University of New Hampshire

University of New Hampshire Scholars' Repository

Center for Coastal and Ocean Mapping   Center for Coastal and Ocean Mapping

2003

Automatic Statistical Processing of Multibeam Echosounder Data

Brian R. Calder
University of New Hampshire, Durham, brian.calder@unh.edu

Follow this and additional works at: https://scholars.unh.edu/ccom

Part of the Oceanography and Atmospheric Sciences and Meteorology Commons

Recommended Citation

Automatic Statistical Processing of Multibeam Echosounder Data

By Brian Calder, Center for Coastal and Ocean Mapping & Joint Hydrographic Center, University of New Hampshire, United Kingdom

This paper presents the CUBE (Combined Uncertainty and Bathymetry Estimator) algorithm. Our aim is to take advantage of statistical redundancy in dense Multibeam Echosounder data to identify outliers while tracking the uncertainty associated with the estimates that we make of the true depth in the survey area. We recognise that a completely automatic system is improbable, but propose that significant benefits can still be had if we can automatically process good quality data, and highlight areas that probably need further attention. We outline CUBE and its associated support structures, and apply it to a dataset from Woods Hole, MA, USA. We illustrate CUBE's output surfaces, show that the algorithm faithfully maintains significant bathymetric detail, and how the algorithm's auxiliary outputs can be used in the decision-making process. Comparison with a selected sounding set shows that CUBE's outputs agree very well with traditional approaches.

Introduction

The Data Processing Challenge

Processing of Multibeam Echosounder (MBES) data is a challenging task from both hydrographic and technological perspectives. There has been an emphasis in the past on improving methodologies and technologies for the collection of data without a corresponding emphasis on new methods for data processing. We are now faced with the situation that we can collect data much faster than we can conveniently process it. With modern shallow water systems running at up to 9,600 soundings/second, data collection at the rate of approximately 250 million soundings/day/system is possible. Processing at that rate using conventional methods is more difficult: it is no longer realistic to continue with the traditional hand-examination processing methodology.

We have to find some acceptable solution to handle automatically as much of the data as possible. Ironically, collecting dense MBES data may be the best solution to the problem of MBES data. Multibeam systems and operating procedures have advanced to the stage where most data is mostly correct most of the time. With suitably dense MBES data, we should be able to construct statistically robust estimates of depth in almost all cases, and use the consistency of the data to indicate areas where there are difficulties that require further attention. An automatic method also provides an objective approach to the problem. Human operators are currently making subjective decisions about every single sounding that they select as 'not for use', with the time burden and quality assurance/control concerns that this subjec-

tivity implies. Regardless of training, experience and dedication, this will eventually lead to mistakes that may be untraceable. An objective automatic method should mean that the operators only have to examine the data that does not correspond to the norm. That is, we should have the operators examine only the data that really needs work, not routinely examine every sounding being gathered. In this way, we reduce the number of subjective decisions that have to be made, reduce operator fatigue and burnout, and facilitate faster processing of data.

The traditional hydrographic approach has been to consider the quality of the component soundings that are represented on the smooth sheet (i.e., the primary archive of the survey). Previous work on automatic processing has maintained this idea, whether attempting to simulate the human operator [Du et al., 1996], nominate dubious soundings by a robust measure of local neighbour properties [Debese, 2001; Debese & Michaux, 2002; Eeg, 1995], or by looking at statistical consistency in an area [Ware et al., 1992; Gourley & DesRoches, 2001] (see [Calder & Mayer, 2002] for a more extensive discussion). However, what this approach answers is the question 'How good is this measurement?' and not the question 'How well do we know the depth at this point on the seafloor?' We contend that this latter question is the one that we should be answering; that is, the processing goal is to determine the depth in the survey area, rather than select soundings. Once we have determined the depth sufficiently well across the survey area to build a suitable surface model, we may make hydrographic decisions on what is significant and what is not.

