Miniature 6-DOF inertial system for tracking HMDs

Eric Foxlin, Michael Harrington and Yury Altshuler

InterSense, Inc., 73 Second Ave., Burlington, MA 01803

ABSTRACT

Current HMD applications are hampered by the limitations of head-tracking technologies now in use. Commercially available magnetic, optical, acoustic, and mechanical head-trackers suffer from various problems such as vulnerability to interference, line-of-sight restrictions, jitter, latency, small range, and high cost. This paper presents inertial-sensor-based hybrid tracking technology that was developed to combat all these problems. Two commercially available products, the IS-300 and the IS-600, are described, both based on the same miniature triaxial inertial sensor device. The IS-300 is a sourceless 3-DOF orientation tracker, using gravimetric tilt-sensing to prevent any gyroscopic drift in pitch and roll, and optional geo-magnetic compassing to prevent any gyroscopic drift in yaw. The IS-600 is a hybrid acousto-inertial 6-DOF position and orientation tracking system. It tracks changes in orientation and position by integrating the outputs of its gyros and accelerometers, and corrects drift using a room-referenced ultrasonic time-of-flight range measuring system. This paper overviews the theory of operation of both systems, and reports bench-testing results designed to evaluate the resolution, accuracy, and latency of each system.

Keywords: head-mounted displays, tracking, inertial measurement, gyros, accelerometers, magnetometers, ultrasonic time-of-flight, virtual environment simulation

1. INTRODUCTION

One of the main technological bottlenecks in virtual environment, augmented reality (AR), and teleoperator systems using head-mounted displays (HMDs) is head-tracking. Head-tracker noise causes the virtual scene to jitter, reducing presence and causing simulator sickness. Latency causes the objects in the virtual world to swim during head movements. Range limitations prohibit the use of VR for applications such as mobile operations training or architectural walk-throughs. Distortions of the head-tracker output mapping cause confusion during haptic interaction with physical objects that should match virtual visual objects, and cause visual misregistration of virtual and real objects in AR displays. Other important factors are environmental interference, degree of user encumbrance, convenience of operation, and cost.

In the past, there have been a variety of mechanical, ultrasonic, magnetic, optical, and gravimetric head-trackers available.[1,2,3] Mechanical trackers are capable of very good accuracy, resolution, and interference immunity, but they have extremely limited range and tend to encumber the user. Ultrasonic time-of-flight trackers can have relatively large range at low cost, but the need to wait for echoes to die out before initiating a new measurement can cause low update rates, particularly when tracking large volumes. The most common technology today is magnetic tracking, which is convenient because it does not have the line-of-sight problems of optical and ultrasonic systems. The biggest problems with magnetic systems are distortions caused by metal objects, and a very rapid decrease in accuracy and resolution with distance. Optical head-trackers have been built using cameras or other position-sensitive optical devices to track IRED beacons, and using laser range measurement techniques. Both methods tend to provide reasonable range, accuracy and resolution, but suffer from line-of-sight restrictions, reflections and very high cost. Gravimetric trackers (typically using a fluid-filled inclinometer and a compass) are usually provided as a built-in component of very low-cost consumer HMDs. These systems measure orientation only, are sensitive to magnetic interference, and suffer from a high level of “slosh”, ie. false orientation readings caused by linear accelerations.

This paper describes two commercially available motion tracking systems, the IS-300 and the IS-600, which are different from previous motion trackers in that they rely on micro-machined inertial sensors for their primary measurements. This technology offers several potential advantages:

- Very low latency
- Prediction based on directly sensed motion derivatives
- Superb resolution / negligible jitter over entire range

The orientation tracker (IS-300) and the orientation outputs of the 6-DOF tracker (IS-600) have the following benefits in addition:

- Unrestricted range
- Immunity to all forms of interference
- No line-of-sight problems
- Portability: no source to set up.

In the next section we provide a theory of operation for these two systems, beginning with a brief background review of inertial navigation systems, followed by discussion of some of the challenges of building miniature inertial systems for motion tracking on a human scale. Section 3 delves into specific discussion of the IS-300 and IS-600 systems, including hardware and firmware configurations and operating mode selection considerations and trade-offs. Section 4 shows test results on the resolution and accuracy of the systems in two of their main modes, and the conclusion is in Section 5.

2. THEORY OF OPERATION

2.1 Inertial Navigation Background

The operating principles for measuring orientation and position of a moving body using only gyroscopes and accelerometers have been well established in the field of inertial navigation systems (INS) [4,5,6]. The original navigation systems were built with a gimballed platform that was stabilized to a particular navigation reference frame by using gyros on the platform to drive the gimbal motors in a feedback loop. The platform-mounted accelerometers could then be individually double integrated to obtain position updating in each direction. Most recent systems are of a different type, called strapdown INS, which eliminates the mechanical gimbals, and measures the orientation of a craft by integrating the angular rates from three orthogonal angular rate sensing gyroscopes (hereafter "gyros" or "rate gyros" or "angular rate sensors") strapped down to the frame of the craft. To get position, 3 linear accelerometers, also affixed to orthogonal axes of the moving body, measure the total acceleration vector of the body relative to inertial space. This acceleration vector can be converted from body coordinates to earth coordinates using the known instantaneous orientation of the body determined by the gyros. Position is then obtained by subtracting off the effect of gravity from the measured acceleration and then performing double integration starting from a known initial position. Figure 1 illustrates this flow of information.

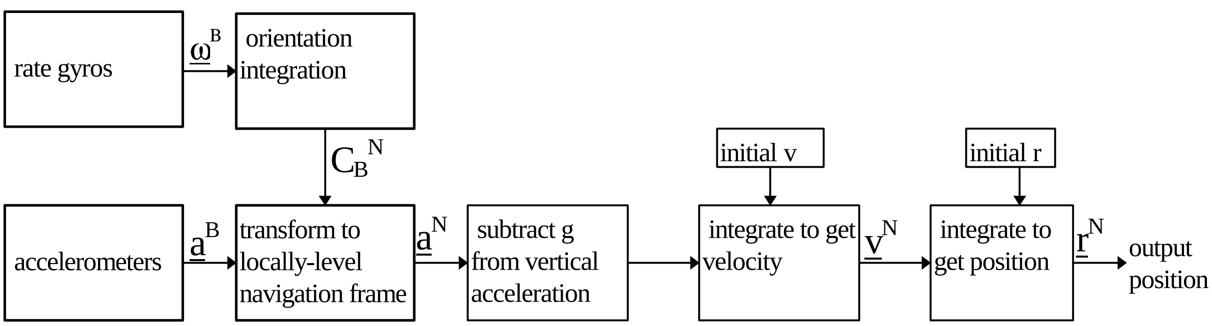

Figure 1: Basic strapdown inertial navigation

Drift in the determination of orientation results mostly from gyro biases, defined as the output produced by a gyro at rest. Fixed biases, if uncompensated, lead to a constant rate of drift after integration. However, the startup biases can usually be measured before take-off and corrected. What matters, then, is bias stability. The typical drift performance of a ring laser gyro (RLG) is about 0.001°/hr, which is sufficient for an INS whose position indication needs to be accurate within about a mile for one hour. Smaller and less costly are fiber optic gyros (FOGs) and dynamically tuned gyros (DTGs), which can achieve drift rates in the .01-1°/hr range, sufficient for short-duration tactical missile flights. Recently, very low cost Coriolis vibratory gyroscopes (CVGs), including micro-machined versions, have been developed for the automotive and remotely piloted vehicle markets. These have drift rates ranging from several degrees per hour to a degree per second.

Drift in the linear position determined by an INS arises from several sources. First, there are accelerometer instrument errors, such as bias stability, scale factor stability, nonlinearity and misalignment. Inertial grade accelerometers, such as those in the 1 mile/hr INS mentioned above, must keep all these errors to a few µg. Considering the 50g maximum accelerations that these accelerometers must contend with, this represents a 10^6 dynamic range! Tactical grade inertial systems can get by with 100µg-class accelerometers. Since position is obtained by double integrating acceleration, a fixed accelerometer bias error a_b results in a position drift error, $x = \frac{1}{2} a_b t^2$, which grows quadratically in time. It is therefore especially critical to accurately estimate and eliminate any persistent bias errors. A second critical cause of error in position measurement is error in the orientation determined by the gyros. Since the INS interprets the direction of the measured acceleration according to the computed orientation of the platform, any error in this computed orientation will cause it to integrate the accelerometers in the wrong direction, thus deviating slightly from the true course of the vehicle. More importantly, the cancellation of gravity will be performed imperfectly by the navigation computer, causing a horizontal acceleration of 1g X sin(error angle) to be erroneously added to the earth-frame acceleration vector. Thus, to take proper advantage of 100μg-class accelerometers, the pitch and roll accuracy must be better than 0.005 degrees for the duration of the flight, which puts a far more difficult task on the gyros than the accelerometers.

