Measurement and calibration of static distortion of position data from 3D trackers

Steve Bryson†

RNR Technical Report RNR-92-011, March 1992

Applied Research Branch, Numerical Aerodynamics Simulation Division
NASA Ames Research Center
MS T27-A
Moffett Field, Ca. 94035
bryson@nas.nasa.gov

Abstract

Three-dimensional trackers are becoming increasingly important as user inputs in interactive computer systems. These trackers output the three-dimensional position, and often the orientation, of a sensor in space. The three-dimensional tracking is often, however, highly distorted and inaccurate. The purpose of this paper is to discuss methods for the measurement and characterization of the static distortion of the position data. When the distortion is constant, various methods can be used to calibrate the data from the tracker to increase accuracy. Several preliminary methods are discussed in this paper, including polynomial and weighted lookup methods. The measurement and calibration methods are applied to the Polhemus electromagnetic tracking system, but are applicable to tracking systems based on other technologies.

† Employee of Computer Sciences Corporation. Work supported under government contract NAS 2-12961

1: INTRODUCTION

1.1: Purpose and motivation

Three-dimensional trackers are becoming increasingly important as user inputs in interactive computer systems [1][2][3][4]. These trackers give the position of a sensor in three-dimensions. If the tracker were perfect, the position returned by the tracker would exactly correspond to the position of the sensor in appropriate coordinates. In reality, trackers fail to be perfect. Distortions are introduced into the tracking data so that the position returned by the tracker only loosely corresponds to the actual position of the sensor, and then only with a limited volume of space. This distortion is typically a function of the sensor's distance from some source, and is dependent on the ambient environment. If this distortion is constant in time, it can be measured and the actual position of the sensor can be inferred from the distorted data. This will be called calibrating the tracker.

The purpose of this paper is to discuss a method of detailed measurement of the tracker output for a known set of tracker positions. These known tracker positions will fill a volume of space with a resolution of 12 inches. Using these measurements, the distortion of the tracker data can be determined, as well as methods of calibration. These measurements have been performed on a Polhemus Isotrack tracker, an electromagnetically based tracking system that provides three-dimensional position and orientation of a sensor. The measurements have been taken twice within the same location and at different locations to measurement the dependance of the distortion on the ambient electromagnetic environment.

While this paper reports on measurements of commercial a system, the results reported are not to be taken as characterizations of these systems. The Polhemus system measured was modified to decrease filtering. One cannot assume that these results represent the performance of currently available commercial systems. Polhemus offers a calibration system for their tracker, and the Isotrack measured in this paper had not been calibrated by Polhemus.

1.2 Magnetic tracking systems

The trackers measured in this paper are based on designs which use orthogonal electromagnetic fields to sense three-dimensional position and orientation [5]. A source contains three orthogonal coils which are pulsed one at a time. The sensor also contains three orthogonal coils which measure (simultaneously) the three components of the electromagnetic signal being given off by the currently pulsed source coil. These components are measured for the three source coils giving nine measurements. These measurements are processed by a microprocessor in the trackers to deliver the three dimensional position and orientation for the sensor. The raw result of this measurement is rather noisy, requiring filtering of the signal. This filtering in the tracker and enhances the static position measuring performance at the expense of the accuracy of the dynamic aspects of the measurement. Due to ambiguities in the calculation of position, the tracking takes place on one side, or 'hemisphere' of the source.

Due to the dependance of the position measurement on the measurement of the local electromagnetic field, these trackers are sensitive to the ambient electromagnetic environment. It is known that large ferrous objects will distort the field and therefore the tracker data. Cathode Ray Tube displays also highly distort the field. Generally, the manufacturers of these trackers suggest that they be used in environments with as little ferrous metal and electrical devices as possible.

This demand of clean electromagnetic environments cannot often be met. Also, even in the case of relatively clean environments the tracker signal is still somewhat distorted. Finally, the tracking depends on the measurement of a field generated by the source, and so is dependant on the distance of the sensor from the source. The Polhemus tracker is advertised to be accurate for sensor-source distances of up to three feet or so.

1.3 The measurement process

The tracker measurements were performed during the author's tenure at the NASA Ames VIEW Lab under Scott Fisher and later Steve Ellis. The VIEW lab is in the Aerospace Human Factors Research Division at NASA Ames Research Center. The measurement protocall was first outlined by Scott Fisher and Doug Kerr. The VIEW lab was moved into a new building during the time of these measurements, affording the possibility of measuring the dependance of the tracker distortion on the environment.