The restatement of the hydrographic question above is intuitively appealing. It is inherently statistical in nature, accepting that our knowledge of the depth may be limited, and subject to update as we gather more data. It implies that we can and should use more than one sounding (if available) to update our information on depth, using redundancy to deal with the noise inherent in each measurement. And it focuses directly on the quantity that we want to measure, aiming to get as close as possible to the 'correct' answer directly, before subsequently applying any safety constraints mandated by good hydrographic practice (see, e.g., [Smith et al., 2002]).

However, it also poses some problems. How do we estimate the errors in the measurements? How do we distinguish normal statistical variations from outliers? How do we utilise information from a set of neighbouring soundings to estimate the 'true' depth? The extent to which we can resolve these problems defines the advantages we can expect from an automatic processing method.

Hydrographic Concerns

As conscientious hydrographers concerned with safety and charting, the notions of 'estimated' depths, surface models and combinations of measurements should raise some concern, if not eyebrows. It is important to point out, therefore, that we do all of these things already. For example, we estimate depth by measuring travel time of sound and converting it, more or less well, into range, and thence through some ray approximation of acoustic refraction into depth and distance. We make an implicit prediction of surface continuity in every chart constructed through the use of selected soundings or contours. We combine measurements from a myriad of systems to make every MBES measurement. Each one of these measurements is in error, and so therefore is any combination of them. Hence, it makes no sense to talk of any one sounding as being the depth – all of the soundings have some error, and this error is not uniform across the swath, between systems, or across all survey environments. Consequently, unless we take account of these errors, we may be deceived about the depth in an area due to noise in the MBES system, in the motion sensors, or in the GPS.

Currently, we deal with data by experience and practice. We expect certain MBES to fail in certain ways; we ask operators to make subjective decisions on what is real and what is not; we strip out data past a certain off-nadir angle, even though it appears to be 'normal', based on a qualitative or roughly quantitative observation that outer beams are more noisy. However, none of these solutions is really adequate as data volumes increase – with a modern MBES survey, can we really affirm that we have inspected every sounding?

We suggest that a statistically justified estimate of depth is not only a reasonable method of proceeding, it is a required method [Smith et al., 2002]. It is certainly a more objective solution to the problem.

The CUBE Algorithm

We propose an algorithm that takes uncleaned MBES data and attempts to estimate the true depth at a collection of point locations arranged in a

network over the survey area. At each point, or node, we maintain an estimate of the true depth and the *a posteriori* variance of this depth, which we update as more data becomes available in the area. In order to deal with noise or outlier data, we implement a monitoring scheme that checks new data against current estimates; if the data is inconsistent (outwith limits based on the expected error associated with the data), then it is modelled and tracked separately. Hence, each node is represented by a collection of potential depth estimates, or hypotheses, each with an estimate of depth and its posterior variance. After all data is assimilated (or on demand), we attempt to choose the most likely hypothesis at each node according to a suitable metric – our goal is to determine the true depth by choosing the hypothesis that appears most likely given, e.g., number of depth soundings which agree on the depth, closeness to neighbouring depths, or consistency of data. We thereby construct a set of point estimates over the survey area, each theoretically representing the best statistically supportable estimate of depth in its location. These point estimates may then be connected into a surface description of the area, which is more readily manipulated and processed. Since the heart of the algorithm is concerned with the estimation of uncertainty in the measurements, we call the algorithm CUBE (Combined Uncertainty and Bathymetry Estimator).

The rest of the paper outlines the CUBE algorithm (for a more detailed mathematical development, see [Calder & Mayer, 2002]), and describes the trial implementation that has been built to test the ideas presented. We then describe a hydrographic survey in Woods Hole, MA, which illustrates the behaviour of CUBE, and the use of diagnostic indicators to guide operator effort.

Method

Estimation at a Point

The basic element of CUBE is an estimation node, defined at a point location with respect to some fixed projected co-ordinate system. We can define the location of a node absolutely, and the node therefore represents a true point in space. An immediate consequence is that the node only has to consider a single depth, since there can only be one seafloor at a point location. Therefore, the node does not need to track horizontal uncertainty (its location is known exactly), but only vertical uncertainty in the true depth at the location. Another immediate consequence of this basic definition is that the estimator we build only has to determine an unknown constant, which makes the estimation task significantly simpler.

