Constellation™: A Wide-Range Wireless Motion-Tracking System for Augmented Reality and Virtual Set Applications

Eric Foxlin*, Michael Harrington, and George Pfeifer

InterSense Incorporated

Abstract

We present a new tracking system for augmented reality and virtual set applications, based on an inertial navigation system aided by ultrasonic time-of-flight range measurements to a constellation of wireless transponder beacons. An extended Kalman filter operating on 1-D range measurements allows the inertial sensors to filter out corrupt range measurements and perform optimal smoothing and prediction, while at the same time using the pre-screened range measurements to correct the drift of the inertial system. The use of inside-out ultrasonic tracking allows for tetherless tracking over a building-wide range with no acoustic propagation latency. We have created a simulation to account for error sources in the ultrasonic ranging system. The fully implemented tracking system is tested and found to have accuracy consistent with the simulation results. The simulation also predicts that with some further compensation of transducer misalignment, accuracies better than 2 mm can be achieved.

CR Categories and Subject Descriptors: I.3.6 [Computer Graphics]: Methodology and Techniques - Interaction Techniques I.3.7 [Computer Graphics]: 3-Dimensional Graphics and Realism- Virtual reality; I.3.1 [Computer Graphics]: Hardware Architecture-Input devices.

Additional Keywords: motion tracking, inertial, ultrasonic, kalman filtering, augmented reality, virtual sets, accuracy, latency, sensor fusion

1. INTRODUCTION

There is an ever expanding set of interactive graphics applications which require smooth and fast free-space tracking of some part of the user's body, or some hand-held object. Head-mounted displays (HMDs) for immersive virtual environment simulations have stimulated a tremendous amount of activity since the early 1990s. Many virtual prototyping systems were developed, often using "goggles and gloves" for interaction. While the media has been distracted by the new phenomenon of the world-wide web, virtual environment technology has made great strides, especially in the area of real-time rendering on affordable hardware, and has been silently catapulted out of the laboratory and into real-world applications.

73 Second Ave., Burlington, MA 01803, ericf@isense.com

To appear in Proceedings of SIGGRAPH 98 (Orlando, Florida, July 19-24, 1998) Computer Graphics Proceedings, Annual Conference Series, ACM SIGGRAPH

Recently, there has been considerable interest in wearable Augmented Reality (AR) systems and virtual set generation for television studios. While these seem to present fairly dissimilar tracking problems (tracking a headset v. tracking a camera), they both require a long-range tracking solution with very high accuracy that will work reliably in an uncontrolled environment full of interference sources. The most immediately promising applications for AR seem to be wearable or mobile computers to assist workers in assembly or maintenance of complex machinery from aircraft [18,19] to buildings [6] to human patients [21]. In the case of assembling wire bundles for aircraft, the workpiece may be over 100 feet long, and the AR tracking system must operate over this span with undiminished performance. Likewise, in the virtual studio it is necessary to track a camera which is being carried about freely in a very large space full of metal, electronics and bright dynamic lighting. In addition to long range and difficult operating environments, both applications share the need for tracking accuracy sufficient for visual registration of computer generated and real objects. Tracking is an urgent unsolved problem for these two applications. This paper is an effort to address it.

With such a plethora of different graphics applications that depend on motion-tracking technology for their very existence, a wide range of interesting motion-tracking solutions have been invented and brought to various stages of maturity over the years. Surveys of the myriad magnetic, optical, acoustic, and mechanical tracking systems are available in [2,7,17]. Many HMD applications only require motion over a small region, and these traditional tracking approaches are usable, although there are still difficulties with interference, line-of-sight, jitter, and latency. We have previously described an alternative solution based on inertial sensing technology with automatic drift correction [9] which overcomes the problems with interference, line-of-sight, jitter and latency. In fact, that drift-corrected inertial tracking system is sourceless and operates over an unlimited range. However, it is only able to track 3-DOF orientation. To correct positional drift in a 6-DOF inertial tracking system requires some type of range or bearing measurements to fiducial points in the environment.

In this paper we present a new tracking system concept, a working system based on this concept, test results and a demonstration of the capabilities of this system in a mock virtual set camera-tracking application. The new concept is an extension of our previous work on inertial orientation tracking technology. The inertial tracker provides a self-contained orientation tracking system with unlimited range which does not suffer from the drawbacks associated with source-based or mechanically-linked tracking systems. It also contains triaxial accelerometers which are double integrated to obtain changes in position, relative to a known starting position. The double integration leads to an unacceptable rate of positional drift and must be corrected frequently by some external source.

