

A vision-based head tracking system for fully immersive displays

A. Hogue, M. Robinson, M. R. Jenkin, R. S. Allison

Department of Computer Science and
The Centre for Vision Research
York University
4700 Keele St., Toronto, Ontario, Canada, M3J 1P3.
(hogue,matt,jenkin,allison)@cs.yorku.ca

Abstract

Six-sided fully immersive projective displays present complex and novel problems for tracking systems. Existing tracking technologies typically require tracking equipment that is placed in locations or attached to the user in a way that is suitable for typical displays of five or less walls but which would interfere with the immersive experience within a fully enclosed display. This paper presents a novel vision-based tracking technology for fully-immersive projective displays. The technology relies on the operator wearing a set of laser diodes arranged in a specific configuration and then visually tracking the projection of these lasers on the external walls of the display outside of the user's view. This approach places minimal hardware on the user and no visible tracking equipment is placed within the immersive environment. This paper describes the basic visual tracking system including the hardware and software infrastructure.

Keywords: Virtual Reality, tracking, optical, lasers.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism – Virtual Reality I.4.8 [Image Processing and Computer Vision]: Scene Analysis – Tracking.

1. Introduction

Immersive displays have become a popular technology for scientific visualization, psychological research, teleoperation, task training/rehearsal, and entertainment. Advances in projection technology have facilitated the development of immersive displays ranging from large single wall projections (e.g. the PowerWall¹⁶), three-wall displays (e.g. the Immersion Square⁸), four-wall displays (e.g. the CAVE^{TM3}), five-wall displays (e.g. the CABIN¹⁷), and more recently six-sided displays (e.g. the Immersive Visual environment at York – IVY¹⁹). The technology is beginning to move out of the lab into the commercial arena, and vendors such as FakeSpaceTM and TAN Projektionstechnologie GmbH have begun to design and build immersive projective displays.

Currently, most existing projective immersive displays are designed with relatively small numbers of walls. Although there are various reasons why the non-fully-enclosed immersive environments have been constructed, the lack of “full enclosure” does simplify a number of design and construction details. As the number of walls increases many of the problems that can be solved “easily” in immersive projective displays with small numbers of walls become much more complex. Entry/egress, projector placement and most importantly, head tracking, become very complex issues. The limiting case of a fully enclosed (six-sided) environment is certainly the most challenging. The primary constraint imposed by fully enclosed environments is that the user is able to look in all directions.




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.*

To date, at least seven six-sided immersive environments have been developed:

1. COSMOS (VR Techno Centre in Gifu Japan)^{24, 6}.
2. HyPi-6 (Fraunhofer Institute for Industrial Engineering)^{10, 18}.
3. The PDC VR-CUBE (Centre for Parallel Computers in Stockholm)⁵.
4. The VR-CAVE (VR-CENTER NORD at Aalborg University)²⁰.
5. The C6 (Iowa State University)¹³.
6. ALICE (Beckman's Integrated Systems Laboratory in the University of Illinois)²¹.
7. IVY (York University)¹⁹

The underlying goal of these devices is to present a fully enclosed visual environment to the user. The IVY (Immersive Visual environment at York) six-sided fully immersive environment is shown in Figure 1. IVY is an 8' cube in which each surface is rear-projected. The rear wall (not shown in Figure 1) slides out of the way to provide entry/egress from the environment. Once the user enters IVY and the rear wall is closed, stereo imagery is projected onto each of the six sides. In order to project the correct images on the six sides, it is necessary to know the location and orientation of the user's head, otherwise the user is more likely to experience discomfort (headaches, nausea, disorientation; symptoms collectively known as cybersickness²²). In displays which contain less than six sides, it is possible to use 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). However, six-sided displays impose a unique constraint: the tracking equipment must be placed outside of the




Figure 2: *Basic approach. The user wears low power laser diodes whose projections are tracked via a camera outside of the display. The head position and orientation is computed from these projections.*

user's field of view to not negatively impact their sense of immersion or "presence." Due to the fully enclosed nature of these displays, the user is able to move around and look in any direction, thus the only reasonable place to position the tracking equipment is outside of the working volume.

