SHIPBOARD ADCP MEASUREMENTS

E. Firing and J.M. Hummon
University of Hawai'i at Manoa School of Ocean and Earth Science Technology, 1000 Pope Road, Honolulu, Hawai'i 96822 U.S.A. e-mail: efiring@hawaii.edu, hummon@hawaii.edu

1. INTRODUCTION

Ship-mounted acoustic Doppler current profilers (ADCP) have been used for over 25 years, and have been available on most research ships for much of that time. They are easy to operate on a routine basis, permitting nearly continuous monitoring of upper ocean current structure beneath each ship. Nevertheless, their potential has not been fully realized for a variety of reasons, including data degradation because of poor system installations; lack of monitoring to detect and correct system faults; lack of attention to data processing and archiving; and lack of clearance to make measurements in foreign EEZs.

The purpose of this document is to provide a starting point, an introduction, for those who may wish to improve the future usefulness of shipboard ADCPs. Additional historical background and detail may be found in the predecessor to this document (WHPO, 1994) and in King et al., 2001. More up-to-date details about the relevant hardware and software will typically be found in evolving documents on the web, some of which will be referenced here. In some cases, the interested reader may need to make personal inquiries in order to get the latest information.

Here is the short-form recipe for the best use of a shipboard ADCP:

1. Select a set of instruments suitable for the ship and for its region of operation (e.g., deep water versus shallow).
2. Design and build a good transducer installation, minimizing noise and bubbles, and maximizing access.
3. Carefully install and check the heading sensors.
4. Operate the system routinely and consistently.
5. Monitor the system for correct operation.
6. Process and archive, in a public location, all measurements.

It is obvious—but somehow seems to get overlooked at critical times—that there is no substitute for getting good raw measurements. Sophisticated processing and personal attention sometimes can ameliorate the effects of data degradation, but at considerable cost, and rarely with results as good as could have been achieved with a good installation, running properly.

2. INSTRUMENTATION

A shipboard ADCP system requires the Doppler sonar itself, plus at least two ancillary data streams: position and heading.

2.1 ADCP

At the time of this writing, there is only one manufacturer of scientific shipboard ADCP systems in common use on medium and large research ships: Teledyne RD Instruments (TRDI). Therefore we will discuss the few specific types of instrument presently available from TRDI. This is not intended to pre-judge any competing products that might appear in coming years. We also note that the one-of-a-kind shear-optimized Hydrographic Doppler Sonar System (e.g. Rainville and Pinkel, 2004) is not within the scope of this document.

The most basic characteristic distinguishing one instrument from another is the operating frequency. Lower frequencies require larger transducers, and yield greater depth range at the expense of vertical and temporal (or horizontal, if the ship is underway) resolution for a given accuracy. To a first approximation, range is inversely proportional to frequency; but other factors, such as the highly non-uniform distribution of acoustic scatterers, make this a very crude approximation.

2.1.1 Instruments

The primary line of blue-water shipboard ADCPs is named Ocean Surveyor. Using phased-array transducers, such that a single flat panel produces all four acoustic beams, it is available in nominal frequencies of 38 kHz, 75 kHz, and 150 kHz (OS38, OS75, and OS150, respectively). Typical effective depth ranges for these three models are 1000-1600 m for the OS38, 600-800 m for the OS75, and 200-250 m for the OS150.

For improved resolution at short ranges (less than 100 m depth), the 300-kHz Workhorse Mariner (WH300) may be used. It has a conventional transducer design with a separate ceramic element for each beam.

Ideally, a blue-water research ship should have at least two ADCPs: either an OS75 or an OS38 for range, and either an OS150 or a WH300 to provide better resolution in the upper ocean. In considering the OS75 versus the OS38, one should bear in mind that the additional range of the latter can provide substantial scientific benefit, while the cost difference, even including the installation cost, is minuscule compared to, say, the annual operating cost of the ship; the cost of the OS38 is also small compared to the cost of the other major sonar typically found on large research ships, the multi-beam bottom mapping system.

