Limits of Absolute Heading Accuracy Using Inexpensive MEMS Sensors

Gregory Tomasch and Kris Winer, Tlera Corporation, *tleracorp@gmail.com*

June 1, 2018

We measure the absolute heading accuracy achievable using two popular MEMs motion sensor suites and a variety of readily-available calibration and sensor fusion solutions. We find that absolute heading error of 2 degrees rms is routinely obtainable and that, in many cases, 1-degree rms error can be achieved. Examination of the residual error shows it systematic and quasi-sinusoidal. We expect heading accuracy can be improved further as clearly one or more fundamental sources of error remain unaddressed. Two areas that seem potentially fruitful for exploration are correction of accelerometer sensor non-orthogonality and misalignment between the accelerometer and magnetometer sensors.

Introduction

Improvements in micro-electro-mechanical systems (MEMS) technology and the advent of open-source, computationally-efficient sensor fusion algorithms have made accurate absolute orientation estimation available at very low cost. While pitch and roll can be estimated with an accuracy of < 1 degree root-mean-square (rms or 1 sigma error) using just a three-axis accelerometer, a three-axis gyroscope, and straightforward sensor fusion methods, absolute heading requires a 9 degree-of-freedom (9 DoF) solution (accelerometer, gyroscope, and magnetometer) as well as more sophisticated calibration and sensor fusion algorithms. We have found that using simple offset bias calibration, open-source sensor fusion algorithms, and 9 DoF MEMS motion sensors, typical absolute heading errors of ~4 degree rms are routinely achievable in a relatively crude 90-degree-turning test. This accuracy is fine for video demonstrations but for real-world applications--like ergonometrics, head tracking, dead reckoning, and autonomous flight control--absolute heading accuracy needs to be ~1 degree rms or better.

Here we estimate the limits to absolute heading accuracy achievable using inexpensive MEMs motion sensors with the following improvements in method:

> **Calibration**- we use more sophisticated methods of accelerometer and magnetometer calibration for the off-board fusion algorithms and input sensor offset/scale corrections into the EM7180;
>
> **Fusion algorithms**- we use both open-source and proprietary fusion algorithms for comparison to PNI Corp's "SpacePoint™" adaptive fusion algorithm embedded in the EM7180 motion co-processor;
>
> **Test method**- we calculate the EM7180 and off-board orientation solutions simultaneously and use an accurate 2D rotation stage to estimate absolute heading errors.

Despite the increased sophistication, the calibration and fusion methods are readily available and straightforward to implement for those whose applications require the best available absolute orientation accuracy.

General Approach

EM Microelectronic's EM7180 sensor fusion co-processor does an impressive job of taking data from low-cost MEMS sensors with no additional calibration and converging to a stable and accurate absolute orientation solution. This is done by internally comparing the gravitational and magnetic orientation solutions and adjusting the sensor signals according to a set of assumptions in the adaptation algorithm. However, the EM7180 cannot perfectly compensate for all sensor imperfections in all circumstances. Occasionally, for an individual set of sensors the EM7180 struggles to converge to a stable solution and might also suffer from residual heading error. We have found that estimating the accelerometer and magnetometer sensor bias offsets and axis scale error corrections and inputting these into the EM7180 at startup can overcome this limitation in most cases.

However, the observation that there are practical limits to the degree of sensor non-ideality an adaptive algorithm can fix raises several questions:

1. What practical heading accuracy level can the user reasonably expect?
2. What are the relative contributions of sensor quality, sensor calibration and fusion algorithm to heading estimate accuracy?
3. How "Good" does sensor calibration have to be? Do more sophisticated calibration methods yield significantly better results?
4. For well-calibrated, high-quality sensors, how important is the actual fusion algorithm?

In order to start answering these questions, we have undertaken a *simultaneous* performance comparison between a baseline case and various options. The baseline case is using the EM7180 and its internal adaptive sensor calibration and proprietary fusion algorithms "as-is"; with no pre-calibration of the sensors. The various absolute orientation solution options include one or more of the following:

- Inputting accelerometer/magnetometer offset bias and scale error corrections into the EM7180 prior to algorithm startup;
- More sophisticated accelerometer and magnetometer sensor calibration methods for use with alternative fusion algorithms;
- Our own proprietary fusion algorithm;
- The open-source Madgwick 9 DoF fusion filter.

To do this, we let the EM7180 calculate its orientation solution as usual *and* simultaneously output the raw sensor data stream. We applied accelerometer and magnetometer calibration corrections on the host MCU and processed the data into parallel orientation solutions using our own proprietary fusion algorithm or the open-source Madgwick fusion algorithm. Note that both methods were being applied to the same sensor data at the same time, eliminating all questions of trial-to-trial sensor variability.

Sensor Platforms

We consider Invensense's MPU9250 (MPU6500 accelerometer/gyroscope with embedded AK8963C magnetometer) 9 DoF motion sensor and ST Microelectronic's LSM6DSM combination accelerometer/gyroscope coupled with ST Microelectronic's LIS2MDL magnetometer. Absolute orientation accuracy depends mostly on the quality of the underlying sensor data and how well the sensors are calibrated; no amount of sensor fusion sophistication can make up for jittery data with large offset biases. These two sensor suites have proven to be stable, low-jitter sensors capable of excellent

heading accuracy as we will show. We pair the sensors with EM Microelectronic's EM7180 motion co-processor. The EM7180 is an I2C master to the slave sensors and manages configuration of and data reading from the sensors. The EM7180 provides scaled or raw sensor data and orientation quaternions to the MCU (Cortex M4F 80 MHz STM32L4) host. The EM7180 uses PNICorp's "SpacePoint™" proprietary adaptive fusion algorithm and in most cases provides absolute heading accuracy of 2 degrees rms or better with these sensors.