2.2 Mechanization for Small Scale Applications

The scale of the head-tracking problem is vastly different from that of global navigation. The area over which tracking is desired is only tens or at most hundreds of feet. Velocities of humans are trivial compared to planes and rockets. The temperature range is only the small variations in room temperature. But the most significant simplification of the problem is the time scale. Uninterrupted mission times of 10 hours are not required; in fact, most HMD users would be happy if they could have high-quality head-tracking operating for 10 minutes before they had to reset the tracking system. Given these facts, one might quickly proclaim the head-tracking problem trivial.

However, there are other scale differences that make inertial head-tracking far from trivial. To guide a plane from one airport to another, an INS only needs to report position within a mile or so (radio-navigation aids can take over within this range). Military applications may have somewhat tighter requirements of hundreds or perhaps even tens of meters, but only for short durations. A VR head-tracker needs a repeatability accuracy of an inch or two, with a resolution of less than a millimeter! For an AR head-tracker the accuracy and resolution must both be under a millimeter. The size and cost, and therefore the performance, of the gyroscopes and accelerometers must also be scaled down tremendously. Micromachined gyroscopes and accelerometers being developed for the automotive industry are the most practical contenders for VR applications, and they only offer bias stabilities of about 0.1-1°/sec and 1-10 mg, respectively. The challenge of using INS technology for human motion tracking is to achieve far tighter positional accuracy using far less accurate sensors, but only for a short time interval.

2.2.1 Design for a 3-DOF Orientation Tracker

A reasonably good solid-state angular rate sensor bias stability of 0.1°/s translates to drift similar to a minute hand. After 10 minutes, the indicated pitch, roll, and yaw values could be wrong by as much as 60°. We have previously presented a solution to this orientation drift problem[7]. A two-axis fluid inclinometer was used to correct drift in pitch and roll, and a two axis fluxgate magnetometer was then used to correct drift in yaw. This basic idea has also been used for mobile robot attitude estimation[8] and for autonomous underwater vehicle navigation.[9] In those two systems, data from the inclinometer and compass was continuously fed back into a Kalman filter attitude estimator. This technique has the drawback that linear accelerations corrupt the performance of the inclinometer, and the Kalman filter unknowingly folds the corrupted data into the orientation estimate. In contrast, our approach took advantage of the burst-like nature of head motion to implement a simple but very effective technique. The inclinometer and compass were ignored during head motion, and consulted during natural head-motion pauses to cancel any orientation drift which has occurred up to that point. This simple algorithm was demonstrated to provide very satisfactory performance when used with high-quality rate gyroscopes having bias stability of about 1-2°/minute. In subsequent work, a concerted effort has been made to get good performance using much smaller gyros. The sophistication of the drift correction algorithms have been increased and calibration techniques have been developed to obtain all the performance possible out of the miniature sensors. The result of this development effort is the IS-300 described in Section 3.2.1.

2.2.2 Design for a 6-DOF Orientation and Position Tracker

In a high-performance INS, the angular error accumulation in the first minute after startup might be only 0.001° X 1/60 = 0.000017°. The error in gravity cancellation would be only sin(0.000017) = 0.29μg, so the dominant error in the short-term determination of position would be accelerometer bias. For our miniature solid-state inertial measurement unit (IMU), the exact opposite is true. If the orientation accuracy is about 0.25° (which is better than most existing head-tracking systems and quite adequate for HMD applications) the erroneous horizontal acceleration will be an unacceptable 4 cm/s^2.

One might ask whether a pure inertial system could ever be manufactured that would be sufficient to solve the 6-DOF head-tracking problem. The best accelerometers under ideal conditions should be able to maintain 1 cm accuracy for about 1 minute. There is an additional source of error that could prevent this performance even if perfect accelerometers existed—gyro drift. Since the orientation angles measured by the gyro subsystem are used to perform the coordinate conversion of the acceleration readings prior to gravity compensation, a pitch error of $\delta\theta$ will cause a component $\sin(\delta\theta)$ of gravity to be subtracted off of a_x and create a false horizontal acceleration in the output of the inertial navigator. To achieve 1 cm accuracy for a minute (assuming perfect accelerometers), the gyros cannot drift more than 1 μrad/minute = 0.003°/hour, which is just achievable by the best gyros. To achieve 1 cm accuracy for 1 minute would then require a state-of-the-art inertial navigation system much heavier than a human head.

Since pure inertial 6-DOF tracking will not be practical for HMDs in the foreseeable future, we set out to develop a complimentary system which can provide position aiding updates while making the least possible compromise to the performance and convenience of inertial tracking. Several factors were considered in selecting an aiding technology:

- Large range is very desirable
- Immunity to persistent distortions like metallic interference is more important than resistance to transient noise
- The system need not provide orientation - we already have highly accurate orientation from the inertial system
- The system need not provide high update rate - the inertial system only needs position aiding about 4-5 times/sec
- Ease of use, ease of setup, and low cost are important

Based on these criteria, time-of-flight (TOF) ultrasonic ranging was selected as the most appropriate option. Section 3.2.2 describes the hybrid acousto-inertial tracking system we developed based on this analysis.

3. SYSTEM DESCRIPTION

3.1 Sensing modules

3.1.1 InertiaCube™ integrated inertial instrument

Both the sourceless orientation tracking and the hybrid 6-DOF tracking product configurations make use of an ultra-miniature smart sensor module called the InertiaCube. The InertiaCube is a monolithic part based on micro-electro-mechanical systems (MEMS) technology involving no spinning wheels that might generate noise, inertial forces and mechanical failures. The InertiaCube simultaneously measures 9 physical properties, namely angular rates, linear accelerations, and magnetic field components along all 3 axes. Micro-miniature vibrating elements are employed to measure all the angular rate components and linear accelerations, with integral electronics and solid-state magnetometers. The geometry and composition of these elements are proprietary, but the functional performance of the multisensor unit can be understood sufficiently by reference to the equivalent diagram in Figure 2.

Figure 3 illustrates the basic physical principal underlying all Coriolis vibratory gyros. Suppose that the tines of the tuning fork are driven by an electrostatic, electromagnetic or piezoelectric drive to oscillate in the plane of the fork. When the whole fork is rotated about its axis, the tines will experience a Coriolis force $\underline{F} = \underline{\omega} \times \underline{v}$ pushing them to vibrate perpendicular to the plane of the fork. The amplitude of this out-of-plane vibration is proportional to the input angular rate, and it is sensed by capacitive or inductive or piezoelectric means to measure the angular rate.

