New compass calibration procedures for Seagliders

Jim Bennett
University of Washington
School of Oceanography
29 June 2012

Background and results

Following the observations that strong magnets near Seaglider hulls--deliberate or accidental--can invalidate a pre-deployment compass calibration, we investigated the requirements for, and accuracy of, an in-flight procedure for calibrating a Seaglider carrying a Sparton model SP3003D compass.

Experiments determined that a shallow (~100m) two-dive sequence, requiring two roll angles and two pitch angles, suffices for accurate calibration. Comparing observations gathered during a standard calibration and simulated in-flight 'dives' on the same calibration stand without wanding off the glider show that these calibrations are as accurate as the standard pre-deployment compass calibrations using our usual stands and routines. Both calibration procedures typically yield a heading RMS error of 1.2 degrees (N=10). We verified their acceptability for navigation by commanding gliders to steer a given heading in both still water and water with known mean current, then comparing the reported against actual heading based on GPS readings at start and end of the dives.

We also established, using calibrated angle gauge-blocks and several compasses (N=6), that the accelerometer-based 'raw' pitch and roll values reported by the Sparton SP3003D compasses are accurate to at least 0.5 degrees throughout their range. Therefore the use of the original TCM2 Seaglider pitch and roll correction procedure (and resultant coefficients) is unwarranted. (Indeed, its use can be detrimental to accurate heading correction, especially at higher rolls.)

Taken together, these observations recommend the use of a new on-shore ('whirly') compass calibration procedure, which duplicates the in-flight procedure, which is much simpler, faster, and less prone to data-collection errors. It does not require an accurately calibrated fixture and can be performed wherever there is a 'quiet' magnetic field (e.g., a parking lot, or similar).

The in-flight procedure then allows pilots to recalibrate the compass at the beginning (and during) the deployment in the case where the pilot suspects the glider may have been exposed to strong magnetic fields during power cycling, transit or storage.

No change in the on-board glider compass correction code is required.

New versions of the MATLAB compass calibration scripts are available separately. Descriptions of both procedures are included below.

The new on-shore and in-flight calibration procedures, like the standard calibration procedure, assume that changes in the HV pack position produce negligible variation in the vehicle's hard- and soft-iron fields. Previous measurements indicate this is the case for Seagliders. Deepglider, however, owing to the closer proximity between its HV pack and compass, will require different calibration and correction schemes; this is work in-progress.

Finally, for those gliders retaining the original magnetic reed switches for on/off control, we recommend an improved wanding technique suggested by Fritz. Using the grey Seaglider magnet-on-a-stick, gently bring the magnet to the side of the hull near the switch, and then slowly rotate the magnet along the stick axis once--like a lollipop. This rotation of the permanent magnetic field appears to reliably engage the reed switches without excessive waving about and likely minimizes hard iron changes in the forward battery cladding.

Thanks to Fritz Stahr, Mike Johnson, Kirk O'Donnell, Karl Kunkle, Tor Bjorklund, Charlie Eriksen, Geoff Shilling and Jason Gobat for many ideas, code, and patience during these experiments.

The calibration procedures, an overview

The goal of the calibration procedure is to characterize the glider's hard-iron (permanent) and soft-iron (metallic and electronic) distortion of the background magnetic field of the Earth. Once characterized, these distortions can be inverted (hence removed) and the remaining magnetic vector, rotated to the X and Y plane, indicates direction of the magnetic North Pole, hence the glider's heading. Since the hard and soft iron fields of the glider remain constant in the frame of the glider, as we rotate the glider through different pitch, roll and heading orientations, we can distinguish the glider's fields from the constant background Earth field. Thus we need to sample a distinctive set of pitches, rolls, and headings in order to tease apart the glider's contribution to the overall measured field.

Empirically, we have determined that continuously sampling heading during a complete 360 degree turn at two different roll-angles, one to port and one to starboard, and at least two modest pitch-angles (e.g., 15-deg and 30-deg) both up and down suffice to create an accurate calibration. (More distinct pitches and rolls are better, but not practical during flight.)

The calibration procedures for on-shore and in-flight operation are described below. In addition to the software provided, you will need: 1) a terminal program with log/capture-file utility that you know how to use, 2) an installation of Matlab, version 2009a or better, which includes the Optimization Toolbox and, 3) for the in-flight calibration, an internet connection when computing the calibration. You should add the directory containing the calibration software to your Matlab path.

In the following discussion substitute your glider id for GGG.

