


AnSlope NBP04-02

LADCP Operation Cruise Report



Lowered Acoustic Doppler Current Profiler (LADCP) during AnSlope Cruise NBP04-02

(A. Thurnherr, M. Visbeck, and B. Huber)

The LDEO LADCP system comprises 2 RDI WH300M ADCPs in deep pressure housings mounted on the CTD frame, one looking up and one down, connected to an external battery housing designed and fabricated at LDEO. Power is provided by a single pack of 35 alkaline D cells in series, manufactured by Mathews Associates. The two heads are operated in master/slave mode, with the downlooking head serving as master. The synchronization signal, communications lines, and power lines are served to the heads and battery pack via a breakout cable designed at LDEO and manufactured by Impulse. Preliminary processing was performed immediately following data download using the LDEO LADCP processing software built and maintained by M. Visbeck.

The LADCP system underwent several revisions and upgrades during the cruise. Initially the precast setup and postcast data download were handled by RDI DOS-based software and DOS batch files running on a Dell Latitude C610 computer with Windows 2000 operating system. The post processing software is Matlab-based, and was initially run on the Dell using Matlab 6.5R13 under Windows2000. The netcdf library installed evidently has some problems, so the post processing output was provided as .mat files rather than netcdf files. As software revisions were made by M. Visbeck, the stations were reprocessed on his Powerbook, and a cruise netcdf file was produced and updated throughout the cruise. The processed data was placed on a web site accessible from the NBP intranet web.

1.1 Visbeck LADCP processing software version 8b April 2004 (M. Visbeck)

All AnSlope-1 and AnSlope-2 LADCP data were processed with a MATLAB based software written mostly by Martin Visbeck and can be obtained from:

<http://www.ldeo.columbia.edu/~visbeck/ladcp>

Version 8b has several minor changes and a few more significant differences between the flavors of version7. The “common” m-files are grouped together, while “local” m-files are organized under a separate heading that need to be adjusted for each particular cruise.

Furthermore, the weighting of the inverse solution has been altered after a number of helpful discussions with Andreas Thurnherr. The most significant change was that version 8 uses the bottom track data to constrain the U_ctd velocities. All previous versions used to directly constrain the U_ocean estimates over the range that bottom track data are available. Significant changes were also made to the implementation of the drag model, but it still should be considered experimental and is turned off in the new default settings. A new constraint was added that allows one to restrict the energy projected to any integer vertical mode (ps.smallfac). For very deep profiles this reduces “runaway” shears due to the expected random walk of shear errors. This constraint is enabled in the default mode at very low influence. The data loading function (loadrdi.m) was changed in several small ways. The most significant change is that the construction of bottom track data from the water profile data was moved to a dedicated routine (getbtrack.m) which also saves all bottom track data for later consistency check. Finally, version 8 has full support of RDI beam coordinate data.

The other significant change is a more elaborate scheme to detect any time offset between the CTD and LADCP data (loadctd.m).

The structure of the LADC processing was as follows:

A wrapper m-file was constructed that allows processing of the whole cruise (ladcpbatch.m). This means that it is not anymore necessary to have one m-file for each station and enables reprocessing of all profiles.

First the input files names for each cruise were assigned in the f-structure: One for each up/down looking ADCP, a CTD-time series file, a navigation time series file (same as the CTD time series file in our case), a CTD-processed profile file and the SADCP profile file name.

Second some cruise specific default parameter were given, such as the vertical resolution of the output (ps.dz=10 [m]).

Then the main LADCP processing was initiated by calling LAPROC. The following main task are performed:

- 1) LOADRDI: load ADCP data and merge up-down looker into one matrix
- 2) GETSERIAL: assign serial number to instrument from either the deployment log file or a locally provided lookup table that decoded the CPU board serial number.
- 3) LOADNAV: local file that encodes ship's navigation and provides the position as a function of time. Note this function assumes that the ADCP and Navigation time are "the same". The mean station position gets calculated and the magnetic deviation calculated and applied if not explicitly given prior to the call of LOADRDI
- 4) GETBTRACK: checks the RDI provided bottom track and makes a secondary bottom track based on the water track data. The primary bottom track to be used can be chosen and both are saved for later diagnostics.
- 5) LOADCTDPROF: local file to load the processed CTD data file and to calculate the sound speed profile as well as the ocean stratification.
- 6) LOADCTD: local file to load the CTD time series. Then a W_{ctd} gets calculated from dz/dt and gets compared to that of the ADCP. A number of different tests are performed to determine the best time lag between CTD and ADCP profiles based on the W time series comparison. Then the ADCP time is adjusted to match the CTD time series. This choice is justified if the navigation data are provided via the CTD time series. If the navigation is external one might want to adjust the CTD time stamps....
- 7) GETDPTHI: calculates the water depth from bottom track distance and ADCP depth time series (either based on time integration of W or preferably the CTD pressure/depth time series).

At this point a plot of some of the engineering data is made and LAPROC calls the second batch of processing files given in PRESOLVE:

- 1) PREPINV: condenses the raw data into a largely reduced number of "super ensembles" which we chose to collect all consecutive data over a 10m depth interval. Prior to the averaging both velocity records get rotated to a common heading base which was selected to be the average heading between down and up looking instrument.

- 2) LOADSADCP: load sadcp data that were provided by the on-board real time processed hull ADCP data (150kHz). For this cruise a special set of commands were implemented to generate a MATLAB file automatically every hour.
- 3) LANARROW: is a set of calls to give a first estimate of the solution and then remove 1% of the most inconsistent data
- 4) A second call to PREPINV adjusts the up/down looking instrument further. First the preliminary estimate of the ocean velocity profile is removed from each raw ensemble. What remains are nbin realizations of U_ctd which should be the same for both instruments and all bins. For each instrument one range averaged mean offset is calculated and removed from the raw data. This reduces the offset between the two heads and could be interpreted as a tilt/compass error. Then the super ensembles are recalculated.

At this point all data are load and screened and ready for the final processing steps collected in RESOLVE:

- 1) GETINV: sets up the inversion and performs the solve. A number of changes have been implemented with regards to the weighting of the various constraints. Also a table of the strength of each of the possible constraints is produced.
- 2) CHECKINV: graphically displays the relative contribution of each constraint to the two solutions (U_ocean and U_ctd). It also lists the expected certainty and the actual difference between the constraint and the solution.
- 3) CHECKBTRK: performs an assessment of bottom track biases between the RDI and water track derived bottom track.
- 4) GETSHEAR2: performs the classical shear based solution and matches the vertical mean flow velocity with that of the full inverse solution.
- 5) GETKXPROF: is an experimental implementation of the Polzin-Gregg-Haney method to compute vertical diffusivities from the internal wave shear spectra. This method is not fully tested and should be take as “experimental”
- 6) BATTERY: outputs the best guess for the battery voltage. Serial number specific calibrations are allowed.
- 7) SAVEARCH: makes MATLAB, ASCII and NETCDF output files
- 8) SAVEPROT: saves some of the most salient information about the processing

The cruise batch file ends with a call to CRUISE which reads all individual NETCDF files and generates one single NETCDF file for the whole cruise as well as series of html index files and table to allow fast browsing of the results.

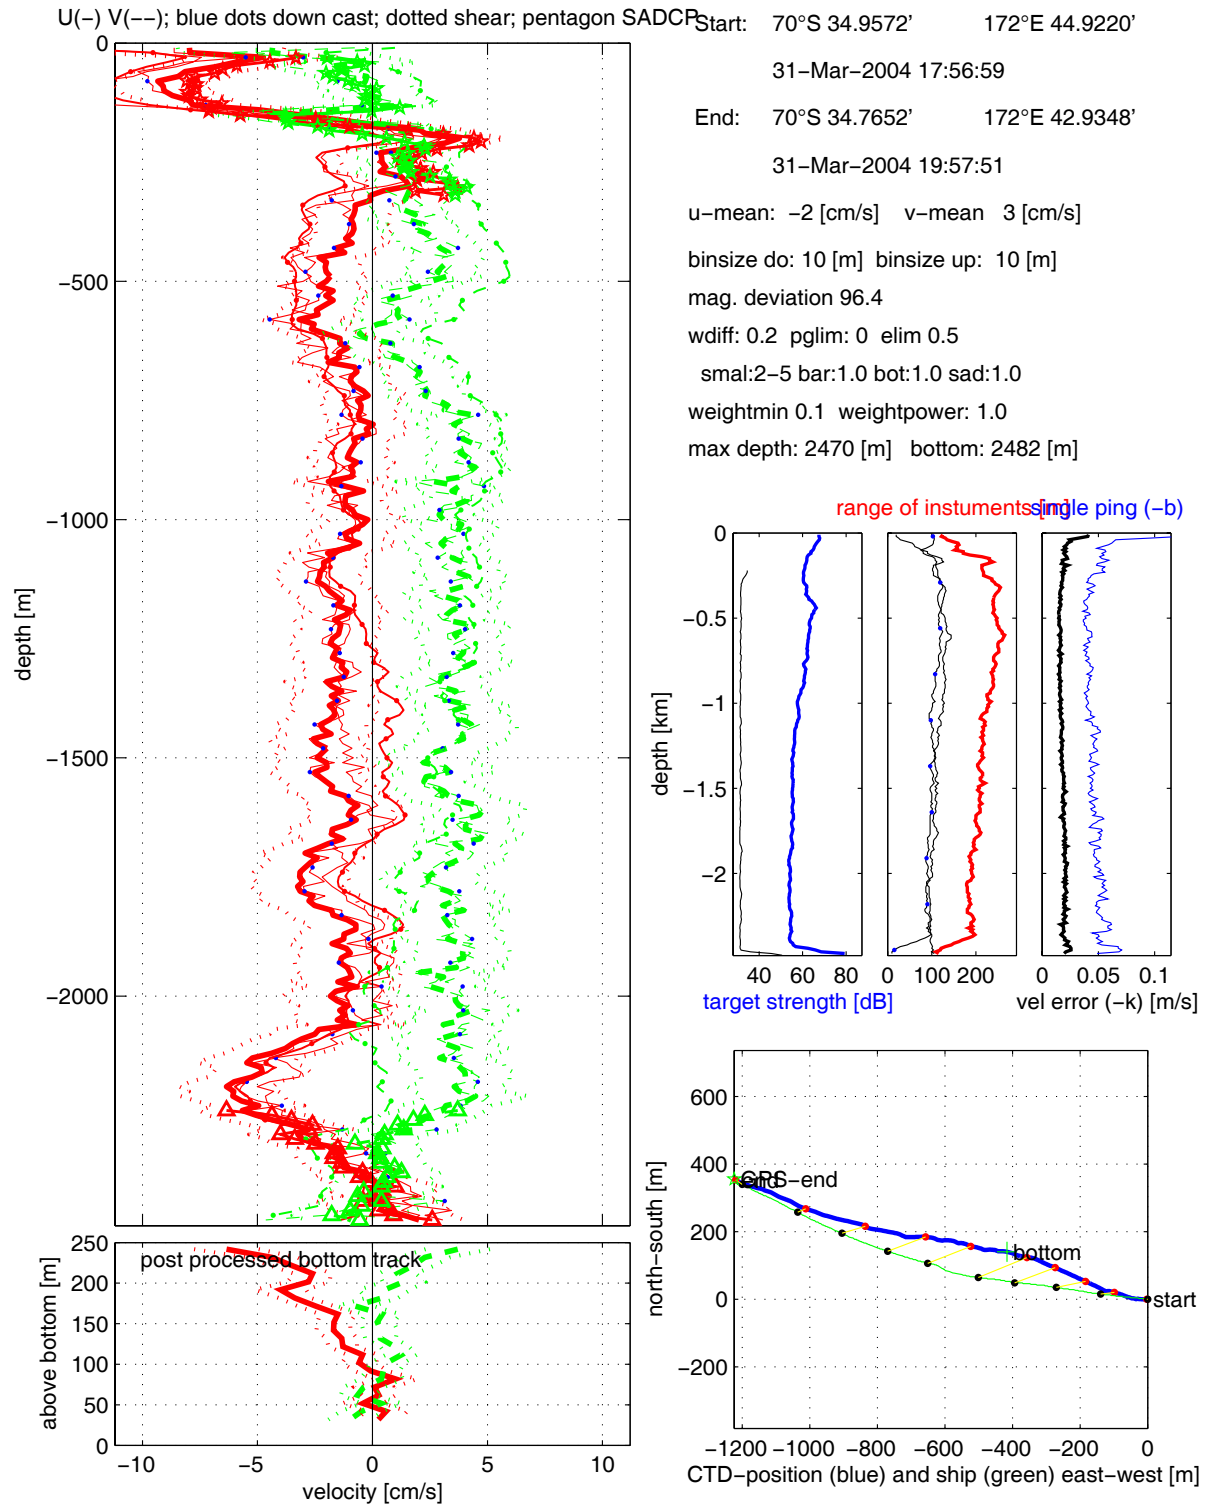
Finally we saved 13 plots for each cast: (Figs 1-10, 12,13)

- 1) Main ocean velocity profile plot plus a number of additional plots: velocity error, range, target strength and CTD trace position.
- 2) Raw data plot. Version 8 had now the correct ADCP voltage and computes ranges for each beam.

