Methods for the Evaluation of Orientation Sensors

K. Waller, A. Hoover
Electrical & Computer Engineering Dept.
Clemson University
Clemson, SC 29634-0915

E. Muth
Psychology Dept.
Clemson University
Clemson, SC 29634

Abstract *In this paper we describe a benchmark suite that evaluates the performance of orientation sensors under different types of motion. Sets of movements are implemented on a spinning platform to which orientation sensors can be attached. The actual position of the platform along with the headings produced by the sensor are recorded, and used to calculate useful statistics and performance metrics for a sensor. Three orientation sensors, the Honeywell HMR3300, the InterSense InertiaCube3, and the Artis MOCOVE were evaluated using the developed benchmark suite. This work has application to man-machine interfacing, robotics, virtual and augmented reality.*

Keywords: tracking, orientation sensing, performance evaluation

1 Introduction

An orientation sensor is a device for measuring physical orientation. Originally developed for navigation, orientation sensors are now being used for applications in robotics, virtual reality, man-machine interfacing, and tracking. These applications exhibit different types and ranges of motions than do the navigation problems for which orientation sensing was originally developed. In addition, the size and cost of orientation sensors have been greatly reduced as these sensors find mass-market applications. These trends have resulted in a number of competing products, with little or no comparative data available to researchers and consumers. With no standard for evaluation, it is difficult to compare the performance of one sensor to another. It is also difficult for researchers to identify gaps in currently available technologies, in order to guide sensing and tracking research. This paper examines the problem of orientation sensing independent of application, and proposes an evaluation framework in which these problems can be addressed. We also evaluate three commercially available sensors and compare their performance in our framework.

Early orientation sensors were often referred to as inertial measurement units (IMUs). IMUs have been in use on both spacecraft and aircraft navigation for many years [7]. These large and heavy units usually consist of ring laser gyroscopes and require delicate calibration using a mechanical platform before being used [4]. Due to their use on manned equipment, the accuracy of these sensors is critical, as well as their robustness. Thus research has been done to evaluate these large systems [8], and failure detection systems have been implemented along with failure recovery systems to ensure the safety of the people onboard an aircraft [5]. Redundant sensors have been developed to ensure the continued operation of the sensor under a failure [6]. A benchmark has been produced by the U.S. Federal Aviation Administration (FAA) to ensure that all IMUs aboard aircraft meet the same minimum standards for performance and safety. These sensors tend to be expensive and are produced in smaller numbers. However, with the development of inertial sensors that are implemented at the microelectromechanical system (MEMS) level, several new markets have arisen. Since these sensors are both small and light, they do not get in the way and can thus be used to track human limb movements [13]. They can be used on head mounted displays (HMDs) for virtual reality and augmented reality applications [1]. They can be used by themselves in robotics application for navigation of mobile robots [3], or they can be used in conjunction with other sensors such as GPS [9].

These applications produce a new set of problems and evaluation criteria. Failure detection, recovery, and system redundancy are no longer as important because the applications are usually not life critical. The sensors are often repositioned, moved from subject to subject, even from one environment to another, which makes delicate, highly sensitive calibration techniques not applicable. The acceleration forces and angular rates that are being sensed are usually orders of magnitude different than that of what an aircraft could undergo. Due to these differences, the evaluation methods used to measure the performance of navigation inertial measurement units are not applicable to this new type of low-cost MEMS orientation sensor. In the absence of an evaluation or comparison standard, vendors producing these sensors typically publish performance data using whatever metrics help sell the product. This practice produces data sheets that are not easily comparable with those of other sensors, and which may not contain the conditions under which testing was performed. In research, the lack of a consistent testing and evaluation framework prevents a better understanding of the gaps and problems with existing technologies and techniques. There are three basic types of orientation sensors: accelerometers, gyroscopes, and magnetometers. Each technology has advantages and disadvantages, in terms of accuracy, stability, and repeatability. The sensors can be combined, for example by placing an accelerometer, gyroscope, and magnetometer on each axis, and combining the sensor readings using filtering techniques. This type of sensor is sometimes called a magnetic angular rate gravity (MARG) sensor[2], and has become a popular implementation for orientation sensing. However, questions still remain as to the best methods for combining sensing modalities, for physically implementing the sensing technologies, and for combining and filtering the sensor readings. Research addressing these issues would be helped by a consistent testing and evaluation framework.