In the VIEW lab, the tracker sources are mounted just below the ceiling, using the hemisphere below the tracker for the

volume of operation. A standard 6'x4' pegboard on a stand was constructed. This pegborad was used to determine a set of actual positions for the sensor. The floor of the VIEW lab was marked with positions for the pegboard so that a regular 8'x8'x6' volume could be measured. Certain holes were marked on the pegboard at 12" intervals as the measurement sights. At each measurement sight 60 measurements of the tracker were taken with the sensor at a fixed orientation. These 60 measurements took approximately one to two seconds. The 60 measurements were then used to compute an average position for the time of the measurement and a standard deviation over that time. Both the position measurement and the standard deviation are three-dimensional vectors.

At each location of the VIEW lab, called the old and new VIEW labs, the above set of measurements was taken twice on the Polhemus Isotrack, to look for any time dependance of the distortion. A set of measurements was taken of the Polhemus Isotrack at the new VIEW lab offset 6" along all three axes from the standard locations. This allowed the study of calibration methods at points other than those used to generate the calibrations.

2: MEASUREMENT RESULTS

2.1 Representations of results

Tracker distortions can be viewed in several ways. Fundamentally, they are functions from three-dimensional space (the actual sensor position) to three dimensional space (the position returned by the tracker). The clearest way of representing this distortion is with an error vector, defined as the returned position minus the actual position. This will generate a set of error vectors, one at each point of measurement. This may be drawn on a three-dimensional graphics device and clearly indicates the distortion. On the printed page, however, the three-dimensional drawing of this vector field gives an ambiguous picture. Another characterization of the error is to graph the magnitude of the error vector as a function of distance. This lends itself well to drawings on a two-dimensional page, and several results will be printed with this method.

2.2 Three-dimensional Structure of Tracker Distortion

For the measurements in the old VIEW lab, the Polhemus Isotrack shows a distinct axially symmetric bowing of the data, as indicated in figures 1 and 2. Figure 1 shows the error vectors in a perspective rendering for the old VIEW lab. The error vectors that are longer than 20 inches have been filtered out as they cluttered up the screen. The dots show the locations of the measurement sights as described in section 1 above. The Polhemus source is above the center of the measurement volume. The distortion seems to be axially symmetric in the sense that the error vector field can be approximately generated by taking the error vectors on a plane perpendicular to the floor passing directly under the tracker source and rotating it about the axis defined by the vertical line directly under the source.

Figure 1: A perspective view of the error vectors in the tracker distortion measured in the old VIEW lab. The dots are the measurement locations whose error vectors are larger than 20 inches. The measurement locations are on a cubic lattice with 12 inch spacing.

Figure 2 shows two y-z slices from the same data set one at x = - 6 inches and the other at +6 inches, so they are on either side of the tracker source. In this figure the error vectors have been projected onto the y-z plane. The bowing phenomena is apparent, and the three-dimensional structure of the distortion can approximated by rotation about the center vertical axis of each figure. Figure 3 shows two planes of error vectors from the new VIEW lab. Note that the bowing is present but is much flatter, with a more severe upturn far from the tracker source. The distortion in the new VIEW lab seems to contain a constant z offset, which suggested a systematic error in the data collection and/or analysis. Careful measurement, however, confirmed that the constant z offset is truly present in the data.

Figure 2: Two y-z planes of the data of figure 1. These data are taken to six inches to each side of the tracker. The error vectors are projected onto the y-z plane, and the measurement spacing is 12 inches.

Figure 3: Two y-z planes of the error vectors from the new VIEW lab in the same format as in figure 2.

2.3 Tracker error as a Function of Distance

The magnitude of the error vector for the old VIEW lab data set shown in figure 1 above is plotted as a function of distance in figure 4. It is shown both with all data points, and in closeup showing those data points whose measurement points are within 50 inches of the tracker source. This plot shows all data points. It is apparent that the error vectors in the old VIEW lab were extremely sensitive to distance. Out to about 50 inches, the relationship between error and distance follows something of a curve implying that out to this distance we can expect some success at calibration. Beyond 50 inches the scatter in the relationship between error and data becomes very large.