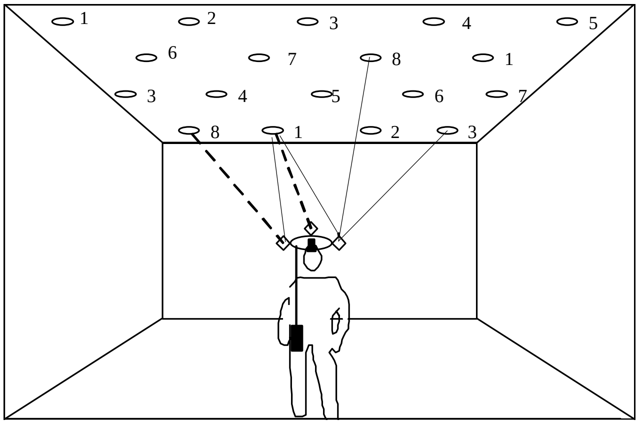

Figure 1: General idea of the Constellation™ system

The CONSTELLATION tracking system is similar in its basic principles of operation to an aided inertial navigation system (INS), except that it operates indoors, has much finer resolution and accuracy, and uses acoustic rather than RF technology for range measurements. Figure 1 illustrates the system, configured for tracking an HMD in a wide-range VR or AR application. The HMD is equipped with an integrated inertial sensing instrument called the InertiaCube™ and, in this example, 3 ultrasonic rangefinder modules (URMs). The rangefinder modules communicate with a constellation of transponder beacons which may be mounted at any known locations in the environment. The beacons are activated one-at-a-time by infrared trigger codes emitted by the rangefinder modules. As each beacon receives its own unique code, it responds by emitting an ultrasonic pulse. The rangefinders count the time-of-flight (TOF) until the pulse arrives, and use the speed of sound to convert the TOF into a distance. These range measurements are fed into an extended Kalman filter (EKF) which makes small adjustments to the position and orientation trajectory which is being update at a high rate by the strapdown INS. At least 6 range measurements, connecting between at least 3 HMD-mounted microphones and at least 3 fixed transponder beacons, are required to completely determine the position and orientation of the HMD. Figure 1 shows an example of 6 suitable ranges, which illustrates that multiple nearly simultaneous measurements from each triggered beacon can be used if available, but are not required. Two degrees of freedom can be resolved by stabilizing with respect to gravity, so only 4 of the myriad potential lines-of-sight need to be open to continue tracking indefinitely, and fewer than 4 can be sufficient to sustain reasonable tracking for a while.

We believe this new tracking system architecture has several compelling advantages:

- It is simple and practical compared to other systems with scaleable-range capabilities
- It is possible to wear the whole tracking system, including all of the sensors and the computational unit. This results in a tracker that is completely untethered.
- It is inertial sensor-based, conferring high update rates and superior smoothness and predictive capability. It can withstand the loss or corruption of a large portion of its acoustic range measurements without a significant degradation in performance.
- The acoustic ranging system is inside-out compared to other acoustic trackers. Since the sound waves propagate spherically *from* the fixed beacons *to* the moving target, the TOF recorded at the moment of detection represents the instantaneous radius measurement with no latency.

1.1 Previous Work

We are not the first to brave the design of a scaleable-range tracking system. A system called the optical ceiling tracker has been in development for many years at UNC-Chapel Hill [22]. It uses a cluster of head-mounted cameras looking at an array of computer-controlled infrared-emitting diodes (IREDs) mounted in ceiling tiles. Although it is an optical tracker and ours is a hybrid acousto-inertial tracker, both systems are based on an array of fiducial markers on the ceiling and designed to offer the same advantages of high accuracy, potentially limitless range, and relative immunity to occlusion through redundancy. Further, both use extended Kalman filtering algorithms to process single measurements at a time [23]. Another optical constellation-based approach was recently proposed [14] which makes use of quadcells instead of lateral effect photodiode cameras. Quadcells are extremely simple and inexpensive optical direction-sensors which eliminate the need for lenses and the weight and optical distortion they introduce. However, quadcells detect the direction to a light source based on the ratios of the illumination received on each of four photocells, and these ratios may be affected by both diffuse and specular reflection of the LED beacon strobes off of various surfaces.

There are a variety of reasons why we chose to employ acoustic range-finding instead of optical bearing-angle measurement to correct the positional drift in our system:

- It requires no head-mounted cameras, only a few tiny ultrasonic microphones, leading to lower weight, power consumption and cost.
- The orientation is already available from the inertial system. The cumbersome head-mounted camera approach was developed to achieve superior orientation resolution. The simpler acoustic and outside-in optical trackers are sufficient for tracking position even though they are not very precise for orientation.
- The mathematics are simpler. Three range measurements pin down the position. Six bearing angles (normally measured two at a time) are required to solve for position and orientation.
- Microphones are available with very wide fields of view compared to cameras. Thus it is possible to use fewer beacons in the constellation and still be sure there will be several redundant lines of sight available.

In addition to the two aforementioned constellation-type tracking systems, there has been much previous work on inertial and acoustic technologies. At least three authors have exploited the motion derivatives provided by inertial sensors to add prediction capability to HMD tracking systems [1][4][15]. In the navigation arena, the aided inertial navigation approach used in this paper has been well known, and a wide variety of radio-frequency navigational aids have been used, including LORAN, OMEGA, radar, GLONASS, and GPS for maritime and aviation applications, as well as star-trackers for space navigation.

Finally, ultrasonic time-of-flight ranging techniques have been used in numerous commercial products for 3-D motion tracking (Logitech 6-D Mouse, Mattel Power Glove, Lipman VSCOPE, Kantek Ringmouse). In particular, the Lipman VSCOPE and Kantek ring-mouse have wireless infrared-triggered transponders. Also, [20] describes a large-volume extension of the Logitech device in which, based on the current position of the tracked object, the nearest of a number of switchable reference triangles is automatically selected and used.

1.2 Contribution

This paper contributes the following new concepts and results in motion tracking for interactive graphics:

- A novel acousto-inertial hybrid tracking approach and a working system. (Demonstrated on video performing an uninterrupted tracking sequence spanning several rooms.)
- The first TOF motion tracker with latency less than the flight time of the ranging signals. This is possible due to the unique inside-out configuration of the transmit-receive pairs, which in turn is possible because the use of inertial tracking allows for processing of non-simultaneous range measurements.
- An example of the usefulness of single-constraint-at-a-time Kalman filtering for designing robust sensor-fusion based motion-trackers.
- An analysis of the tracker's Geometric Dilution of Precision (GDOP) and simulation results to understand its sensitivity to systematic error sources.