By way of comparison, a conventional inertial measurement unit (IMU) senses 6 of these properties using 6 separate instruments (3 rate gyros and 3 linear accelerometers) each of which by itself would typically be larger, heavier, and more expensive than an InertiaCube. Unlike conventional rate gyro and accelerometer instruments, which must be

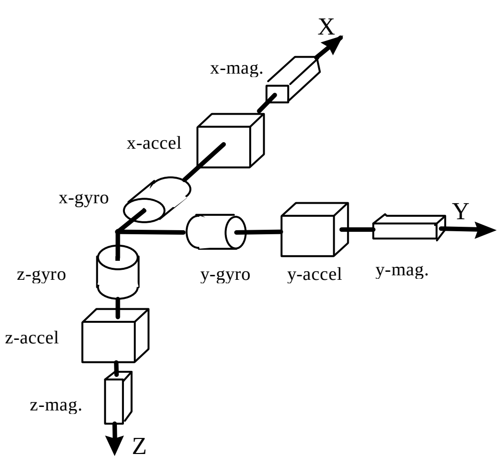

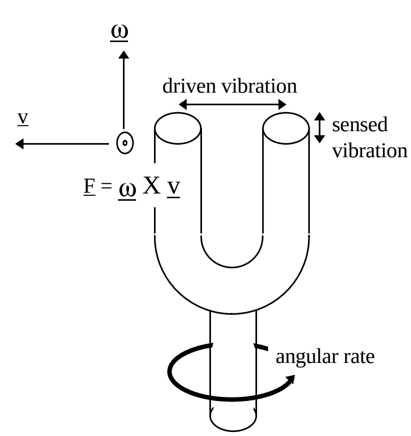

Figure 2: Functional diagram of InertiaCube **Figure 3: Principle of Coriolis vibratory gyroscope**

carefully aligned on a precision machined triaxial mounting block, the InertiaCube is a monolithic device with its orthogonal outputs factory calibrated to precise alignment. Being a digital device, the InertiaCube cabling and connectorization is relatively non-critical, and the cables can be extended to 30 feet or more without fear of contaminating sensitive analog signals. The power consumption of the InertiaCube is 30 mA at 9V, which makes it suitable for prolonged operation from a small battery in future wireless applications. Figure 4 shows an InertiaCube next to a floppy disk for scale.

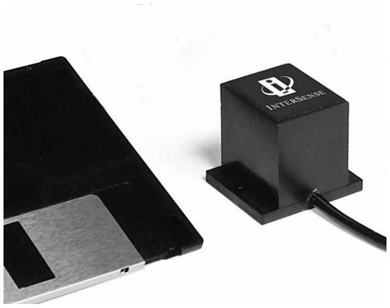

Figure 4: InertiaCube

Figure 5: SoniDiscs

Figure 6: L-Bar

3.1.2 SoniDisc™ remote-triggered acoustic pulse transmitter

The IS-600 uses an acoustic TOF ranging system to prevent positional drift. For maximum accuracy and resolution, acoustic range measurements are made with unidirectional TOF measurements from the SoniDisc transmitters to the ReceiverPods. Transmit/receive mode ultrasonic ranging systems normally require a wire cable connection between the location of the transmitter and the receiver to allow for the synchronization of the pulse emitter and the receiver timer. In order to make the SoniDiscs convenient to use on a human, a wireless remote synchronization method is employed in which the ReceiverPods broadcast an infrared trigger code which uniquely identifies one of up to 8 differently coded SoniDiscs. The selected disc responds immediately with a 40 KHz ultrasonic pulse. Separate timer counters are started in each of the ReceiverPods at the instant of the IR broadcast, and halted by the arrival of the pulse at the ReceiverPod microphone. Making use of the speed of sound (which is calculated from the measured temperature), range measurements are obtained to the three ReceiverPods, and used to compute position as described in Sections 3.3.3 and 3.3.4 below.

The SoniDisc is a battery powered wireless transponder which receives infrared (IR) signals and transmits ultrasonic pulses in response. Both the IR reception sensitivity and the acoustic transmission beamwidth are adjusted to be as wide-angle as possible, and they approximate a hemispherical coverage area so that the SoniDisc can be at an oblique angle of up to +/- 80 degrees from the line of sight to an R-Pod and still communicate with the R-Pod in both directions. The achievable range rolls of gradually with off-axis angle, and at +/- 60 degrees it is approximately halved. Figure 5 displays a photograph of some SoniDiscs.

3.1.3 ReceiverPods and L-Bar support

The SoniDisc's acoustic pulses are detected by three ReceiverPods mounted at the vertices of a large triangle in an L-Bar support structure. The L-Bar provides a fixed frame for positioning the ReceiverPods at precise 39" spacing, as expected by the firmware. It is normally hung from the ceiling over the desired tracking workspace. To achieve a horizontal alignment of the receiver coordinate reference frame, the L-bar can be leveled using adjuster hooks and a spirit level as shown in Figure 6. For the 6-DOF Fusion Mode described in section 3.3.4, this step is required. To increase the flexibility of this system, a future version of the IS-600 is being designed with the ReceiverPods removable from the L-Bar support. The user can then have the option to position the ReceiverPods in any convenient manner, including inside of confined spaces such as automobiles and cockpit simulators. Of course, when this option is used, the user must make the extra effort after installation to measure the cartesian x,y,z coordinates of all the pods and enter this information into the system.

3.2 Product Configurations

3.2.1 IS-300 Sourceless Orientation Tracking System

For some applications, 3-DOF orientation tracking is sufficient. The classic example is the immersive fly-through application, in which a user navigates through a virtual world by looking in the direction in which they wish to move and giving a manual or verbal command to move forward. Many games, visualizations, and vehicle simulators are of this type. 3-

DOF orientation tracking is most effective with large scale virtual worlds, where the objects to be traveled amongst are usually some distance away. When dealing with small scale objects that the user tries to manipulate at arm's distance, the lack of head-motion parallax can feel unnatural.

For applications which can work well with 3-DOF orientation tracking, there are some compelling advantages due to the "sourceless" nature of the Gyroscopic Earth-stabilized Orientation Sensing (GEOS™) paradigm described in section 3.3.1. First, since the sensing is based upon self-contained inertial angular rate sensors, with drift containment based on the earth's gravitational and magnetic fields, the effective range of this tracker is most of the surface of the earth. Secondly, there are no significant interference sources to corrupt the accuracy of gravimetric vertical sensing. Even large mountains only deflect the gravitational vertical by a few arcseconds. Pitch and roll accuracy are therefore immune to any form of environmental interference. This is notable because pitch and roll errors can be perceived in an immersive HMD application, and are therefore more important than yaw errors, which cannot. Due to it's global range, the IS-300 is easily portable. An HMD with an InertiaCube in it can be quickly moved from one workstation to another without the need to also move and set up a source.

In order to realize the benefits of easy portability, substantial effort was invested to miniaturize the processing electronics unit for the IS-300 and IS-300 PRO. Figure 7 illustrates the simple configuration of these products. The dotted lines indicate that the IS-300 PRO can be configured with one to four InertiaCube sensors, whereas the standard IS-300 only supports one.

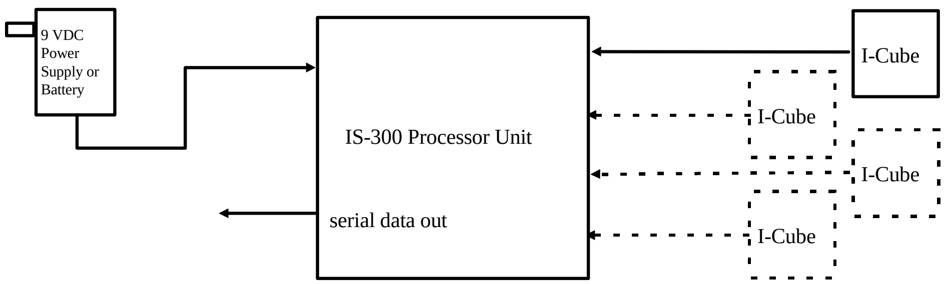

Figure 7: IS-300 HW diagram

3.2.2 IS-600 Acousto-Inertial Hybrid 6-DOF Tracking System

Figure 8 illustrates the configuration of the hybrid acoustic/inertial tracking system. The IS-600 hardware is essentially a superset of the IS-300, adding an ultrasonic range measurement system on top of the basic IS-300 PRO components. The drawing illustrates the IS-600 being used to track a helmet to which are attached an InertiaCube and two SoniDiscs. The rear SoniDisc has just received an infrared trigger code matching its internal ID, and is in the process of emitting an omnidirectional acoustic pulse of 40 KHz energy. The IR trigger code was broadcast widely from 3 IR LED banks located in the ReceiverPods. Only one SoniDisc should be activated at a time.