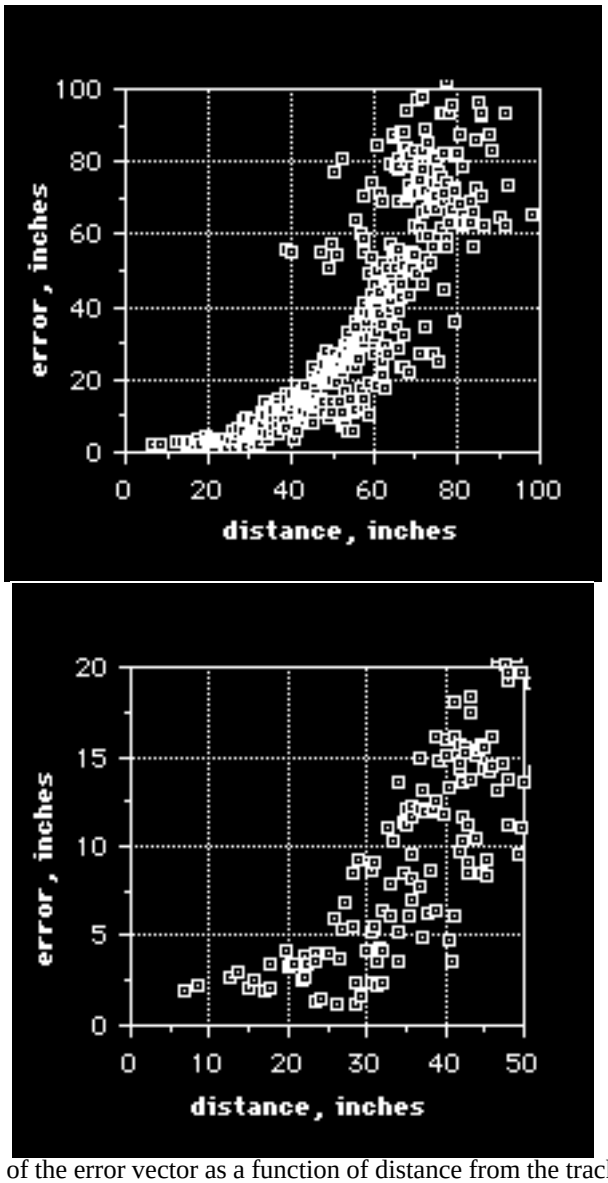


Figure 4: The magnitude of the error vector as a function of distance from the tracker source for data taken at the old VIEW lab. Left: the entire data set. Right: closeup of those measurements within 50 inches of the source.

Figure 5 shows the same data for the new VIEW lab. In this new environment it is apparent that the tracker is much less sensitive to error out to about 50 inches, but has a constant error throughout this range. In the overall graph of figure 5, we see that there is an interesting band pattern in the error at 40 inches and beyond. This phenomenon is not understood.

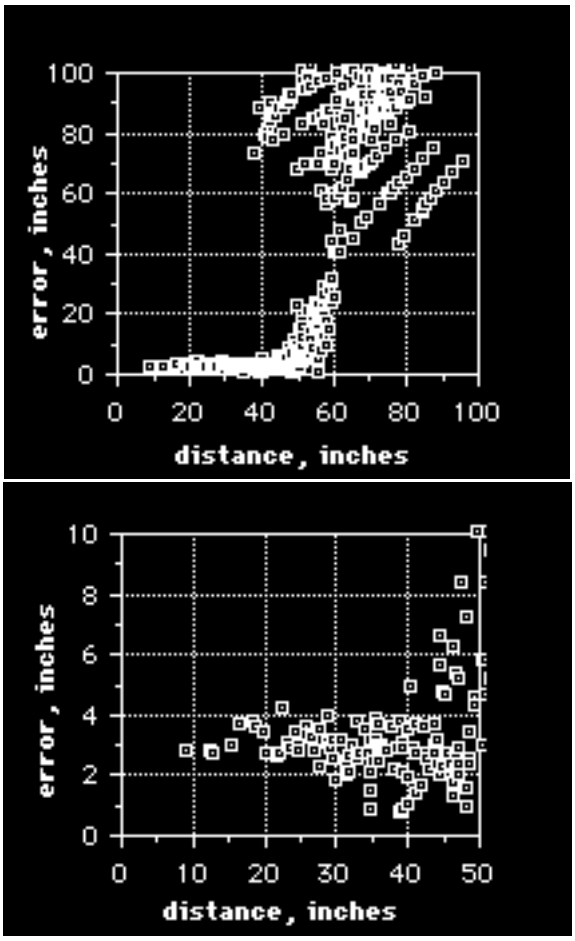


Figure 5: The magnitude of the error vector as a function of distance from the tracker source for data taken at the new VIEW lab. Left: the entire data set. Right: closeup of those measurements within 50 inches of the source.

2.4 Tracker Noise and Repeatability

As described in section 1, tracker noise was measured by taking 60 data points at each measurement location. Tracker noise is defined as the standard deviation of these 60 measurements. Figure 6 shows the tracker noise plotted as a function of distance to tracker source for the new VIEW lab, again with whole data set and in closeup. For distances less than 50 inches the standard deviation grows with distance, but is small in this range, less than 0.5 inches. Beyond 50 inches the noise can be quite large, with standard deviations of tens of inches. This implies that we cannot expect calibration to work beyond 50 inches.