2. SYSTEM DESCRIPTION

2.1 Hardware Overview

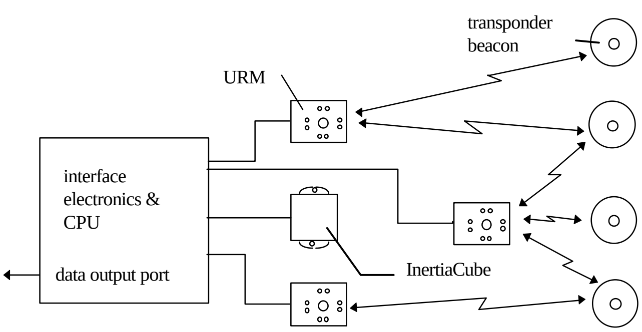

Figure 2:Schematic overview of hardware

Figure 2 illustrates the main hardware components of the tracking system. Just as GPS has a space-based constellation of satellites and a vehicle-born receiver with antennae, this system has a ceiling-based constellation of transponder beacons and a camera- or person-worn tracker unit with ultrasonic rangefinder modules (URMs) and an InertiaCube™ inertial sensing device.

Figure 3 shows a diagram and a photograph of the InertiaCube integrated inertial sensing device manufactured by InterSense for this and related applications [11]. The InertiaCube senses angular rate about and linear acceleration along each of three orthogonal body axes, as illustrated in Figure 3. A portion of a floppy disk is visible in the photograph to highlight the InertiaCube's compact dimensions: 2.7 X 3.4 X 3 cm.

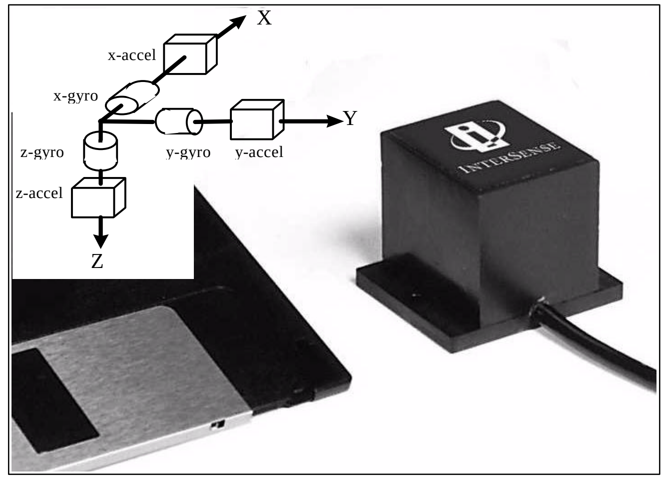

Figure 3: Schematic and photograph of InertiaCube

Each URM consists of a 40 kHz ultrasonic microphone, 4-8 infrared emitting diodes (IREDs) and the necessary electronics, as illustrated in Figure 4. It is not necessary for the IREDs and microphones to be physically mounted together, but it makes logical sense since a blocked line of sight between a beacon and a microphone makes it futile to trigger that beacon.

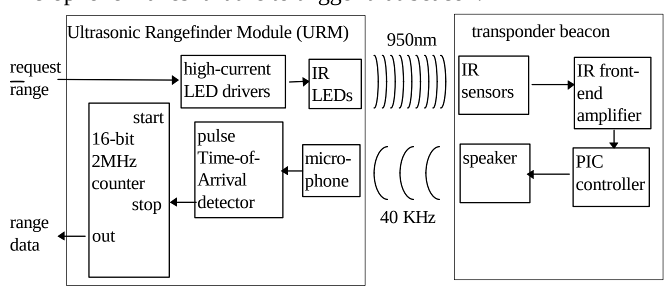

Figure 4: URM and transponder block diagrams

2.2 Software

The tracker system software has two major tasks: acquisition, and tracking. Acquisition occurs whenever the powered-up tracker enters or re-enters a room that has transponder beacons. The purpose of acquisition is for the tracker to determine its initial estimate of position and orientation (a.k.a. pose) so that the tracking algorithm can begin the process of recursively refining and updating the pose estimates. There are currently only eight differently coded beacon types (we are working on increasing this). The acquisition algorithm works as follows:

1. Identify the nearest 4 beacons.
2. Search throughout the entire constellation (which must be pre-known to the tracker) for places that have this combination of beacons in close proximity. Test each such hypothesis to see if a self-consistent trilateration solution can be found using the actual range measurements with the hypothesized beacons .

3. If only one combination of beacons passes the test, use the starting pose determined in this self-consistent trilateration of all three microphones and move on to tracking.
4. If there are multiple 4-tuples in the constellation which are consistent with the initial set of range measurements, try to use range measurements to other beacons to resolve ambiguities.

Note that with only 8 different beacon codes, there are only C(8,4) =70 different combinations, which means that a large constellation would have a lot of repetitions of the same group of four adjacent codes. To overcome this we are increasing the number of beacon codes to 16, which would provide 1820 unique 4-tuples. Even larger constellations would require a different scheme using zone codes and specific beacon codes, because the acquisition time to sequence through more than 16 beacons would be too long.