Since the user is fully enclosed, state-of-the-art optical trackers and acoustical trackers which require a line-of-sight to the user are inappropriate. Currently, magnetic tracking is the technology of choice for fully immersive displays. The COSMOS, PDC VR-CUBE and the VR-CAVE all use the Polhemus FASTRAK[®] magnetic tracker¹¹ while the C6 and ALICE both employ Ascension Technology's MotionStar Wireless[®] magnetic tracking system². The Polhemus FASTRAK[®] requires the user to be tethered to the base station with long visible cables, while the MotionStar Wireless[®] system uses a magnetic field emitter outside of the working volume with an extended range of influence. Unfortunately, magnetic tracking systems have a number of disadvantages¹⁴. They are dependent on the local ambient electromagnetic environment and thus are subject to distortion and noise when used in close proximity to metallic objects or stray magnetic fields. A number of six-sided immersive displays are constructed out of wood to reduce this interference. The quality of the magnetic tracking measurements is a function of the magnetic signal strength. Thus, as the user moves further away from the magnetic field emitter the precision decreases. This implies inconsistent tracking throughout the working area and is illustrated in the work of Kindratenko¹⁵ where a comparison is given between a hybrid inertial-ultrasonic tracking system from Intersense¹² and a magnetic tracking system.

2. Basic Approach

In order to overcome the limitation of existing magnetic trackers, we have developed a novel "outside in"




Figure 3: User with helmet (a) and tracking device on rotational stage (b). Note that the lasers are mounted behind the user so that the laser beams strike walls outside of the user's view and cannot be seen.

vision-based tracking system (briefly described in ¹⁹) for tracking users within a fully enclosed projective immersive environment. The optical tracker utilizes commercial cameras and computers, is capable of obtaining 6 DOF pose estimates of the user within the environment at 15-20Hz, and is designed to be used as either a standalone tracking system or primarily as part of a hybrid optical-inertial tracking system⁹. A tracking system should not interfere with the normal motion of the operator, it should have a fast update rate, low latency, and be accurate. As the enclosed nature of IVY limits the applicability of existing tracking technologies, a new approach was required (see Figure 2 for an illustration). Even though the enclosed nature of the device makes existing approaches inappropriate, it does enable alternate approaches. The projective surfaces outside of the view of the user can be used as part of a system to track the user within the environment. A fixed arrangement of low power laser diodes is attached to a helmet that is worn by the user (see Figure 3). Cameras are positioned behind the screens such that they can view the entire projection surface. The current implementation uses off-the-shelf FireWire® (IEEE 1394) digital video cameras to acquire the images. By tracking the projections of the laser beams on the surfaces of the display in each image, we are able to compute and track the user's head pose. Arranging the laser diodes in a known geometry enables us to constrain the pose of the device based on these projections and allows us to compute the final correct pose. An Extended Kalman Filter is employed to predict, gate, smooth and obtain the final pose estimate.




Figure 4: Basic laser geometry. The four lasers are established so that lines drawn through their beams would intersect at a common point P_0 , and $P_3P_0 \cdot P_1P_2 = P_4P_0 \cdot P_1P_2 = P_3P_0 \cdot P_4P_0 = 0$.

Various configurations of laser diodes could be used to localize the user. Our implementation uses a simple arrangement of four laser diodes in the geometric configuration shown in Figure 4. Two of the laser diodes are arranged so that they project in opposite directions along a line, and the other two diodes are arranged so that they project orthogonal to each other and orthogonal to this line. The projection directions of all four laser diodes intersect at a single point, P_0 . Given the projections of the four laser diodes on the exterior walls of the environment it is possible to obtain strong constraints on P_0 and to define a 3D coordinate frame aligned with this point.

