (on or off) of each individual laser diode allows the tracking system to easily identify which diode is being tracked in each frame simplifying the computational requirements of the algorithm.

Various configurations of laser diodes could be used to localize the user. In the previous approach discussed in [11], a simple arrangement of four laser diodes was used. This provides a simple technique for computing the pose of the device, however it requires that all laser projections be visible by the camera system at each update. Thus, the system cannot handle occlusions (when lasers shine into corners or are occluded by an object within the physical space, e.g. the arm of a user).

The Hedgehog has been designed to overcome the limitations of the four laser geometry by placing redundant lasers on the user. This allows a larger range of motion to be tracked by utilizing all of the lasers in an intelligent manner. Since the Hedgehog contains more than four lasers, it can function in more extreme conditions within the display than the previous method. The redundancy of the lasers also allows the system to achieve a higher accuracy. Each laser contributes to the overall pose of the device through the use of a state estimation scheme.

2.1 State estimation and intelligent laser control

State estimation is typically performed through the usage of a recursive filtering scheme such as the Kalman filter [17]. Many variations on the Kalman filter have been proposed to accommodate different motion models. The most widely used variation being the extended Kalman filter [27, 9], the iterated Kalman filter [1], and the unscented Kalman filter [16]. These types of filtering schemes normally assume that all measurements that come from a particular sensor are taken simultaneously. This is never the case with today's electronics and computers. Digitization takes some amount of time, thus if several sensors must be polled or digitized for a particular measurement then each sensor will be measured sequentially. This is not normally a problem if the measured motion is small between successive measurements or if the tracked object is stationary, however it becomes an issue when the tracked object is moving significantly. In this scenario, the measurements might not accurately reflect the motion of the object. The SCAAT [26] method allows partial measurements to update the state estimate. This improves

the robustness of the system and increases its temporal responsiveness. SCAAT can be applied to various types of state estimation filtering schemes, however in this paper we chose to employ the SCAAT algorithm in the form of an extended Kalman filter. We decided to use SCAAT for the Hedgehog simply because we have a definite mapping of a source and sensor pair, i.e. single laser diode and a single camera constitute a sensing pair.

Using a control mechanism for selecting the visible laser diodes allows the system to identify and account for occlusions in the display volume. Both static occlusions (corners of the display, or improper camera placement) and dynamic occlusions (user raising their arm in front of several laser diodes) can be accounted for. Using a predictive state estimation filter allows the algorithm to predict where each laser will shine in the next frame. By modeling the areas which are not visible in each camera sensor, we estimate the optimal set of lasers that should be activated to view the maximal number of lasers. Since the lasers are labelled correctly, we are able to determine the visibility of each laser as it is activated. If the measurement is updated, then we give more credibility to that laser diode and less credibility to diodes that have not been seen in many measurement updates. This allows the system to detect areas which contain occlusions. Thus, we can activate only the lasers that should not shine into the occlusion barriers. This novel idea has not yet been employed in an optical tracking system.

3 The SCAAT Filter

The single-constraint-at-a-time (SCAAT [26]) approach to state and parameter estimation is employed to enable measurements from a *locally unobservable* system to estimate a *globally observable* system. Applying this method to an extended Kalman filter [27, 9], it is possible to update the state of a system using a single sensor/source measurement. The main idea of the SCAAT filter is that each measurement contains some information about the state of the system. By predicting what the observation should be, the error between the actual observation and the predicted measurement can be computed. The error is used to compute the Kalman gain which updates the predicted state and covariance of the filter.

Another key notion used in the SCAAT method is the idea of separating the sensor parameters from the source parameters. The isolation of the measurements from each type of sensor allows an elegant method to perform autocalibration online.