- 3) Detailed results of the inversion: left panel is the un-attributed velocity as a function of bins and super ensemble velocity. This plot should be “white noise” with an rms of the super ensemble error. Any structure points to poor performance of the inversion. Middle panel plots each super ensemble ocean velocity estimate (color) as a function of depth and time. The detected bottom is given by black dots. One expects this plot to show horizontal stripes of one color. Vertical stripes point to poor performance of the inversion (tilt, compass, bottom track issues). The right panel is similar to the middle panel except that it plots a dot for each single ocean velocity bin (black up looker, blue down looker). The green lines give an estimate of the uncertainty. The red line is the results from the inversion, the black line is the result from the shear profile (if available).
- 4) Top panel is the beginning / end of the cast in time/depth space. Blue dots represent surface detection. Middle panel is the bottom of the cast with the black line the time integrated W , the red line the best estimate of $Z(t)$ (CTD-depth is available) and the blue dots are the distance of the bottom. The black line is the best estimate of the bottom at the deepest point of the CTD. The lower left panel shows the CTD depth as a function of time if CTD – pressure is available. The lower right panel shows a few sample W time series after the time adjustments between CTD and ADCP time series have been performed.
- 5) Shows the difference between the up/down looking compass. Top line is the compass “adjustment” applied. Middle panel is the difference between the compasses. Bottom panel is the rotation needed to best match the reference velocities between each instrument.
- 6) Heading, pitch and roll difference between the up/down looking instrument as a function of the down looking heading/pitch/roll.
- 7) Top panels time series of U_{ctd} , U_{ocean} at depth of CTD and U_{ship} (from ships navigation). If $ps.dragfac > 0$ a red line shows the expected U_{ctd} from the drag model. Left middle panel horizontal distance of CTD from ship. Middle right panel W -CTD. Bottom panel show position of CTD and ship relative to starting position.
- 8) Vertical diffusivity profile (not fully tested)
- 9) Top U/V ship board ADCP profiles within the CTD cast time. Bottom ships position and where SADCPC profiles were taken.
- 10) Top Mean U velocity offset applied to reduce up/down looker difference. Middle same as top for V velocity. Bottom: implied tilt error if the difference was due to false projection of W into the U/V component.
- 11) Brief summary of the most disturbing errors/warning encountered during the processing. Meant to guide the data collection and notify the operator about potential issues. (not shown)
- 12) Weights used for the inversion: Top panel weight used to constrain U_{ocean} . Bottom panel weight used to constrain U_{ctd} . One would like to see mostly velocity data and only a few extra constraints.
- 13) Performance of the bottom track: Upper right U velocity. Black dots represent $U_{adcp} - U_{ctd}$ from solution for profiles where bottom track data are present. Red dots mark bin

range that was used to make water track data. Below that green histogram of $U_{brk_RDI} - U_{ctd}$ that should be think an with a 0 bias. Middle same for bottom track made from water pings (own). Bottom super ensemble bottom track data (could either be own or RDI [default]). Upper left same as upper right for V component. Left bottom same as Upper left except for W component. Right bottom plots $abs(W)$ normalized by the reference layer W and shows the expected low bias for bins “below” surface due to the expected larger beam angle for bottom returns of the center of the beams.


Station : NBP0402-224 Figure 1



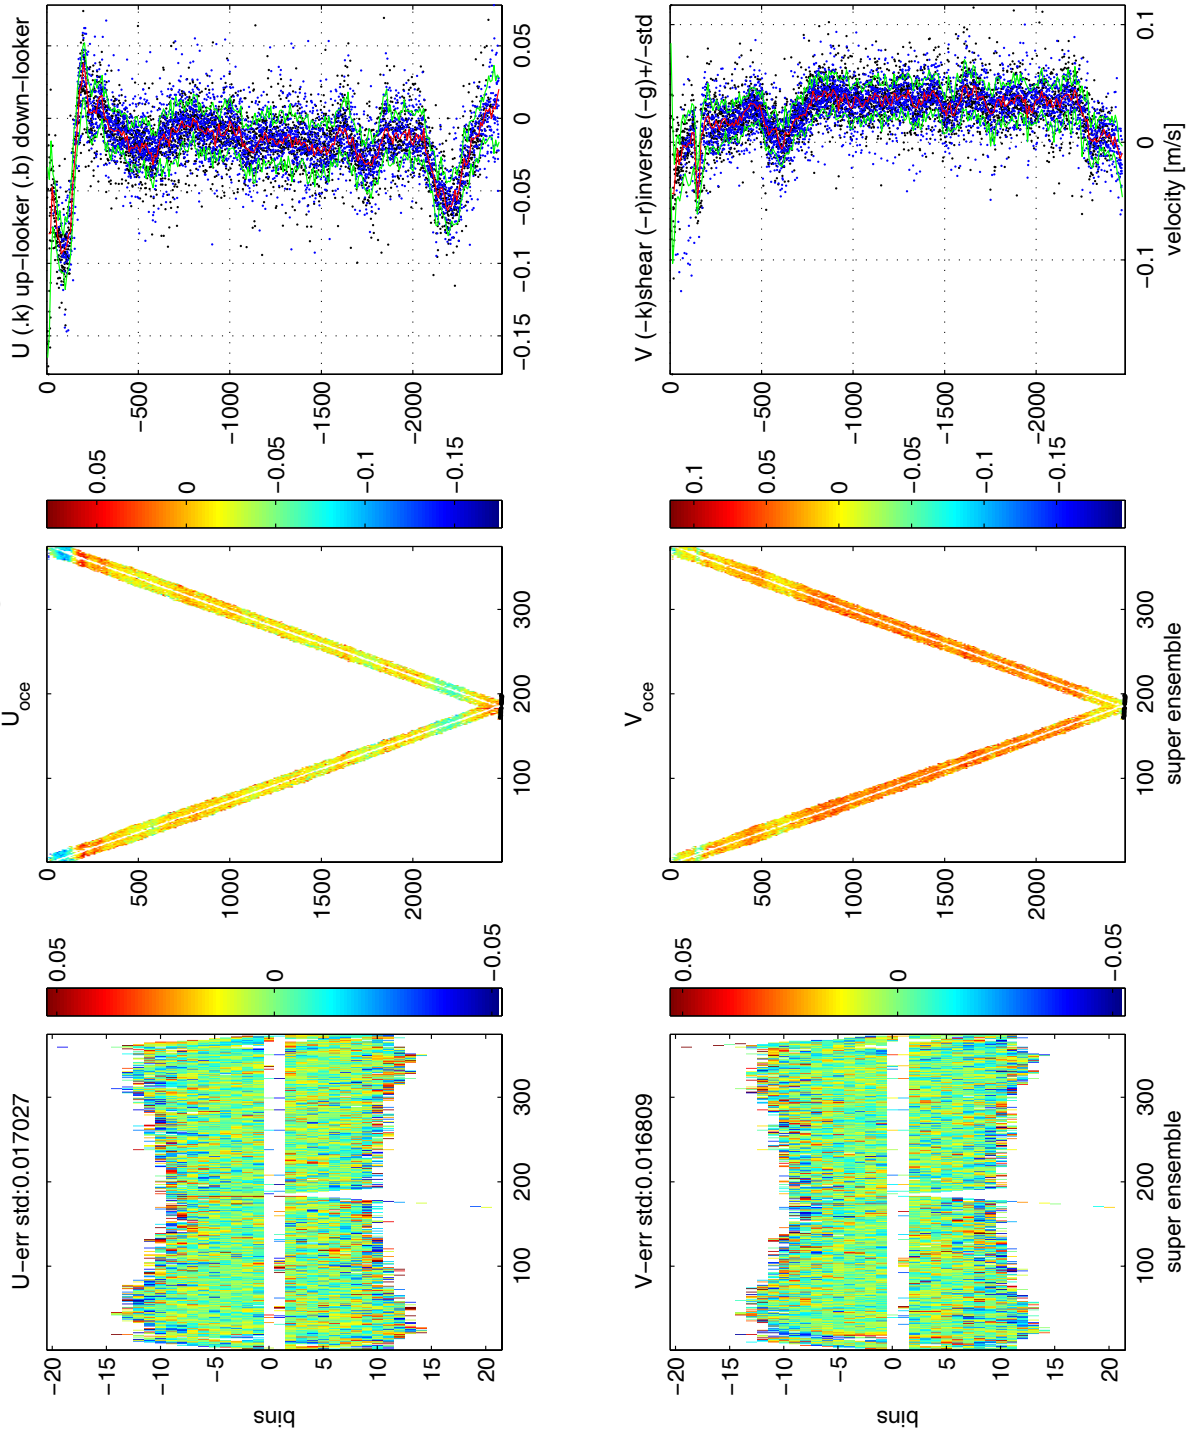
LDEO LADCP software: Version 8b: 2 April 2004

