Distortion Correction of Magnetic Fields for Position Tracking

Gabriel Zachmann
Fraunhofer Institute for Computer Graphics
Wilhelminenstrasse 7
64283 Darmstadt, Germany
email: zach@igd.fhg.de

Abstract

Electro-magnetic tracking systems are in wide-spread use for measuring 6D positions. However, their accuracy is impaired seriously by distortions of the magnetic fields caused by many types of metal which are omnipresent at real sites. We present a fast and robust method for "equalizing" those distortions win order to yield accurate tracking. The algorithm is based on global scattered data interpolation using a "snap-shot" of the magnetic field's distortion measured once in advance. The algorithm is fast (it does not introduce any further lag in the data flow), robust, the samples of the field's "snap-shot" can be arranged in any way, and it is easy to implement. The distortion is visualized in an intuitive way to provide insight into its nature, and the correction algorithm is evaluated in terms of accuracy and performance. Finally, a qualitative comparison of the suceptibility of a Polhemus and an Ascension tracking system is carried out.

Keywords: scattered data interpolation, distortion, virtual reality, cave

1. Introduction

The possibility of measuring the position and orientation of a point in space is one of the enabling technologies of virtual reality (VR). Electro-magnetic trackers have become the most wide-spread devices used in today's virtual reality systems ([12, 1]) in order to track the position and orientation of a user's hands and head, or to track instruments such as an endoscope and scissors. They're also being deployed in real-time motion capture systems to track a set of key points and joints of the human body. Commercial optical tracking systems are getting more mature; however, they're still much more expensive than electro-magnetic systems, but are not yet quite as robust in terms of drop-outs.

Unfortunately, there is one big disadvantage of electro-magnetic trackers: the electro-magnetic field itself, which gets distorted by many kinds of metal. Usually, it is impossible to banish all metal from the sphere of influence of the transmitter emitting the electro-magnetic field, especially when using a long-range transmitter: monitors contain coils, walls, ceiling, and floors of the building contain metal trellises and struts, chairs and tables have metal frames, etc. While tracking systems using direct current seem to be somewhat less susceptible to distortion by metal than alternating current systems, all ferro-magnetic metal will still influence the field generated by the transmitter.

A distortion of the magnetic field directly results in mismatches between the tracking sensor's true position (and orientation) and the position (orientation) as reported by the tracking system. Depending on the application and the set-up, mismatches between the user's eye position (the *real viewpoint*) and the virtual camera's position (the *virtual viewpoint*) impair more or less the usability of VR. For example, in assembly tasks or serviceability investigations, fine, precise, and true positioning is very important [6]. In a cave or at a workbench, a discrepancy of 7 cm (3 in) between the real viewpoint and the virtual viewpoint leads to noticeable distortions of the image[1], which is, of course, not acceptable to stylists and designers. For instance, straight edges of objects spanning 2 walls appear to have an angle (see Figure 8), and when the viewer goes closer to a wall, objects "behind" the wall seem to recede or approach (see [10] for a description of some other effects). Mismatches are most fatal for Augmented Reality with head-tracking in which virtual objects need to be positioned exactly relative to real objects [14].

In order to overcome these adverse effects, we have developed an algorithm which can correct these distortions of the magnetic tracking field. The algorithm is based on measured data which relate the true position to the position reported by the tracking system at several points within the

[1] This is just a rule of thumb, of course. The threshold at which a discrepancy between the real and the virtual viewpoint is noticeable depends on many variables: expertise, distance from the cave wall or projection screen, size of the cave, etc.

