

A reassessment of vessel coordinate systems: what is it that we are really aligning?

John E. Hughes Clarke
Ocean Mapping Group
Dept. Geodesy and Geomatics Engineering
University of New Brunswick
jhc@omg.unb.ca

Abstract

As the accuracy of component sensors in a swath sonar system have improved, the error budgets have increasingly been biased towards integration imperfections. A particular problem is confident alignment of ship coordinate systems and component sensors.

Whereas inter-sensor alignment is well understood and can be easily tested for dynamically (the patch test), alignment of any one sensor with the ships coordinate system, especially for cases where installations take place underway, is not currently adequately addressed in standard dynamic survey calibration procedures. Failure to address this can result in static position biases due to incorrectly reporting the lever arm offsets. Particularly now that RTK positioning is available the importance of these small biases will be of concern.

The use of dual-antenna GPS as a means of dynamically aligning sensors with the ships reference frame is proposed. By sequentially occupying well-established location pairs within the ships reference frame, the time series of azimuth and tilt of that pair may be directly compared with heading and motion sensors. Subsequent inter-sensor alignment (usually the sonar with the motion sensor) can now proceed with the confidence that both sensors may be properly referenced w.r.t. the ships reference frame.

Introduction

A critical component of calibrating an integrated swath sonar system has always been the proper relative location and alignment of all the integrated sensors. The alignment is both relative (one sensor with respect to another) and absolute (each sensor w.r.t. the ships reference frame). Two approaches have been traditionally used: measurement at time of installation, and subsequent field alignment.

Whilst alignment and offset measurement is normally attempted at the time of installation, no subsequent field survey is ever conducted without a field alignment test (normally a conventional patch test). The results of the field alignment are then normally entered as correctors into the integration matrix. The big question is where should those detected misalignment angles go?

The concerns about angular alignment have traditionally focused on the resulting error in the sonar-relative beam vector (Hare et al., 1995) which is normally far larger than the internal ship lever arms. For small vessels where physical offsets of sensors are minimised, the consequences of applying the misalignment angles to the lever arms in different coordinate systems has traditionally been inconsequential. Similarly in large vessels where the sensors offsets (whilst much larger) are generally tiny with respect to the sonar ranges involved (i.e. deep ocean surveying) again the consequences were inconsequential.

But now that centimetric kinematic positioning is a reality in both the horizontal and vertical axes, and sonars are starting to appear that can measure centimetres, we have to be much more careful about what we are aligning. Offsets measured in a ship coordinate system are only good as long as rotations around that coordinate system are reported. If a patch test result is used to alter the alignment of a motion sensor that was previously reporting the orientation of the ships coordinate system, all the previous offsets are now invalid. Unless the coordinate systems are properly understood much of the potential advantages of kinematic positioning cannot be realized.