Once there is a successful acquisition, the state and covariance matrix of the EKF get initialized and tracking begins. Figure 5 illustrates the tracking algorithm. The most important point to note is that the integrated inertial sensors have direct feed-through to the outputs, which insures low latency. The angular rates measured by the gyros are integrated once to obtain orientation, which is output directly. The orientation is also used to transform the accelerations measured by the accelerometers in the constantly changing body-referenced frame into a steady and level navigation frame (hereafter "nav-frame" or N-frame) with its z-axis vertical. The unwanted effect of gravity on this virtual z-accelerometer is first canceled, and then the nav-frame acceleration is double integrated to obtain position, which is output directly. The EKF uses the range measurements to estimate the amount of accumulated error in the orientation, gyro biases, position, and velocity. It applies these error estimates immediately to the appropriate integrator outputs as tiny corrections which prevent the accumulation of error and insure that the EKF is always linearizing about the most accurate possible state. Complementary Kalman filtering is discussed in [3], and the details of our complementary EKF approach are provided in [10].

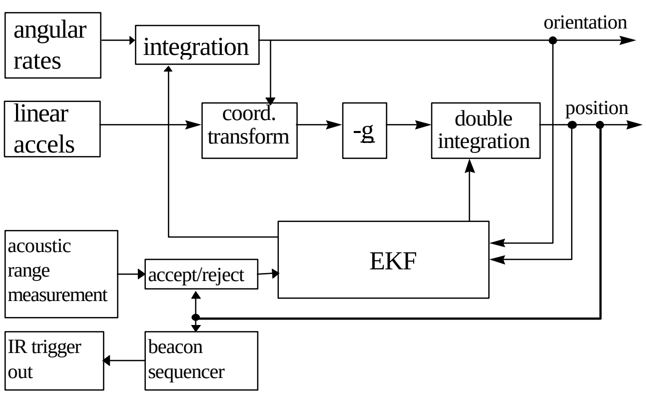

Figure 5: Tracking algorithm flow chart

The selection and utilization of the range measurements is very interesting and deserves some elaboration. First of all, the tracker makes immediate use of individual range measurements as they come in, rather than saving up measurements from 3 beacons, performing a trilateration and feeding the computed position into the Kalman filter as a measurement vector. This technique of processing several scalar measurements instead of one vector measurement is called sequential update Kalman filtering, and it is known to be both faster and numerically more robust because it avoids the matrix inversion in the Kalman gain update step. When measurements containing only partial information about the state are applied during different update cycles, the process has been called single-constraint-at-a-time (SCAAT) tracking [23]. In non-inertial tracking systems, this allows the tracker to have partial updates at a higher update rate, resulting in lower latency and jitter. In our aided inertial design, the pose output gets essentially complete updates (with a little drift) at a high rate of about 500 Hz, but it is still more convenient to make partial drift correction updates immediately upon receiving each range measurement, because at this time an accurate measurement residual can be formed by differencing the measured range and the predicted range computed using the most recent inertial state update.

Secondly, when the tracker receives a new range measurement, it already knows where it is, and it also knows, based on the diagonal elements of the error covariance matrix, approximately how much uncertainty there is in this self-position estimate. Since it knows the location of the beacon that sent the pulse, it can predict what the range measurement should be. If the range measurement doesn't match within the tolerance computed from the covariance matrix, it can be rejected. This is an extremely useful feature in an acousto-inertial tracker. Acoustic range measurement devices always detect the first arrival: a pulse is sent and a counter is started. Since the direct pulse arrives before its echo, the counter is stopped by the first detected pulse at the receiver. Unfortunately, there is occasionally a random background noise or an echo from a previous sampling period which arrives before the real pulse and stops the counter. In our system, we know when to expect the real pulse and can gate the receiver open only during the window of time when the returned pulse is expected. This can likely prevent over 90% of premature pulse detection problems. Because we use the diagonal elements of the covariance to dynamically adjust the acceptance time window, if the tracker misses some measurements it will widen the window and accept subsequent measurements to bring it back on course, instead of becoming completely lost.

3. Constellation Geometry and Error Sensitivity

The constellation may be set up in many geometrical configurations in order to adapt to different types of surroundings. This invites the questions 1) What geometry will result in the highest tracking performance? , 2) What will the performance be for some particular geometry? and 3) How many transponders are really needed? In this section, we develop a simulation to evaluate the sensitivity of the position and orientation calculations to all the known error sources, both random and systematic.

A standard metric used in GPS and other range-based position location systems to evaluate the effect of geometry on positioning uncertainty is the Geometric Dilution of Precision (GDOP):

$$GDOP \equiv \frac{\sqrt{\sigma_x^2 + \sigma_y^2 + \sigma_z^2}}{\sigma_r} = \frac{\sqrt{tr(P)}}{\sigma_r}, \qquad (1)$$

where P is the error covariance of the position solution [16]. This expression is a function of position and it describes, at a given point (x,y,z), how much positional uncertainty there will be if all of the range measurements have the same uncertainty of σ_r. The GDOP is useful for estimating the amplification of random noise in the range measurements, but there are other systematic error sources which may be present:

1. error in the beacon positions
2. temperature error
3. constant time-delay errors in beacons (due to part-to-part variation or electronics drift)
4. constant time-delay errors in URMs (due to part-to-part variation or electronics drift)
5. transducer angle related errors

We have developed a simulation in MATLAB to probe the sensitivity to all these error sources for any desired geometry. The simulation allows the user to enter magnitudes for all of the above systematic error sources, and any desired constellation geometry, then it computes the resulting systematic error and GDOP at a sampling of points within a user-defined test volume.