To demonstrate this, we break the problem down into two parts. The first is to determine P_0 and the coordinate system aligned with P_0 given that one can identify the three-dimensional position at which the beam from *specific* diodes strike the various walls, and the second is to determine which laser spot on a wall corresponds to which laser emitter. (Note that this second problem can be avoided by using laser diodes of different frequencies, or pulsing the lasers.)

For the remainder of this discussion, P_1, \dots, P_4 are the 3D positions at which the laser beams from the respective laser diodes strike the walls of the environment. P_0 lies at the intersection of P_1P_2 with a perpendicular line that passes through point P_3 . This point can be found quite easily by noting that P_0 lies along the line defined by $P_1 + \lambda(P_2 - P_1)$ and $P_1P_2 \cdot P_0P_3 = 0$. Solving these equations for P_0 yields

$$P_0 = P_1 + \frac{(P_3 - P_1) \cdot (P_2 - P_1)}{\|P_2 - P_1\|^2} (P_2 - P_1)$$

This defines the origin of the frame, P_0P_3 defines the forward direction vector for the frame, and the normal of the plane is defined by points P_1, P_2, P_3 ; $\vec{n} = P_0P_1 \times P_0P_3$, which determines the direction of the up vector. Although P_4 is not required in order to compute this

frame (provided that the assignment of laser spots to diodes is known), P_4 will prove useful by providing an additional constraint in terms of determining the appropriate mapping from laser spots to emitters. In terms of the geometry it is important to note that P_0P_4 is perpendicular to the plane defined by $P_1P_2P_3$.

These calculations assume that we know the correspondence between each laser diode and each laser projection. In practice this may be accomplished using different wavelengths, or by pulsing the lasers at known times. In our current implementation we take a more algorithmic approach and use the geometry to place constraints on the finite number of possible configurations and choose the correspondence that minimizes an error function. We must determine the appropriate labelings of the tracked laser projections P_i , P_j , P_k , and P_l with the actual laser points P_1 , P_2 , P_3 , and P_4 . There are 24 possible assignments of the laser points to the emitters. Of all 24 possible assignments, only four are consistent with the geometry of the emitters⁹. Figure 5 shows examples of the possible labelings and the impact this has on the pose computation.

Although there are four configurations that are consistent with the geometry of the laser diodes, the three incorrect assignments are sufficiently distant from the correct pose to be easily disambiguated using temporal coherence. If the correct assignment is $(P_i, P_j, P_k, P_l) \rightarrow (P_1, P_2, P_3, P_4)$, then the three incorrect assignments are

1. $(P_i, P_j, P_k, P_l) \rightarrow (P_2, P_1, P_4, P_3)$. This configuration has the same P_0 as the correct configuration, but is rotated by 180 degrees. With a 15Hz sampling rate, the user would have to rotate at roughly 1350 deg/sec before this configuration can be confused with the correct one.
2. $(P_i, P_j, P_k, P_l) \rightarrow (P_3, P_4, P_2, P_1)$. This incorrect assignment and the final remaining assignment have a different P_0 , and an orientation change of at least 90 degrees. This configuration, like the following configuration, is extremely unstable and can only occur under extremely unusual conditions⁹. With a 15Hz sampling rate, the user would have to rotate at roughly 675 deg/sec before this configuration can be confused with the correct one.
3. $(P_i, P_j, P_k, P_l) \rightarrow (P_4, P_3, P_1, P_2)$. This incorrect assignment is similar to the one above. It has a different P_0 as well as at least a 90 degree orientation change.

A simple temporal tracking system coupled with gating is used to discard these incorrect assignments. Although these constraints allow us to keep a consistent pose, there is still an issue of estimating the initial pose. In our current implementation, the initial

correct assignment is chosen manually. Limiting the tracked rotation to less than 500 deg/sec will easily eliminate the three incorrect assignments.