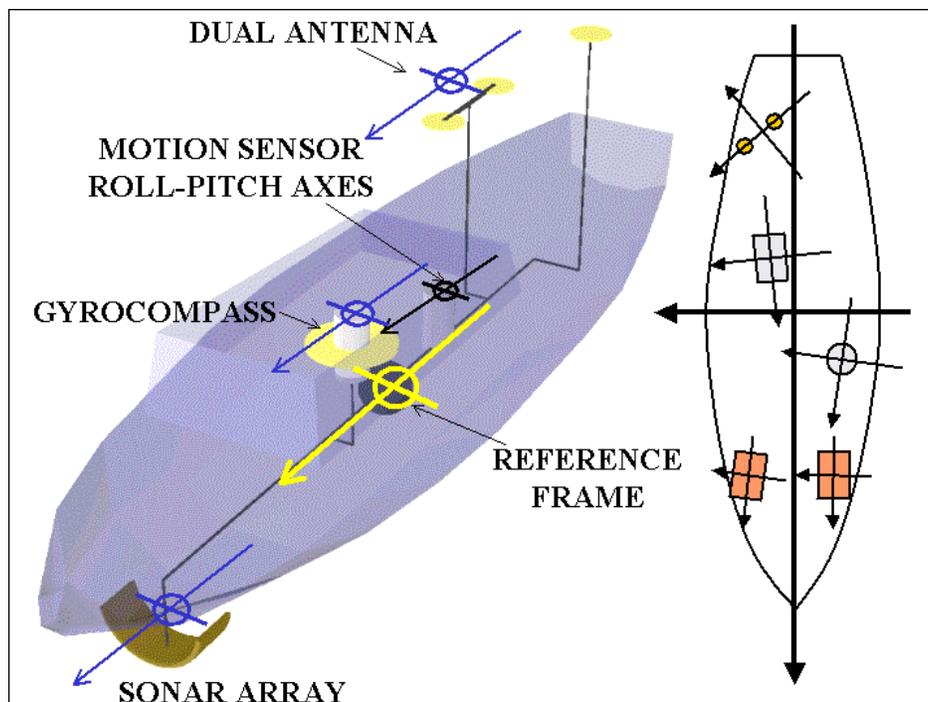


Fig 1: showing the offsets and coordinate systems of a typical suit of swath sonar integrated sensors.

Installation Measurement

Ideally, all instrumentation would be in place on the vessel in dry-dock or on a rigid trailer. As the vessel is immobile, a static measure of all relative locations and sensor orientations can take place.

Lever Arms

Before defining locations within a ships coordinate system the reference frame must be established. The definition of the X-Y axis plane and the origin of that plane (the reference point (RP)) needs to be specified. Ideally for this permanent reference markers must be established within the hull. The X-Y plane need not be parallel either to the geoid surface or the natural trim of the vessel when afloat (as the boat may be sitting at an incline out of the water). In practice, a local-level coordinate system will be established outside the boat and all sensors within surveyed in to that coordinate system by conventional static range and angle measurements (usually a theodolite type instrument). As the ship reference frame (SRF) markers are surveyed in at the same time, all measurements can be transformed from the external local level system to the internal SRF coordinate frame for use at sea.

Mounting Angles

Once the orientation of the XY axes (and by inference the mutually orthogonal Z axis) are defined with respect to the local level and north (or other reference azimuth), static tilt and azimuth of sensors can be measured and those angles again transformed to the ship reference frame. A caveat here is that there be precisely machined mutually orthogonal surfaces on the sensor from which orientation may be adequately measured.

All of the above, however, assumes that the sensors are in place at the time of the survey. If not, the intended position of the sensor may be adequately surveyed (with markers in the vicinity so that small offsets in mounting may be accounted for once installed). But the angular orientations are much harder to obtain once the platform is out of drydock. Even if the original sensor is in place at time of static survey, sensors (particularly the motion sensors) are often changed throughout a survey season. Thus we are left with an operational reality that angular alignments must be estimated dynamically.

Operational Alignment

The Patch Test

All of the above discussion depends on having the luxury of a truly static opportunity and that all sensors are already in place at that time. As these assumptions are removed the reality is that dynamic alignment becomes necessary. The most common method is the patch test (Herlihy et al., 1989, Hillard and Rulon, 1989, Godin, 1997), which is generally regarded as a sufficient test. But the patch test actually only solves for inter-sensor alignment (specifically the sonar against the source of roll, pitch and heading) and one inter-sensor time offset (sonar to positioning clock). Significantly, the patch test makes no measure of any of the sensors w.r.t. the ships reference frame. It has always been assumed that these alignments are already adequately known.

The three angular misalignments are only an estimate of the sum residual misalignment of the relevant angle sensor and the sonar. It is perfectly possible that the neither of these sensors are aligned with the ships reference frame (SRF) (other than grossly, Fig. 1).

When we detect a relative misalignment between two sensors, both of which we have measured approximate (but not precise) alignment with the SRF. How do we implement these extra angular shifts?

For example, when aligning the motion sensor (source of pitch and roll) and sonar, do we assume that one of the sensors is out of alignment with the SRF and the other is o.k. ? Or do share the misalignment offset between the two sensors (if so on what basis do we make this sharing assumption?).

For the purposes of estimating the resultant beam vector (the intersection of transmit and received steered cones) in the local level coordinate systems, the patch test is adequate as it solves for the sonar to motion sensor misalignments. But without a proper alignment of either (and thus by inference both as they are only relatively aligned) sensor w.r.t .the SRF, the lever arm offsets cannot be applied correctly.

For depth solutions where the resultant distance from the sonar to the seafloor is large w.r.t. the sonar to positioning sensors error, this is generally an inconsequential error. Inversely then, where surveys are performed in shallow water with larger inter-sensor offsets, the failure to have proper alignment with the ships coordinate system can result in errors magnitudes that are dominated by the sensor offset errors rather than acoustic range and bearing solution errors.