A final consequence is that, under the null hypothesis that all of the depth soundings in the area are unbiased (i.e., on average, report the true depth), then it does not matter in what order we process the data. That is, we can take it all at once or one point at a time, and in any order. We can, in particular, sequence the data by the order in which it is recorded. Each node thus receives a sequence of data points representing the soundings in its immediate vicinity. The estimator then has to determine the best estimate of true depth from this sequence, and we may treat the problem from the perspective of time-series estimation.


Figure 1: Typical error performance of an MBES system in shallow water. These graphs show performance for a typical MBES on a small survey launch using differential GPS for basic positioning and a high-accuracy attitude sensor. Target depth is 25 metre

Error Models, Information Propagation and Optimal Estimation

CUBE's estimator starts with a quantitative estimate of the errors associated with each sounding. For each data point, we determine the predicted horizontal and vertical error using the model of Hare et al. [1995], which utilises a propagation of variance argument to convert errors in the MBES itself and those of its auxiliary sensors (GPS, IMU) into a predicted error for each sounding. The model is detailed, requiring many properties of the systems in use to be known (e.g., sample rate, accuracy of attitude measurement and patch test, etc.) The configuration is also specific to the survey platform in use since it depends on the offsets between the various instruments; once configured, however, the computation is straightforward. A typical error model response is shown in Figure 1, although this will of course change with MBES in use among many other factors, and should only be taken as illustrative.

The error model provides the basic error measurements required, and is the heart of the rest of the system, but it only provides information about the errors at the nominal location of the sounding, and contains both horizontal and vertical components. Since we are using a set of fixed nodes that have no horizontal errors, we must propagate the information implicit in the soundings to each estimation node location, and combine the vertical and horizontal errors. Our propagation of information method is based on a local bathymetric model that assumes that the local surface consists of at worst a constant slope; as long as we only use soundings that are sufficiently close to the estimation node, this is a reasonable assumption (Figure 2). To ensure that we do not use soundings inappropriately, we also increase the uncertainty associated with a propagated sounding as a function of distance through which the sounding has been propagated. This is implemented by scaling the vertical uncertainty associated with the sounding by a factor that increases quadratically with distance (Figure 3a). To incorporate the horizontal uncertainty, we assume that the sounding could be up to a fixed fraction of the horizontal uncertainty associated with the sounding further away from the estimation node than the nominal location (Figure 3b). Augmenting the distance by this fraction factors in horizontal uncertainty in a reasonable manner: the higher the uncertainty, the larger the distance scale factor and hence the higher the reported uncertainty at the node. Indeed, the scaling process provides many desirable features: soundings with higher initial vertical uncertainty are given lower weight; soundings farther away are given lower weight; soundings with higher horizontal uncertainty cause the uncertainty to scale


Figure 2: Propagation of information. Estimation at a point implies that we need to know the depth there; soundings, however, occur essentially at random. Hence, we must propagate the information to the location of the estimation nodes, taking care to model an increase in uncertainty associated with the fact that we are using the sounding at some distance from the nominal location

Figure 3: Uncertainty in propagation. The uncertainty associated with a sounding must increase the further we move from the nominal resolved location; in this case, it is modelled as a quadratic function of distance. Horizontal uncertainty is taken into account by assuming that the sounding may be up to the maximum likely distance away, rather than at the nominal distance. The difference is a linear function of the estimated horizontal uncertainty

faster, and hence have lower weight.