3. Implementation Details

In the following section, we describe in more detail how the tracking system is implemented. We first describe how we acquire and track the 2D laser points from the cameras, followed by a short discussion on calibration of the system and a description of the method used to discard invalid labellings of the laser dots.

3.1. Laser Point Acquisition

Digital cameras situated outside the immersive display aimed at each of the rear-projection screens allow us to track the multiple laser projections. Each camera is equipped with an optical wavelength bandpass filter with a peak response at 650nm (the laser diode wavelength). This allows us to greatly simplify the image processing routine speeding up the response of the tracking system as a whole. To find the centroid of the laser dot in each image, we employ a sub-pixel peak detector⁴ implemented using a weighted sum of the image intensities corresponding to the image region containing a laser dot. IVY utilizes eight projectors to provide video to its six sides (two projectors are used for each of the floor and ceiling in order to reduce the physical footprint of the device). Eight cameras, connected to a single standalone Linux PC, are used to cover IVY's six sides. This PC performs all of the image processing and pose computation and provides "events" of updated pose estimates to the application.

3.2. IVY screen calibration

Calibrating the optical system is of utmost concern and is performed offline using a multi-stage method, first determining the intrinsic parameters (radial/tangential distortions and principle point) of the cameras using the calibration toolbox for Matlab¹, and the extrinsic parameters modeled as a 2D planar homography between the image plane and the projection screen surface computed using the Discrete Linear Transform algorithm described by Hartley and Zisserman⁷. The final calibration step is to determine the relationship between each of the screen surfaces to the world coordinate system. This is accomplished by defining the world coordinate system as the center of IVY and measuring the rigid-body transformation separately for each screen.

3.3. Discarding Invalid Configurations

Once the tracked 2D laser points are available, it is now possible to determine the pose of the person be-