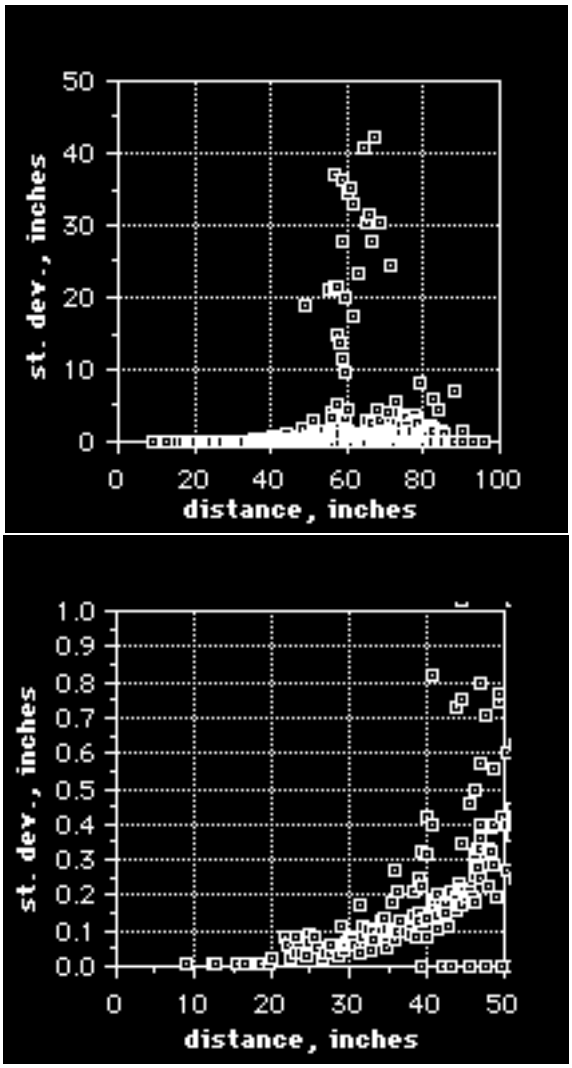


Figure 6: Tracker noise, measured as the standard deviation of 60 consecutive measurements made over a period of about one second, plotted as a function of distance from the tracker source for data taken at the new VIEW lab. Left: the entire data set. Right: closeup of those measurements within 50 inches of the source.

Tracker repeatability was measured by taking data at the same location in the same way on separate days. A measurement of the repeatability is to take the difference between the two day's measured positions for each measurement location. The magnitude of the resulting vector is plotted as a function of distance for the new VIEW lab is shown in figure 7. There is a difference beyond that accounted for by noise that starts in the range of 1-2 inches and rises slowly with distance to about 4 inches at 50 inches distance. Consistent with the noise measurement, beyond 50 inches the difference vector gets quite large. The repeatability measurement implies that calibration cannot be expected to correct for tracker distortion to better than 1-2 inches.

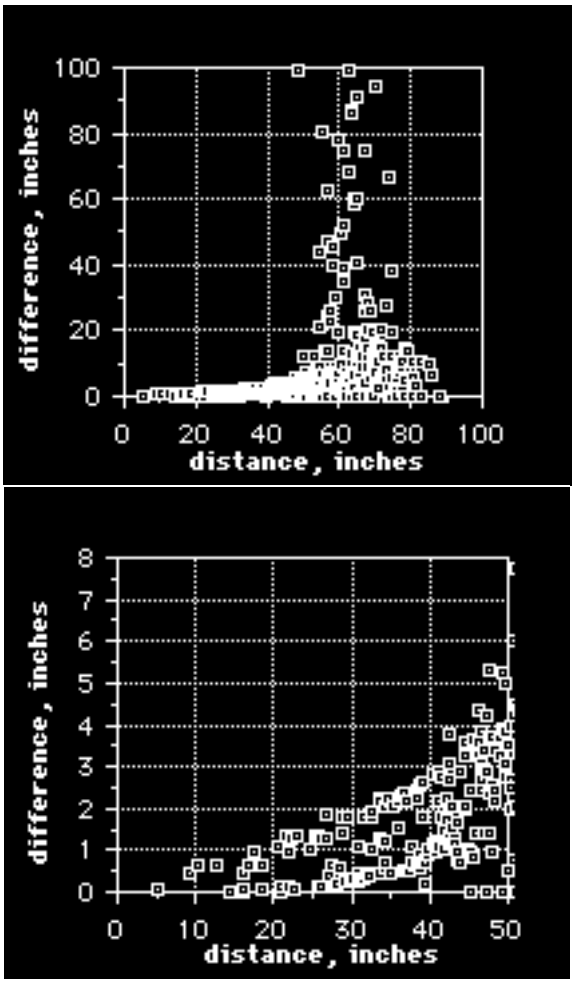


Figure 7: Tracker repeatability, measured as the magnitude of the difference vector of measurements made on different days, plotted as a function of distance from the tracker source for data taken at the new VIEW lab. Left: the entire data set. Right: closeup of those measurements within 50 inches of the source.