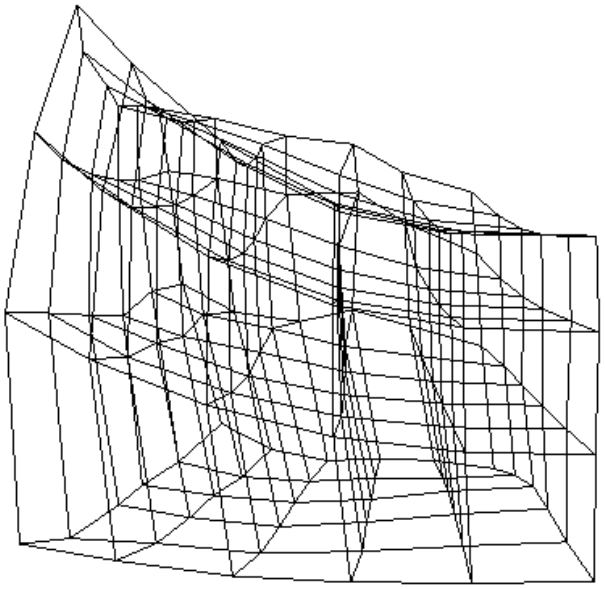

Figure 1. Visualization of the distortion of the field in our cave: the nodes of the lattice are the measured points of a uniform lattice. The discrepancy can be as much as 40-50cm (15-20in).

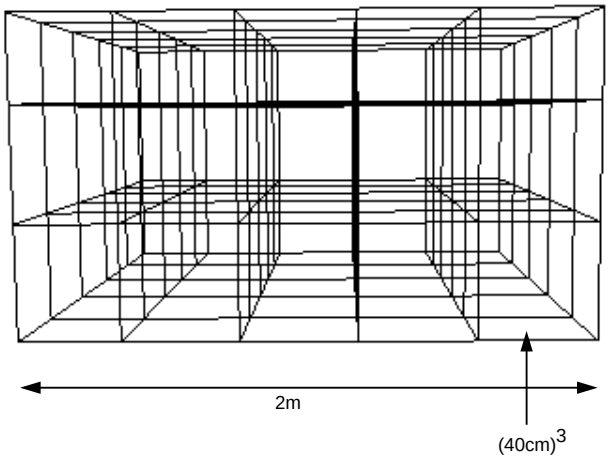

Figure 2. The lattice of sampling points. At each of those points, the tracking sensor's value has been averaged and recorded, which produces the field's "snapshot" on the left.

volume of interest. At run-time of a VR session, the algorithm interpolates these *a priori* measured values with the currently reported position of the sensor.

Since the distortion of the magnetic field is captured by a set of points, a fundamental assumption of our method is that the field does not change over time. Fortunately, in our labs this seems to be true – we could not find any significant changes (see Section 4). Of course, if the set-up is changed, then the magnetic field has to be measured again; the field changes, for instance, when a nearby projector is moved, speakers are installed, or the whole set-up is moved to another place.

Our algorithm has several desirable qualities which makes it very suitable for VR systems. First of all, the algorithm is fast, so it does not introduce any latency into the VR system. Second, the set of measured points can be chosen arbitrarily (it does not need to have a lattice-like arrangement), and more points can be added to it at any time. Last but not least, it is very easy to implement.

While this paper reports on measurements carried out using commercial systems, the results reported are not to be taken as a characterization of these systems. Except where otherwise noted, the Polhemus Fastrak with a long range transmitter was used to generate all the measurements.

The next section discusses earlier work. Section 3 describes the method in detail, Section 4 reports on some results. Finally, Section 5 draws some conclusions and comments on further work.

2. Related Work

To our knowledge, little work has been done on the specific problem of examining and correcting the errors of magnetic tracking devices.

[3] has carried out some experiments on tracker error and noise. Three algorithms for correction of those errors were presented: polynomial approximation and local interpolation with two different weight functions. Although all three algorithms are evaluated, it is not clear whether the two local interpolation methods define continuous functions. The error of those correction algorithms is in the range of 2-10cm (1-5in). The evaluation seems to indicate that accuracy decreases with increasing distance from the transmitter.

A problem similar to the one considered here is the mapping from physical space to computational space and vice-versa, which occurs frequently in scientific visualization. This is usually done by linear local interpolation schemes based on shape functions [2, 4]. These algorithms are well suited for very large data sets. However, they have a few drawbacks: most of them provide only C^1-continuity, additional data structures storing topological information is needed, and there is always a little overhead in order to look up the appropriate shape function.