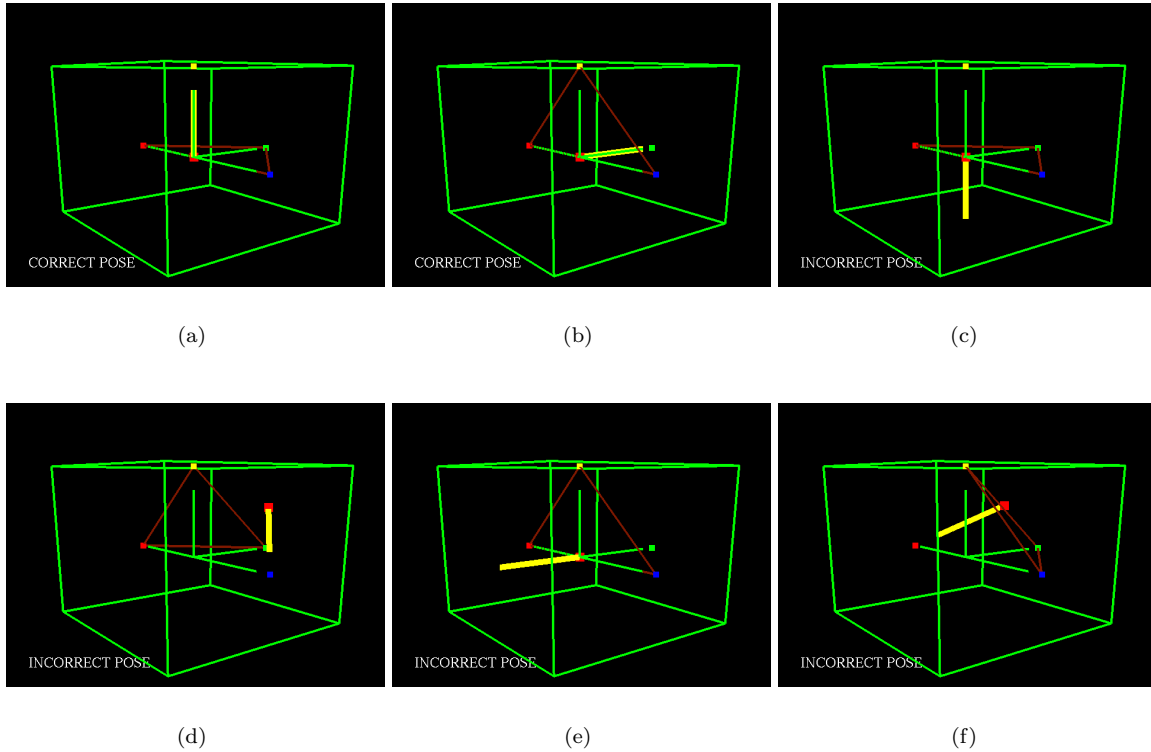


Figure 5: Six examples of the 24 possible labelings and their associated computed pose. Shown here are screenshots from a simulator designed to test the configuration constraints on the laser geometry. The surrounding cube is an analogue of IVY while the smaller dots on the sides of the cube are the associated laser projections. The thick line (shown in yellow) is the computed Up vector, and the computed position is the large dot (shown in red). The connecting lines between laser points indicate which lasers were used to compute the plane for orientation. Each image is labeled with the text “CORRECT POSE” or “INCORRECT POSE” which is automatically computed using only static constraints. (a) is the correctly computed pose while (b) is incorrect but cannot be distinguished using only static constraints. Note that (b) is actually the correct pose rotated by 180 degrees around an oblique axis. A simple temporal mechanism is needed to distinguish between these two solutions.

ing tracked. This is done using the above computation for each possible labeling of the four laser points (24 possibilities). In order to determine which is the correct labeling, we impose geometric constraints on the solution and compute an error function, $\epsilon(i)$, that will allow us to determine which labeling is the correct one. Using three constraints, we are able to determine the correct solution (up to a reflection, see Figure 5(b)) that corresponds properly to the pose of the device.

$$\epsilon(i) = \epsilon_{\perp}(i) + D(i) + \mathcal{F}(i) \quad (1)$$

where i is the current permutation of laser points, $\epsilon_{\perp}(i)$ (the perpendicular error) is the sum of the dot products of vectors that should be perpendicular in this configuration (we take the absolute value of each dot product to ensure an increasing function), and

$D(i)$ is the shortest distance between the computed position P_0^i and $P_0^{i'}$ where $P_0^{i'}$ is computed using $P_4^i P_2^i P_1^i$ rather than $P_3^i P_2^i P_1^i$. This eliminates many of the incorrect labelings since P_0^i and $P_0^{i'}$ will not coincide if the plane normal is computed with the incorrect points. $\mathcal{F}(i)$ is a binary function that is 1 only when our computed plane normal is in the same direction as P_4^i . The perpendicular error, $\epsilon_{\perp}(i)$ is defined as

$$\begin{aligned} \epsilon_{\perp}(i) = & |(P_0^i P_4^i \cdot P_0^i P_1^i)| + |(P_0^i P_4^i \cdot P_0^i P_2^i)| + \\ & |(P_0^i P_4^i \cdot P_0^i P_3^i)| + |(P_0^i P_1^i \cdot P_0^i P_3^i)| + \\ & |(P_0^i P_2^i \cdot P_0^i P_3^i)| \end{aligned} \quad (2)$$

$D(i)$ is defined as

$$D(i) = \|P_0^i - P_0^{i'}\|^2 \quad (3)$$

and

$$\mathcal{F}(i) = \begin{bmatrix} 0 & \vec{n} \cdot (P_4^i - P_0^i) \geq 0 \\ 1 & \text{otherwise} \end{bmatrix} \quad (4)$$

After evaluating $\epsilon(i)$ for each possible labeling, the results are sorted in ascending order according to this error function and the first 2 solutions are taken as the correct pose and reflection. Although there exist four possible solutions, two of these do not occur in practice as they correspond to extreme head positions/orientations within the environment, and are extremely unstable. It is still necessary to distinguish between the final two solutions.