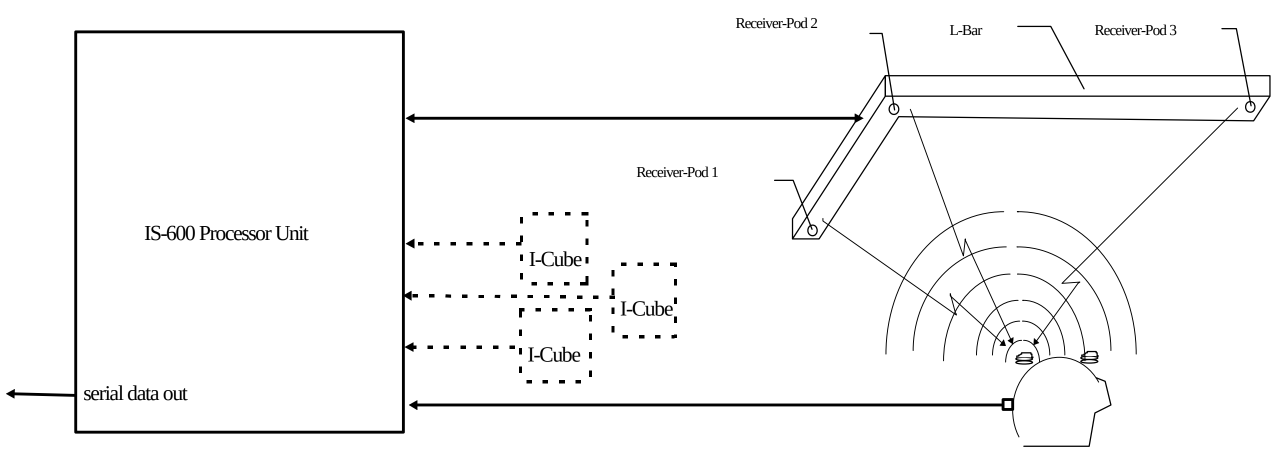

Figure 8: IS-600 HW diagram

The IS-600 has expansion capability up to 4 InertiaCubes and 8 SoniDiscs. The fact that these are packaged as separate modules means that the user has the flexibility to configure the system for a very large variety of tracking tasks. The IS-600 also has several additional operating modes in addition to the GEOS mode of the IS-300. These modes have different strengths and differing configuration requirements, and the following section is an attempt to clarify this potentially confusing aspect of the system.

3.3 Tracking Modes

For a motion tracking system, a **station** means an output data stream corresponding to one particular rigid object being tracked. Some manufacturers refer to such a data channel as a **sensor**, which is perhaps more intuitive for a tracking system in which one physical sensor unit is always used to track one moving object. However, for an IS-600 tracking in a 6-DOF mode, it is necessary to attach multiple physical sensing devices (an InertiaCube and one or more SoniDiscs) to a single tracked entity. We therefore prefer the term station to reduce confusion.

For an IS-600, each station may be configured in one of the four tracking modes described below. Which tracking mode will be used for a particular station depends on what hardware devices are assigned to that station:

hardware assigned to station:	resulting tracking mode:
1 InertiaCube	GEOS mode (3-DOF orientation tracking)
1 SoniDisc	PULSAR mode (3-DOF position tracking)
1 InertiaCube AND 1 SoniDisc	6-DOF Dual mode
1 InertiaCube AND 2 or more SoniDiscs	6-DOF Fusion mode

Table 1: Hardware assignments for each tracking mode

Recalling that the maximum complement of hardware for an IS-600 is up to 4 InertiaCubes and 8 SoniDiscs, an IS-600 could be configured for up to 12 3-DOF stations (4 GEOS and 8 PULSAR), or 4 6-DOF dual mode stations and 4 PULSAR stations, or 4 Fusion mode stations. Since an IS-300/Pro has no SoniDiscs, all of its stations will be configured in GEOS mode by definition.

The following subsections will explain the algorithms of each mode and discuss the characteristics of the mode and trade-offs that can be adjusted with options to that mode.

3.3.1 3-DOF Gyroscopic Earth-stabilized Orientation Sensing (GEOS) mode

Figure 9 shows the processing which is used to compute orientation using this sensor configuration. The basic computation of orientation from gyroscopic angular rates (in the top line of boxes) provides the very rapid dynamic response and high resolution of the system. The accelerometers and magnetometers are used to stabilize the orientation to the earth's gravitational and magnetic fields, thus eliminating the gradual but unbounded accumulation of gyroscopic drift errors. The Kalman filter uses an ever-evolving adaptive algorithm to discard the portion of the accelerometer measurements which are due to actual motion instead of gravity.[10] This is a very important step, because otherwise horizontal accelerations would result in very large transient pitch and roll errors known to technical people as "slosh". The low cost sourceless trackers used in today's consumer HMDs are inclinometer/compass devices, and are thus intrinsically slosh-prone to the point of being uncomfortable to use unless you make a conscious effort to move your head slowly.

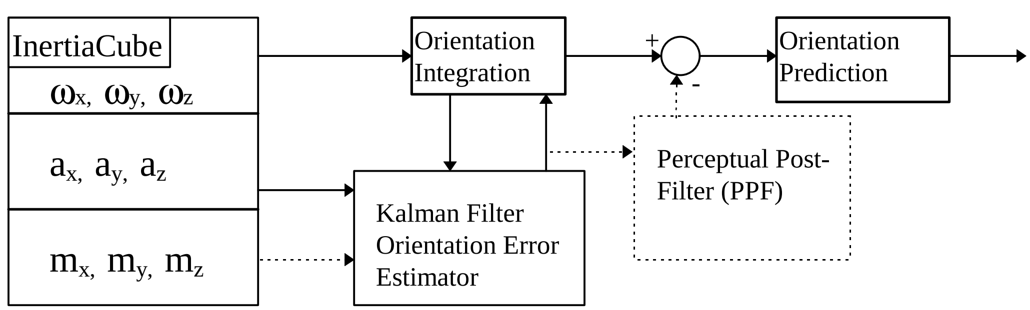

Figure 9: GEOS mode tracking algorithm

In GEOS mode, the reference frame (hereafter referred to as Navigation frame or Nav frame or N frame) is the locally-level geographic frame with its x-axis pointing north, y-axis east, and z-axis down. The Euler angles reported by the tracker can be described as a sequence of rotations applied to the InertiaCube starting with its body axes initially aligned with the Nav frame axes and resulting in the current orientation.

The line from the magnetic field sensor outputs of the InertiaCube to the Kalman filter is a dotted line to indicate that the use of the magnetometers may optionally be disabled. The accelerometer measurements are sufficient to correct all the drift in pitch and roll, and the geomagnetic compassing function is only used to correct drift in yaw. In many fly-through applications absolute yaw referenced to magnetic north is not important and relative yaw tracking is sufficient. This is the case when the user can turn to face an object or rotate the virtual world to bring that object into view. In these situations it may be desirable to turn magnetic yaw compensation off if there are large variations in the direction of magnetic north over the tracking area. With the compassing turned off, the yaw value will drift a few degrees per minute. This drift is too slow too notice while it is happening, but the cumulative yaw error may eventually become noticeable if the user is seated in a fixed chair, and then it may be necessary to send a Yaw Reset command. When yaw compensation mode is disabled, the Nav frame axes are aligned instead to pseudo-north, pseudo-east, and down, where pseudo-north is simply the direction the InertiaCube x-axis was facing on power-up or after a Yaw Reset command.

The Perceptual Post-Filter (PPF) indicated with dotted lines is an optional filter provided for HMD tracking and similar applications which is designed to make the output data minimize perceivable errors rather than minimize mean square errors. Since the human observer is more sensitive to jitter and drift when still or moving slowly, the PPF uses adaptive filtering to preferentially suppress these effects as the head slows down. Interestingly, the PPF only filters the corrections to orientation made by the error estimator, and not the orientation signals themselves. Therefore, the trade-off is increased dynamic error, and not increased latency. If a head at rest suddenly makes a rapid movement, there will be no additional latency imposed by the PPF in the all-important head-motion-to-visual-feedback sensorimotor loop. The PPF is a constantly evolving family of fuzzy rule-based algorithms which have been quite successful in eliminating most perceivable jitter and drift without introducing any latency.

The GEOS mode update rate is the fastest of the four modes. It provides complete orientation updates at 500 Hz on an IS-600 or an IS-300 PRO, and at 150 Hz on an IS-300.