3. Equalization of the Field

The problem we are presented with can be stated as the well-known interpolation problem in 3-dimensional space: Given two sets of points $\mathcal{P} = P^i \subset \mathbb{R}^3$ and $\mathcal{Q} = Q^i \subset \mathbb{R}^3$, we are looking for a function $f : \mathbb{R}^3 \to \mathbb{R}^3$ such that $f(P^i) = Q^i$. Furthermore, we want the function to be sufficiently "smooth": it should be at least $\mathcal{C}^1$-continuous, and it should interpolate "intuitively nice", in particular, there should be no oscillations.

We will call $\mathcal{P}$ the *measured points* in *tracker space*, or, a *snapshot* of the magnetic field; $\mathcal{Q}$ will be called the *true points* in *true space*.

All interpolation schemes can be classified into two categories: global and local. Global methods take all P^i into account, while local methods consider only some P^i within a certain neighborhood. One might also consider the approximation problem. However, we opted for interpolation, since there is a satisfactory solution in our case (see below).

Some local interpolation methods (such as the ones described in [3] or the ones used in scientific visualization) involve a look-up of the closest points or the enclosing cell. This needs some additional data structures if the look-up is to be exact or if it is to be fast. One way or another, there is always some computational burden imposed by the look-up itself, which is another reason why we have chosen a global method.

3.1. The correction algorithm

Because of our considerations above, we implemented *Hardy's Multi-Quadric* method (HMQ), which is a scattered data interpolation scheme [11, 8]. It can be used to construct an interpolation function $f : \mathbb{R}^n \to \mathbb{R}^m$, with arbitrary m, n.

The general form of the interpolation function for $m = n = 3$ is

$$f(P) = \sum A_i \omega_i(P) \quad , \qquad P, A_i \in \mathbb{R}^3$$

$$\omega_i(P) = \sqrt{(P - P_i)^2 + R^2} \quad , \quad R > 0$$

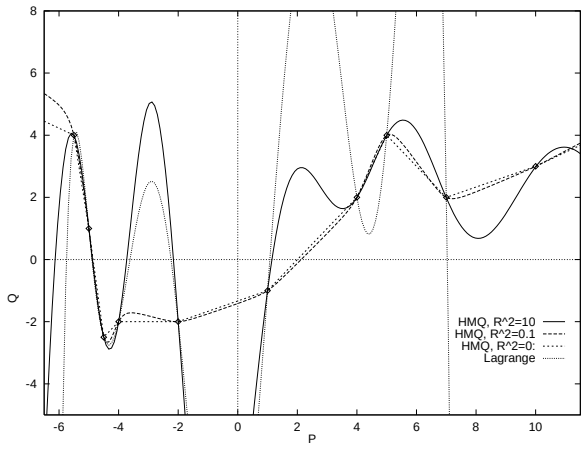

Figure 3. 1D examples of HMQ interpolation functions through 10 points with various R^2 parameters. For $R^2 = 0$ the function is piecewise linear. The Lagrange interpolant exposes fatal oscillations.

Requiring $f(P_i) = Q_i$ leads to three sets of linear algebraic equations with a symmetric matrix. All three matrices have size $N \times N$, $N =$ number of measuring points.

We cannot use the Cholesky decomposition because the matrix is not positive definite, since the upper-left 2×2 sub-determinant is negative:

$$\begin{vmatrix} |R| & \beta_{12} \\ \beta_{21} & |R| \end{vmatrix} = R^2 - \beta_{12} < 0$$

with $\beta ij = \sqrt{(P_i - P_j)^2 + R^2} > R^2$, provided $P_1 \neq P_2$. LU decomposition [15] is, therefore, a natural choice for solving these equations, which is what we do.

The HMQ method does not tend to oscillate as polynomial interpolation schemes do (e.g., Newton or Lagrange interpolation), since the degree of the interpolating function does not depend on the number of sample points. Instead, the smoothness of the HMQ interpolation function depends on the parameter R^2. It can be shown, though, that $f \in \mathcal{C}^\infty$ for $R^2 > 0$ [11]. Figure 3 compares the HMQ method to Lagrange interpolation (in 1D), and shows the effect of various R^2's.

There are other scattered data interpolation functions such as *Shepard* interpolation, or natural Hermite spline interpolation [9]. However, all of these do not seem to be better than HMQ (sometimes much worse) while being much more involved [13].