The SCAAT filter as presented in [26] begins by estimating a state with a position vector (\vec{p}), velocity vector (\vec{v}), incremental orientation angles (ϕ, θ, γ), and their derivatives ($\dot{\phi}, \dot{\theta}, \dot{\gamma}$) and is denoted as

$$\hat{x} = [\vec{p} \quad \vec{v} \quad \phi \quad \theta \quad \gamma \quad \dot{\phi} \quad \dot{\theta} \quad \dot{\gamma}] \quad (1)$$

An external quaternion is maintained after each update step using the incremental orientation angles denoted by

$$\hat{q} = [q_w \quad q_x \quad q_y \quad q_z] \quad (2)$$

The filter first performs a *time update* step using a linear process model ($A_{\delta t}$) to predict the state (\hat{x}^-) using the current δt between measurements. The covariance matrix is also predicted (P^-) with the process model and the process noise estimate ($Q_{\delta t}$) is applied. The time update equations

are

$$\hat{x}^- = A_{\delta t} \hat{x}_{t-\delta t} \quad (3)$$

$$P^- = A_{\delta t} P_{t-\delta t} A_{\delta t}^T + Q_{\delta t} \quad (4)$$

The measurement update equations are where the SCAAT method differs from conventional Kalman filtering. The measurement function ($h_{\sigma}(\bullet)$) has as its input the predicted state (\hat{x}^-), the source parameters (\vec{b}), and the sensor parameters (\vec{c}). This function computes what the measurement should be, i.e. projecting the state onto the source/sensor pair. In essence it predicts a noise-free response of each sensor and source pair given the systems' current state. For each measurement function, a corresponding measurement Jacobian (H) is also computed. The measurement update equations are thus:

$$\hat{z} = h_{\sigma}(\hat{x}^-, \vec{b}_t, \vec{c}_t) \quad (5)$$

$$H = H_{\sigma}(\hat{x}^-, \vec{b}_t, \vec{c}_t) \quad (6)$$

where $H_{\sigma}(\bullet)$ is defined as

$$H_{\sigma}(\hat{x}^-, \vec{b}_t, \vec{c}_t)_{i,j} \equiv \frac{\partial}{\partial \hat{x}_j} h_{\sigma}(\hat{x}_t^-, \vec{b}_t, \vec{c}_t)_i \quad (7)$$

The next step in the measurement update is to compute the Kalman gain

$$K = P^- H^T (H P^- H^T + R_{\sigma,t})^{-1} \quad (8)$$

The residual error is computed between the predicted observation (\vec{z}) and the actual measurement ($\vec{z}_{\sigma,t}$) from the sensor

$$\Delta \vec{z} = \vec{z}_{\sigma,t} - \vec{z} \quad (9)$$

The state (\hat{x}) and error covariance (P_t) are then corrected

$$\hat{x}_t = \hat{x}^- + K \Delta \vec{z} \quad (10)$$

$$P_t = (I - KH) P^- \quad (11)$$

and the external orientation quaternion is updated with the incremental orientation angles estimated in the state

$$\Delta \hat{q} = \text{quaternion}(\hat{x}[\phi], \hat{x}[\theta], \hat{x}[\gamma]) \quad (12)$$

$$\hat{q} = \hat{q} \otimes \Delta \hat{q} \quad (13)$$

The final step in the measurement update is to zero the incremental orientation in the state for the next iteration

$$\hat{x}[\phi] = \hat{x}[\theta] = \hat{x}[\gamma] = 0 \quad (14)$$

3.1 The Hedgehog SCAAT Filter

The implementation of the SCAAT method for the Hedgehog tracking system follows directly from above. Our process model $A_{\delta t}$ simply integrates the position and orientation using a standard constant velocity model

$$\vec{p}_{t+1} = \vec{p}_t + \vec{v}_t \delta t \quad (15)$$

$$\phi_{t+1} = \phi_t + \dot{\phi}_t \delta t \quad (16)$$

$$\theta_{t+1} = \theta_t + \dot{\theta}_t \delta t \quad (17)$$

$$\gamma_{t+1} = \gamma_t + \dot{\gamma}_t \delta t \quad (18)$$

We must specify the source/sensor pair for the Hedgehog measurement function. In the simplest scenario, the sensor

is defined as the camera that views the display walls. Our algorithm requires that the measurement be in the coordinate frame of each wall, i.e. in real world units and relative to one corner of the screen surface. Thus, our sensor is actually a *virtual* camera that provides measurements in 2D with real-world coordinates and is attached to each wall. The source is defined as the laser diode which emitted the beam of light. The sensor parameters \vec{c} are twofold

1. a 2D translation, $t_{wall} = [t_x \ t_y]^T$
2. an affine 2D transform, $P = \begin{bmatrix} a & b \\ c & d \end{bmatrix}$

The laser diodes are parameterized by a unit direction vector, $\vec{L}_i = [x \ y \ z]^T$ with the origin of the laser frame being the currently estimated position of the Hedgehog device.

