

USING DGPS TO MEASURE THE HEAVE MOTION OF HYDROGRAPHIC SURVEY VESSELS

by Peter KIELLAND¹ and John HAGGLUND²

Paper presented at the Institute of Navigation's National Technical Meeting "Navigating the 90's: Technology, Applications and Policy", Anaheim California, USA, 19-20 January 1995.

Abstract

One quite significant error source encountered by hydrographers is wave induced vertical motion of their survey vessel (heave). In heavy swells, uncorrected heave noise will degrade the accuracy of the surveyed soundings upon which mariners rely for safe navigation. Heave motion can be measured using inertial technology thus enabling the raw surveyed soundings to be corrected to calm water conditions. Unfortunately, the high cost of inertial heave compensators has prohibited their widespread use. This paper documents a test carried out by the Canadian Hydrographic Service in which very accurate relative position derived from GPS phase observations were used to determine heave corrections for a hydrographic survey vessel. The algorithm is simply a high pass filter acting on the unused DGPS vertical position record already being observed on the vessel. An inexpensive pitch and roll inclinometer is used to correct for the lever arm effect between the GPS antenna and the sounder's transducer. The experiment indicated that decimetre heave compensation accuracy was obtained.

INTRODUCTION

Hydrographic data plays a major role in assuring safe navigation. From a navigator's perspective, the safety intrinsic to any surveyed seafloor model is its

¹ Canadian Hydrographic Service, Ottawa, Ontario.

² Nortech Surveys Canada, Calgary, Alberta.

fidelity with respect to the real seafloor. If a seafloor model defined by hydrographic data perfectly models the true seafloor then that data is perfectly safe. All bathymetric seafloor models are less than perfect and some are very imperfect. This continuum of degrading spatial fidelity is critical knowledge to mariners. When deciding how close they can safely approach shoal features portrayed by a seafloor model, mariners need to know just how reliable that picture is. This requirement for safe navigation defines the two prime directives for hydrographers:

- 1) Minimize the spatial imperfections in the surveyed seafloor model.
- 2) Estimate whatever level of spatial infidelity remains after all data collection and modeling is complete.

This paper deals with only a small part of the first of these two tasks. If we wish to minimize the spatial imperfections in a seafloor model, there are two types of error sources that must be controlled. The first is the interpolation errors caused by merely sampling discrete point soundings over the continuous seafloor surface. Reducing interpolation errors can only be accomplished by observing depth soundings at closer intervals over the seafloor being surveyed. As a rule: the rougher the seafloor being modeled, the denser the depth samples must be. The desire to totally eliminate interpolation errors has given rise to multibeam echo sounders which increase the efficiency of collecting extremely dense soundings. For sparser data sets: after data collection is complete, gridding techniques based on geostatistics can be used to estimate the probable magnitude of interpolation errors at every location on the interpolated seafloor surface model [1].

The second error source that can be minimized is instrumental errors which contaminate either the position or depth coordinates of the sampled data points. There are a host of factors which contribute to instrumental errors. Increasingly sophisticated signal generation, signal timing and signal propagation modeling have been developed to minimize them. A good example of this trend towards sophisticated instrumentation is GPS. There are many other examples of electromagnetic and acoustic survey systems in which instrumental error sources have been efficiently minimized. One instrumental error source specific to hydrography is the wave induced vertical heave motion affecting bathymetric data. A low cost means of reducing these heave errors is the subject of this paper.

HEAVE ERRORS

In order to compensate for the effect of changing water level (due to tidal effects or other environmental factors), all measured depth soundings must be reduced to a vertical sounding datum. This reduction to datum enables navigators to later make use of the surveyed soundings by adding their current local water level elevation onto the charted depths. The water level elevations can be obtained from published tide prediction tables. The sounding datum elevation is established by first analyzing a history of water level elevations. A statistical value is then determined for "Lower Low Water Large Tides": an elevation which the water level very rarely falls below. During a hydrographic survey, a tide gauge keeps track of the water

level elevation above this datum and the water level elevations are subtracted from all acoustically measured depths.

This whole data reduction process is essentially a water level transfer. It requires a noise free (calm) water surface for projecting the tidal elevations observed at the gauge site out to the location of each surveyed sounding. Any uncorrelated noise (such as wave action) that exists between the tide gauge location and the depth sounder location adds directly to the overall uncertainty of the charted soundings. If for example the survey vessel is sounding in 1 metre swells, then all of its measured depths will have a +/- 1 metre uncertainty. For many critical survey situations these instrumental errors must be reduced.

CONVENTIONAL HEAVE COMPENSATION

There are two traditional methodologies for reducing heave noise in the data.