The on-shore 'whirly' calibration procedure

Find a location with a known (or suspected) 'quiet' magnetic field away from large metal objects. If you have a access to one of the standard calibration locations, use the usual calibration jig. Otherwise, you can use a typical Seaglider shop jig on a relatively flat and smooth surface (a piece of plywood suffices) that allows you to spin the jig smoothly without excessive jarring or rattling. The surface doesn't need to be strictly level but it helps make it easier when you are doing the spins described below.

Mount the Seaglider pupa in the jig, place the jig on the spinning surface. Start a capture file on your favorite terminal program and start up the Seaglider.

First, if you are NOT at one of the known calibration locations around UW, record the latitude and longitude position and date by getting a GPS fix. (For example, attach an antenna pointed skyward and run the GPS selftest (menu hw/gps/selftest).) The known locations are the OSB highbay parking-lot location (the default location and abbreviated "hb" when asked by the Matlab scripts), the UW Urban Horticultural Center location (abbreviated "hc") , and the NOAA Sand Point location (abbreviated "sp").

Next enter the compass hardware menu and select the 'whirly' option (menu hw/compass/whirly -- NOT the 'whirlraw' option). After warming up, the compass should start streaming data to the terminal, which is being captured by the log/capture-file utility. (If you are using a recent glider software release, you should be able to pause and resume this flow using the space bar at any time. This is sometimes useful, for example, to pause recording data as you adjust the pitch and roll of the vehicle or if trucks/cars pass by your calibration fixture. If you can't pause, don't worry. It doesn't appear to matter as long as you can avoid trucks and other large moving magnets.) Hitting control-Q will quit.

By hand, roll the glider to port around 30 degrees. Watch the whirly output to determine the roll; it does not need to be precisely at 30 degrees. Then pitch the glider down around 30 degrees; again, this does not need to be precise. With the compass streaming data, rotate the jig (and hence the glider) in a 360 degree circle.

The spin direction does not matter. The rotation should take about 1 minute to get good data density (roughly 30 degrees per 5 seconds).

Once done, leave the roll as-is but pitch the glider up by 15 degrees (so now at -15-deg pitch). Take another 1-minute spin. Since direction doesn't matter you can reverse direction to avoid tangling the comms cable. Next set the pitch to positive 15 degrees, take another spin, then finally pitch up to 30 degrees and do one more spin (4 spins total at this roll angle). We usually skip 0 degrees pitch since a glider can't fly at that angle.

Next, with the glider's 'whirly' program still running, set the glider back level and roll it to starboard roughly 30 degrees. Then repeat the full pitch/spin sequence you performed above. (Each different 15-degree pitch spin paired on both port and starboard rolls increases the final RMS heading accuracy by about 0.1 degrees.)

When done, quit the 'whirly' progam with control-Q, then turn off the glider and close the log/capture file.

In Matlab run the "whirlymagcal" script.

```
>> whirlymagcal
```

The program will ask you for the id of the glider, the compass serial number, and the name of the log/capture file. It will then ask you for the calibration location and date. If the calibration was done at one of the standard UW locations, you can use the abbreviations noted above (e,g., hb, hc, sp). Otherwise, supply the recorded latitude and longitude (comma-separated in decimal degrees, negative values for south latitudes and west longitudes) and the program will determine the local geomagnetic field strength via the internet.

The program will also tell you if the data you collected is insufficient or if the field appears to be non-uniform.

Once done, the script creates a file named whirly_tcm2mat.GGG that you should install on the glider as tcm2mat.GGG. The glider can now navigate.

The in-flight calibration procedure

With the tcm2mat.GGG file from the on-shore whirly calibration on the CF card, launch the glider and achieve good pitch, roll and VBD trim during the initial dives as usual. Depending on the glider's magnetic history since the on-shore calibration, it may or may not steer a good course at this point but that doesn't matter for purposes of the in-flight calibration about to be performed. (Clawing off a lee shore might pose other problems but those are usually solvable by using the various navigation modes on the glider.)

To achieve a good in-flight calibration the glider must perform a minimum of two dives relatively nearby one another to roughly 120 meters. One dive forces the glider to spin with a constant roll to port (negative) on dive and climb, and the other forces spins with a constant roll to starboard on another dive/climb. One dive should have a 15 degree pitch on both dive & climb (so C_PITCH must be close to accurate to get symmetry), and the other dive gathers data on steeper, 30 degree pitch.

As with the whirly procedure above, the order and absolute precision of the rolls and pitches doesn't really matter so long as they are all done on dives in the same region. It is important, however, to complete a full 360 degree turn during each phase of each dive. Ultimately you should have 4 dive segments that mimic 4 of those you did ashore: a 30 deg roll with both +/- 15 deg pitch, and opposite 30 deg roll with +/- 30 deg pitches. If successful the data collected should yield a trustworthy and accurate compass calibration for that deployment.