Our measurement function applies the predicted position and orientation to the laser diode vector and intersects it with the plane of the display wall. This produces a 2D point $\vec{z} = [x \ y]^T$ in the wall coordinate frame. The measurement Jacobian is computed numerically by perturbing the state elements by a small ϵ and producing multiple predicted observations. Given the observation predicted from the unperturbed state the residual is computed and the measurement Jacobian is populated with the differences.

3.2 Autocalibration

Prior to operation, a homography per camera is computed which is a projective planar transformation that transforms image coordinates into the wall coordinate frame. This offline calibration is performed manually by measuring laser points and recording its subpixel image coordinates. By measuring at least 4 points, a homography transformation can be computed [10] transforming the 2D image coordinates into a 2D point that lies on the wall plane. Since the wall plane itself is assumed to be rigid, the 3D rigid transformation to the display's origin can be physically measured.

The offline calibration provides a good initial estimate of the wall plane transformation. After applying the homography and rigid transformation, we are guaranteed to have 3D points that lie on a plane that is very close to the real wall plane. We estimate the offset and rotation of the plane with a 2D affine transformation and offset. Obviously, manual measurements are prone to error and the 3D points are not exactly at the appropriate locations on the real wall. The minimized transformation brings the image coordinates onto a *virtual* plane slightly offset and rotated from the real wall.

The calibration parameters are added to the point to acquire the measurement for the filter as

$$\begin{bmatrix} x'_{wall} \\ y'_{wall} \end{bmatrix} = \begin{bmatrix} a & b \\ c & d \end{bmatrix} \begin{bmatrix} x_{wall} \\ y_{wall} \end{bmatrix} + \begin{bmatrix} t_x \\ t_y \end{bmatrix} \quad (19)$$

SCAAT provides an elegant solution to automatic calibration of the camera transformation parameters. The estimated state vector is augmented with the parameters of the source and sensor for each pair of observations, effectively creating a Kalman filter per source/sensor pair. In our system, due to the high accuracy of the construction of the hardware device, the source parameters (i.e. the laser direction vectors) do not need to be autocalibrated. We employ the autocalibration only to compute the parameters of the wall, i.e. the wall normal and translation to the 3D display coordinate frame. Thus, only 6 parameters need to be autocalibrated.




Figure 3: The Hedgehog hardware. 17 laser diodes are arranged in a symmetrical hemispherical arrangement. As shown in (a) the IS-900 tracker is rigidly attached for experiment purposes.

We create an augmented state vector (\hat{x}), covariance matrix (\hat{P}), state transition matrix ($\hat{A}_{\delta t}$), and process noise matrix ($\hat{Q}_{\delta t}$) with the appropriate wall transform parameters

$$\hat{x} = \begin{bmatrix} \hat{x}_t^T & \hat{x}_{c,t}^T \end{bmatrix} \quad (20)$$

$$\hat{P} = \begin{bmatrix} P_t & 0 \\ 0 & P_{c,t} \end{bmatrix} \quad (21)$$

$$\hat{A}_{\delta t} = \begin{bmatrix} A_{\delta t} & 0 \\ 0 & I \end{bmatrix} \quad (22)$$

$$\hat{Q}_{\delta t} = \begin{bmatrix} Q_{\delta t} & 0 \\ 0 & Q_{c,\delta t} \end{bmatrix} \quad (23)$$

The rest of the algorithm proceeds normally using the augmented versions of the matrices, i.e. the Jacobian is numerically computed in the same manner but for the augmented measurement vector. At the end of the measurement update, we collapse the augmented state vector by extracting the elements not pertaining to the source/sensor parameters and save them in the main tracker state. The sensor parameters are applied to our sensor models and the filter is ready for the next iteration. Using this approach, the 6DOF rigid wall transform can be estimated online as the tracker receives measurements.

4 Hardware Implementational Details

Our implementation of the Hedgehog hardware consists of 17 laser diodes arranged in a symmetric hemispherical manner (see Figure 3). The device is 9cm in diameter and has a height of 6cm. Each laser diode is a visible red laser diode in the 645nm range and are individually controlled by a PIC microcontroller through a serial interface. The lasers are embedded in a custom housing machined from a single piece of Delrin® plastic. The lasers are placed at 45 degree angles from each other in all directions and the orientation of the diodes is accurate to within 0.1 degree. A tether is not required when the system is in operation through the use of a Bluetooth® serial cord replacement module from Free2move [6].