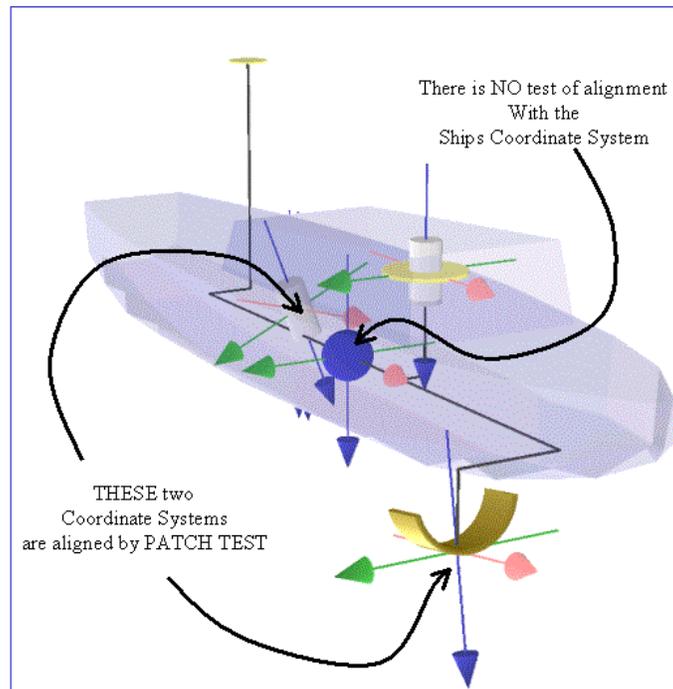


Fig. 2: showing the alignment that patch test solves for, which significantly bypasses alignment of either sensor w.r.t . the SRF (unless the misalignment is gross, for example the offsets are put in backwards).

If the orientation sensor alignment to the SRF is imperfectly known, improper application of lever arms comes about as a result irrespective of how well the lever arms are measured.

As the sensors will be installed and aligned at least by eye, typical misalignments are likely to be no more than a ~ 2 degrees in heading, a degree in roll but possibly up to several degrees in pitch. The heading alignment is usually reasonably easy to guess as it reflects the ships longitudinal axis (something that can be well monitored by surveying in two locations from dockside). The roll alignment however is a tilt over the narrower breadth of the vessel rather than a bearing and thus leveling would be required, a much more tricky problem.

The pitch misalignment has been noted by this author to be the most problematic. This is because for small ship installations, often the sonar is used as the RP, with implicitly the sonar orientation as mounted, reflecting the orientation of the SRF. Also, for hydrodynamic reasons the sonars usually has a +ve pitch alignment from the local level. Quite commonly however, the ships reference frame has been measured from the local level (or at least the boat's usual trim whilst alongside) and thus the two are misaligned in pitch.

Consequences of small mis-alignments with SRF

With the few degree angular misalignments realistically possible in dynamically installed sonar systems, this translates to lever arm position-errors in X, Y and Z that are directly proportional to the length of the lever arm and angular misalignment away from that axis ($\sim 1.75\%$ of the lever arm per degree).

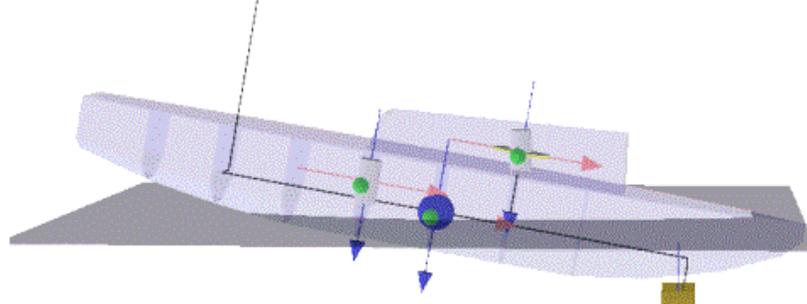
Typical small boat lever arms are a few metres and thus we have a sub-decimetre typical scale of potential error. For horizontal positioning this is rarely limiting as, even if RTK solutions are used, beam horizontal dimensions are larger than this (a 1.5° beam at 45° in 20m of water is $\sim 75 \times 105\text{cm}$) and small-unresolved time delays account for more than this (0.2 seconds at 10 knots is worth 10cm).