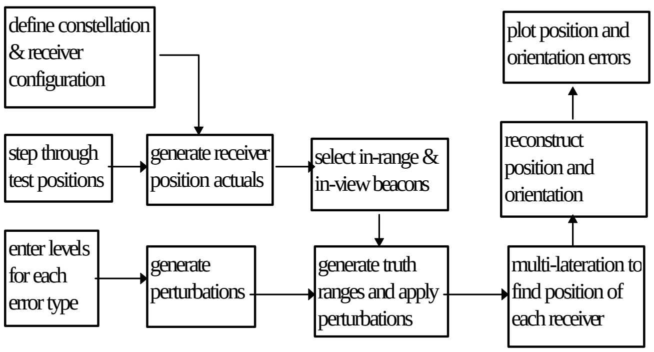

Figure 6: Error sensitivity simulation

Figure 6 shows a block diagram of the simulation. The main flow of the simulation is as follows. First the user is given the opportunity to set up a "trial" constellation for evaluation, and to specify the desired tracking region beneath the constellation. A random horizontal and vertical error are then applied to each beacon in the constellation, uniformly distributed over user-specified intervals. In the main loop of the simulation, the program steps through test positions within the desired tracking volume. At each position it rotates the whole simulated tracker (headset or camera configuration) to a variety of angles, generating the appropriate set of truth positions for all the microphones on the tracker. All beacons which are within range of the tracker are selected using

$$\text{rangelimit} = (5m)\sqrt{\cos(q_t)\cos(q_r)}$$

where $\cos(q_t)$ and $\cos(q_r)$ approximate the off-axis attenuation patterns of our 40 kHz transmitters and receivers. Having selected an active set of beacons, range measurements are calculated, including all appropriate error perturbations, and fed into a multi-lateration algorithm which solves for the positions of all the receivers. There are numerous trilateration and multilateration algorithms in the literature [16][5][13][12]. Although [16] provides an exact closed form solution that is both general and computationally efficient, we chose to employ the classic iterative least-squares approach [12], because it is a closer simulation of the extended Kalman filter used in the tracker. In fact, in the absence of motion, the EKF converges to the same solution as the recursive least-squares approximation used here [12].

Figures 7 and 8 display some simulation results for a constellation which consists of an infinite square array with 2 foot (61cm) grid-spacing and 3 meter height. In all cases, the test volume extended from 0-2.5 meters in the z (height) dimension, and 0-1 foot in the x and y dimensions. Due to the symmetries of an infinite square array, any (x,y) point is equivalent to some point inside of this single-quadrant test region. Therefore, the range of errors displayed in these volumetric visualizations represents the whole range of errors that a tracker would experience over any size workspace, as long as it does not approach too closely the edge of the constellation. (We have not simulated edge effects, but would expect higher errors near the edges). The systematic error levels were set for this simulation run to the following values:

error in beacon positions:	+/-2mm horizontal/ +/-4mm vertical, uniform distribution
temperature error:	0.2° C
beacon variations:	+/- 1mm, uniform distribution
URM variations:	+/- 1mm, uniform distribution
transducer angle related errors:	+/-2.5mm range perturbations at 60° off axis

Table 1: Error source inputs for simulation in Fig. 7-8

These error levels were chosen to reflect what we believe to be the actual systematic error levels in the current setup in our laboratory. They result in a combined systematic positional error shown in Figure 7 which ranges from 2.3-4.7 mm from floor level to 2.5 meters height. By contrast, the positional resolution of pure ultrasonic range measurements, shown in Figure 8, is 0.7-1.5 mm in most of the active volume, increasing to 2.5 mm near the floor. These numbers are obtained via Equation 1 from the estimation error covariance returned by the multilateration algorithm which has been fed individual range noise sigmas corresponding to our hardware test results in Section 4.1.

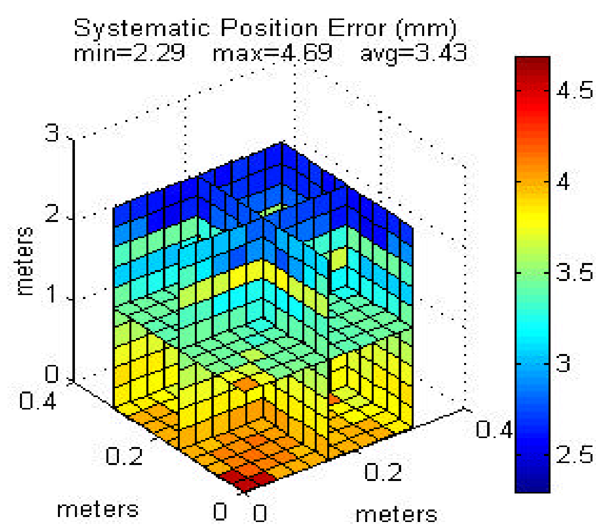

Figure 7: Systematic position errors, assuming error sources listed in Table 1.