2.1.2 Installation

The transducer installation is *critical* to the performance of a shipboard ADCP. Most ships entrain a bubble-laden layer of water under the hull under some sea conditions. When such a bubble layer goes under the transducer it scatters and absorbs the sound, while generating a high level of noise. The result is some combination of biased velocity estimates and reduced range, sometimes making the data unusable. Ideally, the bubble layer problem is solved at the ship design stage. On a SWATH ship such as the *R/V Kilo Moana*, there is generally little or no bubble sweep-down. On a conventional monohull, the transducer location may be selected to minimize exposure to bubble layers; but the most effective measure by far is to mount the transducer in a deep-extending faired housing such as the centerboard on the *NOAA Ship Miller Freeman* or the acoustics pod on the *RVIB L. M. Gould*. Keeping the transducer as far as possible from the bow thruster is also important; when a ship is on station, bubbles and noise from a bow thruster can make velocity profiling impossible.

In addition to the acoustic qualities of the installation, one must consider its robustness and maintainability. If possible, the transducer should be mounted in such a way that it can be removed without drydocking the ship. It is also desirable to have a dry back–that is, to mount the transducer such that the electrical connector on its back surface is accessible in air inside the ship.

Additional information regarding transducer mounting is beyond the scope of this document; suffice it to say that anyone planning to install a shipboard ADCP should poll the community and take every possible measure to optimize the design.

2.2 Ancillary data

Because the ADCP measures the velocity of the water relative to the transducer, ancillary navigational data are essential for calculating the oceanographically relevant quantity, the velocity of the water over the ground. One must know the orientation of the transducer relative to the earth, and the velocity of the transducer over the ground. In addition, for instruments (e.g., Workhorse series) with a discrete transducer for each beam, one must know the speed of sound at the transducer face.

2.2.1 Position

With the universal availability of high-quality GPS positioning, the problem of knowing the velocity of the transducer over the ground is solved, provided one can average over an adequate time interval, and provided the position fixes and the ADCP single-ping profiles are recorded with a common time base. Additional discussion of this topic will be deferred to the section on data processing.

2.2.2 Attitude

Although shipboard ADCP data processing might be optimized by using pitch and roll measurements for each ping, in present practice this is not typically done. The danger of data corruption if the pitch/roll correction is not done well, and the added complexity of making this correction, usually outweigh the small gain in accuracy and effective vertical resolution that the corrections might bring. Therefore, pitch and roll measurements will not be discussed further in this document.

For calculating water velocity over the ground, heading must be known to O(0.1-degree) accuracy. Conventional mechanical gyro compasses are very reliable, but their errors are too large; on time scales of tens of minutes to the length of a cruise, various factors can lead to gyro errors and drifts of roughly a degree in low and middle latitudes, and several degrees at high latitudes. Optical compasses can be better, but they are still subject to reduced accuracy at high latitudes, they have not always been reliable, and a gradual loss of accuracy can be hard to detect. Very high accuracy can be achieved by attitude sensors using differential GPS carrier phase measurements with an array of two or more antennas. In principle, the availability of such instruments should solve the heading-accuracy problem for shipboard ADCPs; but in practice it does not, because the GPS-based attitude sensors are often unreliable.

Given the problems with all known heading sensors, the best practice has three components:

- Use *at least* one gyro compass and one GPS-based compass. The former can be mechanical or optical, or both. The latter can be GPS-only, such as the Ashtech ADU series, or it can be an integrated sensor using GPS together with inertial and/or gyro sensors, such as the POS/MV and the Seapath. GPS-based sensors must provide data messages with information about the source and quality of the attitude measurement. For a GPS-only sensor, the quality estimate will be based on the data redundancy and the consistency of the solution. For a GPS-plus-inertial system, it is crucial to know when the attitude estimate is based on a valid GPS calculation, versus dead-reckoning from the inertial sensors during a GPS dropout. Attitude sensors should be chosen based on their proven record in shipboard scientific applications.
- Take all possible care when installing and maintaining all attitude sensors. GPS-based sensors in particular are very sensitive to the quality of their antenna installation and to the accuracy of their survey data, specifying the antenna array geometry; when an inertial measurement unit is included, the location and the orientation of that unit relative to the antenna array are also critical.
- Monitor the performance of all sensors. No one sensor can be trusted unconditionally; all sensor types can and do fail, sometimes in obvious ways, sometimes via subtle, gradual drifts. GPS-based sensors typically have short data dropouts. If the dropouts become more frequent and/or longer, it often means the unit needs to be reinitialized. If that doesn't help, it may need to be resurveyed, or repaired.