But the vertical component rapidly becomes significant. How important, depends on whether you are using RTK for the vertical solution or heave and induced heave. Induced heave calculations only require RP to sonar lever arms, which are invariably shorter than the RTK antenna to sonar lever arms.

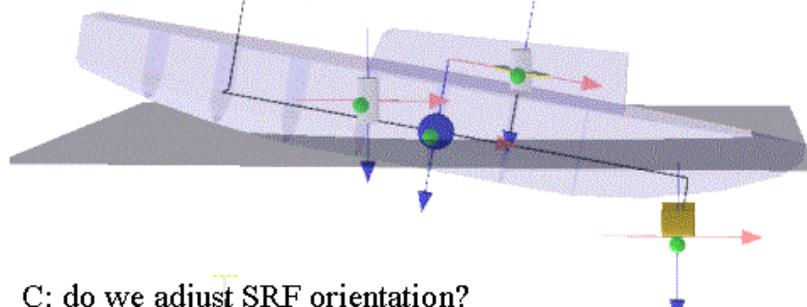
For moderate scale inshore vessels (30-50m), lever arms from RTK-type antennas routinely exceed 10m vertically and fore-aft. To minimize the vertical component of the error, having the RTK antenna nearly over the sonar can minimize the X and Y levers, but this is often not geometrically possible for installation reasons.

If you are using heave and induced heave, one can check the validity of the alignment by reference to the water surface. This is done by measuring both the elevation of the RP, and the draft of the sonar, at dockside. Normally the waterline to RP is the variable that is adjusted daily to reflect loading changes. Therefore the draft of the sonar should reflect the sum of this together with the Z component of the RP to sonar lever arm adjusted for static trim (the induced heave). If the two does not match then one can assume that the motion sensor output is not adequately aligned with the RP.

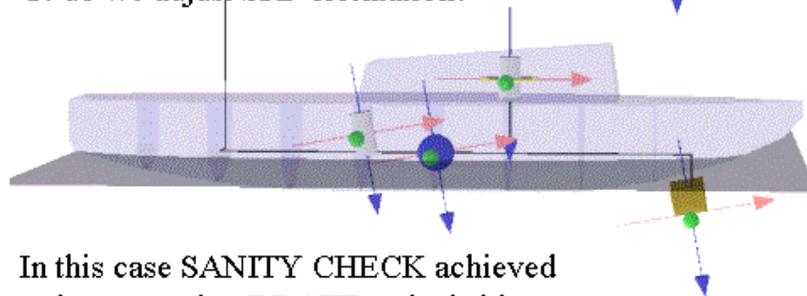
A: we detect a pitch misalignment between IMU and SONAR



B: we adjust IMU output



C: do we adjust SRF orientation?



In this case SANITY CHECK achieved
by measuring DRAFT at dockside

Fig 3: illustrating the problem of applying a pitch misalignment detected through patch test between the sonar and motion sensor. Unless one of the two sensors is already confidently oriented w.r.t. the SRF, the problem is not solvable without an independent check of the elevation at two points in the vessel.

The same sanity check, (looking at the draft alongside) is not however possible if the source of one's vertical is derived from the ellipsoid. Unlike the water surface, the

ellipsoid is not a conveniently level visible plane from which one can make simple relative height measurements (Fig. 4).

Given the limitations of the conventional patch test, it is clear that we need to establish a method for dynamic alignment of the SRF with the motion sensor. The method herein proposed takes advantage of dual antenna GPS systems that are not tied to a fixed baseline.