3.3.2 3-DOF PULSed Acoustic Ranging (PULSAR) mode

The second tracking mode available on an IS-600 is the pure ultrasonic TOF trilateration used to track the position of a single SoniDisc, as depicted in Figure 10:

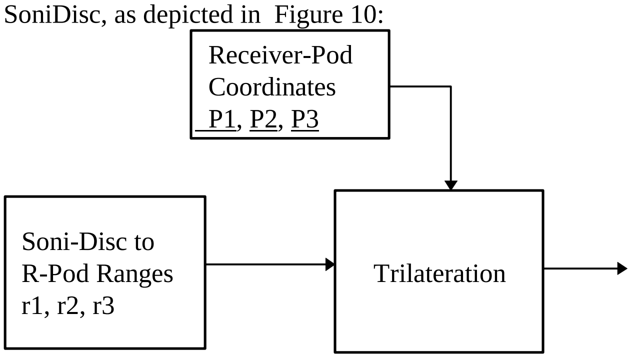

Figure 10: PULSAR mode tracking algorithm

The trilateration algorithm[11] processes 3 measurements to yield a complete recalculation of the SoniDisc location once per pulse. Here lies the major disadvantage of the PULSAR mode; the achievable update rate is limited by the speed of sound (approximately 1 foot per millisecond), the total tracking range and the acoustic reverberation characteristics of the workspace. In most spaces, the maximum update rate without reflection problems is between 20 - 50 Hz. If multiple SoniDiscs are in use, they are time-division multiplexed, and this update rate must be divided by the total number of them. The advantage of this mode is that it is wireless.

The coordinate reference frame for position reporting in PULSAR mode is based on the physical locations of the ReceiverPods, which are in turn governed by the positioning of the supporting L-Bar. The origin is at the location of R-Pod 2, the x-axis points from Pod 2 to Pod 1, the y-axis points from Pod 2 to Pod 3, and the z-axis completes a right-handed triad. For PULSAR mode the L-Bar may be set up in any pose.

3.3.3 6-DOF Dual mode

Six-DOF Dual mode uses the GEOS mode for tracking orientation and the PULSAR mode for tracking position. The InertiaCube and the SoniDisc are tracked completely independently, and inherit all the characteristics described above including the unrelated coordinate reference frames. If the L-Bar is mounted level, with one leg pointing north and the other pointing east, then position and orientation data will be reported in the same coordinate system. This mode is used when you want fast, smooth orientation tracking with prediction, and you also need some positional tracking but position update rate is not critical enough to warrant mounting multiple SoniDiscs on the station.

3.3.4 6-DOF Fusion mode

The 6-DOF Fusion mode works like an inertial navigation system with acoustic range measurements to curtail drift in both position and orientation. Figure 11 illustrates the signal flow in this mode. The same InertiaCube is used as in GEOS mode, but its outputs are processed somewhat differently. The angular rate signals are integrated in a direct manner to obtain orientation, as with GEOS. The linear acceleration signals are used simultaneously for two separate purposes. First, they are transformed from the InertiaCube body frame (B-frame) into a locally-level navigation frame (N-frame). The gravitation constant, g, is then subtracted from the vertical acceleration component in order to cancel the unwanted effects of gravity on the accelerometer triad. Then the remaining acceleration vector is double integrated to track changes in the N-frame position. In a second usage, the B-frame accelerations are also fed into the error estimator to help cancel pitch and roll drift, analogous to the way they are used in GEOS mode. Finally, note that the magnetic field components sensed by the InertiaCube are not used at all in Fusion mode.

Instead, the yaw drift is corrected by the acoustic range measurements simultaneously with the position drift. This is why it is necessary to have at least two SoniDiscs associated with a Fusion mode station. Although what happens inside the extended Kalman filter is actually quite different, the simplified explanation is that the acoustic range measurements can be used to first localize one SoniDisc, then the second, find the vector connecting them and determine a heading direction from it with which to reset the yaw drift. Depending on the baseline separation between the SoniDiscs and the number present, the EKF may also partially or fully override the pitch and roll corrections that the accelerometers make.

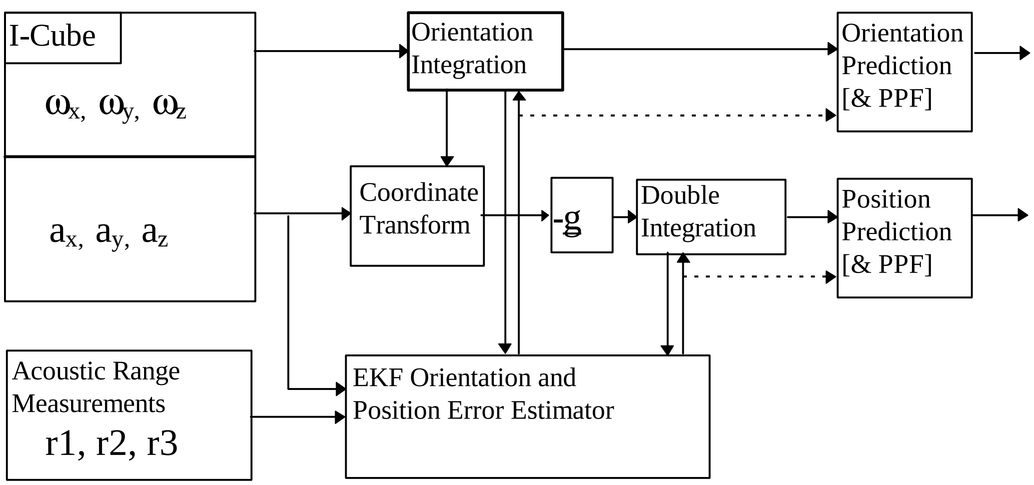

Figure 11: Fusion mode tracking algorithm

The reference frame for both position and orientation tracking is a locally-level frame with the z-axis pointing down. The direction of the x and y axes is determined by the orientation of the L-bar. The L-Bar must be hung level as described in Section 3.1.3 because otherwise there will be conflicting orientation correcting information provided by the accelerometers and the acoustic system, which could lead to output oscillations. If it is not desirable to hang the L-Bar horizontally, then the 3-D coordinates of the R-Pods can be digitized or calculated in a locally level frame N, and downloaded into the tracker. This then establishes N as the navigation frame.

As shown in Table 1, a Fusion mode station must be equipped with an InertiaCube and two or more SoniDiscs. These components are mounted on a rigid object to be tracked, and then the positions of the SoniDiscs with respect to the B-frame axes of the InertiaCube are measured and entered into the tracker. The origin and axes of the tracked station, for the purposes of tracking, will be those of the InertiaCube. At least two SoniDiscs are necessary so that an acoustically-derived yaw direction can be established and used to align the reference frame for the strapdown inertial subsystem with the reference frame of the acoustic range-measuring system. Without this alignment, the accelerometers would indicate translational motion in a conflicting direction from the acoustic subsystem.

Six-DOF Fusion mode takes a little more care to set up, but provides some significant performance advantages:

- Yaw drift correction is accomplished by ultrasonic rather than magnetic means, and it therefore works very accurately even in poor magnetic environments.
- If using 3 SoniDiscs spread out in a sufficiently large triangle, pitch and roll accuracy may be improved compared to GEOS mode.
- Position update rate can be maintained at 150 Hz, even when using multiple SoniDiscs.
- Position tracking resolution and accuracy are superior to PULSAR mode.
- Prediction can be applied to position as well as orientation.
- Position tracking can continue without interruption during brief occlusions of the ultrasonic LOS.

4. PERFORMANCE EVALUATION

In this section we evaluate the performance of an IS-600 system in GEOS mode and Fusion mode. The testing is performed with one particular randomly selected InertiaCube. Therefore this data should not be construed as a specification providing worst-case bounds on the performance of the product, but rather representative sample data.

4.1 Terminology & Methods

A standard measuring system, say a caliper, is characterized by its **resolution** and **static accuracy**. A dynamic measuring system, such as a motion tracker, is additionally concerned with **dynamic accuracy**, and **latency**.

Resolution is the smallest change of the measured property that the system is able to detect. In most systems, the resolution is either limited by the output quantization units, or the output noise level. In the case of most motion trackers, including those discussed in this paper, the resolution is limited by the noise, which is often called **jitter** because of the visual effect the noise has on the computer graphic display. Static accuracy is the amount of output error when the measured properties are held constant. The static absolute accuracy encompasses both non-repeatability and repeatable errors. Non-repeatability is primarily caused by hysteresis and drift. Repeatable error is due to constant input/output distortions such as bias, scale factor error, nonlinearity, and cross-axis sensitivity.