After the soundings are propagated to all nodes in their vicinity, each node has to determine how to assimilate them with the current state of knowledge about the depth in its location. The first stage is to run the soundings through a median ordered queue that implements a permutation of the normal input sequence to ensure that anomalously deep or shallow soundings are delayed before they go to the estimator proper (Figure 4). Since the original sequencing of the soundings is arbitrary, this reordering does not change any significant aspect of the remainder of the estimation, but it does significantly improve robustness by protecting the estimator until it 'learns' about the true depth. The final stage is the estimator proper. CUBE utilises an optimal Bayesian estimator described by a Dynamic Linear Model (DLM, [West & Harrison, 1997]). This estimator is causal and recursive, so that it can start making predictions as soon as the


Figure 4: Permutation of input soundings. Since the ordering of data is not important in CUBE, we can re-sequence the inputs before they reach the Bayesian estimator in order to delay what appear to be outlier points. This is implemented using a moving median window, which delays any soundings that are shoaler or deeper than the rest of the data in the window

Figure 5: Update procedure at a node with a single depth hypothesis. The current estimate is updated with the information implicit in the new sounding. Since the new sounding is believed to be less accurate than the current estimate (i.e., has higher variance), the updated estimate is mostly determined by the current estimate

data starts to arrive, and only requires the current data estimate to assimilate the next data point. This is the basis of the real time implementation of CUBE, and ensures that we do not need to 'backtrack' into the data as each new point arrives. Each sounding that makes it into the estimator is weighed according to its propagated, combined uncertainty against the current state of knowledge of the depth at the node, represented by a depth estimate and measure of a posteriori variance. The weighting factor used balances the variances of the measurement and current estimate so that if the estimate is much more accurate than the measurement, it is only incrementally affected; if the measurement is very accurate, it will have a very significant effect (Figure 5). After the current state is updated, the sounding is no longer required (all of the information implicit in it has been used) and hence it may be discarded; in implementation, it is retained in a backing database for further analysis.

Model Monitoring and Intervention

In CUBE, we have explicitly set up the model to expect a constant depth. In practice, we observe that many soundings are not consistent with this hypothesis: outlier points violate this assumption by implying multiple alternative depths in the same location. Untreated, these points would corrupt the true depth estimate, provoking modelling failure. We use the error estimates of the soundings to provide a calibration point for model monitoring; that is,

under the null hypothesis that the data is consistent with the model, the sounding and the current estimate should agree to within the sounding's predicted error. If they do not (to a statistically significant degree), then we may conclude that there is sufficient evidence to mistrust the sounding (Figure 6). To make this system more useful, we must also observe long-term drifts (i.e., where the data and model drift apart slowly), and sequential failure, where the model is judged as being marginally inadequate for a significant number of samples. All of these may be implemented using the sequential Bayes factor monitoring of West & Harrison [1997]. After failure is indicated, our intervention scheme is to assume that the inconsistent sounding is another potential depth estimate, and to initialise another DLM to represent it. All models are maintained simultaneously and are treated equally until we are required to make a choice as to which one we believe to be the true depth. Maintaining a monotonically increasing list of models gives us some theoretical difficulty, since we have to determine against which model to compare the incoming sounding. We resolve this by choosing the model that is closest to the sounding in a least weighted error sense, with weighting function determined by the predicted error that would result were the sounding to be assimilated. Hence, if the model monitor indicates an outlier, we may safely build a new model track, since the sounding was compared to the best available model and found wanting (Figure 7).

Figure 6: Model monitoring scheme. CUBE predicts that the next depth will be the same as the current estimate, and then uses this as the null hypothesis to test the incoming data (against a simple alternative of a step change in depth). If the null hypothesis cannot be rejected, the Bayesian data assimilation takes place. Otherwise a new depth tack is started

Hypothesis Resolution

Allowing multiple hypotheses provides robustness, but also ambiguity about which depth should be reported. CUBE implements a configurable disambiguation engine to choose a 'best' hypothesis on demand, using predefined metrics on what constitutes 'best' reconstruction.