1) Analog heave reduction

On surveys which still employ analog data capture and processing techniques, the graphical sounder trace must be visually interpreted to select and digitize depths. This presents an opportunity for the hydrographer to identify any oscillations in the seafloor trace that have been induced by heave motion. A regular saw-tooth pattern often indicates the presence of heave induced noise. The reliable identification of heave artifacts during data processing is facilitated if during data collection the hydrographer has annotated the analog trace with sea-state commentaries. If the data digitizer is quite certain that a saw-tooth artifact has been induced by heave noise, then the analog trace can be visually smoothed to simulate calm water conditions. This manual heave compensation technique has obvious shortcomings, particularly if the seafloor is rough. For example, when examining an irregular shoal structure it becomes impossible to reliably differentiate between a spike on the graph that was caused by the survey launch falling into the trough of a wave and an identical spike caused by a dangerous rock on the seafloor.

2) Digital heave compensation

On more modern surveys that log and process data digitally, it becomes feasible to continually measure heave motion and correct all measured soundings. Typically this is done using inertial technology to measure vertical acceleration at the location of the sounding transducer on the vessel. By integrating acceleration twice, the heave motion can be computed in real-time and used to correct the soundings as they are logged from the echo sounder.

The one major drawback of inertial based heave sensing is its high cost. Heave sensors based on a simple triad of accelerometer and rate of turn sensors (such as the TSS 335B) start at about \$30K (US). Damped pendulum based sensors (such as the HIPPY) are also in this price range. More sophisticated gyro based systems (such as the POS/MV) can cost over \$100K. This high capital cost has

limited the use of heave compensation to critical surveys carried out by well funded agencies. Multibeam echo sounders require pitch roll and yaw information (also available from a heave sensor) in order to form and steer the multiple acoustic beams. Heave compensation is therefore considered compulsory only for multibeam sounding operations. While technology has advanced towards somewhat lower cost sensors, inertial heave compensators continue to be relatively high cost survey system components.

A GPS BASED HEAVE COMPENSATION ALGORITHM

Differential GPS presents a third possible method of heave compensation; one which eliminates the need for expensive inertial sensors. By exploiting phase tracking information from GPS receivers we can very accurately measure a change in elevation. Uninterrupted carrier phase tracking can in theory detect changes in elevation of less than a centimetre from one epoch to the next. This characteristic is exploited by the heave compensation algorithm described below.

Heave is considered here to be the height of the vessel's sounding transducer relative to its location "during calm water". Since over time vertical wave movement appears as a random noise, the elevation of calm water should be equal to its average elevation over a period of a minute or two. In this heave algorithm, a moving weighted mean is used to define this "calm water" reference elevation. The weight used for each elevation observation used in the weighted mean is its standard deviation (estimated by the 3D GPS position adjustment). The epoch of each heave value is contained within the sample used to compute the weighted mean reference surface. At each of these epochs, the computed heave value is simply equal to the difference between the instantaneous GPS elevation and its moving mean reference elevation.

The number of measurement epochs in the moving window used to compute the weighted mean is user defined. In this experiment, a variety of window durations between 50 and 250 epochs were tried. One hundred (one second) epochs appeared to adequately model calm water yet still permit the smoothed reference surface to follow trends which might otherwise bias the heave values. Such trends can be caused either by real tidal movement, changes in the trim and draft of the vessel, or any of the many GPS instrumental error sources affecting the absolute elevations. The location of the single heave epoch within the moving data window can also be specified by the user (leading edge, centered or trailing edge). In this experiment we used centered windowing since it provided better heave results. Its only drawback is that, in order to consider both the past and future on either side of the single epoch, centered windowing constrains the algorithm to post processing of logged data. Heave differencing at the windows edge is more susceptible to biasing the heave results but has the advantage of being easily implemented in real-time. Near real-time, centered windowing (heave output delayed by half the window duration) could also be implemented by maintaining appropriate data buffers.

The weighted mean algorithm is implemented as follows:

$$\bar{X} = \frac{\sum_{i=1}^n Wt_i X_i}{\sum_{i=1}^n Wt_i} \quad \text{where} \quad Wt_i = \frac{1.0}{StdDev X_i * StdDev X_i} \quad (1)$$

The traditional formulation of the Standard Error of the Arithmetic Mean for observations with unequal weighing is given by:

$$S_{\bar{x}} = \sqrt{\frac{\sum_{i=1}^n Wt_i (X_i - \bar{X})^2}{(n - 1) \sum_{i=1}^n Wt_i}} \quad (2)$$

to ease the computational burden of maintaining data buffers, this was reformulated as follows:

$$S_{\frac{2}{x}} = \frac{\sum_{i=1}^n Wt_i X_i^2 - 2\bar{X} + \bar{X}^2 \sum_{i=1}^n Wt_i}{(n - 1) \sum_{i=1}^n Wt_i} \quad (2)$$

It should be noted that this Standard Error of the Weighted Arithmetic Mean is actually reflecting the tightness (precision) of the data used to calculate it. It may in fact be very precise, but not necessarily very accurate (due to biasing). As long as any bias is varying slowly relative to the time interval of the window selected, accurate heave estimates can be determined. By definition, whatever bias might exist in the population used to compute the moving weighted mean is also present in the single epoch height value. The elevation bias therefore drops out in the subtraction used to calculate the instantaneous heave value.