We have tried a close variant of HMQ, the reciprocal multi-quadric (RMQ), which is defined by basis functions

$$\omega_i(P) = \frac{1}{\sqrt{(P - P_i)^2 + R^2}} \quad , \quad R > 0$$

Although we expected the RMQ to produce fairer interpolations, we had to learn that this is not true. Also, selection of the optimal R^2 seems to be more critical than with HMQ.

The more general form of multi-quadrics is defined on basis functions

$$\omega_i(P) = ((P - P_i)^2 + R_i^2)^{\mu_i}$$

However, as [9, 7] point out, results are best when $\mu_i = \mu$ and $R_i = R$. We have experimented with $\mu_i = \frac{1}{4}, \frac{1}{2}, 2$, and $||\cdots||$. We found, that interpolation is best for $\mu_i = \frac{1}{2}$. For some of the other exponents a good R^2 is very critical and the coefficient matrix might even be near-singular.

3.2. Measuring the Distortion

The first step of our method is to capture the distortion of the magnetic field so that it can be corrected subsequently based on this data. The measuring needs to be done only once per site and set-up. Obviously, it is still quite desirable that it can be done in a minimum amount of time.

For that reason, we have chosen a regular lattice for the set of measuring points, even though our correction algorithm does not require that. A lattice was chosen because it makes it feasible to generate the true positions of the measuring points automatically, thus minimizing the measuring time. Furthermore, we use four sensors, which further reduces the time. Of course, we take the average of many values (usually 50) when measuring one position sample, because the quality and precision of the data will affect the correction later-on.

The apparatus for moving the sensors at well-defined points in space is a simple wooden boom, which is placed on a small trolley. From its end we hang a plummet on a cord, at which we attach the sensors using small patches of velcro. Thus we can move the sensors about while their height above the ground (the y coordinate) stays precisely the same. In order to position the plummet at well-defined x- and z-coordinates, we fixed paper to the floor (wide wallpaper, for example) and marked those coordinates with a pen.

With a set-up as described above, we are able to measure the field at 144 points ($= 6 \times 6 \times 4$) in about 20-30 minutes. Four positions are recorded at a time, which takes about one second for 4×50 samples (for averaging). Measuring time could be reduced further if more sensors would be used simultaneously.

There are many ways of representing the measured distortion. While [3] uses error vectors, we believe that visualizing the measured lattice provides more insight into the data. One such measured lattice is shown in Figure 1 with the sampling points shown in Figure 2. The long-range transmitter is located near the upper right corner in the front. One can see clearly that the distortion tends to increase with the distance from the transmitter, although there are also regions closer to the transmitter that also have a large distortion. The influence of one of the projectors (which is very close to the cave due to space limitations) can be seen in the back of the lattice.

4. Results

4.1. Accuracy

We have tested our method with very satisfactory results. With a data set of 144 measured points within a volume of $(2.4\text{m})^3$ (see Figure 1), the corrected points are usually within 4cm (≈ 1.5in) of the true points, while some of the measured points (without correction) are displaced by over 40cm (≈ 15in; see Figure 4).

With our algorithm, the error of the corrected points from their true positions does not depend on the distance to the transmitter, nor does it depend on the amount of local "warp" of the magnetic field.

Evaluation of the accuracy of our method was done by moving a sensor along several well-defined straight lines within the measured volume. Then, for all points on the same line, two out of three coordinates of the corrected points should remain constant with a known (true) value. Accuracy is the maximum deviation of these constant values from the true ones.

There are a few factors affecting the quality of the corrected position: accuracy of the measured field data on which the interpolation is based, the sampling density, constancy of the magnetic field, and the parameter R^2 (see below).

We feel that it is very difficult to achieve an accuracy of better than 1cm (0.4in) when acquiring the sample data of the magnetic field, especially when measuring a large volume such as a cave. In order to improve the accuracy of the data, a much more precise way of positioning the sensors within the volume would be needed.