The simplest method chooses the hypothesis that has assimilated the most data points (i.e., which is best supported by the data). This works in most cases, although since it involves no context other than the data points, it can fail under significant noise content (e.g., if there are a burst of errors). Our second method finds neighbouring nodes where there is only one hypothesis, and uses this

certain reconstruction as a guide as to the probable true depth. Then, the hypothesis closest in depth to the guide node depth is used for reconstruction. The final method constructs context using another, potentially lower resolution, surface, constructed either from a previous survey or from the current one. Since this surface is only used as a guide to what the depth is, it does not have to be hydrographically correct and we can take more liberties with its compilation. For example, we can use a simple median bin at low resolution, or interpolate between smooth-sheet soundings from the previous survey, or even from the chart if no other information is available. As long as the surface is in approximately the correct loca-


Figure 7: Model selection for monitoring and test assimilation. Use of a minimum predicted error distance ensures that the 'best' model is chosen, and hence that if the data is found to be inconsistent (see Figure 6), then we can start another depth track since no other model would choose to assimilate the data either
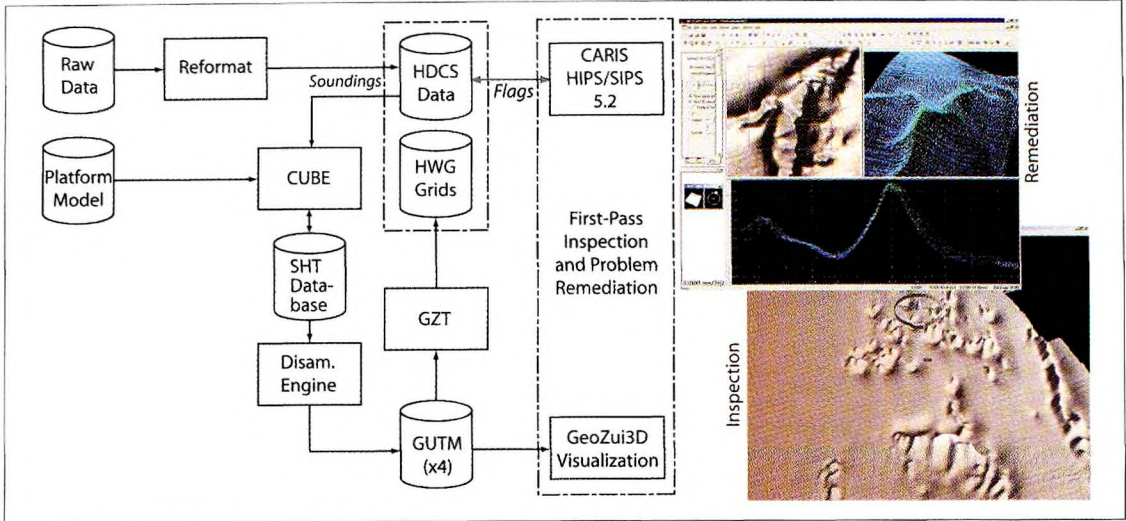
Figure 8: First-pass flow diagram for CUBE processing. We interface to HDCS data so that all normal CARIS/HIPS tools are still available, although for flexibility, we use a separate visualisation suite to display the data, and do the remediation in spatial mode of HIPS 5.2

tion, it should help CUBE, on the average, work out which hypothesis is the correct one. Many other potential solutions exist, and are currently being researched.

Output Products

In addition to the depth, CUBE is capable of providing additional metrics, in particular the uncertainty associated with the depth estimate, the number of hypotheses available at the node, and a measure of how certain the algorithm is about the choice of hypothesis that was made. Each of these is a scalar quantity, and hence may be represented as a surface, or more usefully as auxiliary information on top of another surface (Figures 14-16).

Combinations of these with the depth surface allow the user to see problems in context, and hence make decisions more reliably.

The outputs of CUBE's processing are therefore a set of data vectors per node. It is natural to represent these as separate surfaces, but it is important to note that CUBE's estimates are strictly only estimates at a point, and any interpolation between those points must be considered separately.