This paper focuses on the development of methods for evaluating the performance of this new class of orientation sensor under different types and speeds of motion. We developed a testing apparatus, and sets of motions we call benchmark routines. These routines were designed to evaluate the different characteristics of orientation sensors under different motion conditions. We applied our methods to evaluate three orientation sensors: the Honeywell HMR3300, the InterSense InertiaCube3, and the Artis MOCOVE. These three sensors use different combinations of sensing technologies and filtering, and provide an overview of the range of sensors commercially available today. Our results show that different angular velocities, accelerations, and jerks produce varying accuracies, both between and within individual sensors. To our knowledge, this is the first work to systematically look at the effects of testing methods on the performance characteristics of orientation sensors, outside the navigation community.

2 Methods

In order to evaluate the performance of an orientation sensor, our basic approach is to place the sensor on a motor-controlled turntable and compare the readings of the sensor against the readings of the motor. We assume that a motor can be controlled to a higher accuracy and repeatability than can the sensors being evaluated (and our results show that this is the case). An orientation sensor can be designed to measure orientation around all 3-dimensional axes. However, we assume that the performance will be similar around all axes. Therefore, to simplify the evaluation, we test only a single axis. Our methods consist of three parts: (1) the construction of the motor-controlled turntable; (2) software processes to control, synchronize and record readings from the motor and orientation sensors; and (3) motion routines that are designed to mimic various applications and challenge the orientation sensors in various manners. The following sections describe these methods.

2.1 Motor-controlled turntable

We constructed a spinning platform, or turntable, using a wooden table, a stepper motor and driver from Applied Motor Stepper Drives, a stainless steel shaft and coupler, and a wooden platform. Our turntable can be seen in Figure 2.1 with one of our tested sensors (the HMR3300 sensor) resting on the platform.

The turntable uses a stepper motor because of its high resolution, high accuracy, high torque at low speed, high holding torque, and ease of programmability. The stepper motor driver is the Si2035; it provides 50,800 steps per revolution microstepping and uses 2.0 amps/phase power. It uses the Si command language (SCL) mode which accepts real time commands and queries from a host PC. The motor is a NEMA 23 sized stepper motor that has a holding torque of 74.9 oz-in. Several platform configurations were investigated before settling on the current setup. At the end of our testing, a 6 in. coupled shaft was found to be far enough away from the motor to not adversely affect the sensors. Even so, all sensors were mounted on the edge of the platform instead of over the shaft's pole to reduce possible electromagnetic field (EMF) effects. A 4 in. square wooden platform was found to be rigid enough to not produce vibrations in the platform, and small enough to allow high accelerations.

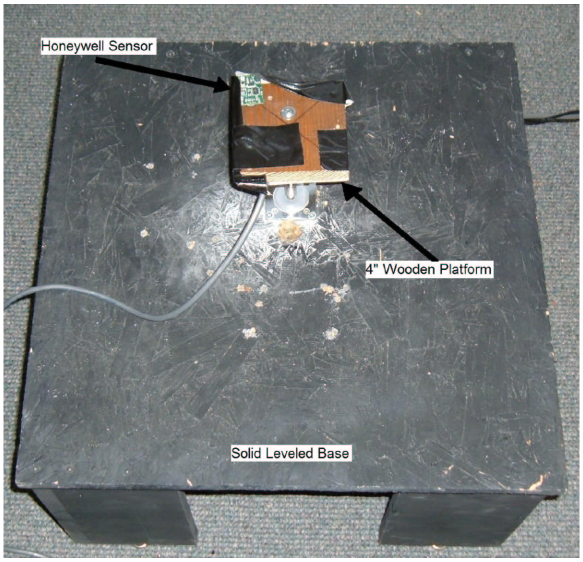

Figure 1: The turntable apparatus, with the Honeywell HMR3300 sensor attached.

2.2 Software

Communication with the stepper motor driver is accomplished via an RS-232 connection. The driver uses the proprietary Si command language (SCL), which consists of a set of ASCII commands and responses. There are two basic types of commands, buffered and immediate. Buffered commands are done sequentially one at a time, whereas immediate commands are executed as soon as they are received, and are done in parallel with buffered commands. The immediate commands allow us to check the position of the motor while it is in motion running a move command. All of the platform and orientation sensor control was done through the use of threads. When the turntable application is first started, it attempts to set up communication with the motor. f successful, a thread is spawned that continually polls the motor driver, asking for the current position of the motor. When an orientation sensor is selected, a thread is created that sets up the needed communication for the sensor, then constantly polls the sensor for a new heading. Starting a motion routine spawns a recording thread, sets up the motor with the proper acceleration and velocity, and runs the routine. All threads interact through a common data structure and control thread, and use mutexes to protect data.