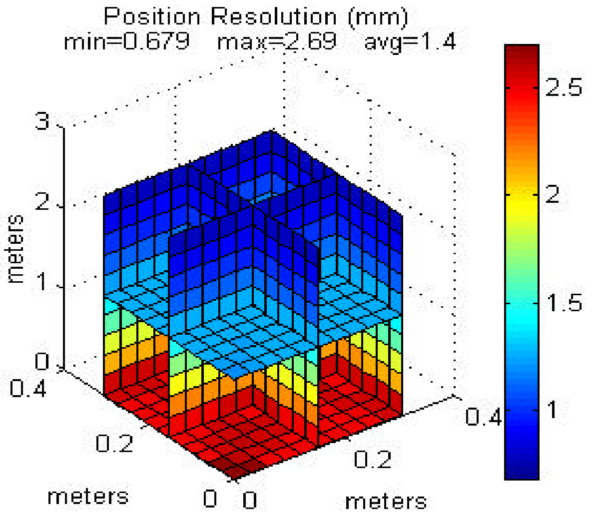

Figure 8: Position resolution

The errors in Figure 7 are largely caused by the transducer misalignment angle errors. We are in the process of developing calibration procedures to better model and compensate for these effects. In Figure 9 we show simulation results predicting approximately 1 mm accuracy when the residual error due to mis-modeled transducer angle effects has been reduced to 1 mm per radian of misalignment, and URM part-to-part variations are measured and compensated out. This excellent accuracy is achieved even in the presence of random beacon placement errors of +/-2 mm horizontal and +/-4 mm vertical, probably because the random errors from the 20-40 beacons participating in the multilateration tend to cancel each other out.

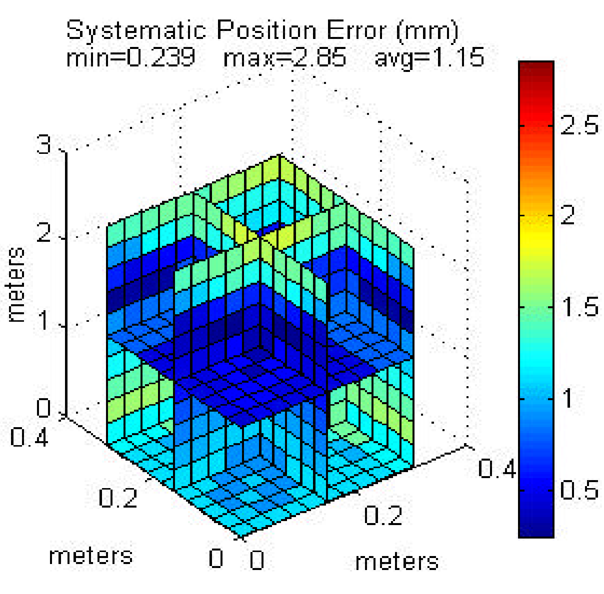

Figure 9: Position errors, assuming improved compensation of systematic error sources

Accuracy of acoustically measured orientation is a function of baseline separation of microphones, so the simulation was run at several different baselines, and the resulting minimum, maximum, and average orientation error throughout a test volume ranging from 1-2 meters below the beacon constellation are plotted in Figure 10. The plotted error is the root-sum-square combination of yaw, pitch, and roll error, with pitch and roll errors truncated to 0.25 degrees because the inertial sensor is able to correct pitch and roll to this level without any ultrasonic aiding [11]. This data is based on the improved error compensation used in Figure 9. As can be seen from the plot, 15 cm of microphone separation, which can be conveniently arranged on an AR head-mounted display, is sufficient to achieve good orientation accuracy. Wider separations, which can easily be arranged on a camera, lead to even higher accuracy.

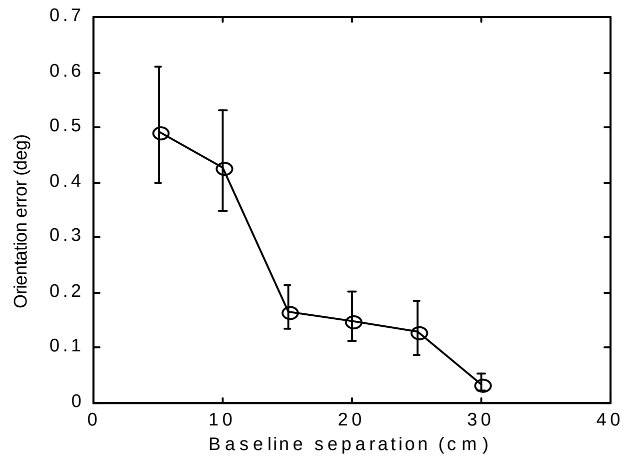

Figure 10: Orientation errors v. mic. baseline separation

4. Test Results

4.1 1-D Ranging Results

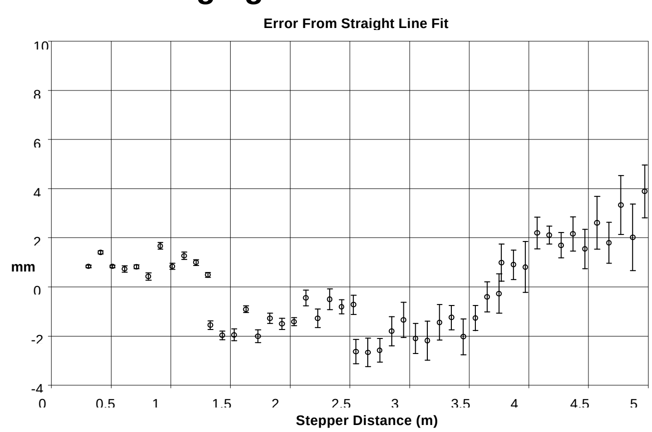

Figure 11: One dimensional ultrasonic ranging results

Figure 11 shows results of testing our ultrasonic ranging hardware prototypes for resolution and linearity. The test was performed using one URM and a transponder mounted on the carriage plate of a leadscrew-driven 4-foot long linear actuator. The rail was moved 4 times to collect approximately 5 meters of data. A single straight line was fit to the entire collective data set, and the residual errors are plotted in Figure 11. The discontinuities in the data were caused by inexact placement of the rail after moving it. Despite these discontinuities, the experiment provides a meaningful assessment of the 1-D ultrasonic ranging performance.