If the standard deviations assigned to the X_i observations truly reflect the uncertainty of the observation, an estimate of the accuracy of the Weighted Mean elevation can be calculated using the weights only:

$$S_x = \sqrt{\frac{1.0}{\sum_{i=1}^n Wt_i}} \quad (3)$$

A previous experiment based on the HPM processing software used in this experiment demonstrated that the estimated standard deviations of the X_i are in fact statistically non biased [1]. Equation 4 is therefore a good indicator not only of the

precision of the elevation of weighted mean reference surface but of its absolute accuracy.

PITCH AND ROLL REDUCTIONS

The raw "heave" from the algorithm described above is merely the high frequency vertical movement of the GPS antenna ... not the heave at the vessel's sounding transducer. Since the GPS antenna must be mounted away from the center of gravity of the vessel, angular pitch and roll movement of the vessel induce a vertical movement of the GPS antenna. This vertical pitch and roll movement contaminates the heave computed for the offset location of the sounding transducer and therefore must be removed. Before using any GPS elevations either for computing the weighted mean height or for differencing the instantaneous heave values, the elevations must first be corrected for vessel pitch and roll effects.

As a rule, the GPS antenna will be located at some considerable distance from the sounding transducer. This forms a lever arm around which the pitch and roll of the vessel add noise to the heave sensed at the GPS antenna. The effect of the antenna's lever arm offsets (defined in the vessel's structure coordinates by D_x , D_y and D_z) can be removed from the heave signal if the angular pitch and roll of the vessel are measured. The lever arm correction to each height is given by (5):

$$Corr_{height} = -\cos(\text{roll})\sin(-\text{pitch})D_{x_{ant}} + \sin(\text{roll})D_{y_{ant}} + \cos(\text{roll})\cos(-\text{pitch})D_{z_{ant}} \quad (5)$$

From equation 5 it is apparent that the horizontal components of the antenna lever arm (D_x and D_y) should be as close to zero as possible. For a vertical lever arm (only D_z), the effect of angular errors on the heave correction is a pure cosine function and therefore has very little effect on the heave values. For example: through a vertical 10 metre lever arm, a 1° pitch and roll error produces only a 0.2 mm heave error. As the inclination of the lever arm increases, the effect of angular error increases as a sine function. Thus, at a rather extreme pitch or roll angle of 20° , the same 1° measurement error over 10 metres would induce an error of 6 cm in the instantaneous heave. In order to minimize pitch and roll induced errors, it is thus very important to mount the GPS antenna as close to the sounding transducer as possible (D_x and $D_y = 0$). The pitch and roll sensor itself should be located close to the center of gravity of the vessel to minimize lateral forces which could increase its susceptibility to overshoot errors. On typical survey launches, the antenna lever arm is less than 10 metres and it is often quite feasible to mount the GPS antenna directly above the transducer.

Since GPS data is already observed for horizontal positioning of the sounding data, any GPS derived heave compensation is essentially free. However, to realize this potential cost-effectiveness, the pitch and roll data required for lever arm corrections must come from an inexpensive sensor. Digital inclinometers with a claimed angular accuracy of $\pm 0.2^\circ$ can now be purchased for under \$0.5K. Since these sensors are currently based on a viscous electrolytic fluid, it is obvious that $\pm 0.2^\circ$ could not be maintained in a high dynamic environment. For this application

however it is not really necessary to use a $\pm 0.2^\circ$ sensor. Occasional 1 or 2 degree pitch and roll errors would still support useful and cost effective heave compensation. Further field testing is required to see how inexpensive inclinometers perform in different dynamic environments. In any event, inexpensive magnetic fluxgate technology is under development which should provide pitch and roll sensing that is impervious to lateral acceleration. Provided the horizontal components of the antenna's lever arm are kept small, it appears at least feasible to use a low cost pitch and roll sensor to provide sufficiently accurate lever arm corrections.

DATA COLLECTION AND TEST METHODOLOGY

Significant instrumental errors were introduced onto the two data sets used in this experiment. Unfortunately, logistical constraints prevented re-observing a "better" data set in time for preparation of this paper. Despite some uncertainties inherent to this test data, the GPS heave estimates derived from it are presented and discussed below. These results provide useful insight into the potential for GPS heave compensation and have helped to identify how to conduct a more conclusive field test in the future. The sub-optimal data sets were actually quite helpful in devising a more robust QC procedure that will be implemented in the production heave algorithm.

The primary objective of this experiment was to ground truth GPS derived heave estimates to see if they are "sufficiently accurate" for use on production surveys. Raw acoustic depths are generally measured with a resolution of 1 dm, however the heave noise in this signal can easily be 10 times that. An optimal level of heave compensation accuracy would therefore be 1 dm, however 2-3 decimetre heave compensation accuracy would often provide a useful improvement over uncorrected soundings. In fact, the accuracy requirement for heave compensation should be dependent on sea state. For example, in 2 metre swells even heave corrections accurate to ± 0.5 metre would still significantly improve the quality of the bathymetry. In calmer seas, applying the same ± 0.5 heave compensation might actually degrade the raw soundings. This sea state dependent accuracy requirement is used in the Quality Control procedure described later in this paper.