2.3 Motion routines

For this research, we wanted to investigate how the orientation sensors performed under different types of movements, as this would help in choosing the right sensor for a particular application. For example, an orientation sensor used as a digital compass for a driver in a car would likely encounter smooth, slow angular rotations, while an orientation sensor mounted on a person would likely see jerky, burst angular rotations, while a sensor mounted on a robot might encounter fast but smooth accelerations. With this in mind, 10 motor routines were developed, each focusing on a certain type of motion. Table 1 lists the developed routines and summarizes each in terms of acceleration and velocity. The acceleration listed is the initial acceleration used to reach the constant velocity listed. Only the smooth-varied routine and the jerky-varied routine have a constant acceleration throughout the routine, which is why the velocity is listed as increasing instead of constant.

The smooth routines consist of three non-stop revolutions in one direction. The first routine is slow, the second is fast, and the third is constantly accelerating. These routines were designed to be the easiest, and thus were expected to produce the best possible results. The different speeds were to see which sensor could handle a higher angular rotation. Jerk is the derivative of acceleration. To produce high jerk in this research, we use abrupt changes in direction, paired with high accelerations. All three of the jerky routines were set up with the same platform path, but one routine was slow, one fast, and one with constant acceleration.

The oscillating routine follows an under-damped oscillation. The platform swings back and forth increasing the distance of the swing by a degree every half period. It was expected that the large number of abrupt direction changes, along with the increasing time between changes, would cause a significant accumulation of drift in sensors implemented using just accelerometers and/or gyroscopes. Magnetometers on the other hand were expected to not be able to keep up with the abrupt direction changes, but would not drift over time. The intermittent routine was created to model the tracking of a conveyor belt wheel that pulls a part along, stops for a moment to allow for assembly, and then pulls another part up. We were unsure what this routine would bring out in our sensors, but included it in our testing since it models a commonly found motion. The whip with pause routine spins the platform around as fast as possible 1080°, pauses for

Routine	**Accel.** $(\frac{rev}{s^2})$	**Vel.** $(\frac{rev}{s})$	**Modeled Motion**
Smooth-slow	1	0.1	Vehicle compass
Smooth-fast	20	3	Winch
Smooth-varied	1	incr. (40 max)	Accel. wheel
Jerky-slow	5	0.1	Human torso
Jerky-fast	20	3	Human arm/head
Jerky-varied	2	incr. (20 max)	Human arm/head
Oscillating	10	3	Pendulum
Intermittent	20	20	Conveyor belt
Whip with pause	20	5	Mobile robot
12-minute	5	0.1	Drift measurement

Table 1: The accelerations and velocities used for the benchmark routines.

3 seconds, then backwards 1080°, and pauses for 3 seconds. It then spins the platform around as fast as possible for 720°, pauses for 3 seconds, then backwards for 720°, and pauses for 3 seconds. This pattern continues for spins of size 360°, 180°, 90°, 45°, and 20°. The goal here was to confuse the sensor, then observe whether it would correct itself during the pause, and if so, by how much. It was also designed to show how big of a spin was needed to confuse the sensor. Lastly, the 12 minute routine is just a collection of jerky-slow movements that take roughly 12 minutes to complete. This routine was designed to have a slow enough angular rotation so that all of the sensors would be able to get good headings. It was used to see how much error each sensor would accumulate over a longer running time.

3 Results

The benchmark routines previously described were run on three sensors: the Honeywell HMR3300, the InterSense InertiaCube3, and the Artis MOCOVE[1]. As an example result, Figure 2 shows a graph of the path of the motor during the smooth-slow routine compared to the measured heading from the HMR3300 sensor. As can be seen, the device is accurate around the 0° position, but is off by approximately 20° at the 180° position. Although 20° may seem like a large error, we would characterize this performance as still relatively "on-target" when compared to other results from our tests.

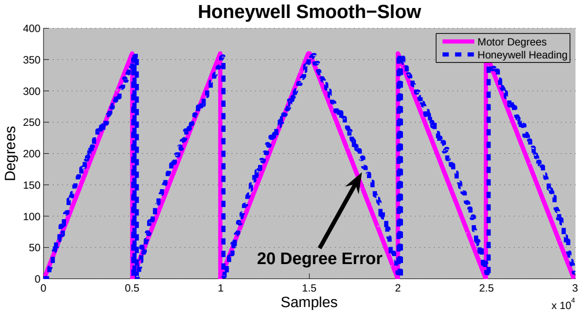

Figure 2: Result for the Honeywell sensor running the smooth-slow routine.