3: CALIBRATION METHODS

3.1 Calibration Methods and their Evaluation

Given a measurement of the tracker distortion, that measurement can be used to estimate the actual position of the sensor given some distorted tracker data. This calibration of tracker data can be done in several ways. We have tried two general methods: polynomial calibration; and weighted table lookup. Both of these methods are mentioned as possibilities in the original publication on this problem [5].

To evaluate the calibration schemes, one can generate the post-calibration error vectors just as one investigated the distortion. The smaller the post-calibration error vectors, the better. This is quite good for a qualitative sense of the success of the calibration via error plots of the same type as in section 2. One must not evaluate the calibrations by measuring their success on measurements taken at the same locations as the calibrating data set. As described in section 1, set of measurements, called the *offset data set*, offset by one-half grid spacing in all three axes was taken for the evaluation of the calibrations.

One can also compute the traditional correlation function from statistics. This is a function that measures the correlation of two sets of data. It varies between zero and one, with one indicating perfect correlation. Let the actual position of the

sensor at the time of measurement is labeled p and the calibrated value of the data is labeled labeled p^. The calibration is considered successful if p^ = p. The accuracy of the calibration was measured with the function

$$1 - \frac{\sum_{k=1}^{n} |p_k - \hat{p}_k|^2}{\sum_{k=1}^{n} |p_k - avgp|^2}$$

where avgp is the average position of all the positions p. This function varies from 0 to 1, with 1 indicating perfect correlation. This function measures the degree to which the calibrated values are more accurate than random guessing, which would center around the average of all measured positions.

3.2 Least-squares fit Polynomial Calibration

The calibration via polynomial fit tries to fit the data via an order n polynomial in x, y, and z. If px (px^), py (py^), and pz (pz^) are the components of the actual (calibrated) tracker locations,

$$\hat{p}_x = poly1(p_x, p_y, p_z)$$
$$\hat{p}_y = poly2(p_x, p_y, p_z)$$
$$\hat{p}_z = poly3(p_x, p_y, p_z)$$

These polynomials are found via least-squares methods. Polynomials of order 1 through 8 were tried, using the new VIEW lab data sets. The measured positions used in the generation of the least-squares fit were filtered to eliminate noisy data, by only allowing those measurements with error vectors of less than 20 inches. These passed 150 data points. Least-squared polynomial fits were generated, and tested using the offset dataset. The resulting correlation for each order of polynomial calibrations are graphed vs the polynomial order in figure 8. The order 4 polynomial has a correlation marginally closest to 1.0 with a correlation of .97. We therefore choose the order 4 polynomial calibration for more detailed comparison with the lookup methods.

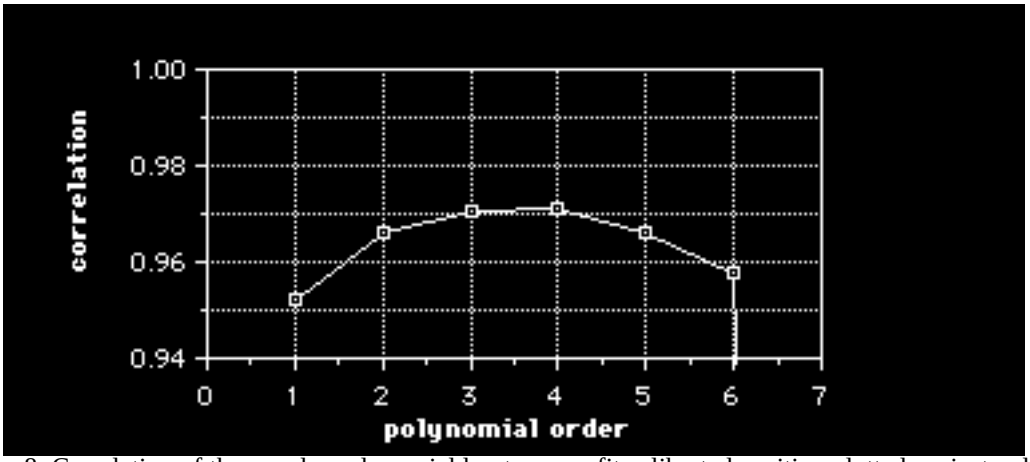


Figure 8: Correlation of the n-order polynomial least-square fit calibrated position plotted against polynomial order.