4.2. The optimal parameter R^2

Admittedly, our approach does have one "magic" parameter R^2 in the basis functions $\omega_i(P)$. It has considerable influence on the "smoothness" of the interpolation function (see Figure 3). The bad news is, no simple and robust formula is known to determine an optimal R^2 [11]. It depends on the number of points, the diameter of their circumcircle, and the values $f(P^i)$. [9] have used $R = 1.25\frac{D}{\sqrt{N}}$, with $D =$ the diameter of the circumcircle of all measured points and $N =$ the number of points. [5] proposed an algorithm which produces near-optimal values for R^2.

Therefore, we developed a program to investigate the effects of R^2 interactively. Thus, given a snapshot of the field,

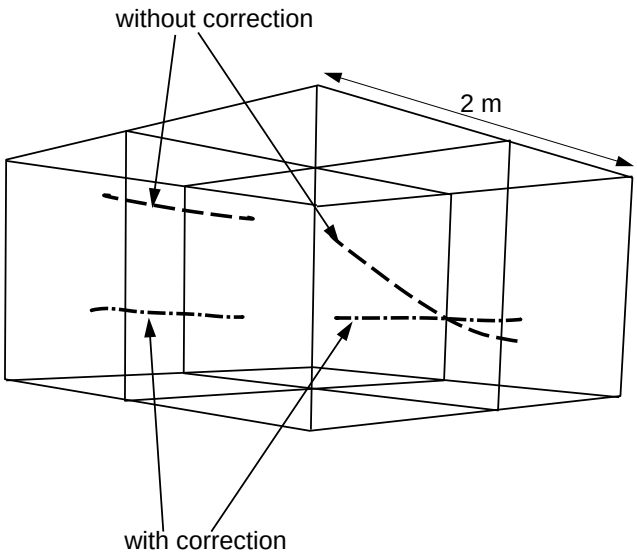

Figure 4. The dashed lines show the locus of points as reported by the tracker when the sensor is moved along a perfectly straight line. The dash-dotted lines show the same lines *with* correction. The error of the corrected positions from the true positions is less than 4cm ($\approx$ 1.5 in). (See also the color plates.)

we can determine an optimal R^2 within a few minutes. The good news is, we found that, with $100 - 200$ measured positions, $R^2 = 10 \ldots 1000$ is optimal, and within that range the exact value has very little impact on the quality of the interpolation.

Of course, it would be quite simple to implement a procedure which calculates an optimal R^2, given a snapshot of the field, a few lines of tracker data, and their true positions.

4.3. Timing

Computation of $f(P)$ is linear in the number of measured samples P_i. Interpolating one 3-dim. position by $f(P)$ with $i = 144$ (i.e., 144 samples) takes 0.4 milliseconds on a 250-MHz R4400 processor.

Since the correction of the field's distortion takes so little time, no additional latency is introduced into the VR system. In fact, a transmission of one complete data record from the tracking system to the host takes much longer.

Solving the three sets of linear equations in order to compute the coefficients A_i is on the order of seconds; for 144 sample points, it takes about 2 sec (250MHz R4400).

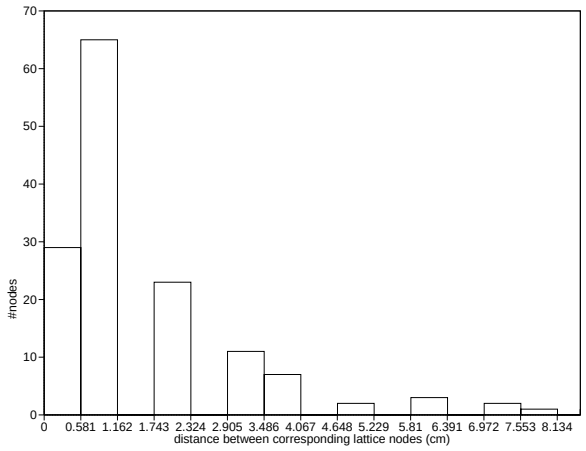

Figure 7. The distribution of the distances between corresponding lattice nodes of two snapshots of the magnetic field on the same site. The second snapshot was taken 6 weeks after the first. Most of the distances are within the measuring accuracy of the data. This verifies the assumption of a static distortion of the magnetic field (at least at our site).