Figure 3 summarizes the average error for each sensor running the 3 different smooth routines, and for running the 3 different jerky routines. Care must be taken when viewing these graphs because when the error is large, it is also a function of the range of the sensor (determination of the quadrant of measurement can be either 2-way or 4-way, depending on the sensor). For these sensors, an error of 45° for the MOCOVE, 90° for the InertiaCube3, and 45° for the HMR3300 indicate that the sensor is lost. Based on this test we see that only the InertiaCube3 can accurately track the larger angular velocities in the smooth-fast routine. It is interesting to note that the MOCOVE sensor actually performs better than the other two sensors at slow speed, but then fails when the speed is increased. The HMR3300 sensor shows a peak accuracy of 19°, even though the data sheet for the sensor claims a 1° accuracy.

Looking at the results for the jerky tests, the InertiaCube3 sensor performed best, although the MOCOVE sensor also performed reasonably well. The HMR3300 sensor tracked the slow jerky motion reasonably well but exhibits a noticeable lag (results not displayed here due to space constraints). Again, the MOCOVE sensor performs better than

---

[1] Artis normally sells this sensor as part of a larger system; they graciously provided a stand-alone sensor for this research

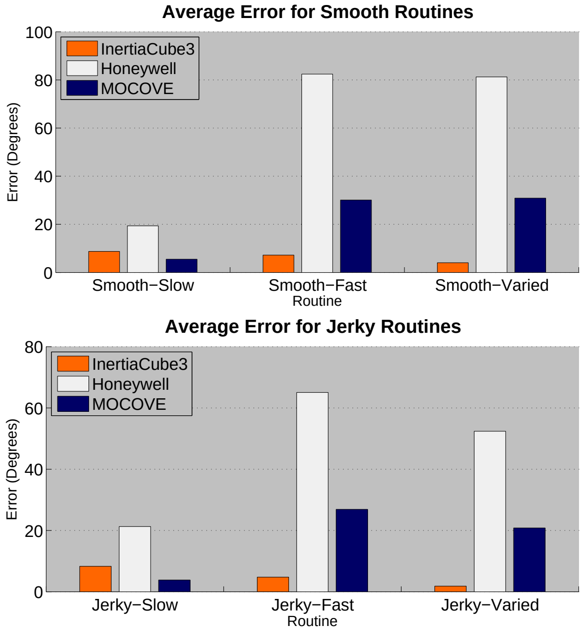

Figure 3: Average error for each sensor on smooth and jerky routines.

the other sensors at slow speeds, but fails once things speed up. Conversely, the InertiaCube3 sensor does better as the speed increases. It is also more accurate during the jerky routines than it was during the smooth routines. The likely cause is that the smooth routines only have two accelerations, the one that gets the platform spinning, and the one that slows it back down, while the jerky routines are filled with accelerations. Accelerometers are a major part of the InertiaCube3 sensor, and by being able to incorporate their measurements into the orientation calculation, a more accurate measurement can be made. It is also worth noting that the performance of the other two sensors remains similar to their performance during the smooth routines.

The other tests, particularly the oscillating test and whip-with-pause test, also produced interesting results that show varying performance between the sensors. Due to space constraints, the interested reader is directed to [11] for more details. After observing the results of the routine tests, we were interested in knowing how fast each sensor could be smoothly spun and still produce a relatively accurate measurement. Each sensor was spun 3 revolutions with an initial acceleration of $20\frac{rev}{sec^2}$, and varying velocities. Figure 4 shows the results of the max velocity test. As previously mentioned, the graph is a little misleading due to the range of each sensor. The important thing to observe in Figure 4 is where the slope increases for each sensor, which indicates the point at which the sensor is starting to fail. It shows that both the MOCOVE and the HMR3300 are accurate to about $0.5\frac{rev}{sec}$, while the InertiaCube3 is accurate all the way up to $5\frac{rev}{sec}$. It is also interesting to note that, in accordance with the routine tests performed, at slow speeds the cheapest sensor (the MOCOVE) outperformed both of the other sensors. The graph also shows rough estimates of how fast a person can spin, turn his or her head, and swing a rifle [10].

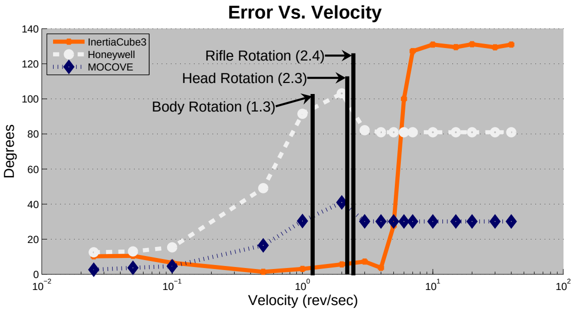

Figure 4: The error vs. velocity plot for all three orientation sensors.