Remediation and Iteration

It is unrealistic to expect that any algorithm will make the correct decision under all conditions. Therefore, it is imperative that there is an operator


Figure 9: Second-pass flow diagram for CUBE processing. This is essentially the same as the first pass, except that we move directly to products from the MapSheet (SHT) database through automatic methods, rather than through some intermediate cartographic extraction. A more detailed description of this process is outlined in Smith et al. [2002]

Figure 10: Woods Hole, MA (H11077), conducted by the NOAA Ship WHITING, 2001 [Barnum et al., 2001]. Both chart and data are reprojected to UTM Zone 19N, WGS-84 ellipsoid. Depth range is (2,30) m, and coverage is approximately 1.7km²

to check the decisions which have been made, and to rectify the problems evident in any area where CUBE either made no decision, indicates that the decision was in doubt, or made what the expert hydrographer believes to be the wrong decision, irrespective of the statistical distribution of data and noise. Our initial implementation uses the traditional data-flagging paradigm to assist CUBE in making decisions where the density of noise is such that the correct depth estimate is not evident to the algorithm. It is also potentially possible to work at the level of CUBE's hypotheses, or in a layered approach (e.g., edit hypotheses, and then data only if the problem is not resolved).

After remediations are made, an iteration of CUBE is required to integrate the modifications with the rest of the data. CUBE is, in this sense, a one-way trapdoor: once the soundings have been assimilated into the estimates, there is no way to back them out except to start again. However, the speed of the algorithm is such that this is not a significant concern. In practice, since the processing is mainly local, we need not re-run the algorithm everywhere – just in regions where modifications have been made. This significantly reduces the computational burden, particularly when the number of modifications is expected to be small.

Implementation

We have avoided, whenever possible, redeveloping tools that are available in COTS software, preferring to interface to available applications for data reformatting, display and manipulation. The essential support requirements for a host system are

that it should have an API for data retrieval, preferably a spatially based one (i.e., that can provide all data within a given radius of a particular point). It should also contain a manipulation system for data so that remediations can be done, and a suitable display system that is capable of displaying multiple surfaces simultaneously. No one system currently available has all of these, so we have built a hybrid system using CARIS/HIPS for data conversion, manipulation and display, GeoZui3D for fast turn-around display of multiple surfaces with overlaid colour-coded data, and Fledermaus/PFM for advanced visualisation, spatially-indexed data retrieval and area-based editing. A combination of bash shell scripts, perl and the GMT package are also used in development and implementation of the various stages of the algorithm and product preparation.

The CUBE process occurs in two passes when used in post-processing mode (i.e., when all data is available). The first pass (Figure 8) generates preliminary surfaces for the user to examine; the second pass (Figure 9) takes any user modifications and generates final product surfaces. We read directly from HDCS data using the HIPS/IO interface libraries, and store CUBE's results in a specialist data structure called a MapSheet (SHT). This intermediate store provides extra flexibility, and allows us to maintain state between data availability.

From the MapSheet, we can generate both HIPS Weighted Grids (HWGs) and GeoZui3D GUTMs. The HWGs are inserted back into a HIPS Fieldsheet, so that they can be seen in conjunction with the raw data; we typically attempt to display the HWGs and data on one screen of a dual-monitor system, and the GUTMs on the other. We have found that it is significantly easier to manipulate data if both representations of the data are available simultaneously, since it is difficult to 'fuse' the information mentally in many cases, and cumbersome to transfer by hand the information from the 3D visualisation, where problems are obvious, to the manipulation system where they can be rectified. It is our experience that getting the implementation of this coupling correct can significantly affect the ease-of-use of a system and hence the potential benefits that can be achieved.