NBP0402-224 Figure 2

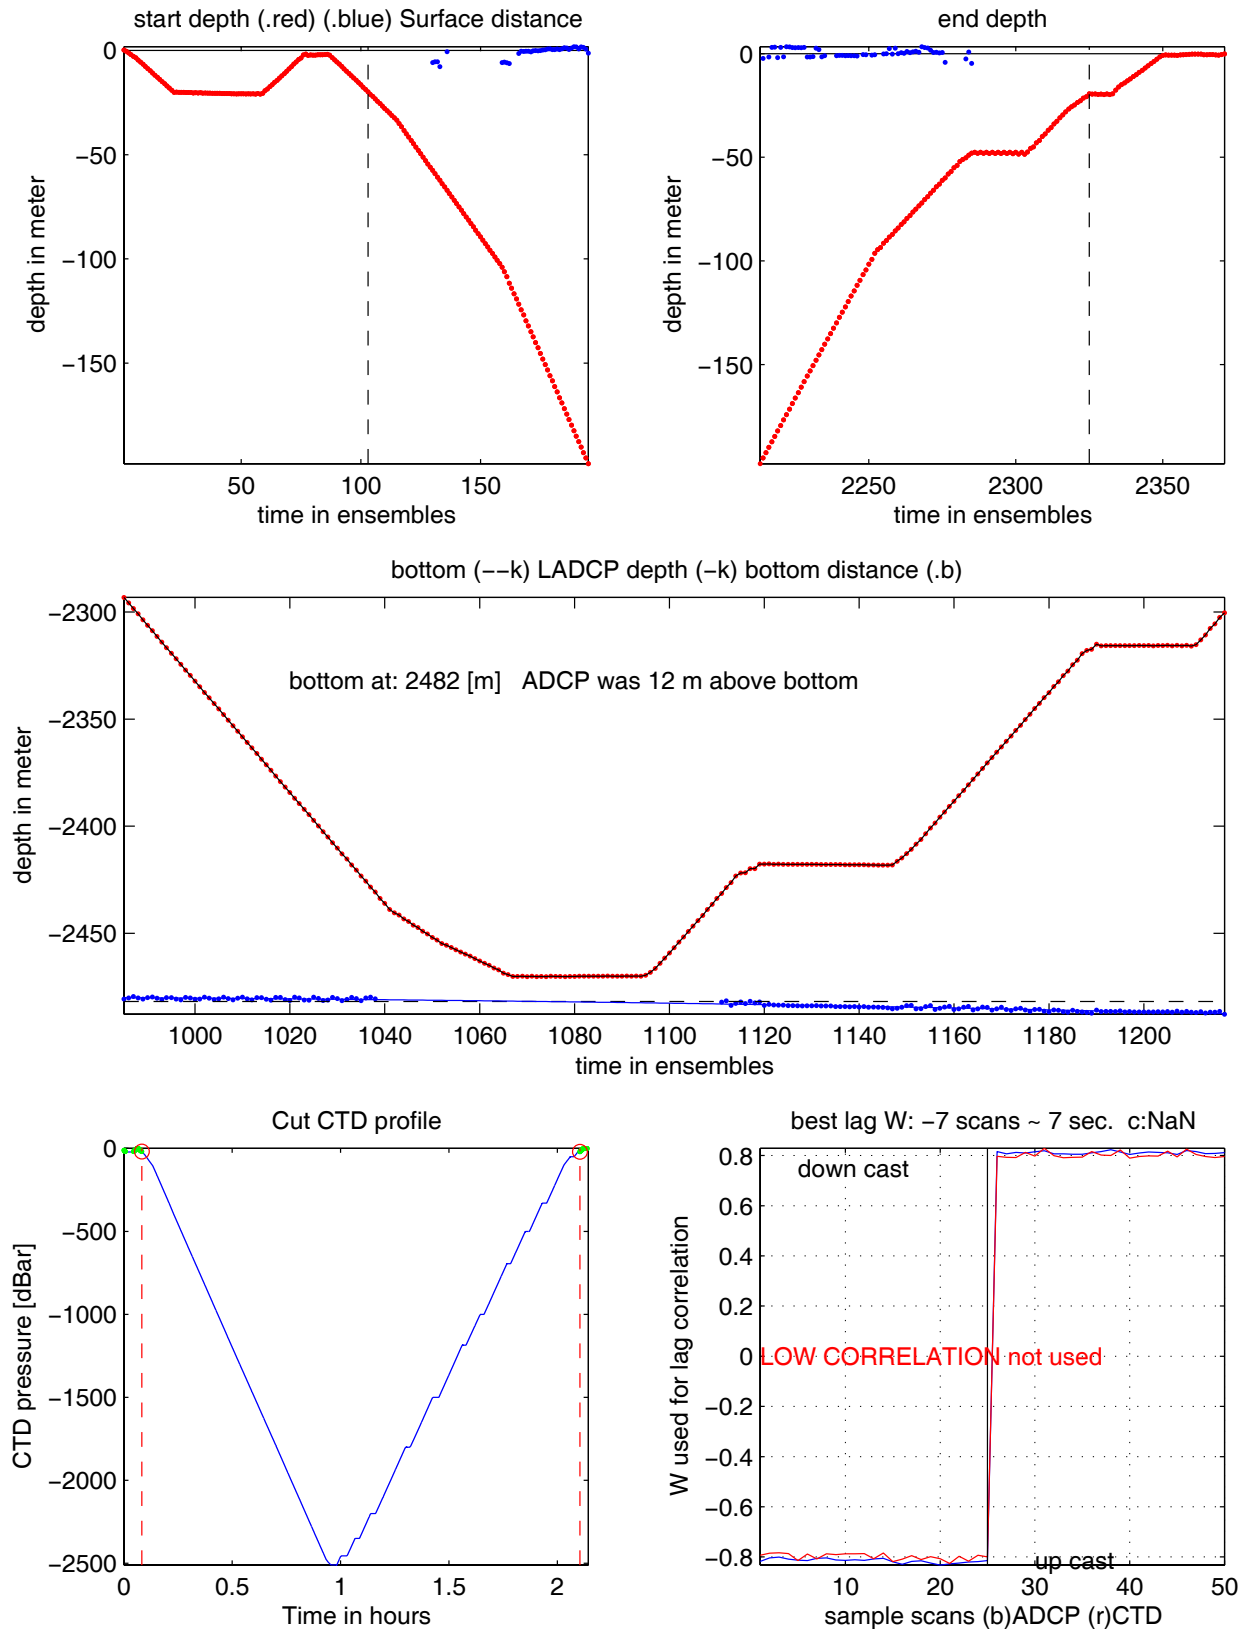
W as function of bindepth and time




NBP0402-224 Figure 3



NBP0402-224 Figure 4




NBP0402-224 Figure 5




NBP0402-224 Figure 6

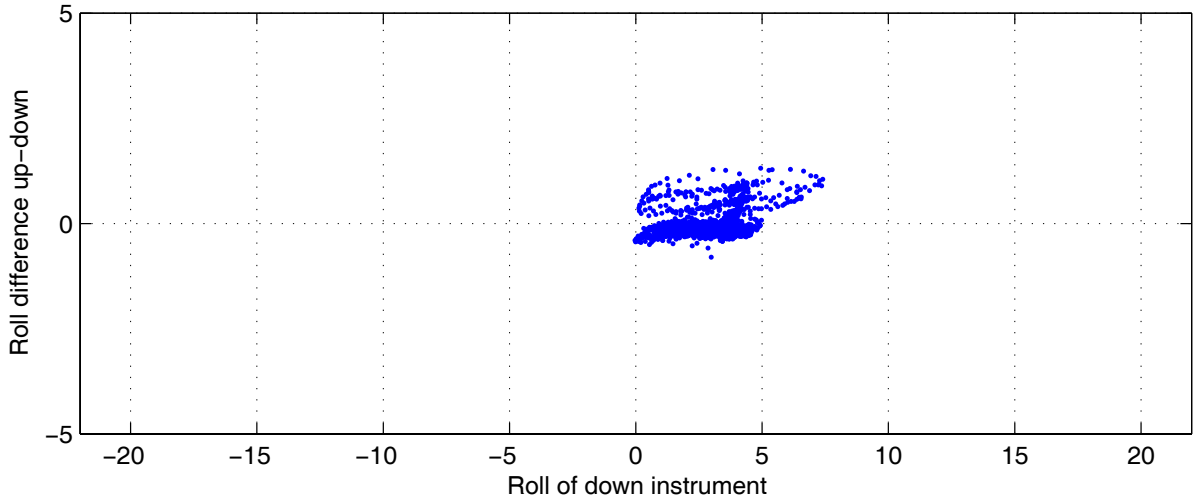
Heading offset : -32.3599



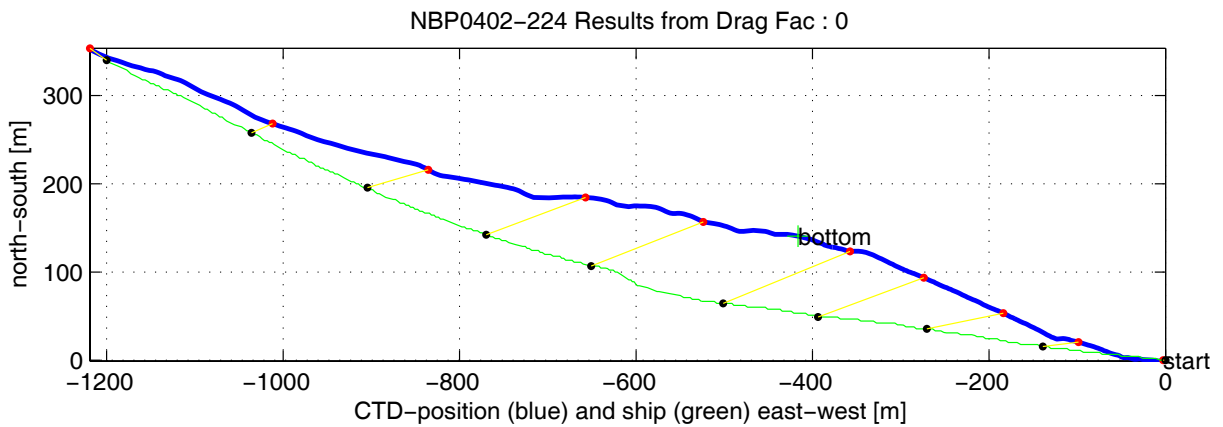

Pitch offset : 2.8751



Roll offset : 0.98822




NBP0402-224 Figure 7




NBP0402-224 Figure 8

Experimental K_z profile from CTD data




NBP0402-224

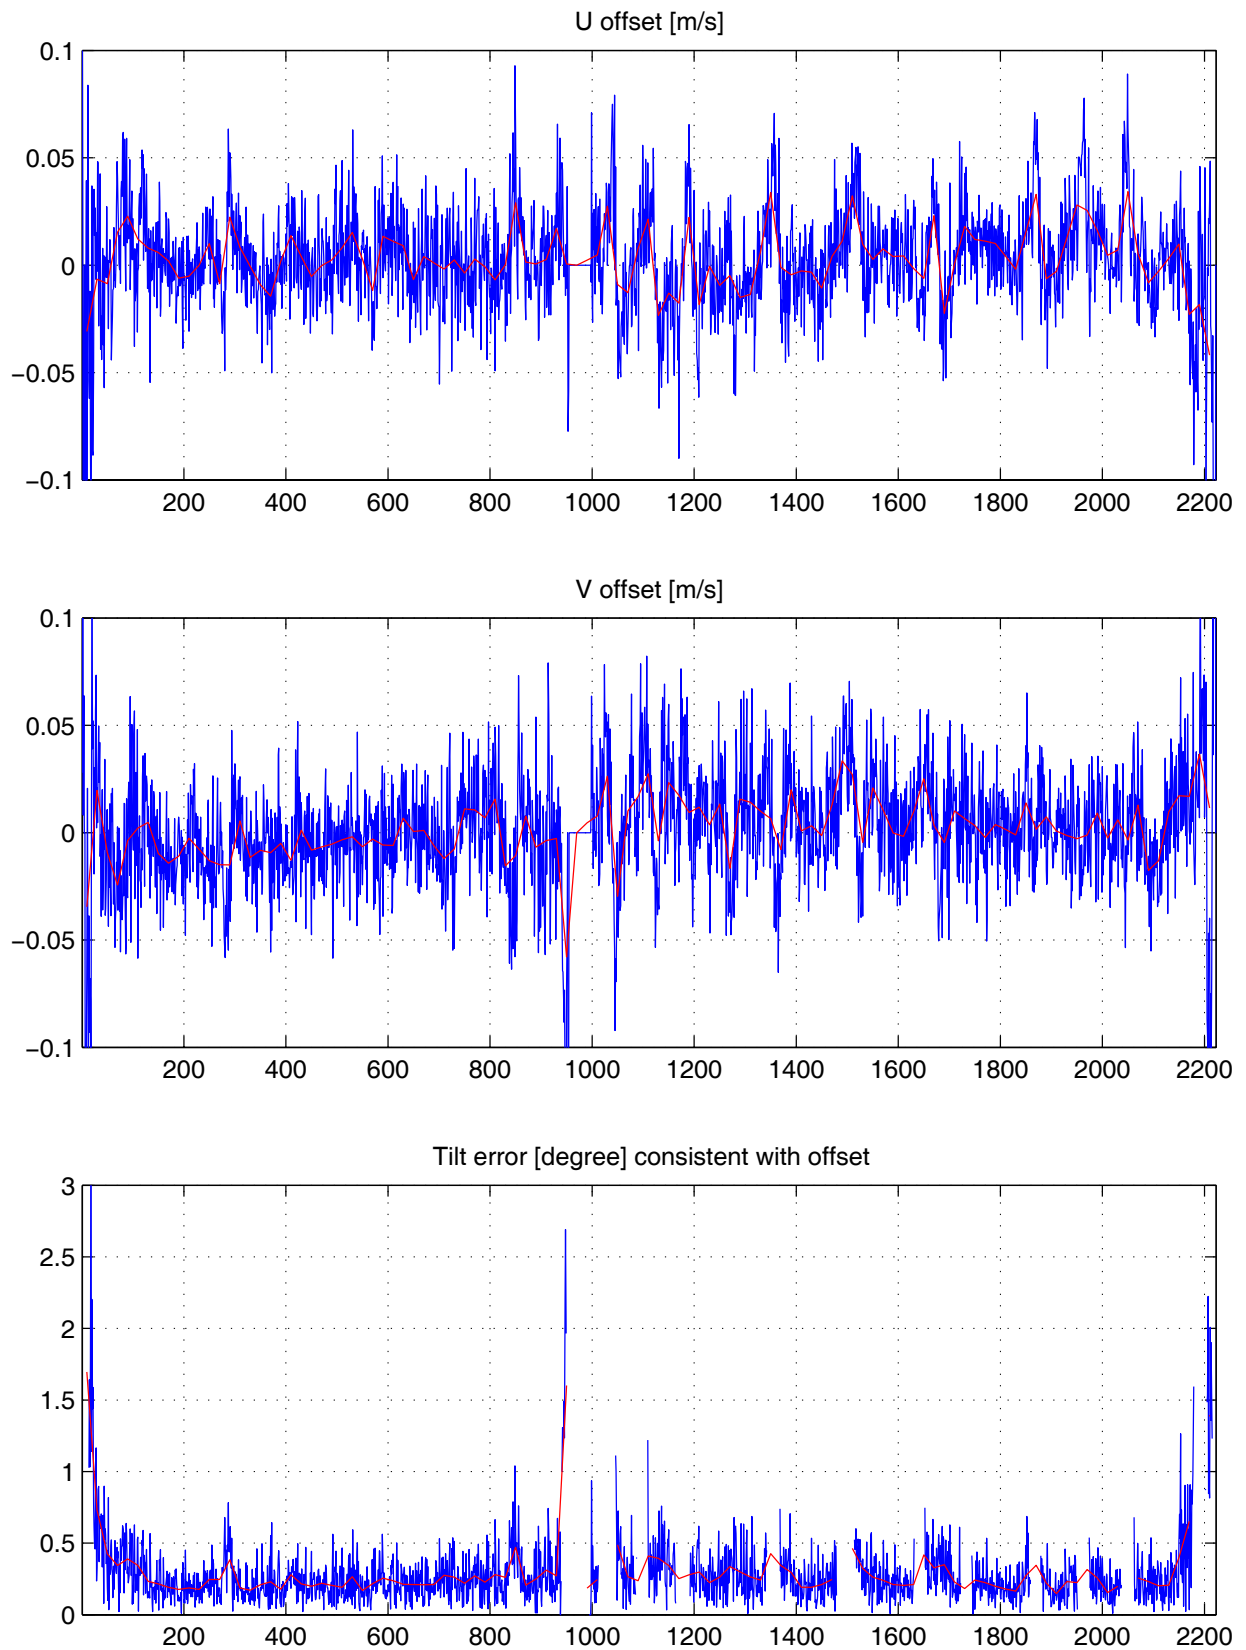
SADCP U(-r) and V(--g) profiles



ship nav (g.) start (bp) end (kp) SADCP (rp)




NBP0402-224 Figure 10




NBP0402-224 Figure 12

ocean velocity constraints



CTD velocity constraints



1.2 Linux Processing Software (A. Thurnherr)

At the start of the cruise RDI Windows programs were used to program the ADCPS and to download the data after each cast. This setup had several disadvantages, the main one being that data downloading had to be slowed down to 57 kbps instead of the maximum speed of 115 kbps. Additionally, the two instruments were connected to the PC using a RS232 switch, which meant that only one instrument could talk to the computer at the same time. Because of good experiences last year on the Aurora Australis A0304 cruise, when full transfer speeds were regularly achieved using a PC running FreeBSD and kermit, it was decided to test a similar setting.

Because no kermit was available on our cruise a replacement for RDI's BBTALK was implemented in perl. The program, called bbabble, is capable of talking to two instruments in parallel (using two separate colors to avoid confusion). The main advantage of having a single terminal program talking to two instruments is that downloading can proceed from both instruments in parallel. Initial tests confirmed that the full downloading speed of 115kbps is achievable under UNIX when downloading from both instruments in parallel and using a laptop. In practice, a speedup factor of 3.5 (compared to the RDI Windows programs) was realized. The program was tested under MacOSX and Linux. Since FreeBSD is more stable than either of these, bbabble will be adapted to run under that operating system as well.

In order to connect the downloading PC to two RS232 devices at the same time, we borrowed two USB-to-RS232 converters from Raytheon. The converters used (Keyspan USA-19QW) worked perfectly both under MacOSX (using the original Keyspan driver) and under Linux (using a driver that is part of Linux). An added advantage of the Keyspan converters is that they have status lights, which helps solve communications problems. We also tried a different USB-to-RS232 converter for which a Linux driver was available. However, this converter did not work --- some tests suggested that it does not seem to be able to send a BREAK, which is required to wake the ADCPs. No USB-to-RS232 adapter is required for bbabble to work, however, as the instruments can also be connected directly to the computer's RS232 port(s).

An additional, unexpected, benefit of bbabble is that it appears more robust than BBTALK in waking up instruments. The instrument used for testing has the property that it often does not respond to the BBTALK wakeup. During several days of development and testing, the same instrument did not fail to respond to a single BREAK sent to it from MacOSX or Linux. The reason may be different BREAK characteristics --- the timing of the RS232 BREAK condition is not well defined and can be chosen in UNIX (using the termios tcsendbreak() routine). The default value (the BREAK condition lasting between 0.25 and 0.5s) worked well for us.

Replacement scripts for the remaining RDI programs (e.g. to erase the memory, send a command file, list the contents of the recorder, etc.) were written using the expect programming language, which is an extension of TCL. While its syntax is somewhat clunky, expect was designed exactly for the purpose of interacting with an interactive system (the RDI workhorses, using bbabble). In particular it is ideally suited to catch error messages and handle timeouts and retries. The new UNIX system was used for somewhat more than half of the casts and appears to be stable.

1.3 Additional Considerations and Findings (A Thurnherr, M. Visbeck)

1.3.1 Firing's Software and Shear Inversion

Processing the LADCP data with Firing shear-calculation software augmented by Thurnherr's shear inversion was not pursued vigorously because it seems that little is to be gained from this. Visbeck's velocity inversion has several inherent advantages and initial inspection of the profiles has shown only few profiles where Visbeck's shear and inverse solutions are significantly different, implying uncertainty. Processing those casts with Firing's software or by using the shear-inversion method does not yield more consistent solutions. The main inherent advantages of carrying out the inversion with velocity (rather than shear) data are that bad velocities (in contrast to a bad shear) do not introduce spurious shear into the solutions. Furthermore, the separation of the measured velocities into CTD and ocean components allows for simpler treatment of the BT data, which directly constrain the CTD velocities. Additionally, the velocity separation allows for a potentially useful additional constraint of the instrument motion using a drag model. The drag model that is currently part of the Visbeck's inversion should be considered experimental, however.