The linearity is approximately 0.1% FS, and the range noise of 1mm per meter of range used in the previous section to generate the GDOP appears justified.

4.2 3-DOF Position Tracking Accuracy

To test the accuracy of the 6-DOF tracking system, we set up a 3 by 3 grid (with one corner missing) of transponder beacons on 2 foot centers on a drop ceiling grid. 1.5 meters below the grid we leveled a table with a 1" grid marked on it, and registered this grid to the constellation coordinates using plumb bobs. A 5-DOF digitizer arm (Immersion Corp.) was placed on this table and registered to its grid with a calibration procedure that involves touching four reference points. The tip of the arm was then attached to a camera-tracker head containing an InertiaCube and 3 URMs separated with horizontal and vertical baseline distances of about 28 cm and 25 cm respectively. The camera tracker was manually moved to 30 locations spaced throughout the test volume reachable by the arm. For comparison to the simulation, Figure 12 plots the root-sum-square (RSS) of the 3 position error components at each point. While the average error is the same as the simulation results in Figure 7, the worst case error is 2 mm larger. This is to be expected because the constellation for this test was equipped with only 8 beacons. When the simulation is run with 8 beacons, it also predicts errors of about 1 to 8 mm.

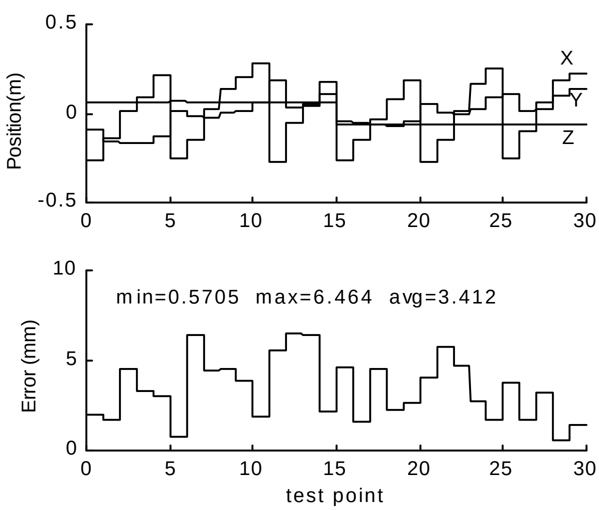

Figure 12: RSS of X,Y, and Z errors at 30 test points

5. Conclusions & Future Work

A new type of motion tracker has been presented which combines inertial orientation and acceleration sensors with ultrasonic ranging devices in a fashion that allows the inertial sensors to filter out corrupt range measurements (ie. echoes and acoustic interference), and perform optimal motion smoothing and prediction, while at the same time using the pre-screened range measurements to correct the drift of the inertial sensors.

Some simulations were performed which indicated that the tracker will be capable of achieving 1-3 mm accuracy levels if used beneath a transponder constellation with 2 ft. spacing that extends 9 ft. beyond the tracking region on each side. Edge effects have not been simulated, but we expect that any dilution of precision near the edges of a room could be overcome by extending the transponder array part-way down the wall.

The prototype ultrasonic distance measuring hardware was tested in a controlled 1-dimensional measurement set-up in order to characterize the error performance and provide meaningful input data for the simulation. The 3-D position locating accuracy was tested in a benchtop configuration with 8 beacons, and found to conform with the simulation's predictions. Finally, the fully functional 6-DOF tracking system was used to record a videotape to demonstrate qualitatively it's resolution, dynamic performance, and range.

The dynamic performance has not yet been tested quantitatively, but in theory there are expected to be no appreciable sources of latency in this system. The range measurements reflect the instantaneous position when received, and the integration of the rate and acceleration data and incorporation of range measurements by means of Kalman filter updates runs at 400-600 Hz. Thus the effective latency of the system is expected to be about 2.5 ms, even though range measurements are received much less frequently.

The simulation indicates that higher levels of accuracy can be obtained by compensating for part-to-part transducer variations. Ultrasonic transducers have an angular dependence which causes a shift of up to 1 mm at 20^o misalignment angle [8], and we find even more at larger angles. We have constructed a computer controlled 2-axis rotation device to characterize this dependency and begun to use the data in the firmware to compensate for the effects.

While the wireless nature of the transponder beacons makes this system easier to set up than other wide-range tracking systems, it is still necessary to accurately measure the beacon locations and download them into the tracking system prior to tracking. This can be time-consuming, and if not done very carefully can become a dominant source of tracking error. In subsequent work, we plan to explore the feasibility of an auto-mapping algorithm that enables a user to install the constellation using the following procedure:

1. A "seed" constellation consisting of 3 rigidly mounted beacons is first hung which establishes the reference frame.
2. The rest of the beacons are hung randomly in any convenient locations. They need not be coplanar.
3. The tracker begins tracking using only the seed beacons. Then it starts trying one additional beacon code at a time until it finds one that responds. The new beacon's position is estimated and entered into the constellation database. As the user walks around the workspace, the tracker finds and auto-installs all the beacons with approximate positions.
4. In a subsequent auto-calibration step, or during normal tracking, there is continuous slow refinement of beacon positions using recursive estimation.