2.2.3 Sound speed

For phased-array transducers such as are used by the Ocean Surveyor series, the scale factor for converting the measured Doppler shift to water velocity relative to the transducer depends solely on the geometry of the transducer. For conventional transducers, however, the scale factor is proportional to the sound speed at the transducer face. If the transducer is exposed to the ocean, or if it is behind a window in a well filled with fresh water, or with water of known salinity, then the sound speed can be calculated based on the temperature of the water near the transducer; the temperature of the transducer itself, as measured by the ADCP, normally is adequate. Ships operating at high latitudes, however, may use a transducer well filled with an anti-freeze solution. Then the sound speed as a function of temperature may be poorly known, so it is best to include a dedicated sound speed sensor in the transducer well, providing a continuous record of sound speed at the transducer face. This method has been used successfully on the *R/VIB Nathaniel B. Palmer* and the *L. M. Gould*, for example, in the wells housing their old 150-kHz instruments.

3. DATA ACQUISITION SYSTEMS

Second in importance to the instrumentation–the ADCP and the ancillary sensors discussed above–is the data acquisition system (DAS). We will briefly discuss the system's essential characteristics, and then introduce the two systems in common use at present, VMDAS and UHDAS.

3.1 Requirements

The DAS has two basic functions: to set the ADCP's operating parameters (depth cell size, number of cells, etc.), and to record the data from the ADCP and from the ancillary sensors. Requirements and desirable characteristics include:

- A user interface that is simple enough to be easy to use, and that makes sensible instrument setup likely, while providing adequate control for the range of conditions that will be encountered. The user interface should make the operating parameters and operating state of the instrument clear, and should provide diagnostic information for monitoring the health of the instruments and of the data acquisition process.
- Flexible and reliable recording of raw data from all sensors, tagged with a common time base. It should be possible to translate from that time base to Universal Coordinated Time (UTC).
- A file naming system that makes it easy to match data from different sensors, and that makes an alphanumeric sort of the file names yield a sequence in temporal order.
- Speed-log functionality adequate to replace the ship's navigation speed log is advantageous for two reasons: it can render the use of a possibly interfering acoustic speed log unnecessary, and it gives the ship's officers an interest in keeping the ADCP running.

In addition to running the instrument and recording the data, a DAS may include automated data processing, display, and access functions. Such functionality facilitates real-time use of the current profiles, such as for adaptive sampling, and general accessibility of the data to the science party throughout a cruise.

3.2 VMDAS

The DAS supplied by TRDI for use with their Workhorse and Ocean Surveyor instruments is called VMDAS. It is a single program running on Microsoft Windows; hence it is easy to install. Although it lacks some desirable characteristics, and is quite limited in its ability to handle ancillary data, it can be used to collect a high-quality data set. Particular care must be taken to avoid serial data overruns in the ancillary data streams, and to avoid any irregularities in the computer's time base. Additional recommendations can be found at http://currents.soest.hawaii.edu/docs/adcp_doc/, and in TRDI's documentation (including the VMDAS help system).

3.3 UHDAS

Many ships in the US University-National Oceanographic Laboratory System (UNOLS) fleet, and a few in the US National Oceanic and Atmospheric Administration (NOAA) fleet, run a DAS developed at the University of Hawaii and therefore dubbed UHDAS. A list of installations is maintained at http://currents.soest.hawaii.edu/uhdas_fromships.html.