Some modifications to Firing's programs were required in order to get them to work with a dual-headed RDI system. (The uplooker code was only in place for Sontek systems.) Additionally, several bugs were corrected. In Firing's design shear editing of the two heads is done separately and the shears are combined when the absolute velocity profiles are calculated. The shear-inversion was updated to allow simultaneous processing of the down- and uplooker shears. This was easy to do but did not lead to marked improvements, as the shear of most casts was well determined by the data from a single head. At least in some of the cases where the down- and upcasts disagree significantly, the disagreement is confined to a single head and data from the other head can be used to process the cast (see also section on wakes below).

One disadvantage of previous versions of the shear inversion was that the GPS data could not be used. It was argued by some users that since GPS data are highly accurate (the largest errors often being related to the lateral offset between the GPS antenna and the CTD wire, and the change of heading of the ship while on station) the GPS constraint should be trusted more than either the BT or the SADCPC constraints. There are problems with this argumentation, however. First, the different data constrain the LADCP profiles in different ways: the BT and SADCPC constraints fix the bottom and top portions of the LADCP profiles, respectively; the GPS constraint, on the other hand, sets the depth-integrated velocities. It is this constraint that forces bad profiles with runaway shear to take on the characteristic X-shape. A second problem with the argumentation that the quality of the GPS data immediately leads to a tight constraint of the barotropic velocity is that in order to calculate the latter the integrated horizontal motion of the LADCP relative to the water during the cast is required as well. Nominally, this uncertainty is small, typically of order 1mm/s in case of our casts. (Since the variance of a sum is the sum of the variances the positional uncertainty grows with the square root of time. The corresponding uncertainty of the mean velocity is therefore approximately the single-ping velocity uncertainty divided by the square root of the number of pings in a profile. With a 1-hour profile, 1.5s ping interval, and a 5cm/s single-ping accuracy the nominal uncertainty is 1mm/s.)

In practice, the uncertainty is significantly larger, however, at least in part because there are gaps in the velocity data (e.g. when the downlooker is close to the sea bed and when the uplooker is close to the sea surface). Having a dual system available allows the uncertainty to be

determined a-posteriori by comparing the barotropic velocities of the down- and the uplookers for any given cast. Figure 14 shows the barotropic-velocity differences between the uplooker and downlooker data plotted against gap length for casts 1--166, excluding stations with instrument problems. As expected there is a correlation between gap length and velocity uncertainty but in the case of our casts the velocity uncertainty never drops below 1cm/s. The average depth of the casts with gap length below 8% is 1900m and the corresponding velocity uncertainty is 1.3cm/s. This is twice the uncertainty of the corresponding GPS velocity uncertainty with an assumed positional accuracy of 30m and a winch speed of 50m/min (no bottle stops).




Figure 14

If it is optimistically assumed that the trend shown in the figure continues to shorter gap lengths, rather than reaching a plateau near 1cm/s, the GPS error does not dominate the barotropic-velocity estimates except for casts deeper than 4000m or so. If a plateau is reached, on the other hand, the GPS uncertainties never dominate the uncertainties in the barotropic velocity estimates and at least the BT constraint is easily as accurate as the GPS constraint.

A further modification to the shear inversion concerns the weight of the BT constraint. In the previous version (the one released after the Aurora Australis cruise in 2003) the BT constraint

was weighted as if there was an independent velocity estimate at each depth in the BT-referenced velocity profiles. (If the BT-referenced profile spanned 100m and the vertical resolution was 10m the BT data were weighted as if there were 10 independent velocity estimates near the sea bed.) This is incorrect, as the BT data really only provide a single velocity estimate, i.e. the package velocity over ground. It is nevertheless useful to impose the BT velocity constraint over a range of depths so as to minimize the influence of the shear uncertainty near the sea bed. Therefore, the BT constraint is applied over the entire BT-profile depth range as before but it is downweighted by the square root of the number of samples in the profile.

Unfortunately, this is still not the correct way of handling the weights, as the standard deviations (and standard errors) of the BT-referenced profiles are dominated by the uncertainties in the WT data and not by the uncertainties of the BT data. (This was noticed when the BT profiles from the casts with the true BT mode were compared to earlier ones and it was found that the standard deviations were the same even though the BT accuracy with true BT pings is nearly an order of magnitude higher than the corresponding accuracy when using BT from WT pings --- see Bottom-Tracking section below).

The changes to the shear inversion are poorly documented, not tested extensively and need significant extra work. No public release is planned at this stage.