Step 4 has been shown to work in the UNC optical ceiling tracker [23]. It is likely to work here as well if systematic errors other than beacon placement errors have been sufficiently compensated such that the system can track its position to an accuracy substantially better than the placement accuracy of the individual beacons. The simulation results in Figure 9 show that even with random beacon placement errors of +/- 2-4 mm, the tracker is able to find its own position to about 1 mm accuracy. This suggests that auto-calibration in this system may indeed lead to successive refinement of accuracy rather than degradation.

6. Acknowledgments

We are thankful to Wallace VanderVelde for many invaluable insights and lessons on inertial navigation and Kalman filtering, to Gary Bishop, Greg Welch, and David Mizell for enlightening discussions, and to our SIGGRAPH reviewers for many insightful and useful comments and suggestions.

7. References

[1] R. Azuma and G. Bishop. Improving Static and Dynamic Registration in an Optical See-through HMD. In *SIGGRAPH 94 Conference Proceedings*, ACM Annual Conference Series, Orlando, FL, August 1994.

[2] D. K. Bhatnagar, Position Trackers for Head Mounted display Systems: A Survey. *University of North Carolina, Chapel Hill TR93-010*, March 1993.

[3] R. G. Brown and P. Y. C. Hwang. Introduction to Random Signals and Applied Kalman Filtering, 2nd ed. New York: John Wiley & Sons, 1992.

[4] S. Emura and S. Tachi. Compensation of Time Lag Between Actual and Virtual Spaces by Multi-Sensor Integration. In *Proc. IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI 94*), pp. 463--469.

[5] B.T. Fang. Trilateration and Extension to Global Positioning System Navigation. *Journal of Guidance, Control and Dynamics*, 9(6), Nov.-Dec. 1986.

[6] S. K. Feiner, A. C. Webster, T. E. Krueger III, B. MacIntyre, and E. J. Keller. Architectural Anatomy. *Presence*, 4(3): 318-325, Summer 1995.

[7] F. J. Ferrin. Survey of Helmet Tracking Technologies. In *Proc. SPIE*, vol. 1456, pages 86-94, April 1991.

[8] J.F. Figueroa. Ranging Errors Caused by Angular Misalignment Between Ultrasonic Transducer Pairs. *Journal of the Acoustical Society of America*, 87(3), Mar. 1990.

[9] E. Foxlin and N. Durlach. An Inertial Head-Orientation Tracker with Automatic Drift Compensation for use with HMD's. In *Proc. Virtual Reality Software & Technology 94*, G. Singh, S. K. Feiner, and D. Thalmann, Eds. Singapore: World Scientific, pages 159-174, August 1994.

[10] E. Foxlin. A Complementary Separate-Bias Kalman Filter for Inertial Head-Tracking. In *Proc. IEEE VRAIS 96*. IEEE Computer Society Press, March-April 1996.

[11] E. Foxlin, M. Harrington, and Y. Altshuler. Miniature 6-DOF Inertial System for Tracking HMDs. In *Proc. SPIE Helmet and Head-Mounted Displays III*, vol. 3362, Orlando, April, 1998.

[12] W. H. Foy. Position-Location Solutions by Taylor Series Estimation. *IEEE Transactions on Aerospace and Electronic Systems*, 12(2), Mar. 1976.

[13] K.C. Ho and Y.T. Chan. Solution and Performance Analysis of Geolocation by TDOA. *IEEE Transactions on Aerospace and Electronic Systems*, 29(4), Oct. 1993.

[14] D. Kim, S. Richards, T. Caudell. An Optical Tracker for Augmented Reality and Wearable Computers. In *Proc. IEEE VRAIS 97*, p.p. 146-151, IEEE Computer Society Press.

[15] U. H. List. Nonlinear Prediction of Head Movements for Helmet-Mounted Displays. *Air Force Human Resources Laboratory , Technical Paper AFHRL 83-45*

[16] D. Manolakis. Efficient Solution and Performance Analysis of 3-D Position Estimation by Trilateration*.* *IEEE Trans. on Aerospace and Electronic Systems*, 32(4), Oct. 1996.

[17] K. Meyer, H. L. Applewhite, and F. A. Biocca. A Survey of Position Trackers. *Presence*: 1(2), pp. 173--200, 1992.

[18] J. Nash. Wiring the Jet Set. *Wired Magazine*, Oct. 1997.

[19] J. Seagull and M. Beauer. A Field Usability Evaluation of a Wearable System In *Proc. International Symposium on Wearable Computers*, Oct. 1997

[20] H. Sowizral and D. Barnes. Tracking Position and Orientation in a Large Volume. In *Proc. IEEE VRAIS 93*, p.p. 132-139. IEEE Computer Society Press.

[21] A. State, M. A. Livingston, W. F. Garrett, G. Hirota, M. C. Whitton, E. D. Pisano, and H. Fuchs. Technologies for Augmented Reality Systems: Realizing Ultrasound-Guided Needle Biopsies. In *SIGGRAPH 96 Conference Proceedings*, ACM Annual Conference Series, pages 439-446.

[22] M. Ward, R. Azuma, R. Bennet, S. Gottschalk, H. Fuchs. A Demonstrated Optical Tracker with Scalable Work Area for Head-Mounted Display Systems. In *Proc. 1992 Symposium on Interactive 3D Graphics*, Cambridge, MA, March 1992.

[23] G. Welch and G. Bishop. Single-Constraint-at-a-Time Tracking. In *SIGGRAPH 97 Conference Proceedings*, ACM Annual Conference Series.