Running on Linux, UHDAS was developed to give the scientific ADCP user community more flexibility and control over data acquisition by providing an open-source code base. Key features include

- the ability to handle any number of ancillary data streams;
- the ability to control any number of ADCPs, and to log their data without duplicating the ancillary data logging;
- support for the original "Narrowband" ADCP from RDI as well as for the Workhorse and Ocean Surveyor series;
- support for interleaved pinging on the Ocean Surveyors, displaying data from both ping types;

- automated data processing, plotting, data and plot access via the shipboard network, and diagnostic emails to shore.

At the time of this writing, UHDAS is an evolving system. It combines elements written in C, Python, and Matlab, with the latter being phased out. Because of its many components, its integration into the ship's network, and its developmental status, it is relatively difficult to install, but once installed it is easy to run. More information about UHDAS is available at http://currents.soest.hawaii.edu/adcp/programs/adcp_doc/index.html. A snapshot of the UHDAS shipboard web page is at:
http://currents.soest.hawaii.edu/uhdas_fromships/example_atseaweb/index.html.

4. ADCP OPERATION

The most general recommendations for operating a shipboard ADCP are:

1. Choose a single, reasonable set of instrument setup parameters, and stick with it; only rarely is it helpful to fine-tune the parameters, or make frequent changes.
2. Use bottom tracking *only* when needed for checking the transducer orientation or other aspects of instrument performance. Routine use in shallow water when leaving or returning to port may be helpful, bearing in mind that bottom tracking degrades the water velocity profile data.[1]
3. Monitor the sensors–especially the heading sensors–and both raw and processed data streams.

4.1 Instrument setup parameters

When choosing the instrument setup parameters, the starting point is always the manufacturer's defaults or recommendations for a given instrument and application. Here we will discuss considerations for only a few of the available setup parameters.

The Ocean Surveyor (OS) instruments can operate in either or both of two modes: "narrowband" (NB), in which the ping is a single unmodulated tone, and "broadband" (BB), in which the ping is modulated via phase reversals to yield several repeats of a code (Pinkel and Smith, 1992). The Workhorse (WH) operates only in BB mode. Compared to NB mode, BB mode provides increased single-ping accuracy for a given vertical resolution, or equivalently better vertical resolution for a given accuracy; the costs include reduced range and robustness. BB mode is more sensitive to noise and to other sonars, and is more likely to interfere with other sonars. As a default for deep-water profiling we recommend using NB mode on the OS, believing the advantages outweigh the reduced vertical resolution.

Inherent in BB-mode Doppler sonars, and in the present implementation of NB mode, is a velocity ambiguity that limits the range of the axial velocity component as measured on each ping by each beam. In the WH, this ambiguity interval is always centered on zero, and the range is set by the "WV" command. The allowable range depends on which of two possible bandwidths (or coding schemes) are being used, as set by the "WB" command. We presently use the default WB0 command (wide bandwidth), in part because it permits a large ambiguity interval; we use WV550 to minimize the likelihood of an ambiguity error. (Note that the axial velocity component may get a large contribution from the wave-induced heave of the transducer.) In the OS, the ambiguity interval is

[1] For small boats in shallow water, lacking accurate heading sensors, routine use of bottom tracking is advised.

not centered on zero, but is biased to fit the typical operation of a ship at sea: the measured horizontal velocity is dominated by the ship's forward motion, at speeds from zero to 6 or 7 m/s, so the ambiguity interval is biased toward the transducer for the forward-looking beam(s), and away from the transducer for the aft-looking beam(s). This requires the "EA" setup command, to specify the orientation of the transducer relative to the keel. High accuracy is not needed—getting within 5 or 10 degrees is adequate—but a bad setting can irretrievably ruin a data set.