When planning the data collection mission for this experiment, it was expedient to make use of a TSS model 335B heave compensator to provide the reference heave values. A TSS 335B had already been purchased by CHS as part of a SIMRAD EM1000 multi-beam sounding system installed aboard the F.G. CREED. The TSS 335B is based on a relatively inexpensive strap-down triad of accelerometers and angular rate sensors. Recent CHS testing has shown this configuration can produce attitude errors during hard turn maneuvers [2]. It has since been replaced with a more costly gyro based inertial platform for providing attitude information to the multi-beam sounder. At the time of data collection for this experiment, the TSS was however deemed sufficiently accurate to act as a ground truth heave sensor.

The DGPS sensor originally selected for this experiment was the Novatel 951 narrow correlator C/A code receiver. The reason for preferring the Novatel is that

all its raw pseudorange and phase measurements are spatially qualified with estimates of standard deviation. The Novatel is the only receiver we are aware of that outputs error estimates for every code and phase observation. A previous CHS experiment documented significant improvements in both positioning performance and position error estimation capability when the Novatel raw data error estimates very fully exploited [1]. The improved performance resulted from using the Novatel raw data error estimates to form a dynamic weighing matrix within the position solution in lieu of the normal practice of using constant a priori weights for each code and carrier measurement.

It was originally planned to use Nortech Survey's (HPM) "Hydrostar Post Mission" PC software to process the L1 only Novatel data. HPM, and its real-time version called HPC are used by CHS for GPS processing and Quality Control. As explained above, the HPC/HPM software had previously been optimized for use with Novatel raw data however raw data from other receivers can be decoded (Trimble, Ashtech and Magnavox). HPC also provides a helmsman's navigation display for following survey lines and can both log and time stamp input from any digital data sensor. The heave estimation algorithm described here was to be developed and tested using the HPM post mission software prior to a possible port to the HPC real-time software. As it turned out, the algorithm proved to be more efficient as a post-mission application.

The first attempt to collect data took place on board a 10 metre survey launch near Victoria BC. HPC was used to log and time stamp all data. The TSS heave data was logged at 20 Hz, all Novatel raw GPS data at 2 Hz and the pitch and roll data from a Trimcube digital inclinometer was logged at 10 Hz. The Novatel raw data at the reference site on shore was logged at 1 Hz. Unfortunately sea state conditions during the time allotted for the heave experiment were flat calm. In order to stimulate heave conditions, it became necessary for the survey launch to perform violent corkscrew maneuvers in the wake of the Vancouver to Victoria ferry. All experimental data was logged in one afternoon and the borrowed equipment immediately returned to its owners.

The Novatel GPS data was then processed through both HPM and the experimental heave filter. When the GPS derived heave values were compared to the TSS reference values, there was a strong visual correlation between the two heave wave forms however very significant wandering biases existed between the two. When the TSS heave record was examined by itself, it became apparent that it had been severely degraded by the violent maneuvers performed during data collection. For minutes on end, the TSS heave values would remain almost completely positive. Obviously real (unbiased) heave values must go both positive and negative as the vessel transits from wave crest to wave trough. It was therefore concluded that the TSS was not outputting usable reference data. It appeared that the violent turn maneuvers of the survey launch were upsetting the turn rate sensors and thus biasing the heave output. This hypothesis was later confirmed by more experienced users of the TSS. It works well when running reasonably straight survey lines but can degrade significantly during hard turns. The west coast data set therefore had to be discarded and new data observed under more realistic conditions of sea state and vessel turn rate. Unfortunately, no time or vessel resources were available to collect new data in time for preparation of this paper.

Fortunately, the University of New Brunswick volunteered to supply a previously observed data set which appeared more suitable for this experiment. This new data set had been collected on day 149 1994 on board the F.G. CREED. The mission took place in the Bay of Fundy approximately 90 km east of St. John NB. The F.G. CREED had been following long straight survey lines so the TSS reference data (logged at 2 Hz) was almost certainly more accurate than what had been collected on the west coast. The F.G. CREED is a SWATH vessel (Small Water-plane Area Twin Hull) riding on submerged floatation pods. The effect of the SWATH hull configuration is to greatly stabilize the sounding platform thus enabling high speed data collection in rough weather. Heave conditions aboard the F.G. CREED were therefore only moderate (about +/- 0.5 m.). Another attribute of the F.G. CREED's hull design is that the spectral content of the residual heave is very distinctive with peculiar horizontal acceleration components in rough weather. The heave data from the F.G. CREED is therefore not very representative of typical survey vessels.

On board the F.G. CREED and also at the reference site at Harbourville, NS., Ashtech Z-XII dual frequency GPS data was logged at 1 Hz. Using the Ashtech Z-XIIs for this experiment had the advantage of providing more processing options than with the single frequency Novatels. Ashtech's PNAV software could now also be used to compute double difference float solutions. PNAV processed L1/L2 data should provide more accurate positions to the heave estimation algorithm than L1 only data processed using HPM's phase smoothing algorithm. This L1/L2 data set, permitted both GPS processing approaches to be tested in conjunction with the heave filter algorithm.