3.3 Lookup Calibration

Lookup calibration uses a table providing an actual position given a measured location. This table is the product of a previous measurement as described in section 1. The basic scheme is: given a measured position, find the nearest pre-measured position via a table search and then look up the corresponding actual position. This method is unambiguous when the input position exactly corresponds to a measured position. When it does not, an interpolation method using several nearby points in the table must be specified. We evaluated two interpolation schemes, both using weighted sums of the eight points surrounding the input position. These two schemes involved weights which are dependant on the distance from the input point

to the pre-measured point. The first scheme, the *linear lookup calibration*, defines each weight as a linear function of distance. The second scheme, the *bump lookup calibration*, defines each weight as an exponentially decaying function of distance.

Given a point to be calibrated, the lookup calibration identifies the nearest point from the calibration set. Then the eight points bounding the octant containing the input point, called *calibration points*, are used to generate a calibration vector for the intput point. This vector is defined as

$$v = Isu(i=1,8, w_i)\ v_i$$

where the index i lables the points bounding the octant. The v_i are the difference vectors between the actual points and the measured points in the calibration data sets. The w_i are the normalized wieghts assigned to each point and are defined differently in the linear and bump schemes. In the bump scheme, w_i is defined by the formula

$$w_i = eSup9(-Bbc((F(r_{max},r_i)-1)Sup9(-1))$$

where r_i where is the distance from the input point to the ith calibration point, and r_{max} is a constant representing the distance for which w_i vanishes. This function is equal to 1 for r_i equal 0. Thus r_{max} must be smaller than the smallest distance between any two points in the calibration data set. In the linear lookup scheme, the product of the other calibration points is used to define the w_i:

$$w_i = Ipr(L(j=1,j\neq i),8, r_j)$$

This function is small for the ith calibration point if the distance from the input point to that calibration point is small. Thus the contribution from the nearest point dominates. Finally, the wieghts are normalized so that the total contribution equals unity:

$$w_i = F(w_i\ ,Isu(i=1,8, w_i))$$

3.3 Evaluation of the Calibration Schemes

The offset data set was used for input to evaluate the various calibration schemes. The uncalibrated errors of this is the data set are shown in figure 5. The error vector magnitudes for the polynomial calibration are plotted as a function of distance in figure 9, for the linear lookup calibration in figure 10, and for the bump lookup calibration in figure 11.

We see from the graphs that the polynomial calibration was successful in reducing the error from about 4 inches to about 2 inches up to a distance of about 40 inches from the source. As expected from the results on noise and repeatability discussed in section 2, the errors at distances of greater than 50 inches were not successfully corrected by calibration. There is a significant amount of error, however, even for those points close to the tracker source.

The linear lookup introduces considerable amounts of error and scatter. It is apparent that the linear lookup fails as a viable calibration method. The bump lookup method is superficially as successful as the 4th order polynomial calibration. In the bump lookup, the error does get smaller for points close to the tracker source, but there is some additional scatter as compared to the 4th order polynomial calibration in the 30-50 inch range.

Figure 12 shows the same y-z plane with uncalibrated data and the three calibration schemes. In this figure it is apparent that the bump lookup method provides the smallest error vectors in the neighborhood of the source. The linear lookup method also has small error vectors in a smaller neighborhood of the source, but greater error farther from the source. The 4th order polynomial lookup has smaller error vectors overall, but has noticeably larger error vectors near the source.