A key aspect of instrument setup is determining the sampling in depth. The primary parameter is the depth cell size, expressed as a projection on the vertical rather than as distance along the beam. Although it is possible to control the length of the transmitted pulse independently of the received cell size, this is not typically done, and the necessary setup parameter is somewhat hidden. The cell size parameter is called "WS" for the WH and the OS in BB mode, and "NS" for the OS in NB mode. Short-term accuracy varies inversely with cell size. Note that the effective vertical resolution is always less–the smoothing in the vertical is greater–than indicated by the cell size alone. Effective vertical resolution is reduced by the spreading of the beams with distance, and by the range of depths swept out by a given cell as the ship pitches, rolls, and heaves. These two factors increase with depth. After setting the depth cell size appropriately for the instrument frequency, operating mode, and application, one sets the number of cells ("WN" or "NN" for BB or NB operation) so that the maximum range usually exceeds the achievable range. The third parameter related to depth sampling is the blanking interval ("WF" or "NF"), the distance from the transducer to the top of the first range cell. With a good transducer installation (little reverberation) and good availability of scatterers in the upper part of the water column, this can be as small as half the cell size. Often, however, this minimum (and instrument default) value leaves noticeable bias in the first depth cell. A larger value–the cell size, or larger–may eliminate the bias. The appropriate blanking interval must be determined for each ship based on review of data obtained under a variety of conditions. With any reasonable value, bias in the top depth cell may still occur occasionally and need to be edited out.

Having set the depth sampling, one must turn to the timing of the pings. Pinging more rapidly permits shorter averaging intervals for a given level of accuracy. It is important to leave enough time between pings, however, so that sound returned from one ping does interfere with the return from its successor. This consideration is most likely to be a limiting factor with lower-frequency ADCPs. For example, if an OS38 pings every 2 seconds in water of 2000-m depth, the bottom-bounce of a ping will arrive at the same time as the 700-m return of its successor, and hence may bias the velocity estimate at that depth. A 3-s ping interval is a reasonable minimum for this instrument. Another consideration in setting ping timing parameters is the potential for systematic interference with other sonars. If there are two or more sonars pinging at regular intervals, and they cannot be synchronized to ping either simultaneously or sufficiently interleaved so as to avoid interference, then they should be set to ping at rates that are not exact multiples; otherwise, the interference will appear at a constant or slowly varying depth.

4.2 System calibration

The velocity measured by an ADCP on a ship underway usually is dominated by the speed of the ship, not the speed of the ocean current. Therefore the estimate of the current is sensitive to small errors in measured velocity; errors in measured speed cause errors in estimating the component of current along the track, and errors in measured direction cause errors in estimating the cross-track component of the current. Speed errors can be minimized by a scale-factor calibration: that is, by determining and applying a small adjustment to the coefficient that multiplies Doppler shift to yield water velocity component along a beam. Direction errors are minimized by accurately determining the orientation of the transducer relative to the heading reference.

4.2.1 Scale factor

In the determination of the horizontal component of velocity, the scale factor depends on the horizontal component of the wavenumber vector of the transmitted sound. (We will assume that the frequency of the sound is known accurately enough, and is stable enough in time, that it can be considered a known constant. We will also assume the transducer is level; taking into account the tilts of a transducer at sea does not change the basic points we are making here.) For ADCPs with discrete transducers, the horizontal component of the wavenumber vector for each beam is a function of the speed of sound at the transducer face and of the inclination of the beam from the vertical. In contrast, for phased array transducers of the type used in the Ocean Surveyors, the horizontal component of the wavenumber vector is determined by the spacing of the transducer elements, and does not depend on the speed of sound.

In practice, other factors may cause small but significant biases in the scale factor, so one wants to estimate it and monitor it. Experience to date indicates that Ocean Surveyor scale factors cluster around 1.003; that is, the estimate of Doppler shift is biased low by about 0.3%. If the calibration procedure indicates a bias exceeding about 0.5%, it indicates that either the data were acquired under adverse conditions and the single-ping editing was not adequate for removing the bias, or there is a problem with the installation.

4.2.2 Orientation (heading)

The heading of the ADCP transducer is that of the master heading sensor, plus the heading offset between the transducer and the sensor. By "master", we mean the most accurate available sensor, typically using an array of GPS antennas. Given good, solid transducer and heading sensor installations, the offset needs to be determined only once, but very accurately. In practice, one wants to monitor routine estimates of this offset—the "heading calibration factor"—so as to detect any shifts that might occur. For example, a heading sensor may degrade in subtle ways, or a GPS antenna array might be bumped out of its original alignment.