The Ashtech GPS data set from UNB did however have three disadvantages:

- 1) The slower (1 Hz) data rate made aliasing of the heave signal more of a problem.
- 2) Unlike the Novatel receiver, the Ashtech Z-XII's data output does not provide error estimates for the code and carrier phase data. We therefore could not exploit HPM/HPC's ability to dynamically weight each observation in the position solution.
- 3) During processing it was discovered that the reference station data set was corrupt. Later tests showed that the data corruption had been caused by the laptop computer used to log data at the Harbourville reference site. The serial port on that computer was not fast enough to keep up with data flow without occasional buffer overflows corrupting the logged Z-XII records. In the data set illustrated below, 84 records were found to be corrupt. While this flaw in the data was not fatal, it complicated interpretation of the results.

No inexpensive digital inclinometer was logged on board the F.G. CREED during this mission. The pitch and roll values needed for lever arm corrections to the GPS data therefore had to be taken from the TSS data record. The lever arm configuration on the F.G. CREED was very sensitive to any errors in the pitch and roll sensor. Instead of being mounted directly above the sounding transducer, the GPS antenna was mounted 7.7 metres aft and 4.6 metres to port of it. The vertical lever arm geometry, it was therefore advantageous to use the pitch and roll angles

from the TSS heave compensator rather than the less accurate attitude data that would have come from an inexpensive digital inclinometer.

Despite the above shortcomings, the F.G. CREED data set offered by UNB was gratefully accepted and processed to extract the heave result illustrated below.

RESULTS

The first round of data processing made use of Ashtech's PNAV software to produce a GPS elevation trajectory. The moving weighted mean algorithm was then applied to this trajectory to extract the GPS heave estimates. PNAV made use of the Z-II's L2 observable in a double difference float solution. PNAV was also used to process the GPS data both as a forward and reverse time series. In addition, PNAV's automatic data filtering option was invoked to help smooth over the corrupt records that had contaminated the reference station's raw data set. This comprehensive PNAV processing insured the cleanest possible elevation file for input to the heave extraction algorithm. After obtaining PNAV results, all data was then reprocessed using Nortech's HPM software. The HPM L1 only, phase smoothed positions were then input to the same heave filter and time series plots of both series of heave results were produced.

Figure 1 depicts 100 seconds of typical heave results. The one caveat here on "typical" is that this 100 epoch data window covers a period during which there were no obvious data dropouts caused by the faulty logging hardware at the GPS reference site. The short time scale of all the figure 1 graphs is also necessary to permit adequate resolution of the predominant 5 to 10 second wave period that the F.G. CREED experienced throughout the mission. Heave amplitudes ranged between 3 and 6 decimetres. In figure 1a, the aliasing effect caused by the low (1 Hz) GPS sampling rate can be seen in the jerky trace near wave peaks. Clipping of the wave peaks is a problem that only a high sampling rate can address. In previous testing, the HPC real-time software has been used with a Novatel GPS data stream of 10 Hz. That rate of sampling and logging should provide adequate antialiasing for a production system.

Figure 1a shows the TSS heave as well as the GPS heave derived from both the PNAV positions and the HPM positions. It is difficult to see all three traces on the graph because the PNAV and HPM results are almost coincident. The TSS heave results are more easily distinguished as a separate trace but are still tightly correlated with the two sets of GPS derived heave. There does not appear to be any correlation between the wave amplitude and the discrepancy between the TSS and GPS derived heave. This would indicate that the same level of heave errors present in this sample of 0.5 metre waves would also be present in heavier swells. That hypothesis cannot however be tested until more data sets are collected in a variety of sea state conditions.

Figures 1b and 1c show a clearer picture of the difference between TSS and GPS heave results. Figure 1b shows the difference between the TSS and the PNAV derived heave records and Figure 1c shows the difference between the TSS and the HPM heave error during the last 5 seconds of the data set. This rise in heave error