4 Conclusion

This paper presented a way to test a MEMS orientation sensor under different types and speeds of motion in an effort to better understand the performance characteristics of small, commercial orientation sensors. These sensors are generally produced and sold with data sheets that outline some aspects of their performance, but lack enough information to make true comparative evaluations. he performance of these sensors cannot be evaluated using the methods developed for large, expensive, strapdown, redundant IMUs, and thus a new evaluation system or benchmark suite needs to be established. In this research, a testing apparatus was developed, and testing routines were designed to evaluate the different characteristics of orientation sensors under different motion conditions. Three orientation sensors, representative of the market today, were evaluated with the benchmark suite. The developed testing apparatus is a turntable that can precisely spin an orientation sensor via a stepper motor, and can record its exact orientation along with the heading read from an orientation device. Sets of movements we call benchmark routines were designed and implemented to test different properties of the sensors.

The results of this research show that the turntable performs correctly, and is a viable way to test orientation sensors. The three orientation sensors tested show that as expected sensors with somewhat similar data sheet specifications can perform very differently. Under smooth, low-speed conditions (0.1 $\frac{rev}{sec}$), the least expensive MOCOVE orientation sensor outperformed the most expensive InertiaCube3 orientation sensor. However, as the angular velocity increased, the error of the MOCOVE quickly increased. The error of the Honeywell HMR3300 increased as well. Only the InertiaCube3 could accurately measure quicker angular velocities (up to 4 $\frac{rev}{sec}$) and in the range found when tracking the movements of human head and arm movements. Also, different amounts of lag and accuracy were uncovered by different motion routines. Overall, this research shows the need for such a benchmark, and will aid in the process of choosing an orientation sensor for a particular application.

References

[1] R. Azuma and G. Bishop, "Improving static and dynamic registration in an optical see-through HMD", in Proc. of computer graphics and interactive techniques, 1994, pp. 197-204.

[2] E. Bachmann and X. Yun, "Design and implementation of MARG sensors for 3-DOF orientation measurement of rigid bodies", in *Proc. of IEEE Int'l Conf. on Robotics and Automation*, Sept. 2003, pp. 1171-1178.

[3] B. Barshan and H. Durrant-Whyte, "Inertial Navigation Systems for Mobile Robots", in *IEEE Transactions on Robotics and Automation*, vol. 11 no. 3, June 1995, pp. 328-342.

[4] L. Camberlein and F. Mazzanti, "Calibration Technique for Laser Gyro Strap-down Inertial Navigation System", in *Proc. of Symp. on Gyro Technology*, 1985, pp. 5.0-5.13.

[5] T. Chein and M. Adams, "A Sequential Failure Detection Technique and Its Application", in *IEEE Trans. on Automaic Control*, vol. 21 no. 5, 1976, pp. 750-757.

[6] S. Cho and C. Park, "A calibration technique for a redundant IMU containing low-grade inertial sensors", in *Journal of the Electronics and Telecommunications Research Institute (ETRI)*, vol. 27 no. 4, August 2005, pp. 418-426.

[7] J. Gilmore and R. Mckern, "A redundant strapdown inertial reference unit (SIRU)", in *Journal of Spacecraft and Rockets*, vol. 9 no. 1, 1972, pp. 39-47.

[8] J. Harrison and E. Gai, "Evaluating Sensor Orientations for Navigation Performance and Failure Detection", in *IEEE Trans. Aerospace and Electronic Systems*, vol. 13 no. 6, 1977, pp. 631-643.

[9] S. Panzieri, F. Pascucci and G. Ulivi, "An outdoor navigation system using GPS and inertial platform", in *IEEE/ASME Trans. on Mechatronics*, vol. 7 no. 2, June 2002, pp. 134-142.

[10] K. Waller, J. Luck, A. Hoover and E. Muth, "A Trackable Laser Tag System", in *Proc. of WORLDCOMP*, June 2006.

[11] K. Waller, "Methods for the Evaluation of Orientation Sensors", master's thesis, Electrical & Computer Engineering Dept., Clemson University, 2006.

[12] E. White and J. Rios, "FAA certification of a MEMS attitude and heading reference system", in *14th Int'l Technical Meeting of the Satellite Division of the Institute of Navigation*, Sept. 2001, pp. 1349-1360.

[13] X. Yun, E. Bachmann, A. Kavousanakis, F. Yildiz and R. McGhee, "Design and implementation of the MARG human body motion tracking system" in *Proc. of IEEE/RSJ Int'l Conf. on Intelligent Robots and Systems*, Oct. 2004, pp. 625-630.