Both jitter and drift are measured simply by placing the tracking sensor set at rest on a non-moving surface and recording the output data stream for a period of time. Jitter is mainly important because of the rapidly shaking imagery it creates, so we exclude the near-DC component of the still-sensor output data stream from the jitter data below. This extreme low-frequency component represents a slowly undulating output error, that we will categorize under the term **stationary drift** or **bounded drift**. The typical drift of integrated rate sensors or accelerometers resembles a random walk process, which tends to grow larger and larger over an extended period of time, and thus must be measured in degrees/second instead of just degrees. This we call **nonstationary drift** or **unbounded drift**, and it is not present in the outputs of our drift-corrected inertial systems, except in the yaw output of GEOS mode when the use of the compass for yaw drift correction has been disabled. The cutoff frequency between jitter and bounded drift is somewhat arbitrary. In keeping with the connotations of the words jitter and drift, we will consider any output motion that is so slow that it cannot be perceived except by waiting and comparing to a memorized previous position to be drift not jitter. We think most readers will agree that angular rotation slower than the minute-hand of a clock (6 degrees/minute) and linear translation slower than 1 cm/minute satisfy this description.

Static accuracy is the root-mean-square (RMS) deviation of reported angles or positions from the true ones when the sensor set is held at a fixed known pose. To measure static accuracy, one needs a "truth" reference that can be relied upon to be considerably more accurate than the system under test. The static accuracy may vary considerably depending on the pose. We therefore desired a truth reference that could be reconfigured to a large variety of known fixed orientations or positions. For angular static accuracy testing, we used a stepper-motor controlled camera pan/tilt unit (PTU) from Directed Perception, Inc. (Burlingame, CA) mounted on a stable leveled tripod. The PTU, shown in Figure 12, is specified to have an accuracy of 0.05°. For position testing, we taped a large sheet of paper with a precise 1 inch square grid plotted on it to the surface of a flat board leveled on a lab bench. This grid was then aligned with the origin and axes established by the L-bar hanging above it to an accuracy of about 0.05 inches by using plumb bobs. A 5-DOF articulated digitizer arm (MicroScribe from Immersion Corp.) was registered to this coordinate frame, and then used for GEOS mode dynamic accuracy testing and for Fusion mode static and dynamic accuracy testing.

Figure 12: PTU for orientation static accuracy

Figure 13: Arm for dynamic accuracy & latency

4.2 GEOS mode performance

Testing was performed on an InertiaCube performing orientation tracking in GEOS mode with the optional PPF turned off (aka jump mode). As shown if Figure 14, this resulted in r.m.s. jitter of 0.021 degrees and a slow bounded drift of about 0.05 degrees maximum in one direction before reversing. Use of the PPF in "smooth mode" would result in substantially less jitter, but the dynamic accuracy would be worsened.

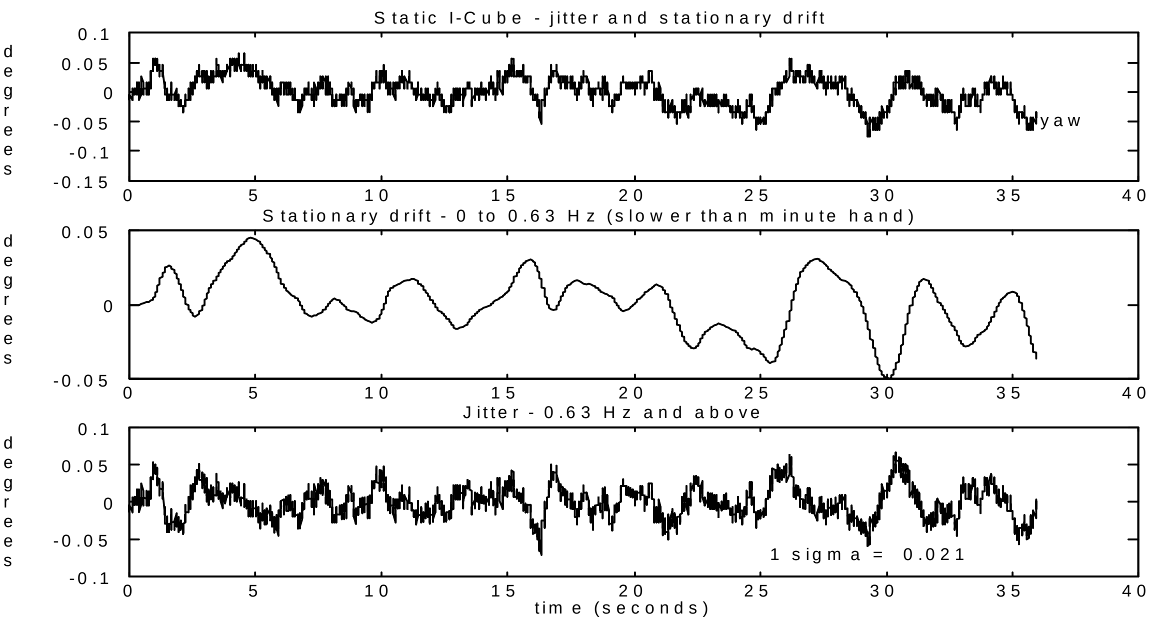

Figure 14: GEOS mode jitter and stationary drift

Figure 15 shows the results of a static accuracy test in which the PTU was rotated to 8 roll angles, and at each roll angle cycled through 6 pitch angles. The InertiaCube was subject to a quick recalibration procedure prior to the test, by placing it on a horizontal surface and removing residual accelerometer biases. The r.m.s. error in both pitch and roll was under 0.25 degrees, averaged over all pitch and roll combinations.

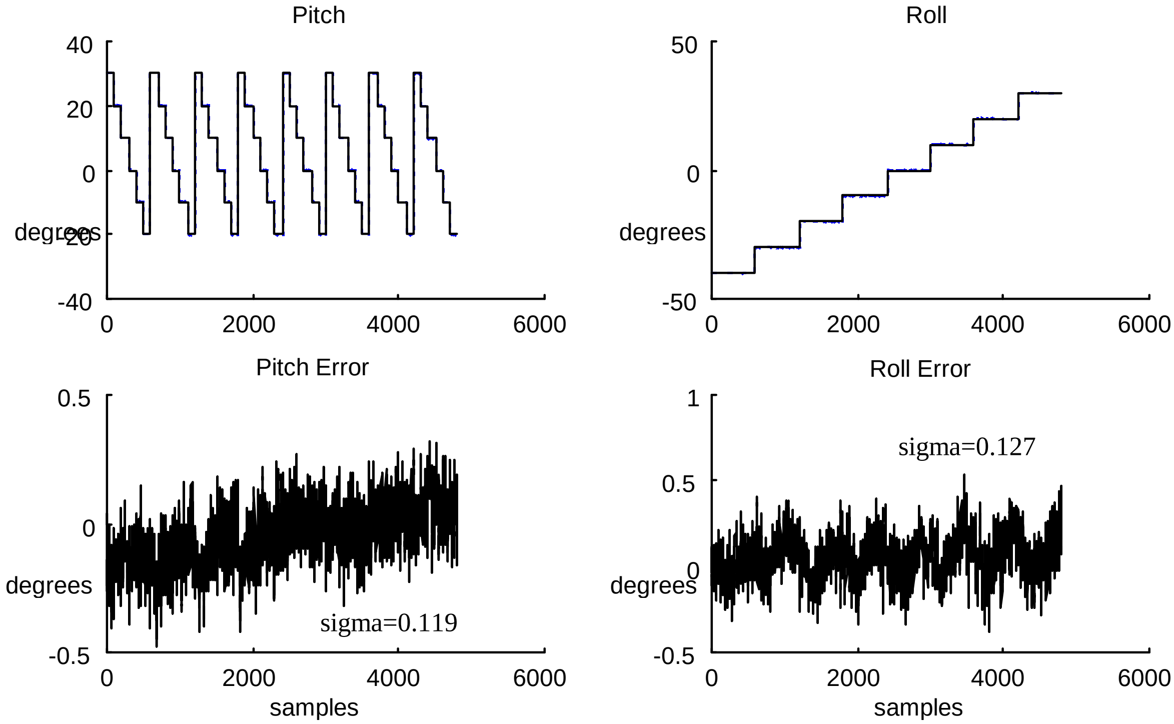

Figure 15: GEOS mode static accuracy

Figure 16 shows the results of an orientation dynamic accuracy test performed by attaching the InertiaCube to the end of the digitizer arm, and waving it fairly rapidly in the air.