The display, IVY at York University, is equipped with 8 digital firewire cameras connected to a single linux PC through 3 firewire buses. Two cameras are required for both the floor and ceiling due to space restrictions in the design of the display. Each camera provides grayscale images at 640x480 resolution at 30fps with an exposure of 5ms. This

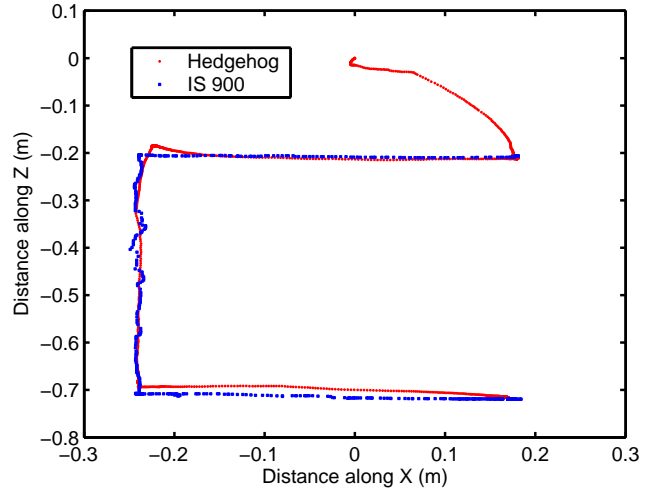


Figure 4: Positional Tracking Experiment 1

ensures that the laser spot is the only bright point in the image and is thus easily identified.

The laser identification and tracking is performed at a 15Hz update rate due to hardware limitations. The hardware limitations involve the camera framerate which at maximum is 30Hz, and the image processing of all eight images per frame on a single cpu PC. This hardware limitation can easily be overcome by utilizing either faster cameras or specialized laser spot detectors, such as position sensitive diodes, per wall which were unavailable at the time of development.

5 Experimental Results

In order to determine the accuracy of the Hedgehog tracking system, we performed several preliminary experiments for position and orientation. The IS-900 from Intersense was used in all experiments to validate our accuracy for position. We placed the IS-900 base transmitter on a rigid frame standing vertically at the entrance of IVY aligned with the X-axis of the display and one of its sensors was rigidly attached to the Hedgehog. The calibration offset was computed separately offline and the measurements were appropriately transformed into display coordinates (relative to the center of IVY). The IS-900 provides a way to determine the Hedgehog's performance relative to the state-of-the-art in commercial systems. In order to accommodate the IS-900, the rear wall of the display was left open approximately 70cm and the camera parameters for this wall were adjusted accordingly.

The following experiments were performed close to the center of the display volume which is used the most in a typical VR application. These experiments show that the Hedgehog is suitable for head-tracking applications, however further experiments will be needed to fully characterize the performance of the system.

5.1 Positional Experiments

The Hedgehog device was placed on a planar rectangular stage of known dimensions. We initialized the tracking system and recorded data from both the Hedgehog and the IS-900 while moving the system slowly around the rectangle