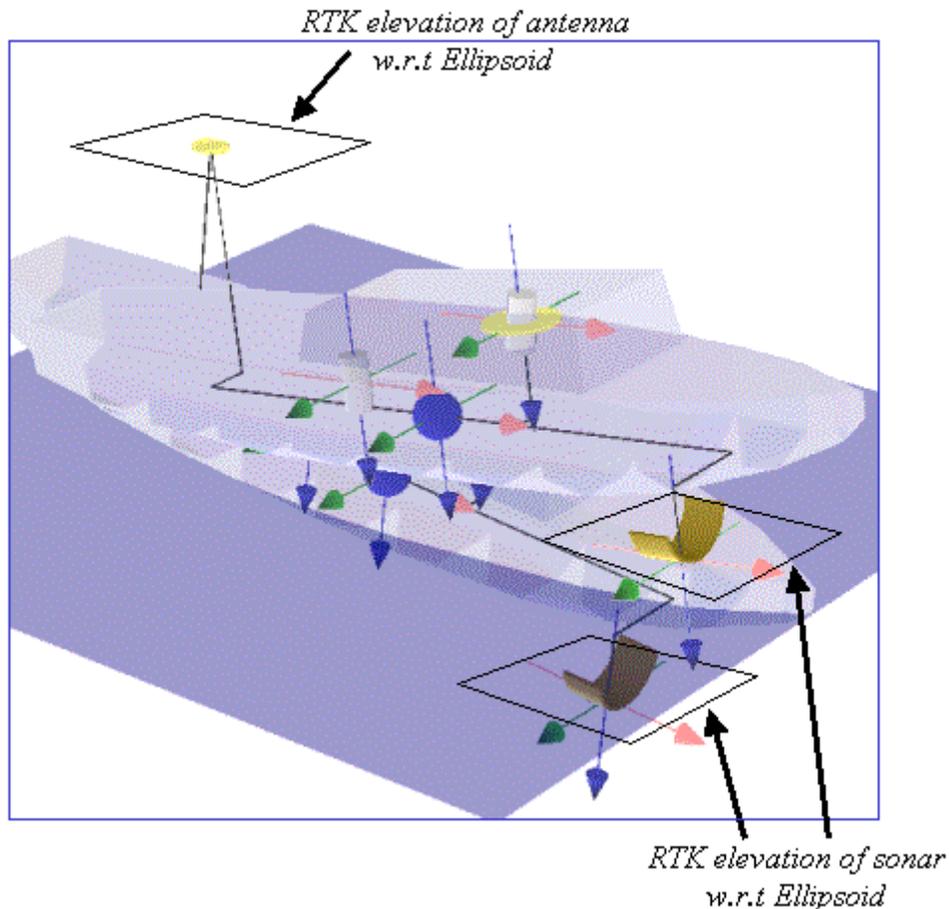


Fig 4: the effect of propagating the position of an RTK fix at the antenna due to a small undetected pitch misalignment between the SRF and the motion sensor. Note the water surface provide no useful confirmation on the validity of the ellipsoid height solution

Multiple-antenna GPS orientation

Multiple antenna GPS was tested as a possible means of providing swath sonar orientation in the early 90's (Lu et al., 1993) and, although the accuracies were shown to be adequate, the data rates and data lag times resulted in their sole use being impractical for real time swath sonar systems. However very rapidly, a subset of the 3-4 antenna

geometry - dual antenna heading, was integrated into a number of inertial sensor systems (POS/MV Applanix 1994; Seapath 200, Seatex, 1996).

The dual antenna systems are now increasingly sold as standalone heading sensors without direct inertial integration (e.g. Trimble MS860). A subset of these systems can also provide single axis tilt measurements, and thus can, if installed along a known baseline in the SRF, provide a measure of the SRF orientation dynamically.

There are two common installation configurations for stand-alone GPS azimuth sensors:

1. Supplied with a rigid factory-measured mounting bar. This reduces the ambiguity resolution requirements, but provides the problem of installing the axis precisely in both in bearing and tilt. The bearing is usually o.k., but as these mounting bars are normally supported at just a single point it is very difficult to be confident about the exact tilt of the mounting bar w.r.t. the SRF if installed at sea.
2. Supplied as two antennas without specific installation requirements. These systems solve for the baseline length dynamically at start up time (usually requiring about 30-90 seconds after power up) and provide a measure of the bearing and pitch of the baseline w.r.t. the ellipsoid. As the antenna can be mounted anywhere, they may be installed on pre-existing mount points previously surveyed in at time of static survey. Thus a confident alignment with the SRF can be obtained even if installed under dynamic conditions.

The second method could potentially be used to survey in the vessel underway and provide a time series of single axis tilts w.r.t. the SRF to compare directly with the dynamic output of the motion sensor. By comparing the two time series, the static misalignment along that axis together with any evidence of cross-talk between the two axes can be estimated dynamically.

Example application

A Trimble MS860 system with the second configuration (i.e. not supplied with a fixed baseline, but rather can be mounted at any separation and orientation) was tested in November of 2002 on CSL Heron. Heron is a 10m survey launch equipped with a Simrad EM3000 and a Seatex MRU-6 motion sensor. The two GPS antennas were mounted on rigid supports, previously surveyed in and known to be in-line fore aft and level with the SRF X axis. The baseline used was 3.04m. The baseline azimuth was used as the prime source of heading for the vessel. Note in order for the azimuth to be useful, the antennas have to be at the same elevation in the SRF, otherwise roll motion when the vessel is pitched will provide a false azimuthal deviation.