This is accomplished by applying a temporal coherence property on possible body rotational velocity. Given the two normal vectors associated with the final two labellings, \vec{n}_1 and \vec{n}_2 and the previously computed normal \vec{n}_p we compute

$$\theta_1 = \cos^{-1} \left(\frac{\vec{n}_1 \cdot \vec{n}_p}{|\vec{n}_1| |\vec{n}_p|} \right) \quad (5)$$

$$\theta_2 = \cos^{-1} \left(\frac{\vec{n}_2 \cdot \vec{n}_p}{|\vec{n}_2| |\vec{n}_p|} \right) \quad (6)$$

and take the solution with the smallest associated value of θ as the final correct pose.

4. Results

4.1. Orientation

The laser diode housing was placed on a rotational stage in the center of IVY roughly 4' above the floor that allowed us to rotate the device precisely at 1° intervals. Figure 6 shows the raw data points for a full 360° rotation on the azimuth at 5° intervals. For each direction vector, points on the unit circle are drawn at both the measured and correct orientation in the same colour (note that due to the accuracy of the measurement these points appear almost coincident). In a cube-shaped immersive display the corners present problems with tracking, and when the lasers shine into the corners, no data can be collected and tracking is lost until the lasers shine onto the screen. Note that these gaps can be filled in by using inertial data⁹.

In a second orientation experiment, rotational data were collected over a 10 degree range at 1 degree intervals on the azimuth. The relative angle, shown in Table 1, was computed between the direction vectors X-Z components and the first reported direction vector.

The mean error of this exercise was approximately 0.1° while the max error was approximately 0.3°.




Figure 6: 360° Raw Orientation Data (taken at 5° intervals). Unit vectors are plotted, with the same symbol, in the recovered and measured directions. The plotted X-axis is the X-coordinate of the unit vector and the plotted Y-axis is the Z-coordinate of the unit vector. Note: the four large holes indicate positions where the lasers were shining into the corners of IVY and thus could not be tracked.

Rotational Stage	Computed Angle
0°	0.0000°
1°	0.9229°
2°	1.9101°
3°	3.2703°
4°	4.1654°
5°	5.0992°
6°	6.2851°
7°	7.0167°
8°	8.3210°
9°	9.1814°
10°	9.8664°

Table 1: Computed angles between the reported direction vectors at 1° increments.

4.2. Position

To estimate the accuracy of the position estimates, we placed the device at 20 known locations (See Table 2) within IVY and recorded the tracker output. The raw data in this test is illustrated in Figure 7. The mean absolute position error was modest at 1.13cm but there were several cases where the error was nearly

Absolute Position (X-Z) (metres)	Reported Position (X-Z) (metres)	Error X (metres)	Error Z (metres)
(0.00, 0.65)	(0.0072, 0.6511)	0.0072	0.0011
(-0.81, -0.65)	(-0.8167, -0.6527)	0.0067	0.0027
(0.81, 0.30)	(0.8504, 0.3016)	0.0404	0.0016
(-0.81, 0.65)	(-0.8175, 0.6559)	0.0075	0.0059
(-0.81, 0.30)	(-0.8175, 0.3059)	0.0075	0.0059
(0.00, 0.30)	(0.0046, 0.3066)	0.0046	0.0066
(0.81, 0.65)	(0.8587, 0.6584)	0.0487	0.0084
(0.81, -0.65)	(0.8586, -0.6546)	0.0486	0.0046
(0.00, -0.65)	(0.0048, -0.6556)	0.0048	0.0056
(0.50, 0.65)	(0.5072, 0.6484)	0.0072	0.0016
(0.50, -0.65)	(0.5070, -0.6531)	0.0070	0.0031
(0.81, -0.30)	(0.8126, -0.3013)	0.0026	0.0013
(-0.81, -0.30)	(-0.8120, -0.3038)	0.0020	0.0038
(0.00, -0.30)	(-0.0056, -0.3045)	0.0056	0.0045
(0.81, -0.30)	(0.8154, -0.3070)	0.0054	0.0070
(-0.50, -0.30)	(-0.5045, -0.3080)	0.0045	0.0080
(-0.50, -0.65)	(-0.5014, -0.6509)	0.0014	0.0009
(-0.50, 0.30)	(-0.5056, 0.3007)	0.0056	0.0007
(-0.50, 0.65)	(-0.5041, 0.6542)	0.0041	0.0042
(0.50, -0.65)	(0.5051, -0.6534)	0.0051	0.0034
(0.50, -0.30)	(0.5040, -0.3026)	0.0040	0.0026
(0.50, 0.30)	(0.5063, 0.3042)	0.0063	0.0042
(0.50, 0.65)	(0.5070, 0.6512)	0.0070	0.0012