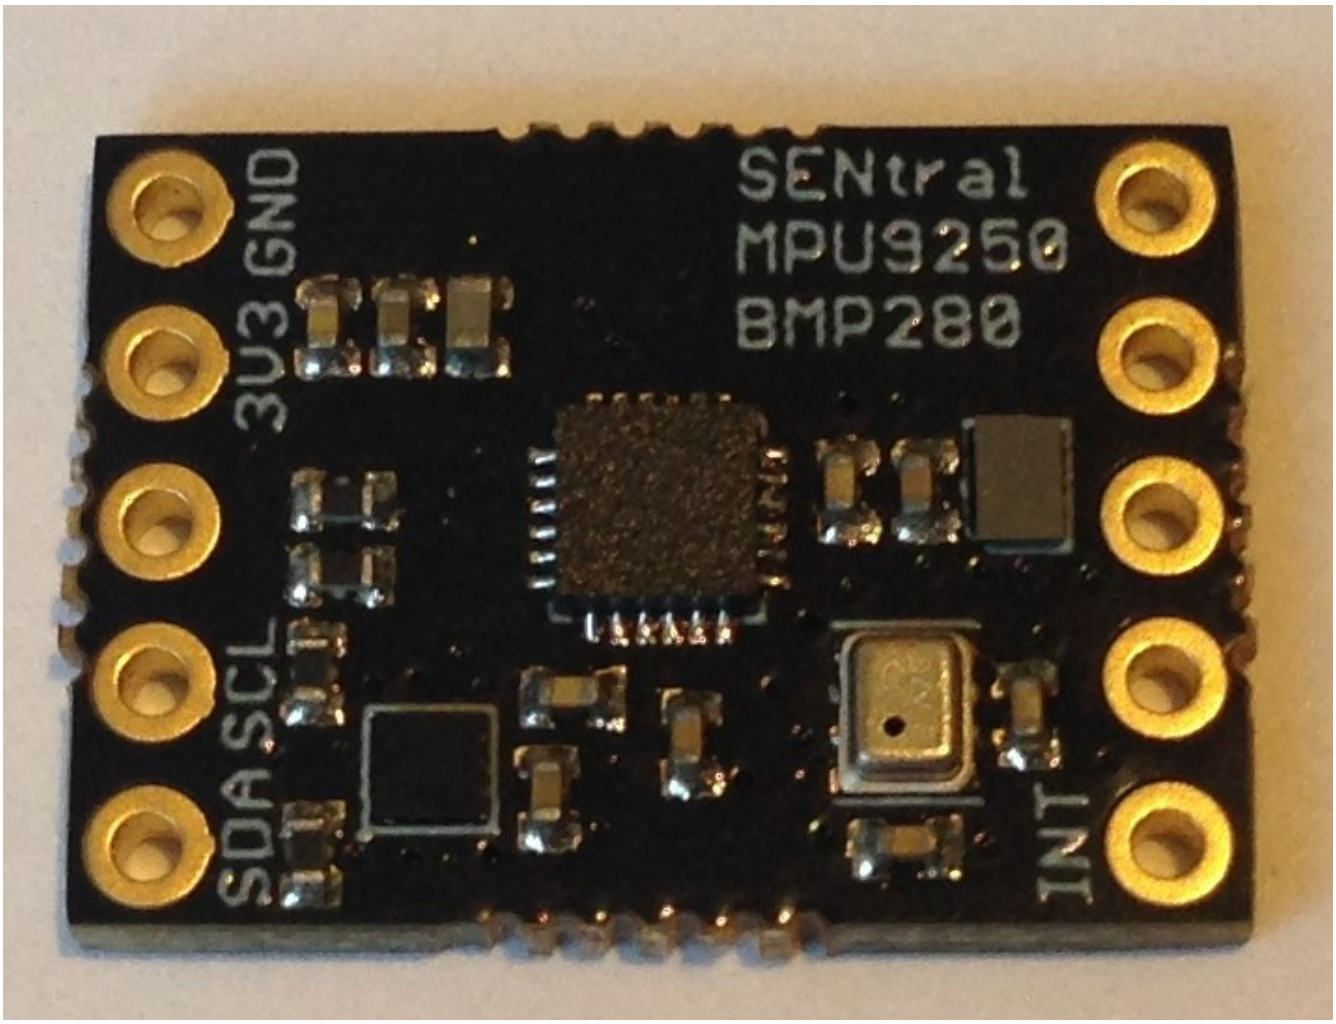

Figure 1a *Production version of one of the two MPU9250 breakout boards we used for these tests. MPU9250 is in the center, EM7180 at lower left.*