The baseline azimuth solutions were generated at 1 Hz and fed into the MRU-6 to be inertially smoothed and re-output at 100 Hz. The magnetic heading of the MRU-6 was not used as the sensor was mounted too close to the alternator and thus strongly affected by spurious induced magnetic fields.

The MS860 outputs a custom Trimble ASCII string (PTNL) which provides both an azimuth and tilt of the antenna pair w.r.t. geographic north and the ellipsoid respectively. The ASCII data, time stamped with UTC was logged asynchronously. The MRU-6 orientation was logged by the Simrad EM3000. The Simrad clock was synchronized at the beginning of the day, but a 1 PPS signal was unfortunately not used. As a result the drift between the two clocks had to be measured by comparing the Simrad clock time stamp on the navigation fix and comparing with the original GGA UTC time (recorded with each fix). All resulting time series have been adjusted to a common UTC time base.

The data were collected during a conventional multibeam bathymetric survey. The vessel was operating at approximately 10 knots at most times, acquiring ~ 10 minute survey lines spaced on reciprocal headings at ~ 60m offset.

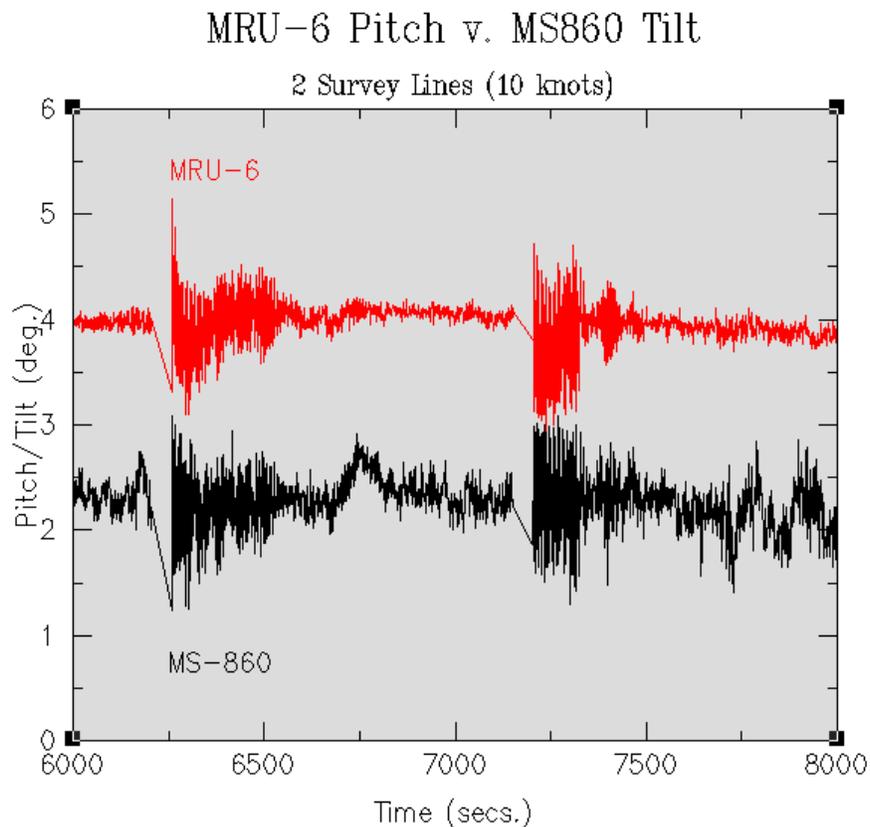


Fig. 5: 2000 second time series plot showing MS860 baseline tilt solutions at 1Hz, compared to MRU-6 pitch measurements at the same epoch (the two instruments were tied together using a common UTC reference). Two inverse-direction survey lines are shown here where the vessel steams through it's own wash at the beginning of each line. Most of the motion characteristics appear identical other than a 1.7° constant offset between the two. The MS860 does occasionally, however, show 100+ second drift events of $\sim 0.3^\circ$ in tilt (at 6750 seconds and 7750 seconds in the above plot). The vessel did not change speed, heading or roll during these events (a straight survey line in near flat calm conditions) and thus the MRU solution (showing no pitching event) is believed to be a more faithful representation of the motion of the vessel.