1.3.2 Instrument Wake and X-Profiles

In a few cases (part of) the upcast of the downlooker is contaminated by shear that is most likely caused by the wake of the CTD platform. The wake contamination is most easily apparent in the left panels of processing software figure 3 where the inversion-residuals are plotted (Figure 15) . The corresponding velocity profiles are of comparatively poor quality, as indicated by the disagreement between Visbeck's shear and the inverse solutions (Figure 16) . In most wake-affected cases this disagreement is noted by the software and a specific warning is output on processing software figure 11. We observed wake problems most often in the downlooker data but there is at least one cast where a portion of the uplooker downcast is contaminated, although without significantly reducing the consistency between the shear and inverse solutions.

Figure 15 suggests that the wake contamination is largely restricted to the first two bins. This assumption is also made in Firing's software, which, in contrast to Visbeck's inversion, contains specific code to handle wake contamination. Strictly speaking, Firing's wake-editing can be configured to remove data from as many bins as required but it defaults to 1 and the recommended value (from the demo) is 2. Unfortunately, the instrument wake does not only contaminate the first few bins but, in our cases, all bins up to the range of the instrument. This was tested by successively removing one bin after the other and re-processing the data --- only when all bins are removed does the shear and inverse solutions become consistent (Figure 17). This is troubling for single-head LADCP casts because it implies that all data from the affected beam have to be discarded and that 3-beam solutions have to be used. Currently, this is not possible either in Firing's nor in Visbeck's software.

p402210 Figure 3

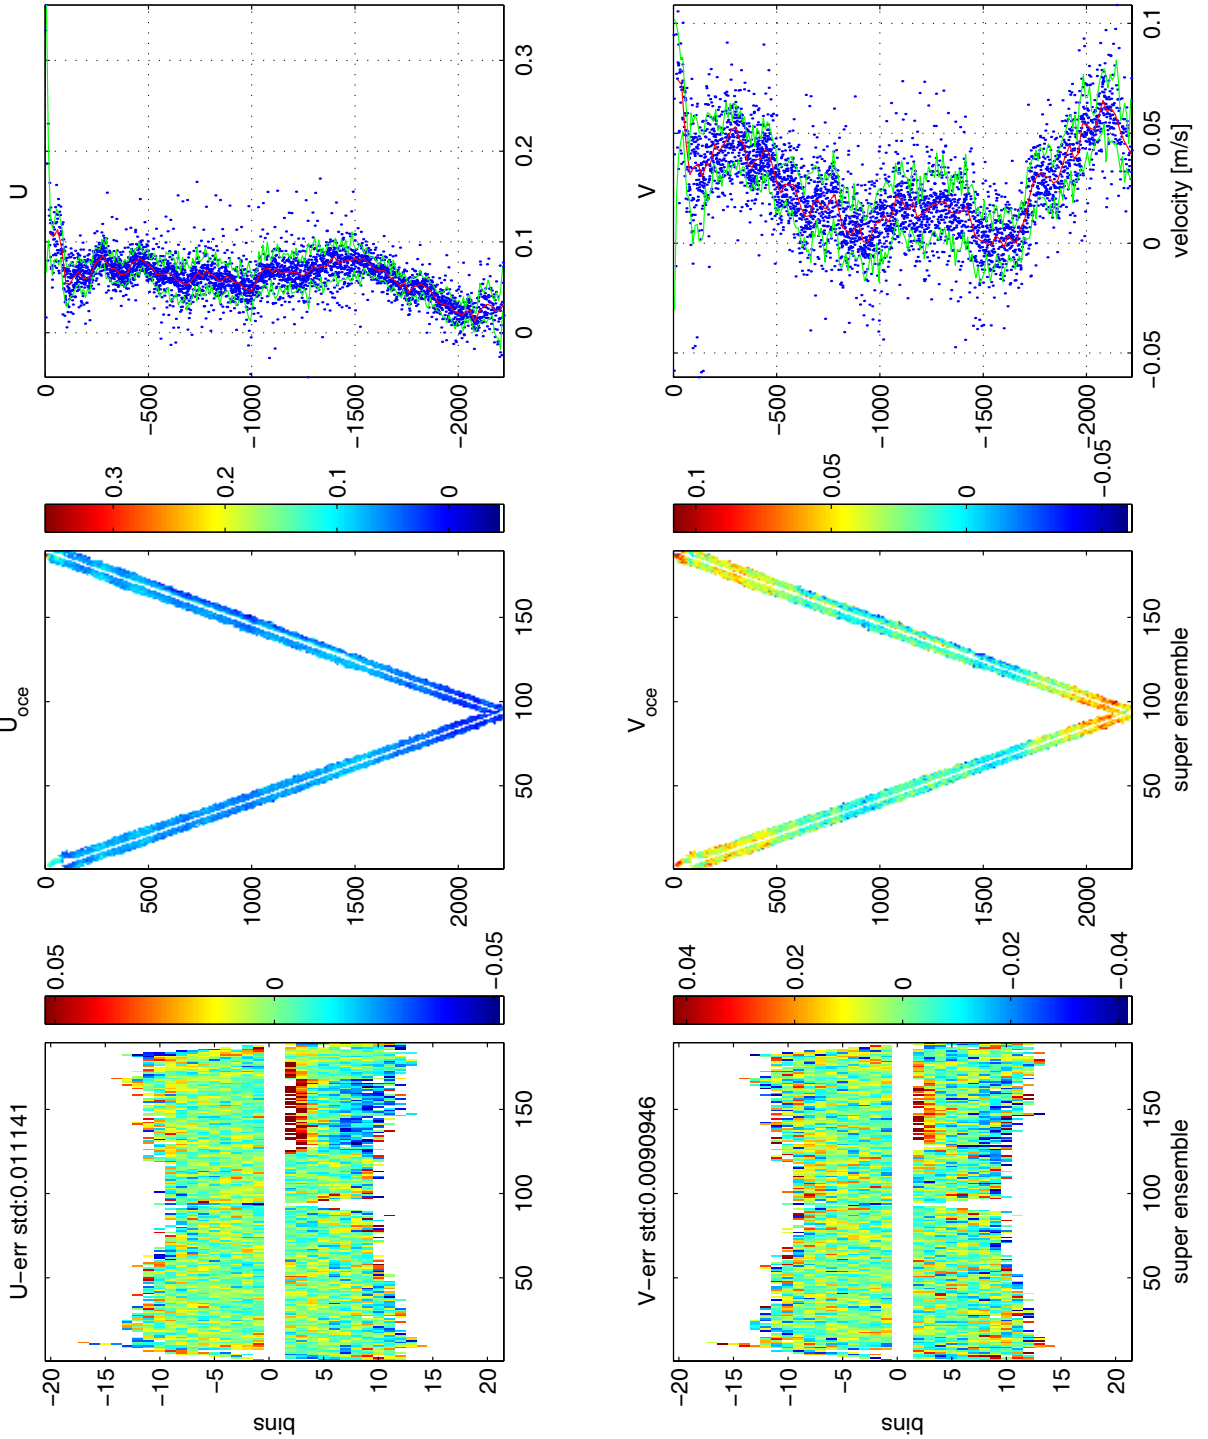
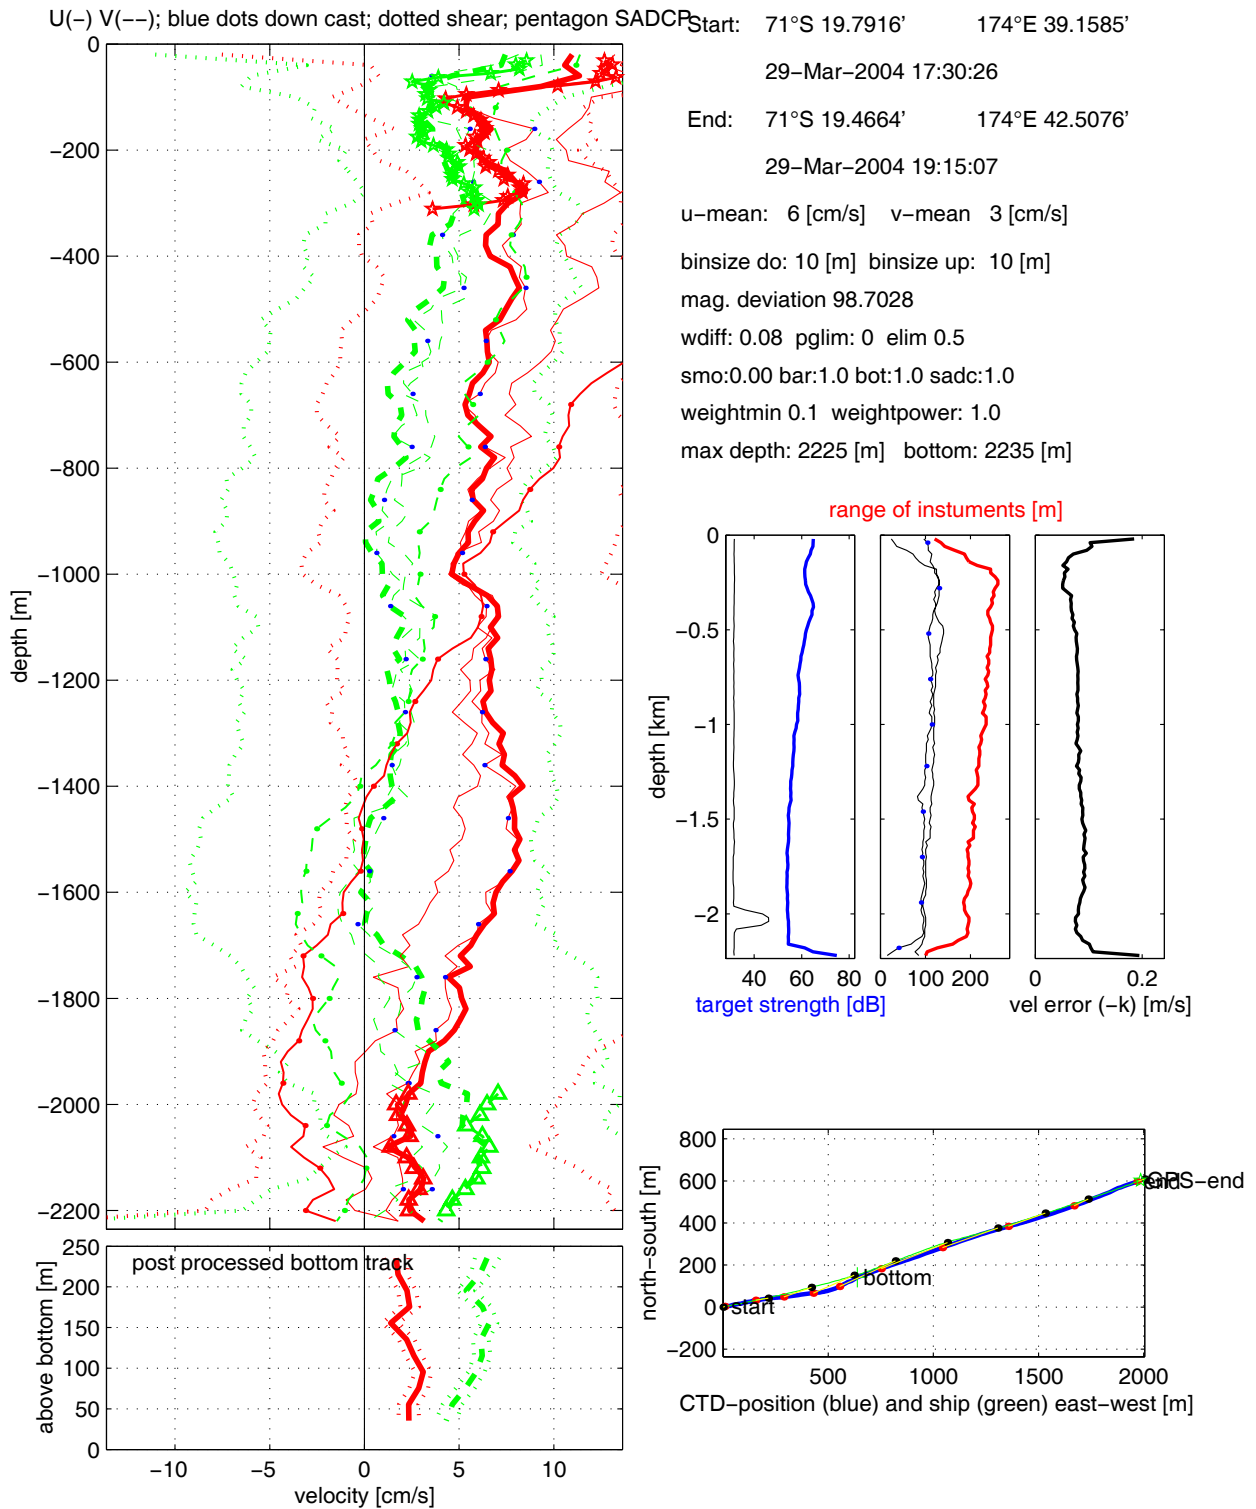