Table 2: Error associated with measured and reported tracker positions. All data points were taken at the same height (*Y*-coord) of 1.35m.

5cm. All errors greater than 1.0cm occurred when $X = 81$ cm. We believe that this is due to the placement of one of the ceiling cameras. Due to space constraints on the physical layout, one camera needed to be placed largely off-axis creating a large perspective distortion in the image.

The noise covariance of each position estimate was also computed and a typical example can be seen in Figure 8(a) using a Linear Kalman filter with variance of 1cm^2 on position. The small covariance (approximately 0.5cm) in the position is attributable to the noise in each laser position estimate due to the limited resolution of the cameras. Since we are acquiring 640x480 resolution images from the cameras, the 2.29m screen is imaged at approximately 500 pixels, making 1 camera pixel correspond to approximately 0.5cm on the screen surface. Using higher resolution images would increase the precision of the tracking system since it would allow us to make better estimates of the laser positions. Since the walls of IVY are fabric walls, and thus vibrate and move slightly when in the presence of large motion within the display, we were concerned how this would affect the position es-

timate. We placed the device in a stationary position and recorded data while violently moving the screen fabric on all walls. The covariance of the estimate can be seen in Figure 8(b). The system reacts well with a spread of approximately 1.5cm even in the presence of large motion of the screen surfaces.

5. Summary and future work

The tracking approach for fully immersive environments presented here has many advantages over existing approaches. The accuracy achieved is not a function of the distance of the user from a base station. The system performance is not degraded by metallic objects or other interference, the user is untethered and is not required to wear a large encumbering device which could compromise their immersive experience. Since the laser diodes are aimed behind the user, their projections do not interfere with the user's visual experience. Also, using off-the-shelf FireWire® digital video cameras allows the tracking system to evolve with the commercial market making it possible to increase the resolution and framerate as new camera technology becomes available. Our current im-




Figure 7: Reported Position of Tracker at known locations in metres. Different symbols are used only to distinguish different measurements.

plementation is limited to approximately 15Hz while work continues on increasing this to 30Hz. Increasing the framerate would have a positive impact on the system performance. It would be easier to disambiguate the invalid pose estimates since it would limit the amount of motion that could occur between updates even more. Work progresses on making the system more robust when the measurements are unreliable or unavailable, e.g. occluded laser dots, lasers projecting into corners make one or more laser projections unavailable and thus increase the delay between updates in this configuration (and increasing the chance of choosing an incorrect pose). This is currently under development by using a SCAAT²³ tracking filter which will allow the user to be tracked consistently even if a laser dot measurement is unavailable for a given frame. Using this algorithm would also decrease the total latency of the system since we would not need to wait until four laser measurements are available to estimate the pose. Although a complete end-to-end latency analysis has not yet been performed, minimum system latency – due to camera capture, initial data processing etc. – is approximately 0.035s.


6. Acknowledgements

We would like to acknowledge Jeff Laurence and James Zacher for helping us build the hardware, and Urszula Jasiobedzka for her help and support. The financial support of NSERC Canada, the Canadian

Foundation for Innovation, SGI, Precarn, IRIS, and the Centre for Research in Earth and Space Technology (CRESTech) is gratefully acknowledged.