4.2.3 Calibration methods

There are two sources of calibration information: bottom-tracking, and water-tracking.

Bottom-track calibration is based on ADCP estimate of the ship's velocity over the ground, as derived from dedicated bottom-track pings. Obviously, it is feasible only when the water is shallow enough that a good acoustic reflection from the sea floor can be obtained. In practice it also requires that the bottom slope be moderate, and that the depth be within a more restricted range, perhaps within 20% to 100% of the instrument's typical water-profiling range. The ship should be underway, preferably with speed and heading fairly constant for at least 30 minutes. Bottom-track calibration is useful primarily for determining transducer orientation relative to the heading sensor. We find that the bottom-track scale factor often differs from that for water-profiling, and it is the latter that we are most interested in.

Water-track calibration takes advantage of the ship's maneuvers in deep water. The scale factor and transducer orientation are calculated so as to minimize the root-mean-square differences between the estimated water velocity over the ground (using GPS navigation) before and after ship accelerations, including major turns as well as stopping on station and getting underway after a station. Each such acceleration provides a very noisy estimate of the calibration factors, so it is essential to average over a large number after eliminating outliers. A typical hydrographic cruise provides a good ensemble; a line with 50 CTD stations may provide 100 independent estimates.

5. PROCESSING

Deriving the desired product — profiles of ocean velocity over the ground — from the recorded ADCP and ancillary data requires several processing steps. Some of these steps are fundamental; others are needed because of deficiencies in the raw data. Detailed description of algorithms is beyond the scope of this publication, but the basic steps are the following:

1. Transform the velocity measurements for each ping from beam coordinates to instrument Cartesian coordinates.
2. Temporally align ancillary data time series with the ADCP time series.
3. Edit the single-ping velocity measurements and the ancillary time series.
4. Estimate the ship's heading for each ping, possibly using a composite based on more than one heading sensor.
5. Vector-average the single-ping velocity profiles in ensembles of a few minutes duration—that is, rotate the velocity vectors into east and north components, and average those components. The averaging should be done after splitting the profiles into the sum of a depth-independent part and a deviation (see Appendix B in Hummon and Firing, 2003). This separation of variables is also used in single-ping editing, as described below.
6. Iterate among editing, navigation, and calibration steps at the ensemble-averaged level:
   - Edit the velocity estimates, typically by inspection after applying automated criteria.
   - Add the ship's velocity based on finite-differences of GPS positions to the water velocity profiles relative to the ship, yielding an estimate of water velocity over the ground. Some light temporal smoothing may be applied.
   - Estimate and apply calibration and heading error corrections.
7. Assemble meta-data, including documentation of processing procedure and any data or processing anomalies.

5.1 Single-ping Editing

Single-ping editing is discussed in some detail by Hummon and Firing (2003); since then we have added some steps, so the procedure will be summarized here.

- Screen for amplitude (acoustic intensity) spikes, such as those caused by interference from other sound sources. For each depth-cell, we calculate a three-profile running median of the received amplitude for each beam. Where the amplitude exceeds this by some threshold, the corresponding velocity estimate is masked, and the mask is extended to include one cell above and one below the spike. The running median estimate of the amplitude is used for all subsequent steps.
- Look for a return from the bottom in each beam. This is done only if a global bathymetry product indicates that the bottom may be within the range of the instrument; otherwise a strong scattering layer can fool the bottom-detection algorithm. Bottom detection is based on the

shape of the amplitude profile. The shallowest bottom depth among the four beams is assigned to the profile as a whole, and the velocity is masked accordingly; because the downward-pointing sidelobe is reflected by the bottom, interference starts at about 85% of the bottom depth when the beam angle is 30 degrees.