Figure 15

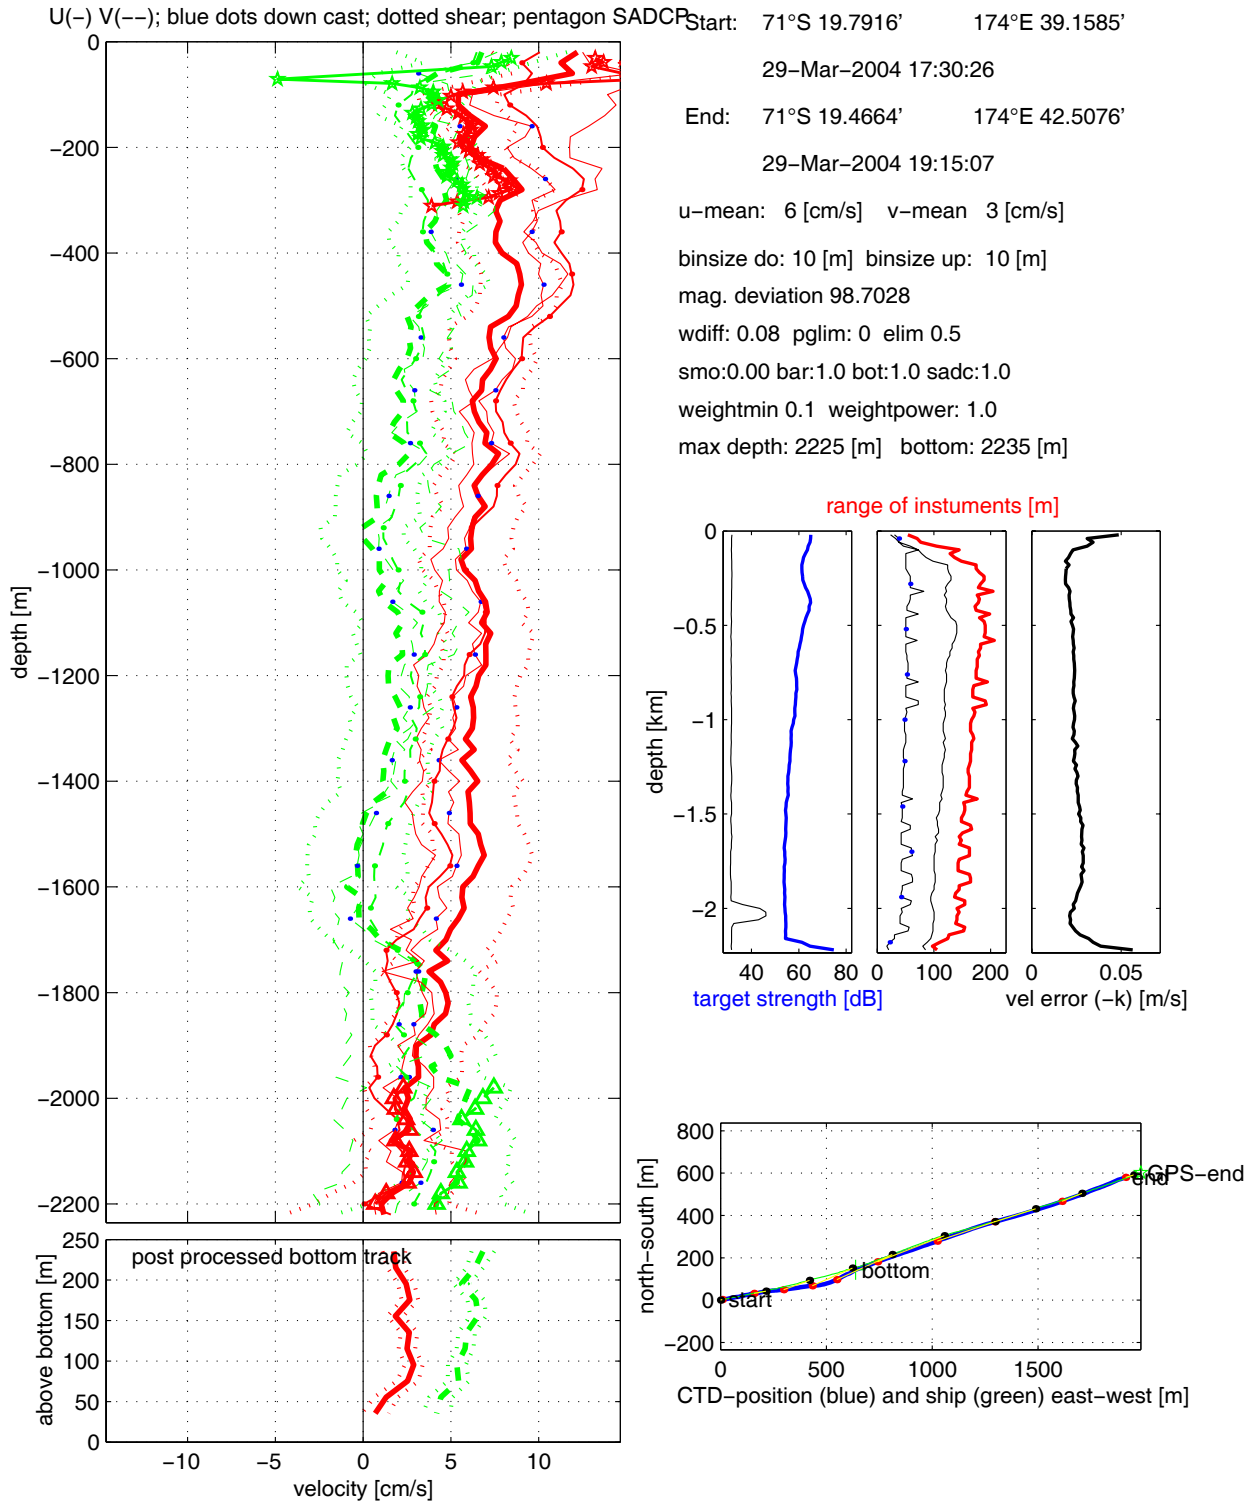
Station : p402210 Figure 1



LDEO LADCP software: Version 8b: 16 Mar 2004

Figure 16

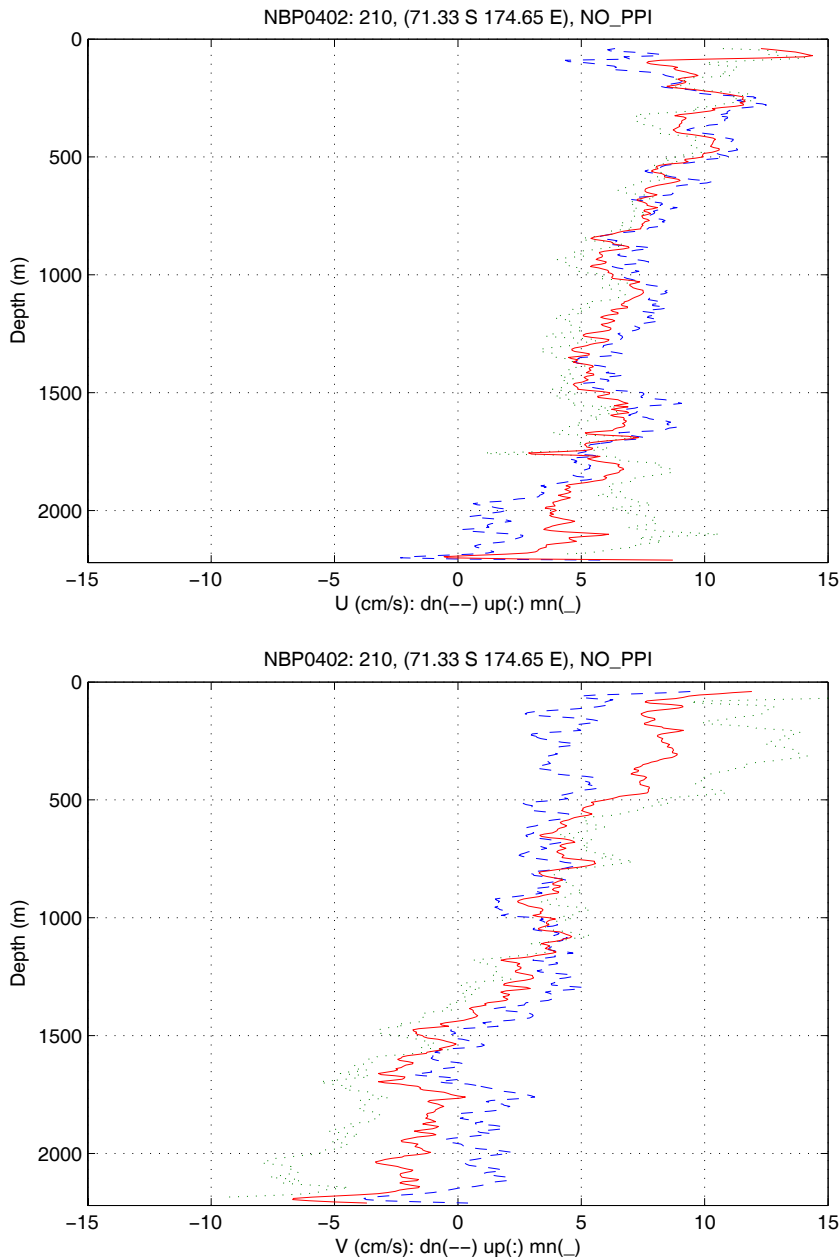
Station : p402210 Figure 1



LDEO LADCP software: Version 8b: 16 Mar 2004

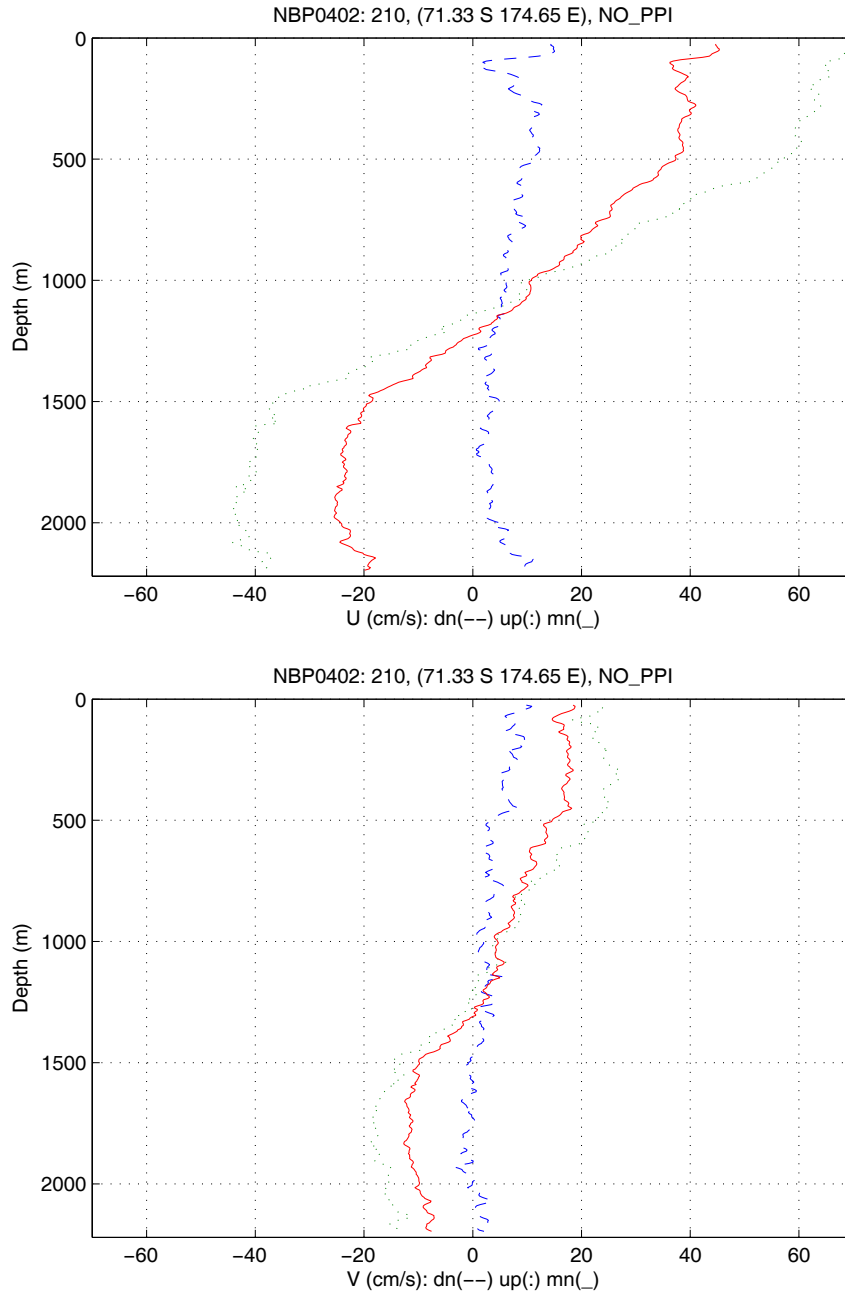
Figure 17

In order to check the effectiveness of Firing's wake editing, the data from a wake-affected station were processed using the recommended wake-editing settings ($\text{min_wake_w} = 0.1$; $\text{wake_hd_dif} = 20$; $\text{wake_ang_min} = 15$; $\text{n_wake_bins} = 2$). While the velocity profile from the uplooker (Figure 18a) is not a particularly good one, judging from the down- vs. upcast consistency, it is nevertheless approximately consistent with Visbeck's inverse solution within the error bars. The downlooker solution (Figure 18b), on the other hand, is a characteristic X profile. This suggests that at least some of the bad profiles observed elsewhere may be caused by wake contamination as well.



2004-4-7 9:32

Figure 18a



2004-4-7 9:36

Figure 18b

1.3.3 Potential Bias in Post-Processed BT Data

There are three methods for getting bottom track data from RDI Workhorse ADCPs (in order of decreasing expected accuracy): using dedicated bottom-track pings, from water-track pings processed by the instrument (RDI BT-from-WT data), and by post-processing regular water-track data. As long as none of these methods has any bias they are equally useful, because the accuracy can be increased more-or-less arbitrarily by adding bottom-tracking stops (as was done

during the Aurora Australis cruise A0304). Assessing whether there are biases is difficult because such biases would most likely be a function of the speed of the instrument over ground. In the case of our data set there are some casts where the instrument was towed near the sea bed (because the ship had to drift with the ice) but those casts are also characterized by comparatively few BT data. In these casts the uncertainties, especially of the post-processed BT data, are too large to determine any bias by direct comparison with the BT-from-WT data.

However, there are indications for potential bias in the BT data determined from post-processed WT data. The top right panel of Fig.2 in Visbeck's inversion paper (JAOT, 2002), for example, shows large (of order 10cm per 10m) shear in the horizontal velocities near the sea bed, implying that the exact location of the seabed (within a bin) can have a significant effect on the resulting BT velocities. Bias in BT data from post-processed WT pings is a problem because a BT-from-WT mode is only available for the 300kHz Workhorse but not for the 150kHz RDI ADCP used by other groups. (Using dedicated BT pings has other disadvantages as discussed below.)