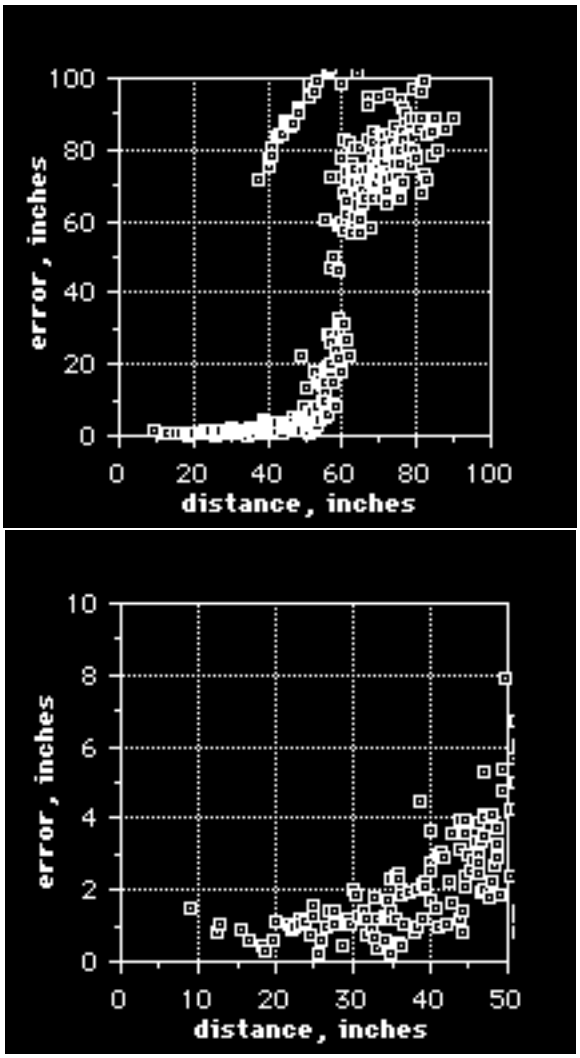


Figure 9: Polynomial least-squares calibration of the data of figure 5. The magnitude of the error vector is plotted as a function of distance from the calibrated tracker source Left: the entire data set. Right: closeup of those measurements within 50 inches of the source.

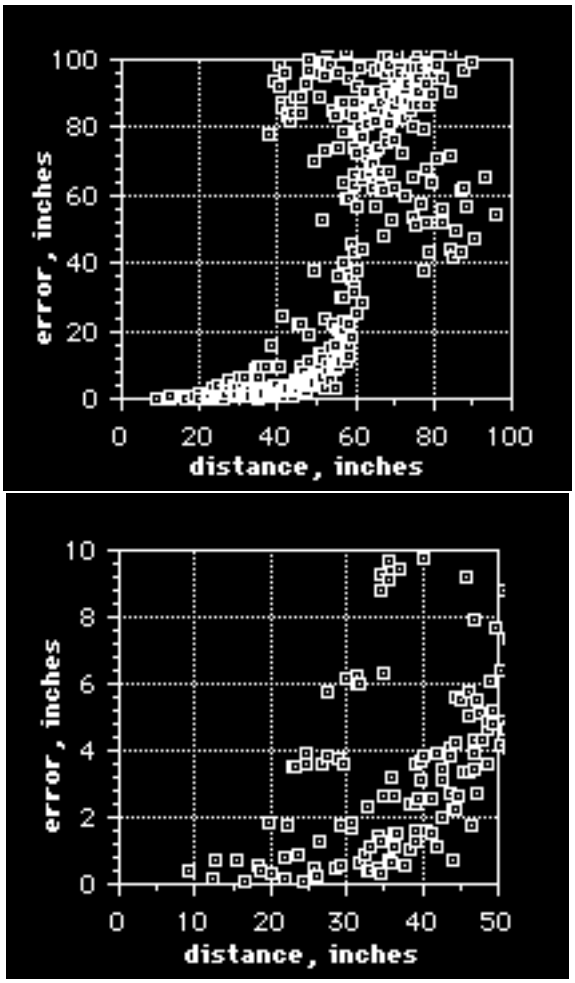


Figure 10: Linear lookup calibration of the data of figure 5. The magnitude of the error vector is plotted as a function of distance from the calibrated tracker source Left: the entire data set. Right: closeup of those measurements within 50 inches of the source.