- Screen the velocity using thresholds for error velocity (the mismatch between the two independent estimates of vertical velocity based on the two pairs of opposing beams) and the correlation recorded by the instrument for each velocity estimate.
- Eliminate "weak profiles", characterized by anomalously few valid velocity estimates near the top of the profile, where a solid return would be expected. This algorithm is particularly powerful in reducing along-track bias under adverse weather conditions or when operating in sea-ice.
- Mask velocity outliers as found in the residual from the separation of variables step used in ensemble-averaging (Hummon and Firing, 2003, Appendix B).

5.2 Editing ensemble averages

Given good single-ping editing, very little editing may be needed after the ensemble-averaging stage. The primary automated step is application of a "percent-good" threshold; e.g., mask out depth cells in which the average comes from less than 30% of the original pings in the ensemble.

The key aspect of editing ensemble-averaged data is visual inspection, looking for the following sorts of problems that may require manual masking of data, or going back to an earlier stage to correct an error:

- The single-ping bottom-detection algorithm does not work in very shallow water, when the bottom is within a few depth cells. Until an algorithm to detect this case is implemented, it must be flagged manually.
- Sometimes a visually-detectable bias occurs despite the single-ping editing. In this case it may be possible to minimize it by changing single-ping editing criteria, or it may be necessary to manually flag the suspect data, if the bias is deemed unacceptable.
- Errors in the heading estimate, either from a poor calibration correction or from problems with the heading data stream, may not be apparent until the fully edited and navigated data are inspected visually.

Underway bias from strong scattering layers presents the ADCP data processor with a dilemma: it is visually obvious that the estimated velocity over the ground is biased in the direction of the ship's motion immediately above the layer, and in the opposite direction below the layer. Should this bias be edited out? Our practice has been to leave it alone on the grounds that the biased measurements are better than no measurements at all. It is left to the user of the data to decide whether and how to apply additional editing.

6. ARCHIVING

For the shipboard ADCP observations to be useful, they must be archived and made available to potential users. There are two levels of archiving: the original data sets, as they are at the end of a cruise, and the data as processed after the cruise.

6.1 Original datasets

It is recommended that the original data sets be archived either locally, typically at the institution operating the ship, or centrally. These original data sets should include everything that would be needed to for later reprocessing, together with the result of whatever automated processing occurred on the ship. Typically this requires saving a single directory per cruise. Note that saving the single-ping data is useful not only in the event that the data need to be re-averaged with improved algorithms or a different averaging interval, but also because useful information (particularly about scatterer distribution) sometimes can be gleaned only from the unaveraged data.

For many ships, the new Rolling Deck to Repository (http://rvdata.us) program now fulfills this need for archiving original datasets.

6.2 Processed data

Although automated processing systems are improving, and in many cases can deliver a product close to a fully-processed dataset, interactive editing and meta-data compilation should always be performed. The result, which we refer to here as "fully-processed", includes only averaged data, so it is two orders of magnitude smaller than the original dataset. This processed version should be archived and served by a long-term facility. A suitable entry point is the Joint Archive for Shipboard ADCP (http://ilikai.soest.hawaii.edu/sadcp/)

7. REFERENCES

Hummon, J.M., and E. Firing. 2003. A direct comparison of two RDI shipboard ADCPs: a 75-kHz Ocean Surveyor and a 150-kHz Narrowband. *J. Atmos. and Oceanic Technol.* 20, 872–888.

King, B. A., E. Firing, and T. M. Joyce. 2001. Shipboard observations during WOCE. In *Ocean Circulation and Climate*. G. Siedler, J. Church, and J. Gould, eds. Academic Press.

Pinkel, R. and J. A. Smith. 1992. Repeat-sequence coding for improved precision of Doppler sonar and sodar. *J. Atmos. and Oceanic Technol.* 9, 149–163.

Rainville, L. and R. Pinkel. 2004. Observations of energetic high-wavenumber internal waves in the Kuroshio. *J. Phys. Oceanogr.* 34, 1495–1505.

WHPO (WOCE Hydrographic Program Office). 1994. *WOCE Operations Manual, WHP Operations and Methods*. WHPO 91-1. WOCE International Project Office, WOCE Report No. 68/91, rev. 1. (unpublished manuscript).