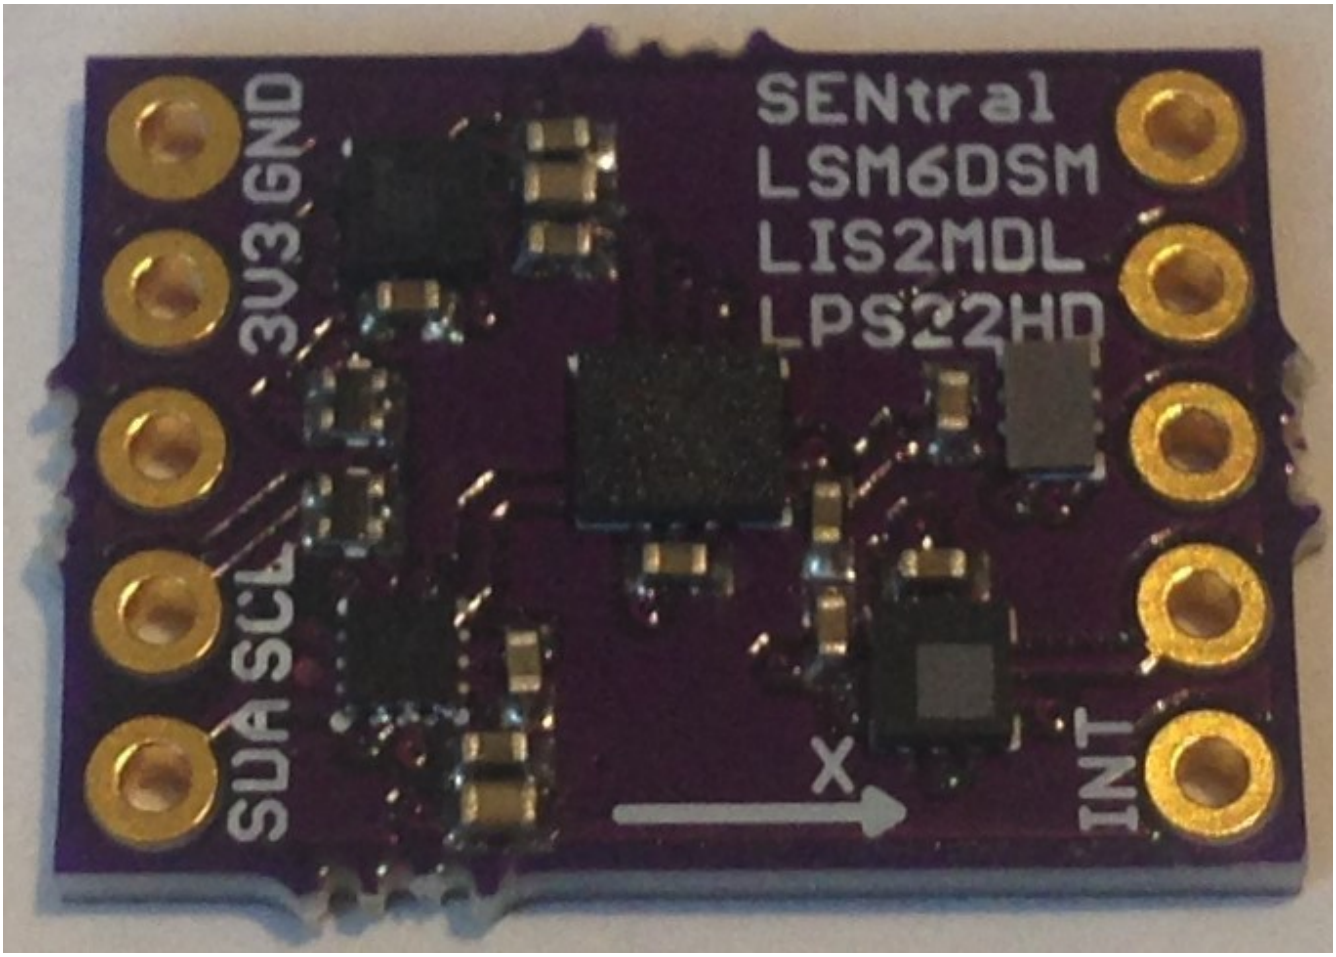

Figure 1 b *Prototype version of one of the two LSM6DSM+LIS2MDL breakout boards we used for these tests. LSM6DSM is in the center, LIS2MDL is top left, EM7180 at lower left.*

Sensor Calibration

Calibrating the gyroscopes is straightforward. We make use of the built-in calibration of the EM7180 which calculates the gyro offset biases when the sensors are at rest. Even without the use of the EM7180, calculating gyroscope offset bias this way, with the sensors flat (pcb normal to gravity direction) and motionless is adequate. One could go further (we have not) by measuring and correcting the differences in axial response.

The goal of accelerometer calibration is to remove both the offset biases as well as the axial scaling differences. A popular simple method for estimating accelerometer offset biases is to place the sensors flat and motionless and average the accelerometer readings (accounting for gravity on the z axis). A better although more time consuming method that we use here is to align each accelerometer axis (+/-x, +/- y, +/- z) with gravity, average readings to provide both offset biases and axes scaling correction.

The goal of magnetometer calibration is also to remove both the offset biases as well as the axial scaling differences. In this case, we need to rotate the magnetometer throughout 3-Space such that we adequately sample the entire 3D response surface. This will provide a cloud of points generally in the form of an ellipsoid with two foci and not a sphere centered at the zero-field origin as we desire. A marginally adequate method for magnetometer calibration is to center the ellipsoid on the origin (estimate simple offset biases) and scale the axial responses. But a better though more computationally

intensive method we use here is to fit an ellipsoid to the cloud of points and create a 3x3 calibration matrix to spherize the response surface and center it on the origin.

Heading Measurement

We purchased a used Zeiss rotation stage from a seller on E-Bay for the heading measurements.