Because of the potential bias problem it was decided to investigate the velocity shear near the sea bed in more detail. In the top right panel of Figure 19 the normalized BT-referenced vertical velocities of cast 224 are shown. (Vertical, rather than horizontal velocities are used because of the comparatively larger signal.) The velocities observed when the instrument was more than 120m above the sea bed are plotted in green, the velocities observed when the instrument was below 70m are shown in red and the remaining velocities are colored blue. At and especially below the sea bed there is a strong shear, which increases with decreasing distance from the sea bed. Velocities below the sea bed are determined by late arrivals, i.e. by sound energy in the "outer" side lobe scattered by the sea floor farther away from the instrument than the main beams. Since the acoustic paths in this "outer" side lobe are angled less steeply than the nominal beam angle the apparent vertical velocities below the sea bed are biased low. In contrast, the energy from the "inner" side lobes arrives early and contaminates the bottom part of each profile (15% in case of a 30-degree beam angle). In some of the profiles (only weakly in the one shown) the contamination from the "inner" side lobes introduces a high bias in the vertical velocities immediately above the sea bed. The solid lines show the expected shear calculated from simple geometric considerations, supporting the hypothesis that the shear is due to sidelobe contamination.

NBP0402-224 Figure 14

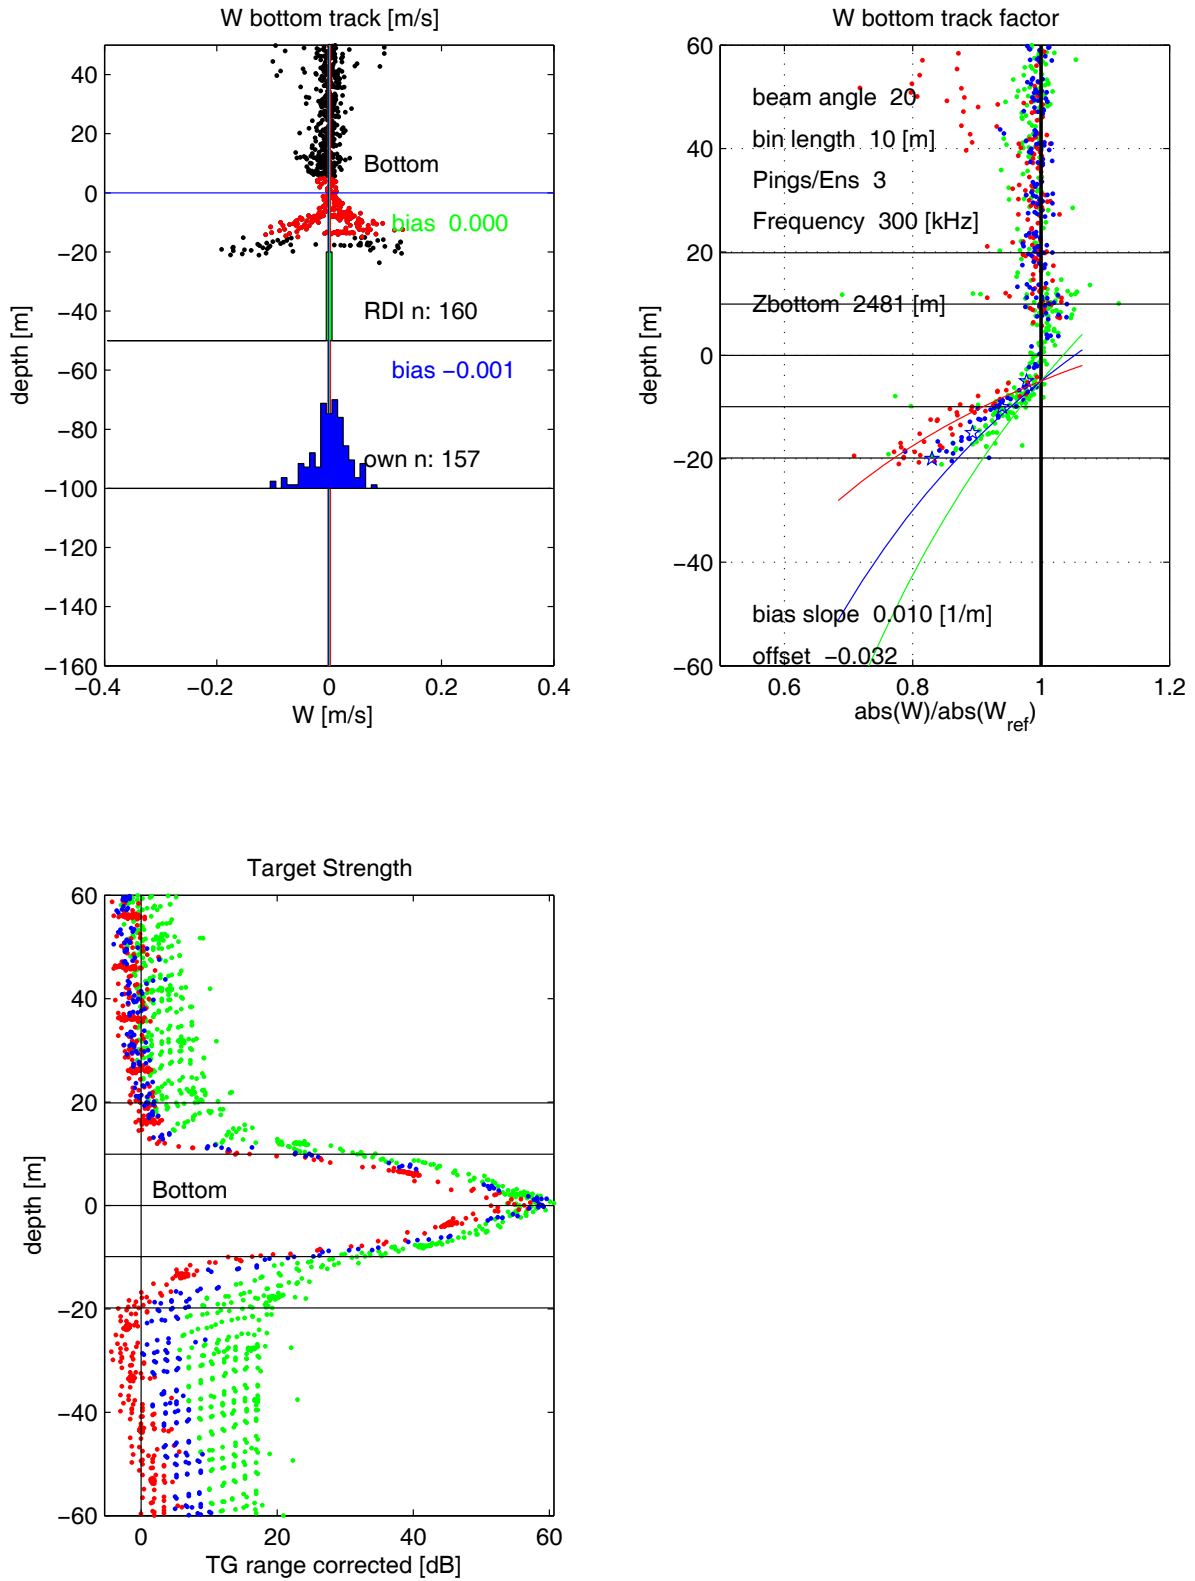


Figure 19

Because of the geometry the side-lobe-related shear in the horizontal velocities near the sea bed is more pronounced than the corresponding shear in the vertical velocities. However, presumably at least in part because of the smaller signal, the pattern is much less clear in the data. It will be noted that the velocities measured at the level of the sea bed are unbiased. If the distance to the sea bed is known accurately, unbiased BT velocities can therefore be derived from WT pings. This is easier to do before the velocity data are binned, which is presumably the reason why the RDI BT-from-WT data are less noisy than the post-processed BT data. In any case, the bias can be minimized if BT data are collected as far away from the sea bed as possible --- it may therefore be advisable to implement a BT stop at some distance above the sea bed, as was done during the Aurora Australis A0304 cruise.