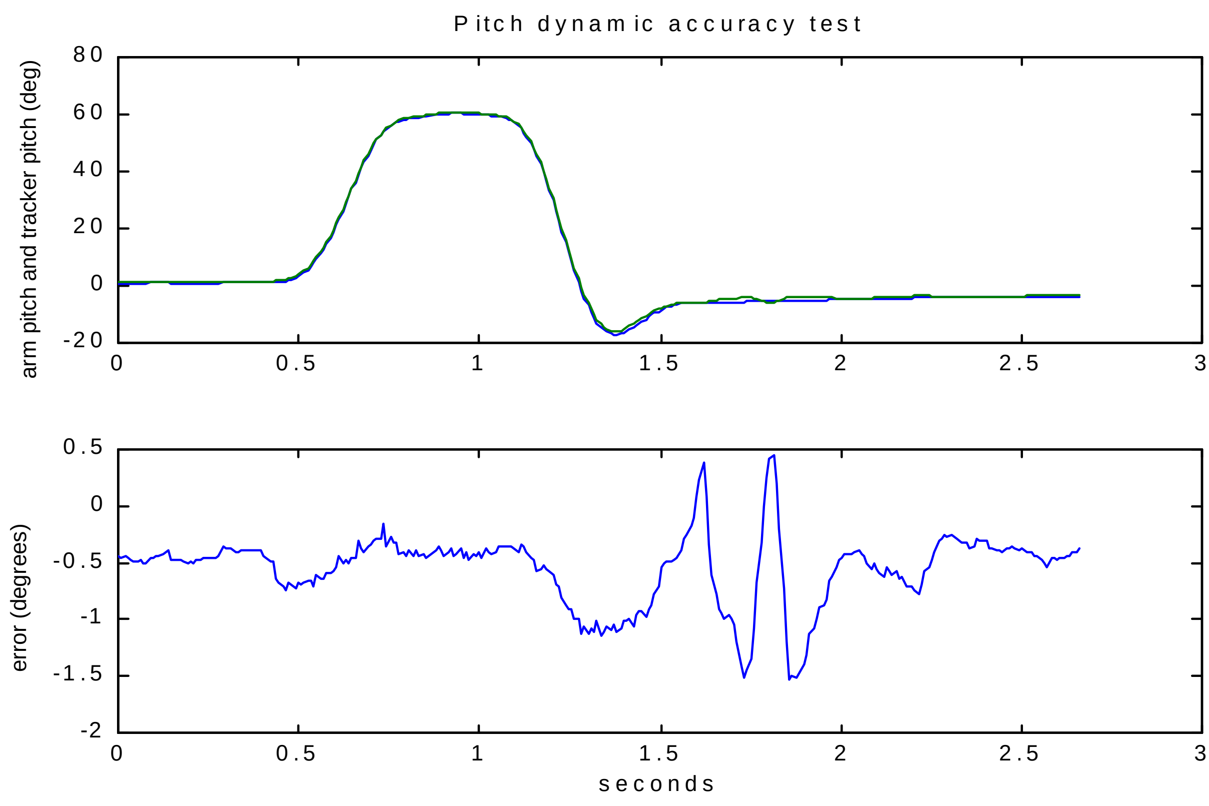

Figure 16: GEOS mode dynamic accuracy

4.3 Fusion mode performance

Fusion mode results are shown for translational degrees of freedom only, since the orientation performance is not that different from GEOS mode. Again both the orientation and positition PPFs were disabled in order to show best dynamic accuracy at the expense of higher jitter. Figure 17 shows the jitter during a 2 second stationary trial, yielding a standard deviation (1 sigma) of about 0.5 mm for X and Y and 0.2 mm for Z.

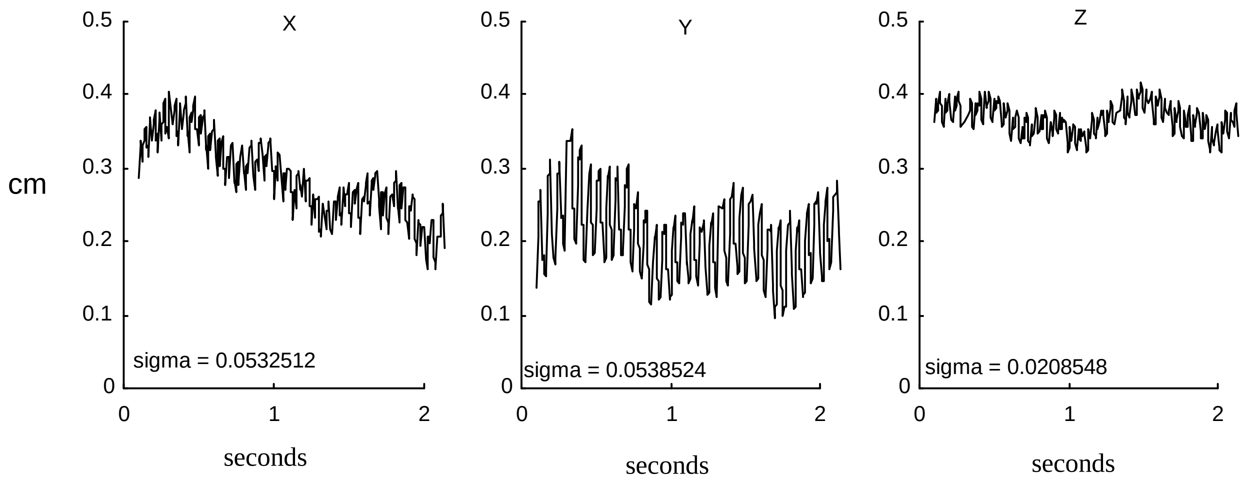

Figure 17: Fusion mode jitter and stationary drift

Figure 18 and 19 show the static and dynamic accuracy of position tracking respectively, collected using a station consisting of 3 SoniDiscs and an InertiaCube mounted on a sheet metal triangle and attached to the end of the digitizer arm as shown if Figure 13. The baseline separation of the SoniDiscs was about 30 cm. The static accuracy was sampled at 10 random X,Y,Z positions, and indicates a deviation of about +/- 5 mm over this span, with an r.m.s. error of 3.3 mm. The dynamic accuracy is a little larger, with peak errors of about 1.5 cm during a fairly rapid movement sequence.