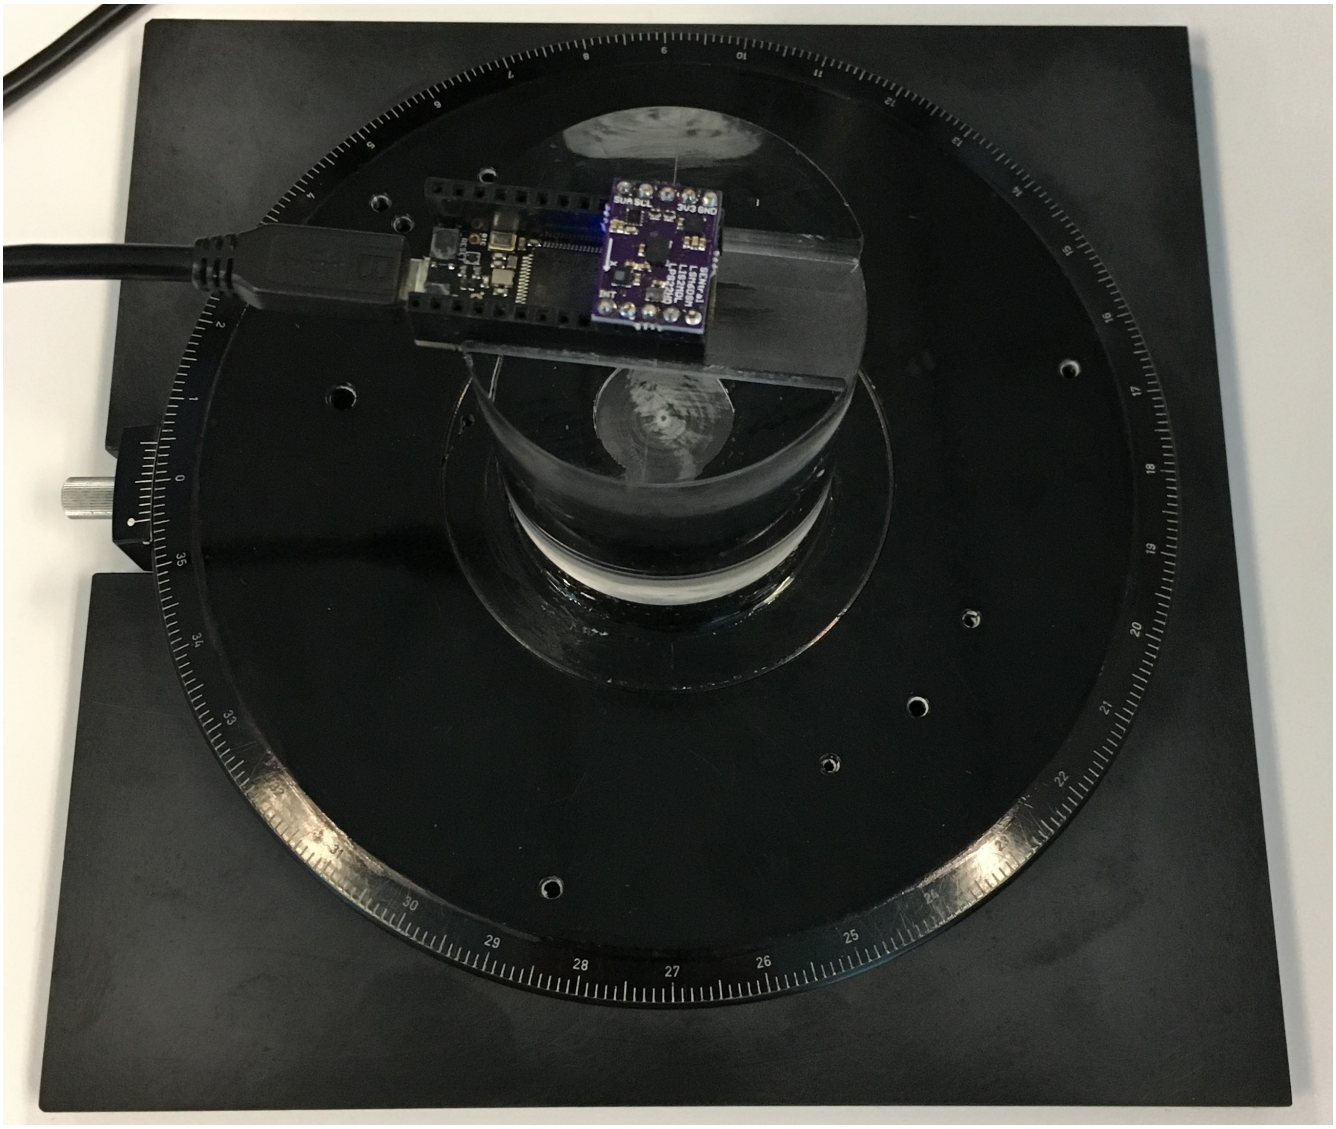

Figure 2 *Zeiss rotation stage used for the heading measurements. Sensor breakout is mounted via machine pin headers onto a STM32L476 development board, which fits into a slot machined into the acrylic cylinder for rigid and reproducible mounting.*

The stage bezel is graduated in degrees and has a micrometer for turning and a stop for secure positioning. The rotation stage is capable of 0.1-degree accuracy.

We added an acrylic cylinder mounted onto the center of the stage with adhesive to serve as a sensor platform. The purpose of the acrylic cylinder is to stand off the sensors from the metal in the rotation stage. Even though we do not believe the metal is ferromagnetic, we wanted to minimize any

interaction. The acrylic makes it easy to check the level of the sensor platform and serves as a thermally stable platform for sensor board mounting.

The measurement procedure was to calibrate the sensors, level the stage on the table, mount the sensor board on top of an acrylic cylinder in the center of the rotation stage, place the stage at a known heading and record an average of 300 readings. Then rotate the stage 15 degrees and repeat the measurement. The data was automatically downloaded into an excel spreadsheet where differences and standard deviations were calculated.

Measurement Results

Typical results for the MPU9250 and LSM6DSM+LIS2MDL using PNI Corp's proprietary fusion algorithms embedded in the EM7180 motion co-processor are shown in Figure 3. In this case we only rely upon EM7180's internal self-calibration capabilities; no bias offsets or scale corrections were sent to the EM7180.

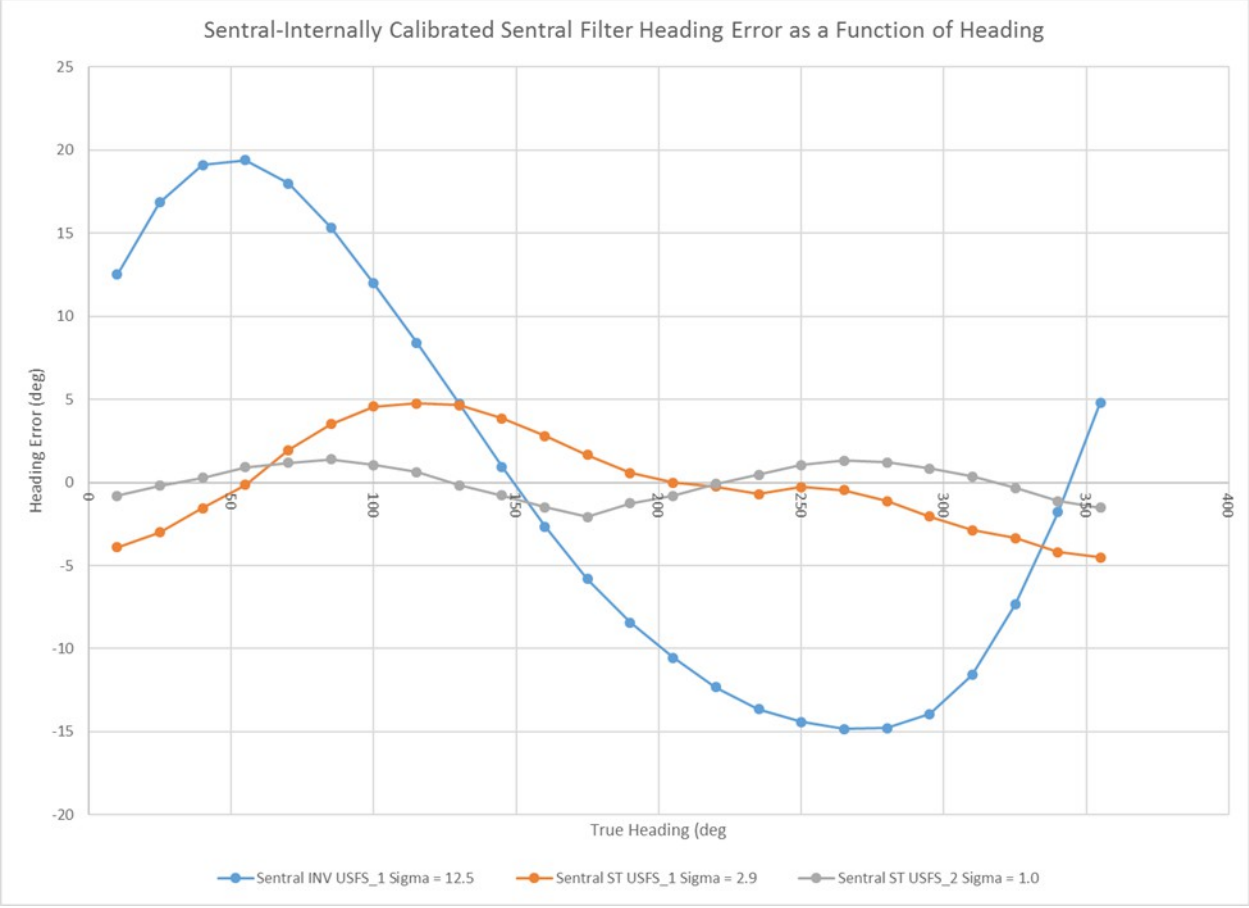

Figure 3 *Absolute heading error as a function of true heading plotted for one MPU9250 and two LSM6DSM+LISMDL breakout boards for the solutions using PNICorp's proprietary sensor fusion algorithms but without sensor offset and scale corrections applied. Blue is the MPU9250, gray and orange are the LSM6DSM+LIS2MDL. Depending upon the nature and degree of sensor imperfections, the residual heading error can vary dramatically*