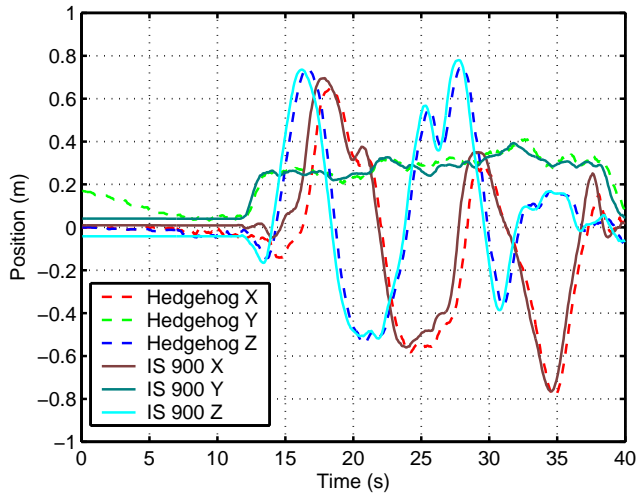


Figure 5: Positional Tracking Experiment 2

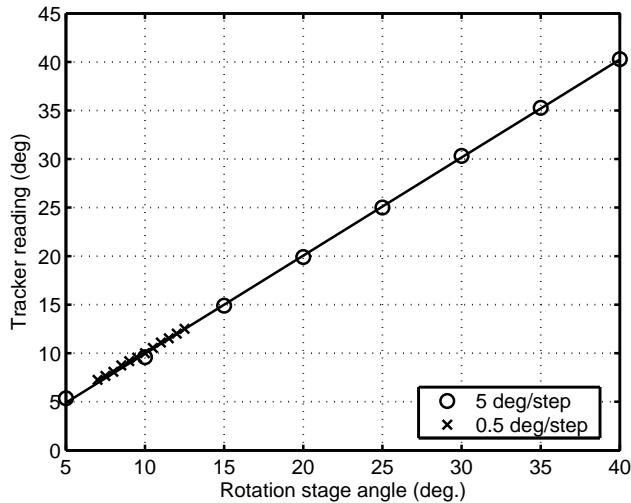


Figure 6: Angular Measurements of the Hedgehog

in a U-shape, see Figure 4.

A second experiment was performed where the Hedgehog was moved at 0.5 cm intervals and 1 cm intervals to determine the resolution, which was found to be 0.2 mm RMS.

A third graph (see Figure 5) shows how the Hedgehog performs dynamically versus the IS-900; we walked the connected devices throughout the entire tracking volume of the display and recorded all measurements taken within a 40 second duration.

5.2 Angular Experiments

The Hedgehog device was placed on a rotational stage with a Vernier scale allowing for angular resolution measurements accurate to 0.5 degrees. We initialized the tracking system and recorded data for every 0.5 degrees on the rotational stage in the first experiment and 5 degrees in a second experiment. Results of these experiments are shown in Figure 6 and demonstrate that the sensor is accurate to better than the accuracy of the rotational stage. Additional examination of the raw data was used to determine the static

angular resolution to be within 0.01 degrees RMS.

6 Summary

In this paper we introduced the Hedgehog optical tracking system. It is capable of tracking the user within a fully-enclosed spatially immersive display with high accuracy and reasonable update rate. Preliminary results reveal that the Hedgehog has a 0.01 degree RMS angular resolution and 0.2 mm RMS positional resolution.

Our approach to tracking in enclosed displays has many advantages over existing technologies:

1. The tracking method discussed is able to reliably track a user within a fully-enclosed display.
2. The hardware is constructed from inexpensive and readily available materials.
3. It performs automatic calibration of the wall transformations making it adaptable to different display types made from planar surfaces (polygonal, single-wall, etc.)
4. The Hedgehog is able to track the user reliably and accurately in the presence of noise.
5. The system is equipped with an intelligent control mechanism to change the state of the laser diodes allowing the system to determine which lasers should be used. This also enables the system to detect and track in the presence of occlusions.
6. The user is untethered and there is no interference with the metallic frame of the display.

Acknowledgements: We would like to acknowledge Matt Robinson for helping to create the software infrastructure, Jeff Laurence and Andriy Pavlovich for helping us build the hardware, Dr. Rob Allison and Dr. Michael Jenkin for guidance and financial support. The financial support of NSERC Canada, the Canadian Foundation for Innovation (CFI) is gratefully acknowledged.

References

- [1] B. M. Bell and F. W. Cathey. The iterated kalman filter update as a gauss-newton method. *IEEE Transactions on Automatic Control*, 38(2), 1993.
- [2] Carolina Cruz-Neira, D. Sandin, and T. DeFanti. Surround-screen projection based virtual reality: The design and implementation of the cave. In *Proc. SIGGRAPH '93*, pages 135–142, 1993.
- [3] Center for Parallel Computing. Primeur: Advancing European Technology Frontiers, World's first fully immersive VR-CUBE installed at PDC in Sweden, 1998.
- [4] E. Foxlin. *Motion Tracking Requirements and Technologies*. Lawrence Erlbaum Associates, 2002. ed. Stanney, K. M. Chapter 8.
- [5] E. Foxlin and L. Naimark. Vis-tracker: A wearable vision-inertial self-tracker. In *Proc. of IEEE Virtual Reality 2003 (VR2003)*, 2003. March 22-26, 2003, Los Angeles, CA.
- [6] Free2Move. Bluetooth serial port replacement module. <http://www.free2move.se>.
- [7] A. L. Fuhrmann. Personal communications.
- [8] K. Fujii, Y. Asano, N. Kubota, and H. Tanahashi. User interface device for the immersive 6-screens display "COSMOS". In *Proceedings 6th international conference on virtual systems and multimedia (VSMM'00)*, Gifu, Japan, Oct. 4-6 2000.