The survey duration was ~4.5 hours during which the two sensors (MRU-6 pitch and MS860 baseline) tilt were recorded at a variety of vessel speeds. Figure 6 below shows a cross plot of the two estimates of the Y axis orientation of the platform for the whole survey day. As can be seen the survey vessel routinely changes trim over 5 degrees as it accelerates from 0 to 10 knots.

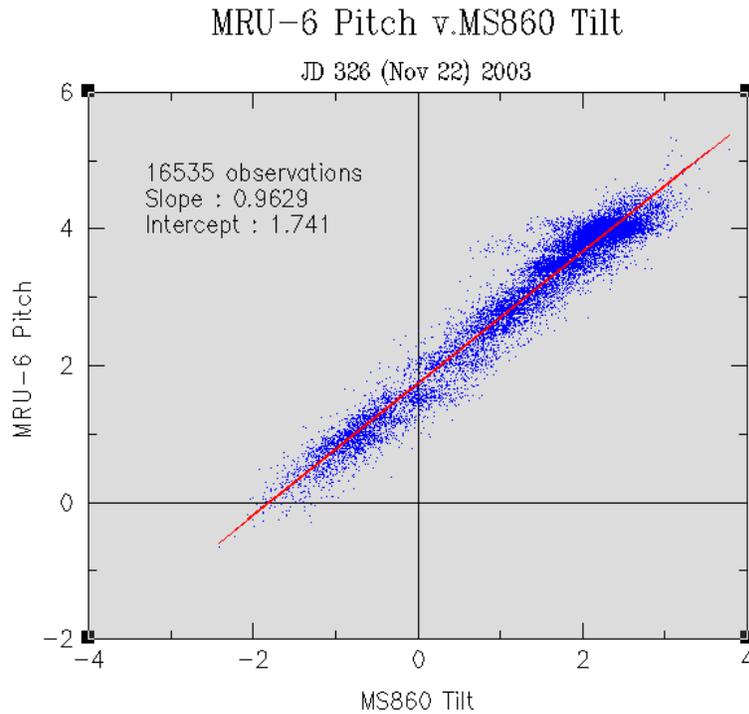


Fig 6: cross plot of all MS860 tilt values (at 1Hz), v. equivalent MRU-6 pitch (down sampled at each epoch). 4.5 hours of collection with abrupt speed and heading changes. The intercept of the regression line on the Y-axis indicates the static pitch bias between the MRU sensor frame and the dual antenna baseline which was exactly fore-aft aligned at a constant elevation in the SRF. The 1.74° shift thus is a robust dynamic indication of the motion sensor alignment w.r.t. the SRF.

From Figure 6 it is clear that, despite a combination of noise in both of the motion sensors, over long time periods (4.5 hours used here) a clear linear trend between the two solutions may be determined providing a reliable dynamic method of measuring the SRF to motion sensor coordinate alignment on one axis.

Interestingly the regression of the pitch v. tilt measurements indicates a slope that is not precisely 1.0. This could reflect non-linearity in the performance of either of the sensors. Another option however is that there is a yaw misalignment of the roll and pitch axes. If so it would require $\sim 15^\circ$ of misalignment ($\cos^{-1}(0.9629)$). This is unlikely and it can be examined by looking for evidence of crosstalk.