Simply relying on the EM7180 to internally compensate for all sensor non-ideality clearly has its limits; many sensor sets yield good to excellent performance (1.0 – 3.0 degrees rms) while occasionally others give an abysmal outcome (12.5 degrees rms in this case). These results are consistent with our early experience using the EM7180. In order to achieve better and more predictable results, we set about developing better accelerometer and magnetometer calibration techniques. We use these to enhance the performance of the SpacePoint algorithm by pre-loading the EM7180 with bias offsets and scale corrections for the individual sensors. Figure 4 shows heading error results for the same three sensors characterized in the previous figure plus an additional Invensense USFS sensor set. Using the calibration enhancements, the results are both very good and dramatically more consistent.

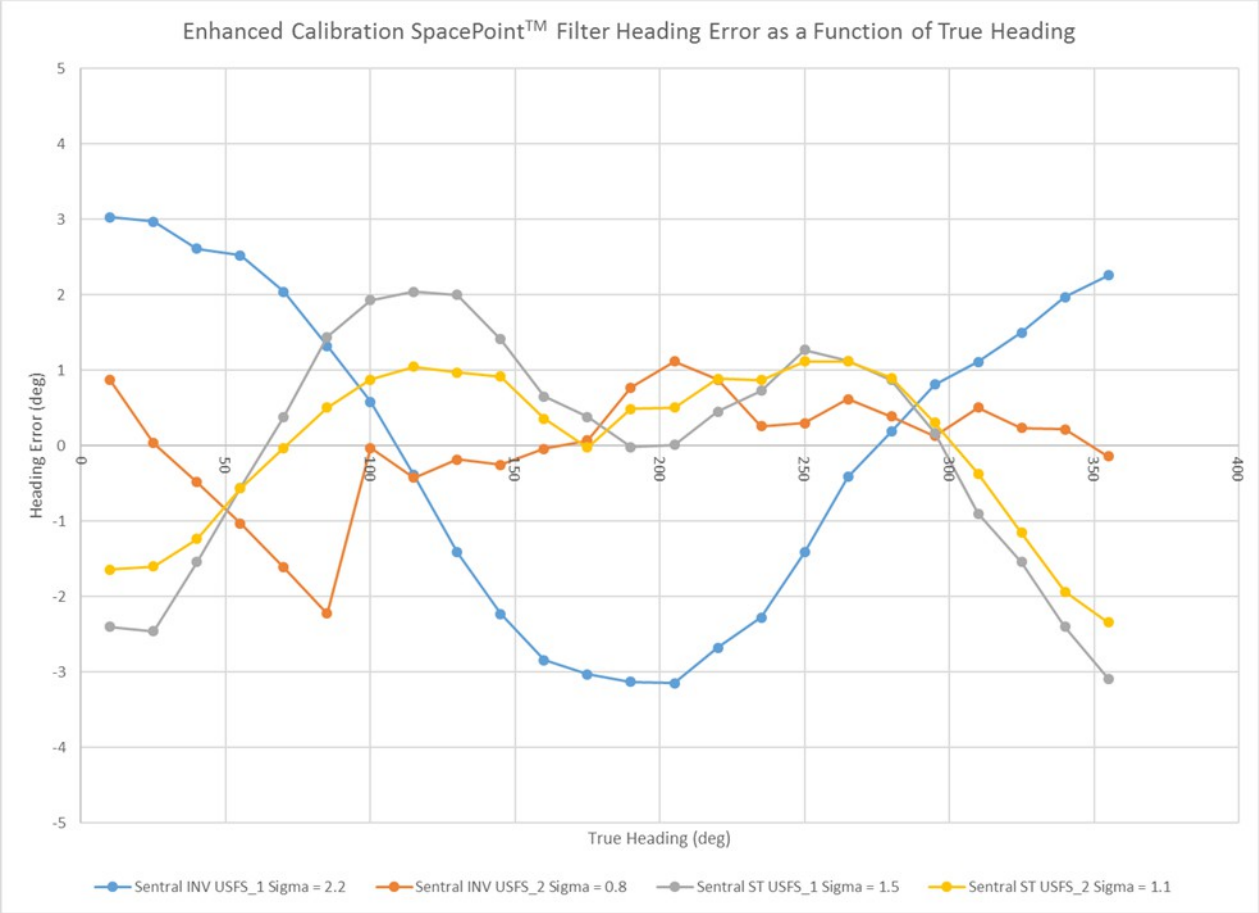

Figure 4 *Absolute heading error as a function of true heading plotted for two different breakout boards each for the MPU9250 and LSM6DSM+LISMDL solutions. Sensor fusion was done using PNICorp's proprietary algorithm and calibration was enhanced by pre-loading the EM7180 with accelerometer/magnetometer bias offsets and scale corrections. Blue and orange are MPU9250, gray and yellow are the LSM6DSM+LIS2MDL. The ST sensor boards are remarkably consistent with rms heading error approaching 1 degree. All data show signs of remaining systematic error.*

Table I shows the rms heading accuracy for the three sensors characterized in Figure 3 with and without enhanced sensor calibration.

Sensor	RMS Heading Error	
	Enhanced Sensor Cal (deg)	EM7180 Internal Cal Only (deg)
Sentral_INV_USFS_1	2.2	12.5
Sentral_ST_USFS_1	1.5	2.9
Sentral_ST_USFS_2	0.8	1.0

Table I *Comparison of PNICorp's proprietary sensor fusion algorithm with and without accelerometer and magnetometer bias offset and scale corrections. Depending upon the particular sensor, the impact of the "Enhanced" sensor calibrations can range from minimal to profound.*

In all three cases, the EM7180's algorithm performed better when accelerometer and magnetometer sensor corrections were applied, as compared to just relying on the EM7180's ability to self-calibrate. In the case of the problematic Invensense sensor (Sentral_INV_USFS_1) the impact of applying the enhanced calibration techniques was quite profound; the rms heading error improved from 12.5 to 2.2 degrees.