To perform these calibration dives, the pilot first needs to prepare a temporary science file that samples sensors every 5 seconds to 200 meters. On the basestation, backup an current operational science file as science.orig, then modify a science file such that the first line lists 200 as the depth bin, and 5 sec as the repeat interval.

Next modify flight parameters in cmdfile to perform the spinning dives for calibration. You should record the operational values of any parameters changed below so you can reset them after the calibration dives. The basic idea of the parameter changes is to force the glider into a static roll (hence spin) on one side, then the other, by forcing the roll centers to approximate min and max roll AD limits.

For the first calibration dive (shallow pitch angles), modify the following parameters in the cmdfile:

// Record magnetometer data on each sample (critical!)

    $COMPASS_USE,4

// Ensure we don't enter an active roll

    $HEAD_ERRBAND,180

// Force a static port roll on dive and climb

    $C_ROLL_DIVE,<value of $ROLL_MIN>

    $C_ROLL_CLIMB,<value of $ROLL_MIN>

// Shallow pitch dive

    $GLIDE_SLOPE,15

    $MAX_BUOY,250

    $D_TGT,150

    $T_DIVE,65

    $T_MISSION,70

When you are ready to perform the first dive, send the modified science file and this cmdfile to the glider. Once the first dive is underway, modify the following parameters in the command file for the second dive (steeper pitch):

// Leave $COMPASS_USE,4 and $HEAD_ERRBAND,180 in place

// Force a static starboard roll on dive and climb

    $C_ROLL_DIVE,<value of $ROLL_MAX>

$C_ROLL_CLIMB,<value of $ROLL_MAX>

// Steeper pitch dive

$GLIDE_SLOPE,50

$MAX_BUOY,70

$D_TGT,150

$T_DIVE,50

$T_MISSION,70

Before the glider returns from its second dive, add $QUIT to cmdfile to hold the glider at the surface while you compute the calibration.

In Matlab, run the "divemagcal" script.

>> divemagcal

The code will ask you for the glider id, the compass serial number, and the dive numbers to use for the calibration (comma-separated). It will determine the location and date of the calibration from the GPS locations in the log files, then compute the calibration itself from the data in the .eng files and produce a file named dive_tcm2mat.GGG.

If the data were collected by a glider using a compass calibration that uses an old pitch and roll calibration (that is, not made with this package), then you must make the onboard tcm2mat.GGG file available to the divemagcal code. The code will use it to determine the original, measured pitch and roll for the purposes of the new calibration. If the onboard tcm2mat.GGG is not available, the code will assume the dive data reflect the measured (uncorrected) compass pitch and rolls.

(NOTE: If you cannot directly access the deployment directory, simply copy the pGGGyyyy.nc file -- or pGGGyyyy.log and pGGGyyyy.eng files -- for the calibrations dives to your matlab directory. These are the only files you need to compute the new compass calibration file.)

If all looks well from the matlab run, rename the created file as tcm2mat.GGG, and have the glider take the following pdoscmds.bat file, replacing GGG below with the id of your glider:

“pdoscmds.bat" file:

// backup old tcm2mat.GGG file

    ren tcm2mat.GGG tcm2mat.old

// install the new tcm2mat file

    xr tcm2mat.GGG

    strip1a tcm2mat.GGG

// use the new file

    menu hw/compass/coeff

If the commands in the pdos file are successful the glider will report the new coefficients you uploaded in the associated pdos log file.

Now restore the science file you want to use for regular data collection by renaming science.orig on basestation to science. In addition, restore the parameters in the cmdfile, especially HEAD_ERRBAND, C_ROLL_DIVE and C_ROLL_CLIMB, to return to normal flight operations. After these files are taken and all seems correct, $RESUME dive operations. The glider will now navigate using the new compass calibration.

The old "standard" calibration procedure

What about the original 60-point calibration, you ask? It still works, of course. However, we recommend you run the "rawmagcal" script:

```
>> rawmagcal
```

to perform the calibration regression and create a raw_tcm2mat.GGG file. It requests the glider id and the name of the tcmGGGxx.raw file you captured on the CF card, in addition to the compass serial number, location, and date. Using the tcmGGGxx.raw file is preferred to using the tcmGGGxx.ave file since the onboard averaging scheme does not take care of all the artifacts it should. Further, rawmagcal uses the compass pitch and roll rather than the asserted frame angles, which are not trustworthy.