1.3.4 Using Dedicated Bottom-Tracking Pings

In order to compare the different methods for determining BT velocities RDI kindly provided a BT upgrade for one of the LDEO instruments. Because of some uncertainty regarding the synchronization of Workhorses without LADCP mode (the LADCP feature has to be de-installed before installing the BT feature) the trials were carried out toward the end of the cruise on stations 225--232. These casts were all executed in heavy ice where it was not possible to tow the instrument, i.e. the BT velocities were too small to determine whether there is a bias in any of the methods. Nevertheless, the experiment proved interesting.

For the tests the uplooker was left with the LADCP upgrade installed. Synchronization between the uplooker and the downlooker (with BT instead of LADCP feature) worked exactly the same as synchronization between two instruments with LADCP features installed. Both instruments used firmware version 16.21.

Figure 20 shows the median fluctuations of successive BT speeds. BT-track data were collected in 3 different configurations: from single-ping WT data using the built-in RDI BT-from-WT mode (red), the same from 3-ping ensembles (blue), and from single BT pings (purple). High median fluctuations are presumably caused by true instrument accelerations. The lower bounds in each configuration are determined by the velocity uncertainties; consistent with this assertion, the three-ping BT-from-WT uncertainties are approximately 1.7 (square root of 3) lower than the corresponding single-ping uncertainties. The uncertainties associated with true BT pings are nearly an order of magnitude smaller than the corresponding single-ping uncertainties from WT data, suggesting that using BT pings is potentially very useful for collecting LADCP profiles.

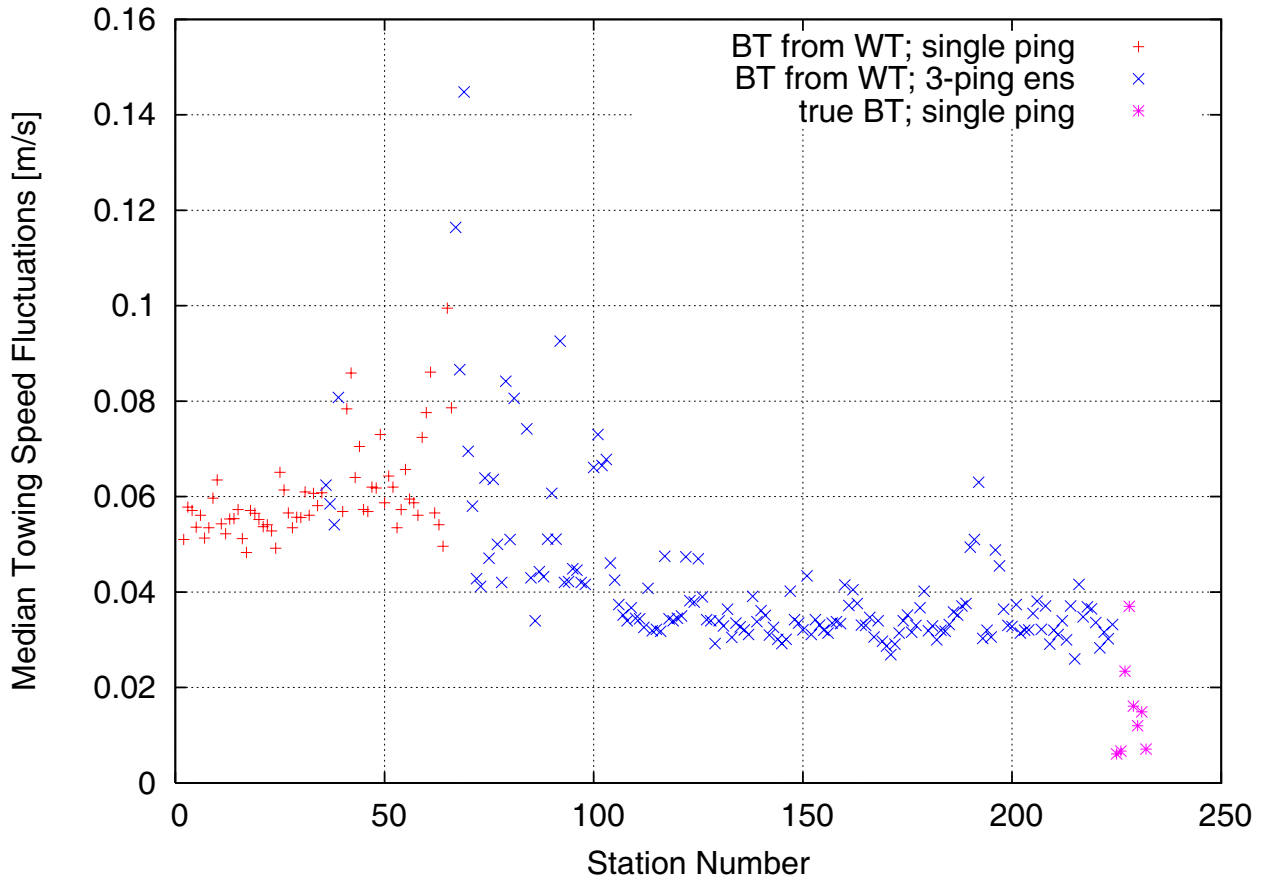


Figure 20

There are three disadvantages to using BT pings, however. First, BT pings increase power consumption --- our data do not allow a detailed assessment of this effect, however. Second, the interleaving of BT and WT pings reduces the available amount of WT data. This can be alleviated to some degree by combining single BT pings with several WT pings in a single ensemble. We used 3-WT+1-BT ping ensembles for all but the first two of our true BT casts. When doing this the time lag between the BT and the WT data can become an issue, however, especially in rough seas. Because our test casts were done in heavy ice it is not possible to assess this effect from the available data. However, lag correlations between w (dz/dt) calculated from CTD and BT data, and from CTD and reference-layer WT data of the single-ping test casts show offsets of 0.3s and 0.8s, respectively. While it may be possible to temporally interpolate the BT data to match the times of the WT data this was not tested.

Another potential way of reducing the number of BT pings without having to resort to multi-ping ensembles consists in using the RDI BD command, which suspends the generation of BT pings for a number of ensembles when no bottom is detected in a given ensemble. While this command is described in the RDI manual it does not appear to be implemented, however --- the LDEO Workhorses return an unknown-command error.

1.4 Cruise-specific issues (A. Thurnherr, B. Huber)

The Dell computer being used developed some problems which we thought were heat-related, manifested by the screen going blank and the keyboard becoming unresponsive (including the power button). When this happened, the only way to revive the computer was to disconnect power, remove and replace the battery to completely hard-reset the computer. We borrowed a Compaq notebook computer from RPSC, but it had no serial port, only USB. We produced a dual-boot disk drive in the Dell, swapped that disk with the Compaq, and ran the system under Linux on the Compaq. The disk swapping was facilitated by having a disk cloning kit manufactured by Apricorn, with several spare hard drives. Thus, we were able to preserve the original Dell disk, clone it to produce a dual boot system, and later, with our third spare drive, produce a Linux-only disk.

The Dell was disassembled to see if we could diagnose the shutdown problem. During the disassembly, it was given a thorough cleaning. Upon reassembly, the shutdown symptoms disappeared. After thorough testing under Windows, the Linux disk was installed and the Compaq returned to RPSC.

Several different instrument setups were used during the cruise. Initially, for casts 1-35, single ping ensembles were used. For casts 36--39, 3-ping ensembles were tried but this led to mutual interference by the instruments and we reverted to single pings for casts 40--66. After solving the synchronization problem (by lengthening the ensemble time) 3-ping ensembles were used for casts 67--224. For casts 225ff the LADCP feature was replaced by the Bottom-Tracking feature in the downlooker. In casts 225--226 ensembles consisting of one BT and one WT ping were used. During the remainder of the casts, the downlooker averaged 1-BT/3-WT pings per ensemble. Because of an oversight, the uplooker remained in single-ping mode (1 ping per 4 seconds), which did not significantly degrade the resulting velocity profiles, however.

A note on timing: during the final tests with the true BT mode it was discovered that after the command TE00:00:03.5 the ensemble time is 3.05s (instead of 3.5s as intended). Using TE00:00:03.50 solves that problem. It is likely that at least some of our initial synchronization troubles are related to this quirk.

Once we entered deeper water some of our profiles began deteriorating because of short instrument ranges. It was noted that in cast 175 most of the downlooker data in bin 1 was rejected. Suspecting ringing, the blanking distance was doubled to 10m for cast 176 and 177 but the same behaviour occurred. We then decided to try a setting first suggested by John Church on the Aurora Australis cruise A0304, namely to set the blanking distance to zero and always discarding the data in the first bin. This solved the problem (i.e. the data in the first bin after this change are as good as those in bin 1), and it appears that our subsequent casts in deep water were less plagued by range problems. Whether this holds up in different locations remains to be seen.

CTD 1: Initial configuration was sn 150 as master, sn149 as slave. The master came back after the cast with no data. During bench test, the unit failed the TRANSMIT sequence of the test, and the beam continuity check for beam 4. The unit was disassembled and found to have a loose cable between one of the circuit boards and the head. Unit was reassembled, tested again. This time is passed the TRANSMIT test, but failed the beam continuity test.

CTD 2: sn 149 slave, sn 754 master. Processing software reported a weak beam 3 for the master. Reinstalled sn150 as master.


CTD 3: sn150 had trouble waking up for cast initialization. Post processing reported weak beam 4 (range of 145 vs 175 for the other 3 beams). Changed heads for subsequent casts: sn 754 slave, sn149 Master. Will return sn 150 to RDI for repairs.

CTD 4: processing software reports possible weak beam 3 on sn 754. This was later determined to be a non-issue, as even the weaker beam had excellent range.

CTD35-67 - experimented with various combinations of ping/sync and ensemble settings to reduce data file size and retain good data. Settled on 3 pings per ensemble, 3.5 sec/ensemble, 0.9 sec per ping.

1.5 Battery consumption (B Huber)

The ADCP reports transmit voltage as part of the Variable Leader Data frame. This was used to track battery consumption, after scaling the integer count to obtain approximate voltage (Figure 21). It was found that one 35 D-cell pack will power approximately 140,000 2-head pings. The battery life depends upon duty cycle; the alkaline packs recover a bit when idle between casts. Frequent, short casts will consume more batteries than infrequent, longer casts of the same number of pings. Seven battery packs were used for the 232 casts of NBP04-02. In retrospect, we probably could have squeezed more out of the packs, especially when casts were expected to be short, but lacking an adequate history of power consumption and battery performance, we chose to be conservative and replaced batteries when the voltage during use went below about 38 volts.



040429 LDEO bah

Figure 21

Addendum to NBP04-02 LADCP report
October 12, 2004
B. Huber

Re: WHM300 sn 150 repairs following the cruise

As noted in the LADCP cruise report, sn 150 experienced multiple problems during the cruise, and was subsequently returned to RDI for diagnosis and repair. The symptoms reported to RDI were failure of beam 4 continuity test, failure of transmit and receive tests, and difficulty waking up (no response to break). Secondly, RDI was asked to evaluate the compass calibration, and perform a new cal if required

Their diagnosis was as follows:

1. Beam continuity, transmit and receive failure, and wakeup difficulty - confirmed, attributed to faulty Receiver, PIO board, and ceramic ground wire.
2. receiver board not properly grounded, beam 4 black wire not grounded. Rcvr board required retuning and upgrade to current rev levels.
3. PIO board requires upgrade
4. DSP board requires upgrade
5. firmware is up to date.

See Repair Report in the Workhorse documents binder for full details of corrective action taken. Summary- all required upgrades performed; compass cal performed; DSP board replaced; receiver board retuned.