We noticed that despite calibration enhancement the heading accuracy achievable with each sensor breakout board still appears to be specific to the board. This was especially true for the MPU9250 solutions. If the MPU9250 magnetometer offset biases were particularly large, then the resulting heading accuracy would suffer no matter how well we calibrated-out the offsets. This makes sense if the bias offsets are due to large assembly strains and the magnetometer is far from its normal operating conditions. We noticed this less for the ST sensors, perhaps because the separate physical packages make such assembly strain less likely or less severe. So perhaps there is an advantage in having a 9 DoF solution requiring two sensor packages despite the added costs in board area and alignment effort. Also, the LIS2MDL is not a Hall sensor, rather it uses anisotropic magnetoresistance so is less affected by temperature changes and seems to offer more stable performance compared to many Hall sensors.

The rms heading accuracy in the four tests above varies between 0.8 and 2.2 degrees rms with these two particular ST sensors performing better overall in terms of both heading accuracy and performance consistency. In all cases there is clearly a quasi-sinusoidal systematic error component. This clearly indicates that our "sophisticated" sensor calibration techniques are still not addressing all sources of systematic error. Our future work will investigate an augmented version of the "Six position" accelerometer calibration technique to generate the full 3x3 correction matrix for the accelerometer sensors. We also acknowledge that the magnetometer and accelerometer axes can be rotated with respect to each other and that any impact from this eventuality is currently unaddressed.

In Figure 5 we show the results of using our own proprietary fusion filter using the same four sensors used earlier. These results were generated on the host MCU at the same time as the data presented in Figure 3. We used the raw sensor data stream from the EM7180 and applied the enhanced calibration corrections on the host MCU. The algorithm generates quaternions from fusing the accelerometer and gyroscope sensor data, compensates the magnetometer data for tilt and then computes heading. The filter is as computationally-efficient as the Madgwick or Mahony filters, so it can be easily performed

on the host MCU but provides heading estimation accuracy similar to the EM7180. Again, for this sampling of sensors the ST solution performs most consistently with low rms heading error of 0.8 – 1.2 degrees. Systematic sinusoidal error is also observed in the results, suggesting that the causes are not directly related to the fusion algorithm. We should also state that this algorithm may have better long-term predictability than the EM7180's SpacePoint algorithm. The behavior of our algorithm is totally dependent upon the quality of accelerometer and magnetometer calibration. It should be static in the absence of sensor drift. If the algorithm delivers bad results, the sensors need to be recalibrated. In contrast, if the EM7180 algorithm gives unanticipated results it is not clear what drifted; the sensors or the SpacePoint algorithm.

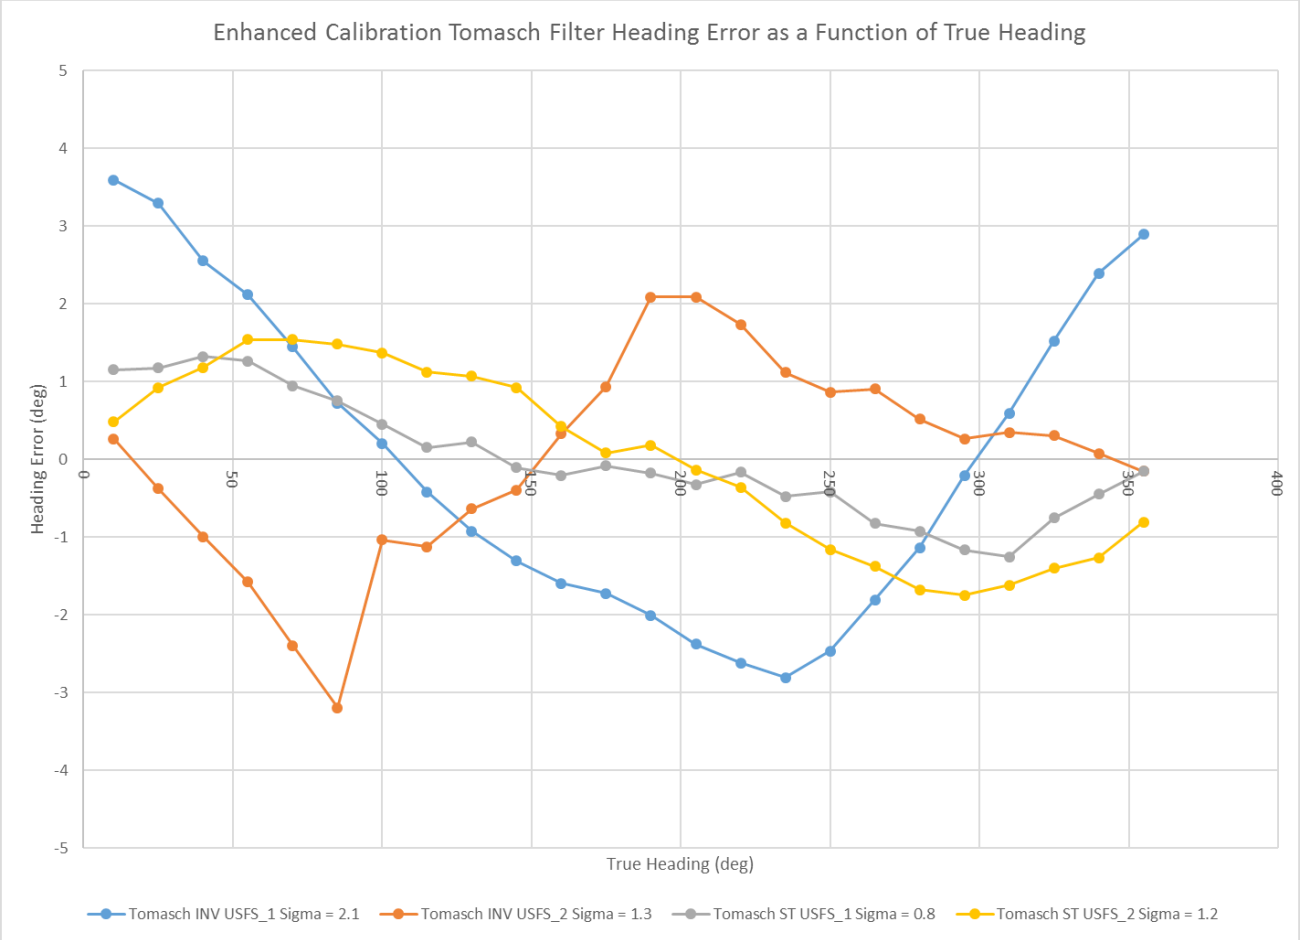

Figure 5 *Results of our proprietary filter using the same two different breakout boards each for the MPU9250 and LSM6DSM+LISMDL solutions. The one MPU9250 board with large magnetometer offset bias is still an outlier.*

In Figures 6 we show results generated using the popular open-source 9 DoF Madgwick fusion filter. In this case we used a different set of production USFS sensors because the original prototype set was no longer available. We used the same simultaneous calculation method to evaluate the performance of the Madgwick 9 DoF algorithm and compare directly to the EM7180's SpacePoint algorithm results, shown in Figure 7 for reference. The rms heading error for the EM7180-based results ranged from 0.82 to 1.2 degrees while the Madgwick-based heading error ranged from 0.76 to 3.0 degrees. So for one instance in four the Madgwick result was as good as the SpacePoint result but in the other three trials the Madgwick result was significantly worse.