Fig. 1a: Heave waveforms from TSS, PNAV and HPM

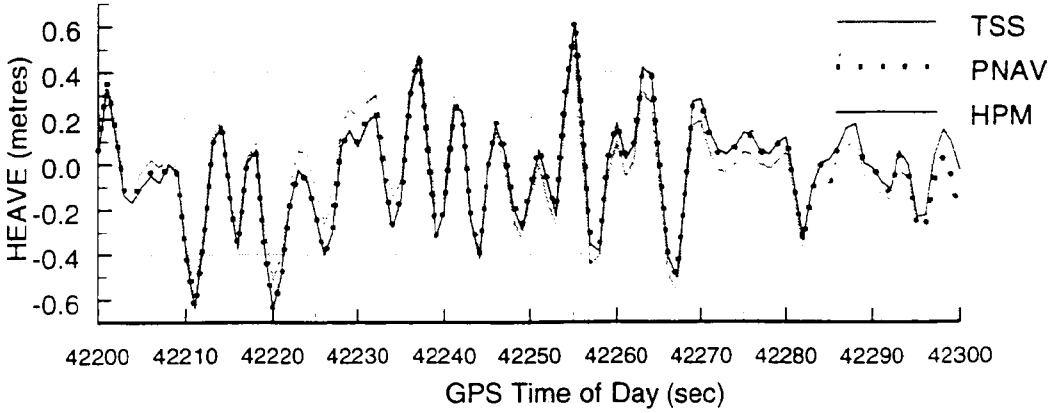


Fig. 1b: TSS Heave minus PNAV Heave

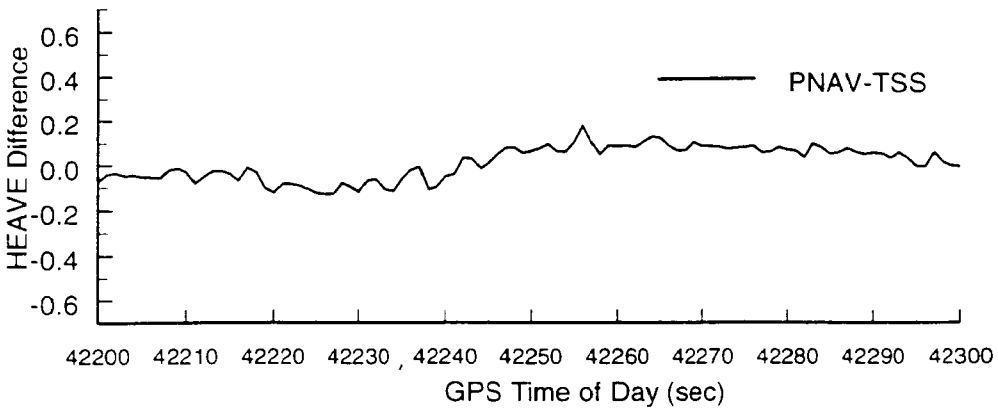


Fig. 1c: TSS Heave minus HPM heave




FIG. 1.- 100 seconds of heave results.

Fig. 2a: TSS Heave Record (20 minutes)

Figure 1 Data Set




Fig. 2b: PNAV Heave minus TSS Heave




Fig. 2c: HPM Heave minus TSS Heave




Fig. 2d: HPM and PNAV Error Estimates




FIG. 2.- 1400 seconds of heave results.

is due to an impending bad reference site record that was later logged at 42340. Its inclusion in the moving data window is starting to degrade the weighted mean reference elevation used to difference those heave values. The 6 worst data logging errors in the reference site GPS data set produced long term biasing affects on the moving weighted mean reference elevations. Those problems will be discussed with reference to Figure 2. In Figure 1b and 1c (except for the last 5 seconds of 1c) both the PNAV and HPM derived heave records are almost identical . Both are within about 1 dm of the TSS ground truth heave record. There is a high frequency (single epoch) noise level of about 5 cm superimposed on a longer term drift that goes from -1 dm to + 1 dm during the 100 second sample.

To provide an illustration of the GPS heave performance over a longer period, Figure 2 illustrates over 20 minutes of results (1400 epoch). Figure 2a shows only the TSS reference values. At this time scale the heave wave form overplots itself and thus provides a noisy looking record of maximum wave heights for the 20 minute period. Figure 2b and 2c show heave discrepancy w.r.t. the TSS for both the PNAV and HPM position record. In Figure 2b, the heave derived from the PNAV processed L1/L2 data maintains a 1 to 2 decimetre agreement with the TSS heave. In Figure 2b, if we consider the TSS heave values to be error free then the mean error of the GPS derived heave is -0.01 metre and the standard deviation is 0.08 metre.

In contrast to the 100 second sample shown in Figure 1, the long term heave results from PNAV processed positions appear very significantly better than those derived for the HPM positions. In Figure 2c the trace of the HPM derived heave accuracy makes 6 sudden jumps to unusable levels of 1 metre or greater. When the GPS data record was examined, it became apparent that all of these jumps correspond exactly to the epochs when severe logging errors occurred at the GPS reference site. On 4 of these 6 occasions the number of SV's logged from the Z-XII receiver dropped instantly from 7 to 3. On the other 2 occasions, the logged record showed zero satellites tracked for 3 or more consecutive epochs. The PNAV, forward/reverse processing combined with its data spike detection filter, removed the effect of the 6 corrupt reference site records.

The 6 severest data dropouts in the reference site's logged data file were responsible for the much poorer long term HPM results. Due to the faulty serial port on the logger, the reference site data was very "dirty". There were 84 logging errors where either no data or only partial data records were written. Seventy eight of these errors were minor dropouts that were successfully bridged over by both PNAV and HPM. However there were 6 severe dropouts containing only partial records which caused some of all of HPM's phase smoothing filters to be reset (i.e. the algorithm decided the satellite was truly lost and would not try to interpolate through the outage).