References

1. Jean-Yves Bouguet. Camera calibration toolbox for matlab. http://www.vision.caltech.edu/bouguet/calib_doc.
2. Ascension Technology Corporation. Motionstar wireless. <http://www.ascension-tech.com/products/motionstarwireless.php>.
3. 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.
4. R. B. Fisher and D. K. Naidu. *A Comparison of Algorithms for Subpixel Peak Detection*. Springer-Verlag, Heidelberg, 1996. <http://citeseer.nj.nec.com/482699.html>.
5. Center for Parallel Computing. Primeur: Advancing European Technology Frontiers, World's first fully immersive VR-CUBE installed at PDC in Sweden, 1998.
6. K. Fujii, Y. Asano, N. Kubota, and H. Tanahashi. User interface device for the immersive 6-screens display "cosmos". In *Proc. VSM'00*, 2000.
7. R. Hartley and A. Zisserman. *Multiple View Geometry*. Cambridge University Press; ISBN:0521623049, 2000.
8. F. Hetmann, R. Herpers, and W. Heiden. The Immersion Square – immersive vr with standard components. In *Proc. Virtual Environment on a PC Cluster Workshop, Provnino, Russia*, 2002.
9. A. Hogue. MARVIN: a Mobile Automatic Realtime Visual and INertial tracking system. Master's thesis, York University, 2003. In Preparation.
10. Fraunhofer Institute IAO. <http://vr.iao.fhg.de/6-Side-Cave/index.en.html>.
11. Polhemus Inc. Polhemus 3space fastrak. <http://www.polhemus.com/fttrakds.htm>.
12. InterSense. <http://www.isense.com/>.
13. Virtual Reality Applications Center Iowa State University. <http://www.vrac.iastate.edu/about/labs/c6>.
14. 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.



(a) Typical Noise Covariance of Stationary X-Z position in metres (filtering with 1cm^2 variance). This shows a spread of 6mm on the Z-axis and 3mm on the X-axis.

(b) Typical Noise Covariance of Stationary X-Z position in metres (filtering with 1cm^2 variance). This shows that even in the presence of large motion of the screen surfaces (screen movement due to fast motion within IVY), the spread of the measurement is 1.5cm on the X-Axis and 1cm on the Z-axis.

Figure 8: Results from collected data. (a) shows the typical noise covariance of a stationary position, and (b) shows how the noise increases when the fabric screens vibrate due to fast motion.

15. V. Kindratenko. A comparison of the accuracy of an electromagnetic and hybrid ultrasound-inertia position tracking system. In *Presence: Teleoperators and Virtual Environments*, 2001. vol.10, no.6, pp. 657-663.
16. University of Minnesota. Powerwall. <http://www.lcse.umn.edu/research/powerwall/powerwall.html>.
17. University of Tokyo. http://www.iml.u-tokyo.ac.jp/facilities/index_e.html.
18. I. Rötzer. Fraunhofer Magazine, Synthetic worlds within six walls 2:2001.
19. 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.
20. Aalborg University. The VRMedialab. <http://www.vrmedialab.dk/>.
21. Beckman Institute University of Illinois, Integrated Systems Laboratory. A laboratory for immersive cognitive experiments. http://www.isl.uiuc.edu/Virtual%20Tour/TourPages/meet_alice.htm.
22. Joseph J. Jr. La Viola. A discussion of cybersickness in virtual environments. In *SIGCHI Bulletin*, January 2000. vol.32 no.1, pp.47-56.
23. G. Welch. *SCAAT: Incremental Tracking with Incomplete Information*. PhD thesis, Chapel Hill, NC, 1996.
24. T. Yamada, M. Hirose, and Y. Isda. Development of a complete immersive display: COSMOS. In *Proc. VSMM'98*, pages 522-527, 1998.