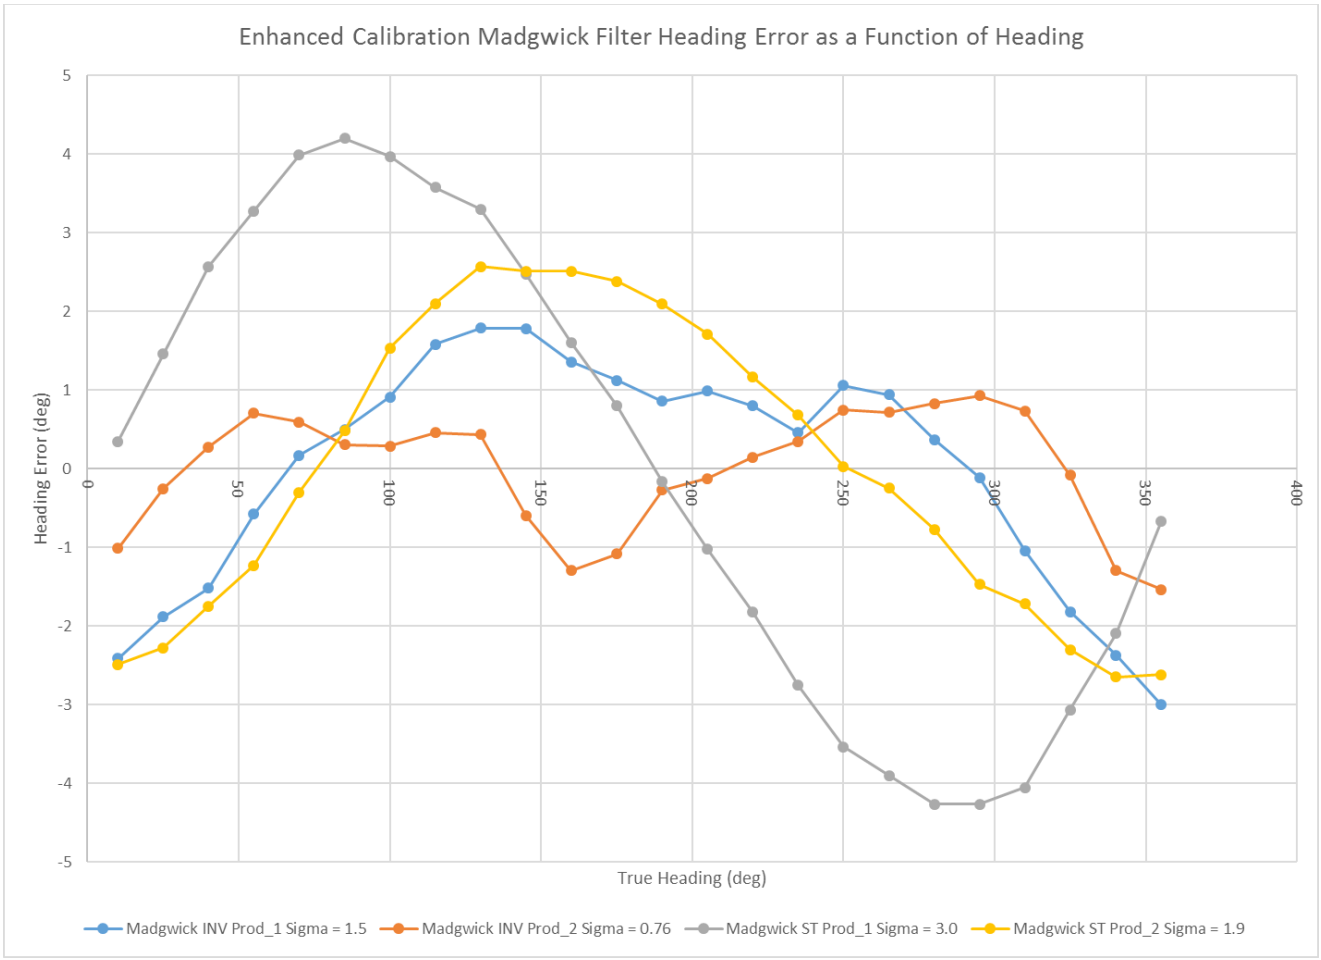

Figure 6 *Results of Madgwick 9DOF filter using two different production breakout boards each for the MPU9250 and LSM6DSM+LISMDL solutions. The heading error is generally poorer than either the EM7180 SpacePoint algorithm or our own proprietary fusion filter (Figures 5 and 7)*