After losing a satellite, an HPM phase smoother is reset to full weight on the noisy pseudorange observations. Phase tracking then regains influence in the solution as a $1/N$ function, where N is the number of epochs. It therefore takes a while for the phase smoothed solution to re-build a high weight on phase observations. During this interim following a satellite reset, the noisy positions will degrade both the single epochs and the weighted mean reference elevations used to compute heave. Once the phase smoothing has converged back to a high weight on

phase, the HPM heave results become identical to the PNAV processed data (as illustrated in Figure 2 between 41850 and 42300). Clearly HPM would have benefitted from the 6 worst records being smoothed out of the corrupt data set in the same manner as the PNAV (forward/reverse) processing.

Figure 2d's upper trace shows a plot of HPM's estimated standard deviation for the elevation trajectory. It illustrates that the 6 bad GPS elevations which caused the degraded heave performance were quite well predicted by the statistics. The two spikes that go off scale were series of "zero SV's tracked" records for which the GPS positions were not differentially corrected at all. For those two periods of single point elevations, the estimated standard deviations went off-scale to 55 metres and 40 metres respectively. The other 4 records that caused less catastrophic SV smoother resets contained corrections for only part of the constellation. In those cases HPM, applied partial differential corrections and left the other SVs uncorrected (a non-RTCM approach based on the logic that partial correction is better than no correction at all). The default HPC/HPC DGPS rule follows the standard "all or nothing" RTCM guideline.

From a Quality Control perspective, Figure 2d would be useful for identifying and flagging the 6 sudden gross errors. HPM's trace of estimated standard deviation is not however as realistic as it could be. Truly non-biased statistics would almost certainly have predicted generally lower error estimates. Recent experience with Novatel receivers indicates that considerable improvements can be gained through "tuning" the weighing matrix to fit the characteristics of the receiver [1]. Prior to this experiment, HPM had not been used with the new Ashtech Z-XII receivers. Its weighting scheme was therefore based on a default a priori weighting model conceived for a "generic receiver" assumed to possess noisier technology than the Z-XII. The position standard deviations therefore tend to be quite strongly biased towards large error estimates. The lower trace on Figure 2d shows the error estimates made by Ashtech's PNAV software. These are generally more realistic (smaller) error estimates based on a more realistic a priori model of the receiver's real noise characteristics. The many small spikes in the PNAV error estimates correspond to the 84 minor dropouts in the corrupted reference site data.

When analyzing the original data set collected on the west coast, it was evident that the TSS "reference" heave estimates were not necessarily always error free. Since the F.G. CREED data had been observed along a fairly straight survey line, the TSS reference heave values appeared to be free of any gross artifacts that can be produced by violent turn maneuvers. However, this TSS reference data still warrants a close look to determine if it is providing the sub-decimetres accuracy required to ground truth this experiment.

What is remarkable about Figures 2b and 2c is the 60 second ripple pattern which dominates the error signal. This regular one cycle per minute artifact has a 1 to 2 dm amplitude and has no apparent correlation to the dominant 5 to 10 second wave periods. What could be causing this error signal? It might conceivably result from a GPS related phenomenon such as multipath. Since the GPS positions were mainly derived from GPS phase observations (which are highly immune to multipath effect), this is not very likely. The often identical heave results obtained from the PNAV and HPM positions also suggests that GPS is not responsible for the 60 second ripple effect. Another possible explanation might be a bug in the

implementation of the test algorithm such as a sign inversion in a lever arm component. Again, this does not appear likely since such a bug would also produce artifacts at the predominant wave period of 5-10 seconds. The most plausible explanation for the observed beat frequency between the TSS and GPS heave signals appears to be that there was some inadequacy in the high pass filtering assumptions used to extract the heave motion from the two sensors. Either the GPS position sensor or the TSS acceleration sensor could be subject to this problem.

Both the TSS firmware and the GPS heave algorithm being tested require a user defined time constant for the high pass filter's bandwidth. Wave periods longer than the chosen bandwidth cannot be sensed. The low end frequency cutoff is required not only to block out the tidal signal but also any other biasing influence such as changes to the draft or trim of the vessel. To produce both Figures 1 and 2 the GPS heave algorithm used a 100 second moving time window for computing the weighted mean reference elevations. During the F.G. CREED data collection, the TSS real-time firmware had been initialized with a relatively short 16 second filter constant. It is obvious that its much longer high-pass filter time constant should enable the GPS heave filtering to detect much longer period waves (if any were present in the vessels real motion).

In Figure 1a the predominant wave period is between 5 and 10 seconds. The question is: were longer period wave phenomenon also affecting the vessel? The F.G. CREED's twin submerged hull configuration might conceivably generate harmonics of the predominant 5 to 10 second waves. The test site in the Bay of Fundy is also subject to very large tides and the test data was collected during quite rough weather. In following seas, any vessel under way has a tendency to surf up and down the waves at a lower frequency than the waves themselves. It is therefore possible that the 60 second ripple artifact apparent in Figure 2 might actually be a real 1-2 dm heave phenomenon that had been blocked by the TSS's 16 second high pass filter but detected by the 100 second filtering done on the GPS elevations.