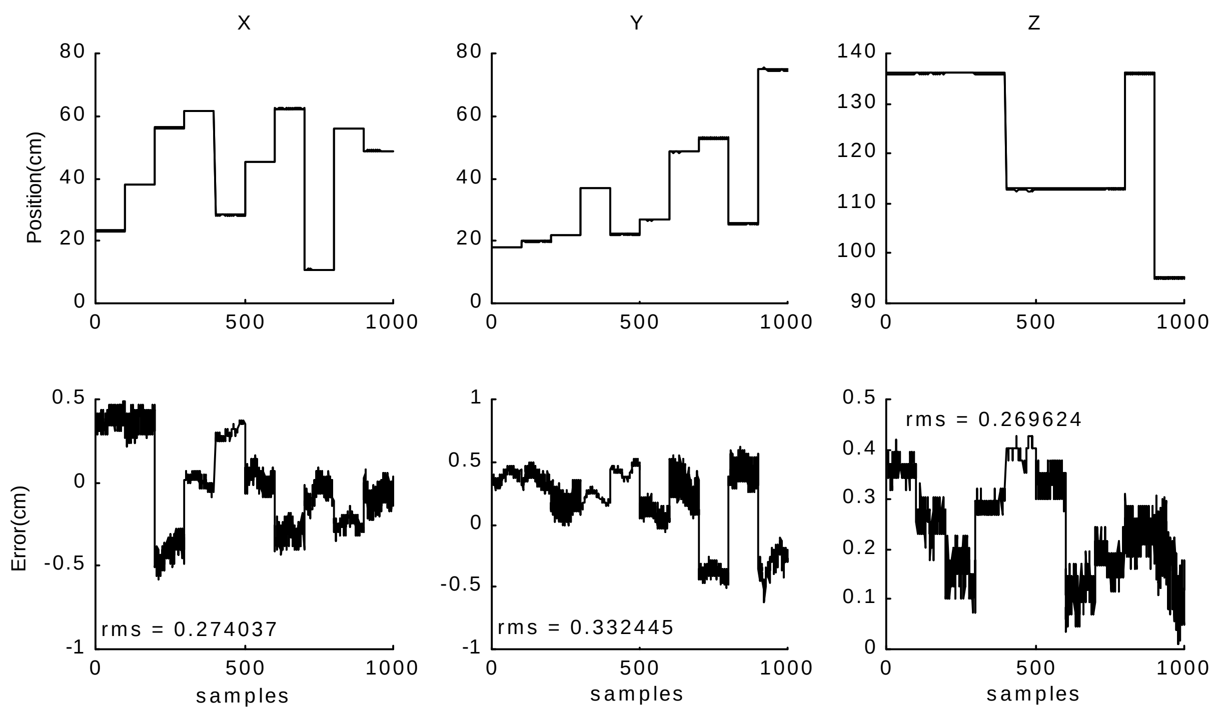

Figure 18: Fusion mode static accuracy

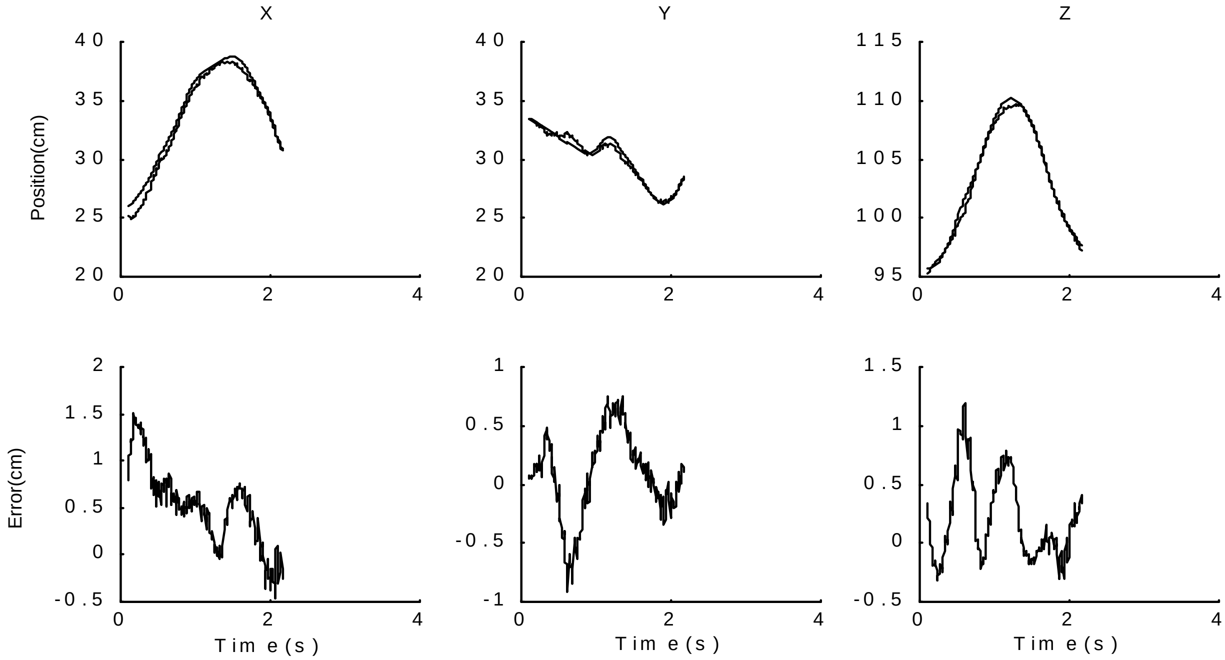

Figure 19: Fusion mode dynamic accuracy

5. CONCLUSIONS & FUTURE WORK

The test results in this paper indicate that excellent HMD tracking performance is achievable with a very small low-cost inertial system using sensor fusion techniques. The hypothesized benefits of low jitter, high update rates, low latency and predictive capability have been realized using commercially available systems specifically designed to be practical for HMD tracking applications. Orientation tracking resolution of 0.02° r.m.s. and static accuracy of 0.25° r.m.s. were achieved while operating in a mode with no postfiltering, so as to reach the optimum performance in latency and dynamic accuracy. The dynamic accuracy was found to be within about 1.5° during a rapid motion for the pitch output. More extensive testing needs to be performed to characterize the relationship between input motion dynamics and dynamic accuracy, and to evaluate static and dynamic accuracy in yaw, with and without geomagnetic heading drift correction. All of these results compare extremely favorably with results that can be achieved under typical operating conditions using other tracking technologies. Although, the testing was performed in a small working volume due to the limited size of the digitizer arm used as a reference, there is no reason to suspect there will be any dropoff in performance of the orientation tracking over any size working volume.

Due to space limitations, position tracking resolution and accuracy were presented only for Fusion mode, which produces superior resolution, accuracy, and update rate as compared to the pure ultrasonic PULSAR mode. The Fusion mode results show a jitter of 0.5mm r.m.s. and static accuracy better than 5 mm. These were tested over a small region as well, but the fusion of the sourceless inertial data with the ultrasonic TOF data gives reason to expect little degradation of resolution in other parts of the working volume. There may be some reduction of absolute accuracy towards the fringes of the ultrasonic working volume due to geometric dilution of precision in the trilateration equations when the lines of sight are at steep angles.

The testing is so far of a preliminary nature, as more exhaustive tests need to be run on a wider range of hardware and configurations (such as multiple station tracking, additional modes, different operating temperatures etc.) Also, latency testing and prediction validation have been postponed due to space and time limitations.

Several enhancements to the inertial tracking systems are currently under development. A version using larger ReceiverPods with higher IR output levels has been produced, which can be configured to track motion over the entirety of a 30’ by 30’ square room. Using the wireless PULSAR mode, this is particularly suited for position tracking of personnel in larger spaces such as museums, training facilites or CAVEs. A version of the L-Bar with removable ReceiverPods is also being prepared to provide greater flexibility for custom installations in tight spaces such as cockpits and automobiles. Finally, we are designing a wireless InertiaCube interface box in order to complement the wireless SoniDiscs and form a completely wireless IS-600 system.

6. REFERENCES

1. D. Bhatnagar, "Position trackers for head mounted display systems: a survey", University of North Carolina - Chapel Hill, TR93-010, 1993.
2. F. Ferrin, "Survey of helmet tracking technologies", SPIE, vol. 1456, pp 86-94, 1991.
3. K. Meyer, H. Applewhite, and F. Biocca, " A survey of position trackers", Presence, Vol. 1, pp. 173-200, 1992.
4. K. Britting, "Inertial navigation systems analysis" , New York, Wiley Interscience, 1971.
5. C. Broxmeyer, "Inertial navigation systems", New York, McGraw-Hill, 1964.
6. R. Parvin, "Inertial navigation", Princeton, New Jersey, Van Nostrand, 1962.
7. E. Foxlin, "Inertial head-tracking", M.S. Thesis, Dept. of E.E.C.S., MIT, 1993.
8. J. Vaganay, M. Aldon and A. Fournier, " Mobile robot attitude estimation by fusion of inertial data", IEEE Intl. Conf. On Robotics and Automation, vol. 1, 1993.
9. E. Bachman, R. McGhee et al., "Evaluation of an integrated gps/ins system for shallow-water auv navigation (SANS), Symposium on Autonomous Underwater Vehicle Technology, Monterey, CA, 1996.
10. E. Foxlin, "Inertial head-tracker sensor fusion by a complementary separate-bias kalman filter", Proc. VRAIS '96 Virtual Reality Annual Intl. Symposium, Santa Clara, CA 1996
11. D. Manolakis, "Efficient solution and performance analysis of 3-D position estimation by trilateration", IEEE Trans. Aerospace and Electronic Systems, vol. 32 no. 4, 1996.