- [9] A. Gelb. *Applied Optimal Estimation*. MIT Press, Cambridge, MA., 1974.
- [10] R. Hartley and A. Zisserman. *Multiple View Geometry*. Cambridge University Press; ISBN:0521623049, 2000.
- [11] A. Hogue. MARVIN: a Mobile Automatic Realtime Visual and INertial tracking system. Master's thesis, York University, March 2003.
- [12] A. Hogue, M. R. Jenkin, and R. S. Allison. In *Proc. of IEEE 1st Canadian Conference on Computer and Robot Vision (CRV'2004)*, 2004. May 17-19, London, Ontario, Canada.
- [13] Fraunhofer Institute IAO. <http://vr.iao.fhg.de/6-Side-Cave/index.en.php>.
- [14] IS900 Tracking System. InterSense. <http://www.isense.com/products/prec/is900>.
- [15] Virtual Reality Applications Center Iowa State University. <http://www.vrac.iastate.edu/about/labs/c6>.
- [16] S. Julier and J. Uhlmann. A new extension of the kalman filter to nonlinear systems, 1997.
- [17] Emil Kalman, Rudolph. A new approach to linear filtering and prediction problems. *Transactions of the ASME-Journal of Basic Engineering*, 82(Series D):35-45, 1960.
- [18] V. Kindratenko. A survey of electromagnetic position tracker calibration techniques. In *Virtual Reality: Research, Development, and Applications*, 2000. vol.5, no.3, pp. 169-182.
- [19] V. Kindratenko. A comparison of the accuracy of an electromagnetic and hybrid ultrasound-inertia position tracking system. *Presence: Teleoperators and Virtual Environments*, 10(6):657-663, 2001.
- [20] J. LaViola. A discussion of cybersickness in virtual environments. In *SIGCHI Bulletin*, January 2000. vol.32 no.1, pp.47-56.
- [21] G. Mantovani and G. Riva. Real presence: How different ontologies generate different criteria for presence, telepresence, and virtual. In *Presence: Teleoperators And Virtual Environments*, 1999. Vol. 8, Pp. 538-548.
- [22] I. Rötzer. Fraunhofer Magazine, Synthetic worlds within six walls 2:2001.
- [23] M. Ribo, A. Pinz, and A. L. Fuhrmann. A new optical tracking system for virtual & augmented reality applications. Technical Report TR004-2001, VRVis, 2001.
- [24] M. Robinson, J. Laurence, J. Zacher, A. Hogue, R. Allison, L. R. Harris, M. Jenkin, and W. Stuerzlinger. IVY: The Immersive Visual environment at York. In *6th International Immersive Projection Technology Symposium, March 24-25, 2002, Orlando, FL*, 2002.
- [25] Beckman Institute University of Illinois, Integrated Systems Laboratory. A laboratory for immersive cognitive experiments. http://www.isl.uiuc.edu/Virtual%20Tour/Tour-Pages/meet_alice.htm.
- [26] G. Welch. *SCAAT: Incremental Tracking with Incomplete Information*. PhD thesis, Chapel Hill, NC, 1996.
- [27] G. Welch and G. Bishop. An introduction to the kalman filter. Technical Report TR95-041, University of North Carolina at Chapel Hill, 1995.
- [28] G. Welch, G. Bishop, L. Vicci, S. Brumback, K. Keller, and D. Colucci. High-performance wide-area optical tracking - the hiball tracking system. In *Presence: Teleoperators And Virtual Environments*, 2001. Vol 10, No. 1, Pp. 1-21.
- [29] T. Yamada, M. Hirose, and Y. Isda. Development of a complete immersive display: COSMOS. In *Proc. VSMM'98*, pages 522-527, 1998.
- [30] S. You, U. Neumann, and R. Azuma. Hybrid inertial and vision tracking for augmented reality registration. In *Proceedings of IEEE Virtual Reality*, pages 260-267, March 1999.