4.4. Repeatability

By repeatability we understand the constancy of the magnetic field's distortion. It is a fundamental assumption in all static correction methods, and the accuracy of the interpolation algorithm can be no better than the constancy of the field.

The good news is that the distortion of the field remains sufficiently constant over time (see Figure 7). On the same site, the magnetic field has been measured by our method, as described in Section 3.2, two times, the second time 6 weeks after the first time. The distribution of the distance between corresponding lattice nodes shows that the distortion of the field does not change significantly over time.

Another measure of constancy is the traditional correlation function. However, it is not clear what a correlation value of $c \in [0, 1]$ maps to in terms of (maximum or average) absolute error (with units of cm or inch).

4.5 Comparison of Polhemus and Ascension

A comparison of the two common commercial magnetic tracking systems, Polhemus's Fastrak™ and Ascension's Flock-of-Birds™, with regard to field distortions was carried out.

We measured the output of a Flock-of-Birds at the same site and at exactly the same grid nodes which have been used for Figure 1. We used the extended range transmitter,

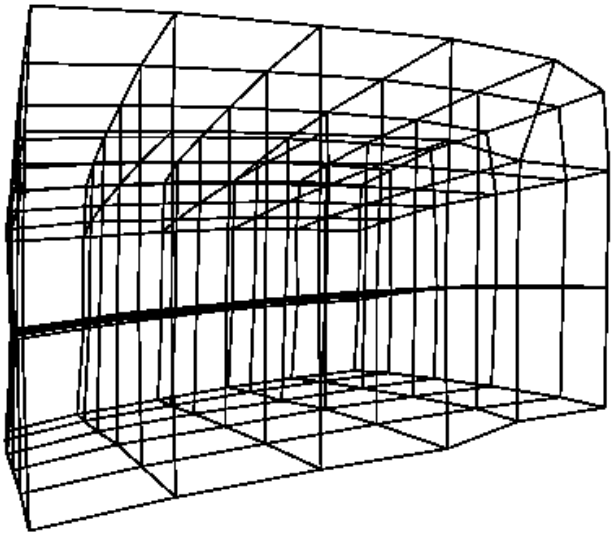

Figure 5. Distortion of the magnetic field using Ascension's Flock-of-Birds. The extended-range transmitter is in the middle of the front right edge.

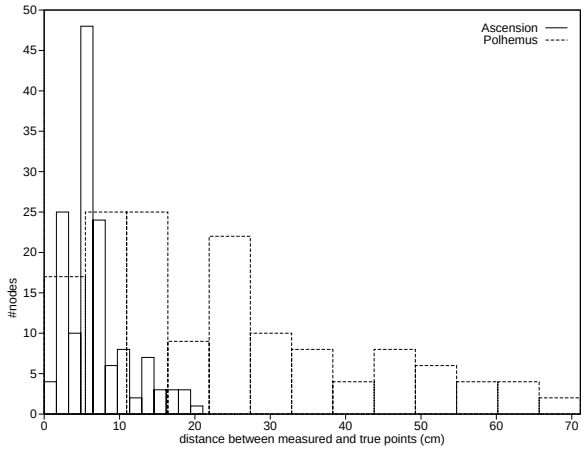

Figure 6. A histogram of the distances of the measured positions from the corresponding true positions, for both Polhemus and Ascension.

however, due to site-specific constraints, it was not possible to position the transmitter at the same location where the Polhemus long-range transmitter was positioned to acquire Figure 1.

The result is shown in Figure 5. Although Ascension's system seems to be less susceptible to ferro-magnetic metal than Polhemus', there is still an unacceptable amount of distortion. While tracking with Polhemus in our cave can be 50 cm ($\sim$20in) off the true position, the position reported by Ascension can be 20cm off. The average discrepancy in our cave is about 20cm for Polhemus and about 8cm for Ascension (see Figure 6.

Due to less distortion with Ascension trackers, the final equalized positions are slightly more accurate than the equalized positions from the Polhemus tracker (typically 0.5-1.0 cm). However, the impact of distortion is almost neutralized when using our correction method.

5. Conclusion

We have presented a practical method to solve arguably the severest problem of magnetic tracking systems: distortions of the (electro-)magnetic field.