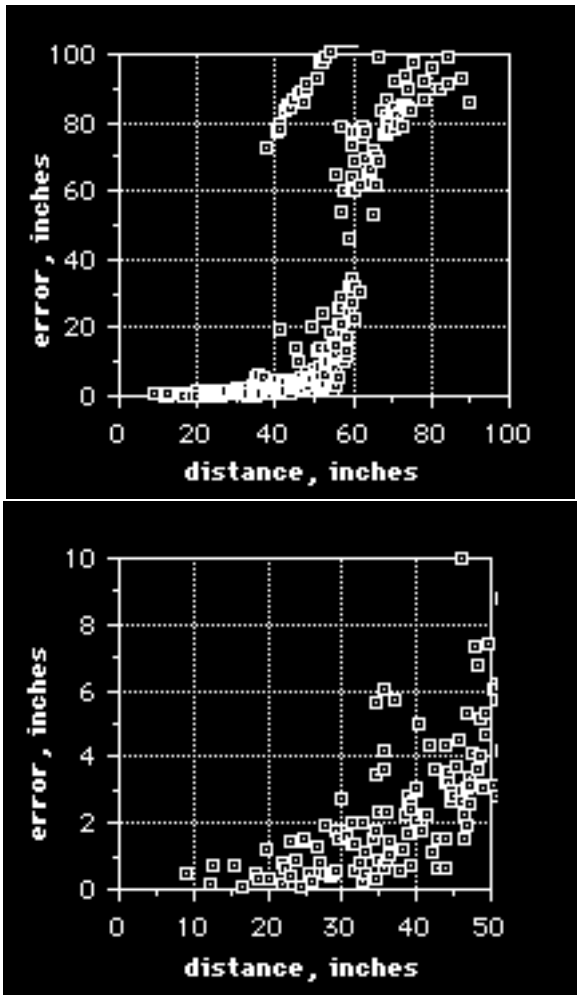


Figure 11: Bump lookup calibration of the data of figure 5. The magnitude of the error vector is plotted as a function of distance from the calibrated tracker source Left: the entire data set. Right: closeup of those measurements within 50 inches of the source.

Figure 12: Error vectors on the same y-z plane of data, uncalibrated and calibrated by three schemes. Clockwise from upper left: uncalibrated data set, 4th order polynomial calibration, bump lookup calibration, and linear lookup calibration.

5. CONCLUSIONS

This paper presents the measurements of the static distortion in the position signal of a Polhemus electromagnetic three-dimensional tracker. It has been found that the distortion is significant, very noisy at distances of greater than 45 to 50 inches, and is very sensitive to location. While these results have been found for a particular tracker, they can be generalized to any position tracker in three-dimensional space.

To some extent, calibration can be used to correct for these distortions. Three calibration methods were discussed in this paper, least-squares polynomial fit calibration, linear lookup calibration, and bump lookup calibration. Of these, 4th order least-squares polynomial calibration and bump lookup calibration performed best in different ways. The 4th order polynomial fit has the best overall behavior, while the bump lookup calibration performed best near the tracker source. None of the calibration methods tried are likely to be optimal.

Study of the noise and repeatability imply limits on calibration success. When tracking greater than 50 inches from the source, the tracking signal is so noisy that no useful calibration can be expected. Inside this distance, repeatability implies a limit of about an inch in calibration accuracy. Both 4th order polynomial calibration and bump lookup calibration perform close to these limits.

Further study in the calibration question can proceed in two obvious directions. The success of polynomial and lookup calibrations suggest that a three-dimensional spline calibration, a combination of global polynomials and lookup tables, should work quite well. Lookup calibration requires more study of the weighting and interpolation question. The bump look-

up calibration method fails for overly distorted data sets and can be refined.

This paper addresses only static position calibration. These trackers also produce orientation data, and the study of distortion in the calibration should be performed. The dynamic signal from these trackers is also significant. Work is proceeding at NASA Ames studying the dynamic distortion in the position data [6]. Calibration methods such as Kallman filters will have to be developed to cover this distortion as well. Finally the possible cross-coupling between these various distortions should be studied.

6. ACKNOWLEDGEMENTS

The author would like to thank and acknowledge Scott Fisher and Doug Kerr for initiating this study and outlining the measurement procedures. Many thanks to Ian McDowall and Vivian Tong for many tedious hours spent assisting in the measurements described in this paper. Thanks also to Rick Jacoby and Steve Ellis of the Aerospace Human Factors Research Division at NASA Ames for useful discussions and assistance.

7. REFERENCES

1 Blanchard, C. et. al., Reality Built for Two: A Virtual Reality Tool, *1990 Symposium on Interactive 3D Graphics*, Computer Graphics, Vol. 24, No. 2, March 1990

2 Bryson, S. and Levit, C., Virtual Wind Tunnel, *IEEE Visualization 91*, San Diego, Ca, Oct. 1991, to appear in *Computer Graphics and Applications*, July 1992

3 Fisher, S. et. al., Virtual Environment Interface Workstations, *Proceedings of the Human Factors Society 32nd Annual Meeting*, Anaheim, Ca. 1988

4 Fisher, S., Virtual Environments, Personal Simulation and Telepresence, *Implementing and Interacting with Real Time Microworlds*, Course Notes, Vol. 29, SIGGRAPH 1989

5 Raab, F., Blood, E., Steioner, T. and Jones, H., Magnetic Position and Orientation Tracking System, *IEEE Transactions on Aerospace and Electronic Systems*, Sept. 1979

6 Bryson, S., Defining, Modeling and Measuring System Lag in Virtual Environments, *SPIE Conference on Stereoscopic Displays and Applications* San Jose, Ca. Feb. 1990, pp 98-109