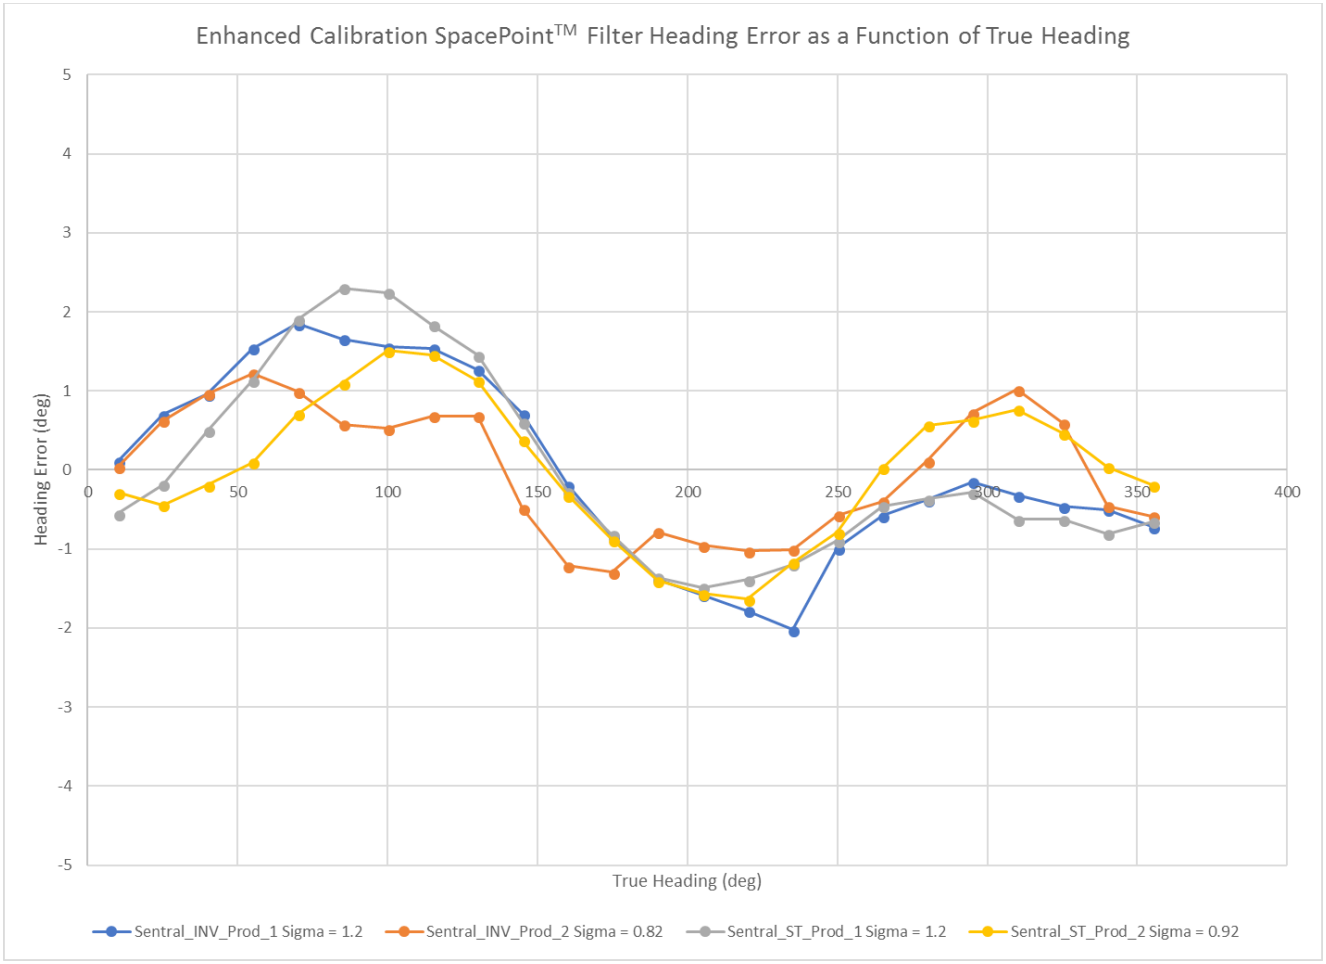

Figure 7 *Absolute heading error as a function of true heading plotted for two different production breakout boards each for the MPU9250 and LSM6DSM+LISMDL solutions. Sensor fusion was done using PNICorp's proprietary algorithm with enhanced calibration. Blue and orange are MPU9250, gray and yellow are the LSM6DSM+LIS2MDL. This data is shown here to provide context for the results presented in Figure 6.*

Table II shows a side-by-side comparison of the SpacePoint and simultaneous Madgwick filter results. Clearly, the SpacePoint algorithm's performance is superior.

Sensor	RMS Heading Error	
	SpacePoint Result (deg)	Madgwick Filter Result (deg)
INV_PROD_1	1.2	1.5
INV_PROD_2	0.8	0.8
ST_PROD_1	1.2	3.0
ST_PROD_2	0.9	1.9

Table II *Comparison of PNICorp's proprietary sensor fusion algorithm with the 9 DoF Madgwick fusion filter.*

Conclusions

We have only examined a small sampling of MPU9250 and LSM6DSM+LIS2MDL breakout boards using these methods, and heading estimation is limited to a 2D plane, but we can still draw some conclusions from such limited testing.

Firstly, proprietary fusion methods such as those embedded in the EM7180 motion co-processor or our own fusion method performed on the STM32L476 host are capable of ~1-degree rms heading accuracy with the right sensors and proper calibration. In terms of the right sensors, either the MPU9250 or LSM6DSM+LIS2MDL will do although we believe the latter offer more consistency and higher accuracy.

Secondly, if a performance-screened sensor board shows very large magnetometer offset biases, we recommend discarding that board for uses where heading accuracy is critical. However, after applying our improved calibration techniques, even our worst performing MPU9250 board with very large magnetometer offset biases showed a heading accuracy of 2.2 degrees rms. This is perfectly adequate for many applications.

Thirdly, in terms of proper calibration, we highly recommend the more complex methods. They will repay the increased effort in calibration with heading accuracy (~1-degree rms vs. ~2-degree rms, for example) that can be twice what could be achieved with the same sensors and sensor fusion but easier (more lax) calibration. Even if the user is relying on the EM7180 and its adaptive algorithm, heading accuracy is both better and much more predictable if accelerometer and magnetometer sensor bias offsets and scale corrections are input into the EM7180.

The results so far show that there is still systematic sinusoidal error in the heading estimation. We hope to understand and eliminate this systematic bias and use this measurement methodology to reach the limit where the statistical sensor noise alone determines heading error. In this case, we would expect to routinely be able to achieve < 1-degree rms heading accuracy, at least with the ST sensor suite.

Our first suspicion regarding a possible root cause for the residual sinusoidal error is inconsistency between gravitational and magnetic measurements. We believe this can be thought of as the general problem of tilt compensation error in a tri-axial magnetic compass. For such a device, it has been

shown that a tilt estimation error of 0.1 degree at a magnetic dip angle of 66 degrees (the magnetic dip angle in Northern California where these measurements were performed is ~60 degrees) results in sinusoidal heading error with an amplitude of ~0.3 degrees. . This being the case, if the gravity-based pitch and roll estimates are off or effectively misaligned to the magnetometer axes by roughly 1 degree, it is easy to imagine a residual sinusoidal heading error of a several degrees in amplitude.

Immediate future work will focus on further improving accelerometer calibration data analysis to generate the full 3x3 correction matrix which will handle any non-orthogonality or cross-talk effects of the accelerometer sensors. We also plan to investigate the effects of residual rotations between the accelerometer and magnetometer sensor axes.

We further note that it would be desirable to repeat these measurements with a wider selection of sensor breakout boards (the ST version is currently in pilot production) as well as use an accurate 3D rotation stage to quantitatively measure the effects of pitch/roll (tilt) on heading accuracy.