Our method is robust, i.e., the error of corrected points does not depend on distance from the transmitter, nor does it depend on the amount of local "warp" in the magnetic field. It does not cause any additional latency in the VR system, and it is easy to implement.

With a very simple apparatus, the field of a $(2.4\text{m})^3$-cave can be captured in about half an hour. For our purposes the process and the apparatus for acquiring the data are reasonably precise.

The proposed algorithm and field measuring method reduce the error of a tracking sensor's position to 2-5 cm (1-2 in), which remedies any distortions in the images projected on our cave walls (see Figure 9).

Based on mathematical experiments with "synthetic" field data, we believe that the algorithm can potentially provide even better correction if the acquired field data are more precise. However, it is quite difficult to position the sensors within the cave with a maximum error of less than 1 cm.

5.1. Future Work

To this point of time, we haven't looked at orientations. In a cave or at a workbench, a skewed orientation of the

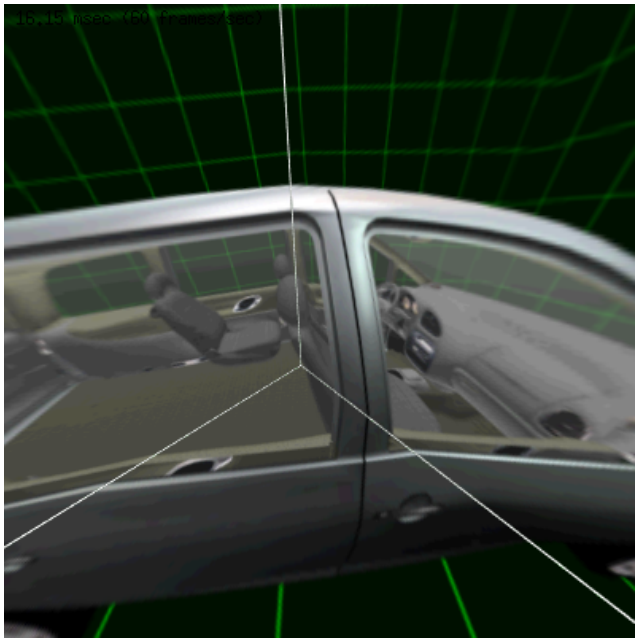

Figure 8. Without correction of the magnetic field tracking, an off-center viewer in a cave will see a distorted image. The effect is even worse when the viewer moves, because objects seem to move, too.

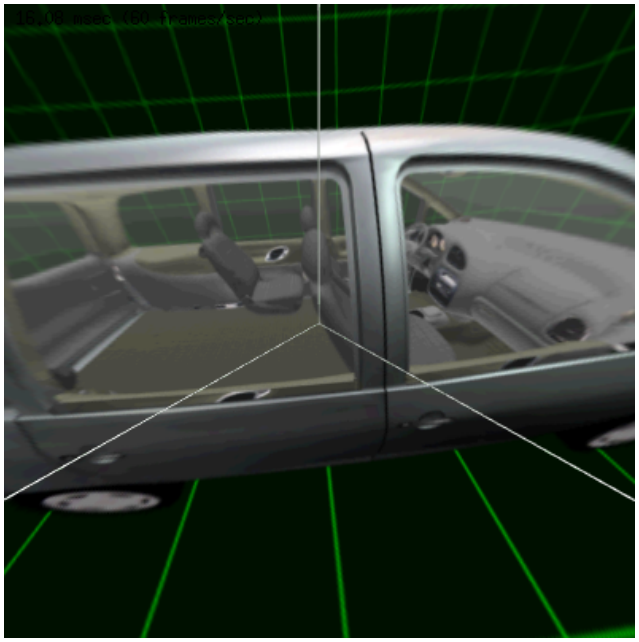

Figure 9. With our method, the perspective is always correct, even when the viewer moves. (See also the color plates.) Data courtesy of Volkswagen AG.

viewpoint is not as fatal as an offset, since the orientation determines only the parallax between the two images for the left and the right eye, respectively However, when tracking the hand of a user who is trying to perform a virtual assembly task, a mismatch between the true and the virtual hand's orientation might lead to confusion and frustration. Therefore, orientations should be corrected as well. This could be done by interpolating positions and orientations simultaneously in 7-dim. space. We still need to investigate whether or not the orientation samples can be obtained with our measuring procedure, or whether they can be derived somehow. Otherwise, sampling the distorted field by position *and* orientation might become a very lengthy and tedious task.