If this hypothesis is true, then the real GPS derived heave errors would be significantly more accurate than the +/- 1-2 dm observed in Figure 1 and 2. In Figure 1 the high frequency (1 to 5 second) noise level of the discrepancy signal is generally less than 5 cm. This higher level of heave sensing accuracy is consistent with the performance repeatedly demonstrated in GPS experiments based on semi-kinematic DGPS [3]. Unfortunately the F.G. CREED data set does not permit us to test this tempting hypothesis. The TSS filtering was done in real-time so its accelerometer data could not be post-processed using a longer filtering parameter than the chosen 16 second real-time value. When planning data collection for a future, more conclusive experiment, this problem will be addressed.

QUALITY CONTROL OF GPS DERIVED HEAVE

The heave filtering algorithm described above is being added to the HPC/HPM software suite used by CHS. The heave algorithm will be incorporated into the HPOST utility used to format an output file containing the depth measurements logged by HPC. Each depth record in this file has its position

interpolated from the neighboring GPS position epochs in the asynchronous logged data stream. The time stamp accuracy of all logged records in the file (depth, position, pitch and roll) is continually calibrated by HPC using the 1 pps time synch pulse from the GPS receiver. The HPOST output file therefore provides an appropriate environment for computing heave corrections and applying them to the logged soundings. A statistical QC algorithm is required to insure that the heave is only applied if it will improve the accuracy of the soundings and that all the final soundings are tagged with realistic error estimates. The QC approach being implemented is as follows:

QC Step 1

Estimate heave corrections as per equation 1. Then estimate the error of each heave correction based on the expected error for the weighted mean reference surface (from equation 4) together with the expected error of the single epoch used to difference the heave value (HPC's estimated standard deviation logged in real-time).

Both the computed heave and its estimated standard deviation are thus computed for each GPS epoch, interpolated for the epoch of each logged depth measurement and then written into the HPOST file. When and if a sounding is corrected for heave (QC step 3 below) the reduced sounding will take on the uncertainty of its heave correction. That uncertainty will also include an estimate of error due to uncertainty in the lever arm correction (lever arm offsets and an estimate of pitch and roll measurement accuracy will be required in the HPOST file's header). Other positioning and acoustic factors affecting a depth's accuracy will not be analyzed by this utility.

A user defined tolerance on the single epoch standard deviation will also reject "bad" elevations. This will cure the poor results derived from this data set due to the logging errors at the reference site.

QC Step 2

As the moving data window computes the "calm water" reference elevations, it will also compute and buffer an "average maximum wave height" within the span of each window. This value to be used for decision making in step 3 below.

QC Step 3

The QC algorithm will then decide whether or not to apply each computed heave correction to its corresponding raw sounding. The decision will be based on the average maximum wave heights observed around the time of the sounding. If the expected error of the heave correction exceeds a user defined proportion of the average maximum wave height, then the sounding will not be corrected (the logic here is that the water is too calm to apply the too uncertain heave "correction"). The uncorrected depth together with its estimated error (the average maximum wave height) is then written to the HPOST file as a separate entry and also flagged as being uncorrected.

If on the other hand the expected error of the heave correction does not exceed the user defined proportion of the average maximum wave height then the sounding will be corrected for heave (the logic being that the uncertainty of the heave correction is low enough to still provide a net gain over not correcting heave at all). The heave compensated depth together with its estimated error (computed in step 1) is then written to the HPOST file as a separate entry and flagged as being corrected.

The statistical details of this general QC algorithm will be worked out during implementation and testing. This HPOST heave utility should be ready for deployment on CHS production surveys during 1995.

Conclusions

This experiment has demonstrated that high pass filtering of GPS elevations can be used to estimate highly cost effective heave corrections for acoustic depth soundings. Preliminary ground truth testing indicates that these heave corrections are accurate to +/- 1 dm. Using L1 only phase smoothed GPS positions to compute the heave corrections can produce virtually the same results as to those computed from dual frequency GPS data and processing. Provided that the GPS antenna is situated directly above the sounding transducer, lever arm corrections for the GPS positions can be made using an inexpensive digital inclinometer. Error estimates for the GPS elevations used in the heave computation permit a robust quality control of the final corrected depth soundings. Further field testing is required to provide a more definitive assessment of the accuracy and robustness of this technique.

Bibliography

- [1] KIELLAND P., TUBMAN T., 1994, On estimating map model errors and GPS position errors: (Applying more science to the art of navigation), International Hydrographic Review, Monaco, Vol. LXXI, No. 2, pp. 47-67.
- [2] HUGHES CLARKE, J.E. and GODIN A., 1993, Investigation of the roll and heave errors present in the Frederick G. CREED - EM-1000 data when using a TSS-335B motion sensor. Ocean Mapping Group, University of New Brunswick, DFO contract FP707-3-5731.
- [3] DELOACH, S., WELLS, D., DOOD, D., PHELAN, R., MORLEY, A., SHANNON, B., 1994, Delineation of tidal datums and water surface slopes with the GPS, US Hydrographic Conference '94, The Hydrographic Society SP No. 32, pp. 214-221.