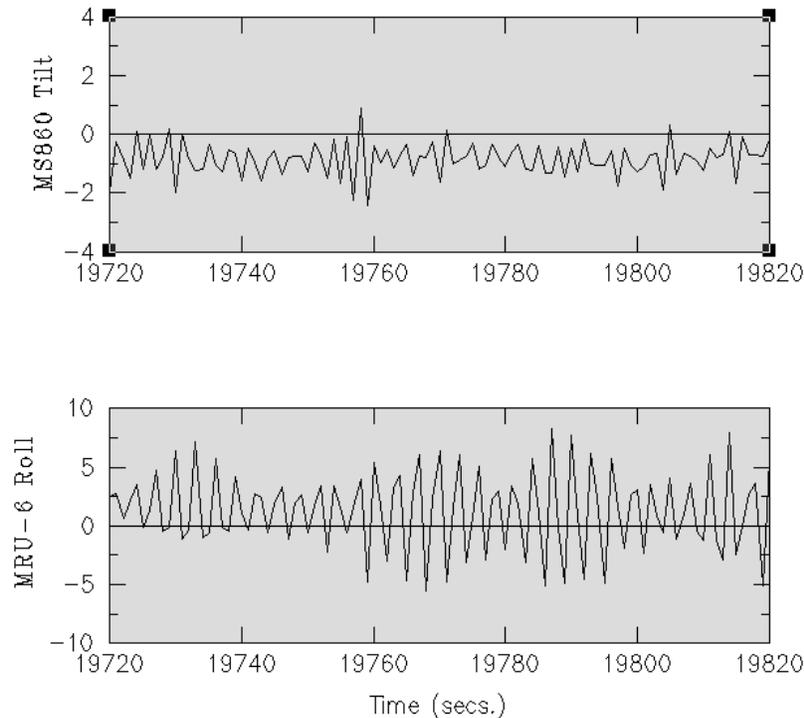


Fig. 7: time series of MRU-6 roll and MS860 Tilt indicating that no cross talk is in evidence indicating that the MRU pitch axis is adequately alignment to the SRF.

To test this, Figure 7, shows the one time period in the day when significant roll was experienced ($\pm 5-7^\circ$ when the vessel steamed along its own wake). As can be seen from the equivalent MS860 tilt time series, none of this visibly leaks through into the tilt measurement (a 15° misalignment would have seen a $\sim 25\%$ leak through).

Although not attempted at the time of this experiment, an across-vessel baseline, using antenna mounts precisely surveyed into the SRF would have allowed measurement of an equivalent Y axis alignment between the SRF and motion sensor. Care should be taken here as the tilt measured is equivalent to the Hippy roll convention as it is the angle from the local level, not the pitched XY plane.

Conclusions

As the accuracy of aiding sensors improve, the requirement for commensurate accuracy on SRF location and alignment must correspondingly improve. Increasingly static surveys are being conducted to establish precise reference frame within a ship and locate sensors. Establishing alignment is a harder problem and compounded by the operational fact that sensors are often not in place at the time of static survey (or are replaced at sea).

Whereas, approximate alignment (within a degree or two) used to suffice for sensor to ship reference frame, this alignment now needs to be almost as good the inter-sensor alignments. Whilst dynamic inter-sensor alignment procedures have been available for many years (the patch test), the procedures do not address SRF to sensor alignments.

Without the luxury of perfect static alignment surveys to meet this new requirement, a method for dynamic alignment with the SRF needs to be established. Herein a method using dual antenna GPS elements mounted sequential at precisely know locations within the SRF is shown to be a viable dynamic method.

Acknowledgements

This work is funded through the sponsorship of the Chair in Ocean Mapping. Current sponsors include the Canadian Hydrographic Service, Kongsberg Simrad Mesotech, U.S. Geological Survey, U.S. Oceanographic Office, CARIS, State Univ. New York, Univ. New Hampshire. The generous loan of the Trimble MS860 was made possible through McQuest Marine Ltd.

References

- Applanix, 1994; POS/MV Operation Manual : Applied Analytics Corporation, technical documentation.
- Godin, A., 1997, The Calibration of shallow water multibeam echo-sounding systems: MEng Thesis, University of New Brunswick, 184pp.
- Hare, R, Mayer, L. and Godin, A., 1995, Accuracy Estimation of Canadian Swath (multibeam) and Sweep (multi-transducer) sounding systems: DFO, CHS Internal Report.
- Herlihy, D.R., Hillard, B.F., Rulon, T.D., 1989, National Oceanic and Atmospheric Administration Sea Beam "Patch Test" Manual: Ocean Mapping Section, Office of Charting and Geodetic Services, NOS, 34 pp.
- Hillard, B.F., Rulon, T.D., 1989, National Oceanic and Atmospheric Administration HydroChart II system "Patch Test" Manual: Ocean Mapping Section, Office of Charting and Geodetic Services, NOS, 35 pp.
- Lu, G., M.E. Cannon, G. Lachapelle, and P. Kielland (1993). Attitude determination in a survey launch using multi-antenna GPS technology. Presented at The Institute of Navigation National Technical Meeting, San Francisco, Calif., 20-22 January, The Institute of Navigation, Washington, D.C., 9 pp. 5121.
- Seatex 1996, Seapath 200 Operator Manual: Seatex Inc., technical documentation.