Our apparatus for measuring the magnetic field should be improved in order to acquire data with an error less than 1 cm.

We have not investigated the effect of small "metallic" changes in the environment on the electromagnetic field, such as a chair would produce. It would be interesting to get at least a "qualitative" impression of those distortions.

Finally, our appraoch assumes a static distortion of the magnetic field. However, in set-ups comprising a Boom and magnetic tracking (e.g., Boom plus glove), the magnetic field is changed by the Boom. Unfortunately, the "shape" of the distortion depends significantly on the position and orientation of the Boom. It is not clear to us whether or not this kind of dynamic distortion is "repeatable" in the following sense: given a certain position (and orientation, maybe) of the Boom, the measured position of a tracking sensor remains constant over time. If the assumption is true, then this kind of dynamic distortion can be corrected by interpolating a tracker's position in 6D (or 9D) space. If the "shape" of the distortion depends not only on he Boom's position but also on its orientation (which it probably does), the process of measuring the field is probably very time-consuming.

6. Acknowledgements

I would like to take the opportunity to thank Prof. Dr. h.c. Dr.-Ing. J. L. Encarnação and Dr. Stefan Müller for the great work environment. In addition, I would like to thank my friends and colleagues for many fruitful discussions and their help, namely Van Do, Alex delPino, Alexander Rettig.

References

[1] P. Astheimer, W. Felger, and S. Müller. Virtual design: A generic VR system for industrial applications. *Computers & Graphics*, 17(6):671–677, 1993.

[2] K.-J. Bathe. *Finite Element Procedures in Engineering Analysis.* Prentice-Hall, 1982.

[3] S. Bryson. Measurement and calibration of static distortion of position data from 3D trackers. In *Siggraph '92, 19th International Conference On Computer Graphics and Interaction Techniques, Course Notes 9*, pages 8.1–8.12, 1992.

[4] P. G. Buning. Numerical algorithms in CFD post-processing. *van Karman Institute for Fluid Dynamics*, Lecture Series:1–20, 1989-07.

[5] R. E. Carlson and T. A. Foley. The parameter r^2 in multiquadric interpolation. *Computers & Mathematics with Applications*, 21:29–42, 1991.

[6] F. Dai, W. Felger, T. Frühauf, M. Göbel, D. Reiners, and G. Zachmann. Virtual prototyping examples for automotive industries. In *Proc. Virtual Reality World*, Stuttgart, Feb. 1996.

[7] N. Dyn. Interpolation of scattered data by radial functions. In F. I. U. C. K. Chul, L. L. Schumaker, editor, *Topics in Multivariate Approximation*, pages 47–61. Academic Press, 1987.

[8] T. A. Foley and G. M. Nielson. Comparing methods of interpolation for scattered volumetric data. *Eurographics State of the Art Reports*, pages 39–59, 1994.

[9] R. Franke. Scattered data interpolation: test of some methods. *Math Computation*, 38:181–200, 1982.

[10] L. F. Hodges and E. T. Davis. Geometric considerations for stereoscopic virtual environments. *Presence*, 2(1):34–43, winter 1993.

[11] J. Hoschek and D. Lasser. *Fundamentals of Computer Aided Geometric Design.* A K Peters, 1993. ISBN 1-56881-007-5.

[12] N. Magnenat-Thalmann and D. Thalmann, editors. *Realism in Virtual Reality*, pages 189–210. Wiley & Sons, 1994.

[13] G. M. Nielson and J. Tvedt. Comparing methods of interpolation for scattered volumetric data. *Siggraph '94, Course Notes 4*, pages 99–123, 1994.

[14] T. Oishi and S. Yachi. Methods to calibrate projection transformation parameters for see-through head-mounted displays. *Presence*, pages 122–135, winter 1996.

[15] W. H. Press, B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling. *Numerical Recipes in C.* Cambridge University Press, 1988.