In real-time mode, it is not sufficient to have this 'once-through' model of processing, since we want to be able to work data incrementally as it is being gathered, typically on a daily cycle. We currently resolve this by maintaining two MapSheet struc-


Figure 11: Reconstructed bathymetry in southwest corner of Woods Hole data, looking west. The main sand ripples have amplitude approx. 0.5m, and wavelength approx. 10m, although they are overlaid with sand ripples of smaller wavelength and amplitude. The rougher texture to the right is thought to be a dumping area overlaid on the sand ripples

Figure 12: Man-made objects. Thought to be the remains of a floating dock and a floatplane, these objects occur in the northwest corner of the survey, just west of the WHOI dock. The many small features on the area around the dock are probably mooring blocks or rocks


Trackline oriented holidays

Figure 13: Example of track-line oriented artifacts that are only obvious at high resolution, but which are symptomatic of a problem with the data acquisition system. Feedback like this from CUBE's outputs as the survey progresses could help with the early detection and remediation of such problems in the field, where the cost of correction is significantly less

Figure 14: Number of hypotheses at each estimation node colour-coded over the reconstructed bathymetry; hot colours indicate that several self-consistent hypotheses were formed. From the pattern of hypothesis clusters, it is immediately obvious that these were caused by pilings for the associated dock structure. This is not obvious from the bathymetry alone

tures, one for 'daily' use and one for 'cumulative' use. At the start of each day, the cumulative Map-Sheets are used to initialise the daily set, and the current day's data is then assimilated. Once any changes to the data have been made based on the intermediate results, the second pass of CUBE is used to assimilate the day's data into the cumulative MapSheets. In this way, the cumulative Map-Sheets should always represent the 'best available' information on the survey. Working in this incremental mode saves considerable time in processing, although the cumulative MapSheets can always be re-constructed at any time simply by re-running the data from the start.

Example: Woods Hole

During the 2001 field season, the NOAA Ship WHIT-ING conducted hydrographic survey operations around Cape Cod, including Woods Hole, MA (41°31'N 70°40'W, registry number H11077), from Great Harbor to Vineyard Sound, Figure 10. Over approximately five survey days, the WHITING's multi-beam survey launch covered approximately 1.7km² in depths from 2m to 30m with full coverage from a Reson 8101 MBES. In total, 37.9x10⁶ soundings were gathered and archived. A POS/MV 320 was used for attitude measurement and positioning was derived from a Trimble DSM212 differential GPS receiver (corrections: Chatham, MA). All of the data was archived in XTF format and then converted into CARIS/HIPS for processing. Corrections for static and dynamic offsets, refraction and tides were made, and the resultant HDCS data was provided as the starting point for CUBE's processing. The data archive contained edit flags, but these were removed from the test set before starting automatic processing. We used a depth gate of (2,30)m to avoid gross outliers, although we allowed all beams to be used rather than applying the standard angle gate of ±60° per the Data Acquisition and Processing Report (DAPR) for the survey [Barnum et al., 2001]. This provided more coverage in very shallow areas hence allowing for a more stable reconstruction, although we did encounter more multiple hypothesis areas because of this decision, and hence have taken more time to work the data than we otherwise might.

We bootstrapped analysis of the data by constructing a 5m median bin using all of the data. This is inadequate for hydrography, but provides a suitable

reference for slope corrections and dynamic depth ranges. We utilised a blunder filter to remove any soundings more than 25 per cent deeper than the median estimated depth (with a minimum depth difference of 1m), and then processed all of the data at 0.5m resolution in order to ensure that small shoals were reliably estimated, and to provide the highest possible resolution surface for the area. The resultant surface was inspected and remediations made by flagging the original soundings. The CUBE algorithm was then iterated to complete the processing. The non-interactive processing took approximately 60 min. per pass on commodity PC hardware; the interactive time was approximately 240 min., although much of that time was spent investigating the many small lumps in the harbour area rather than actually editing data. It is important to note that the robustness of CUBE's estimation algorithms allows us to be a little more cavalier about editing, in the sense that we do not have to remove every single anomalous sounding in the set, simply enough to give CUBE a head start in estimating the surface, i.e., to improve the signal-to-noise ratio.

Therefore, we may remove just the obvious outliers, and allow the algorithm to process those close to the 'true' surface appropriately. This was used to preserve the objectivity of CUBE's estimates.

We found that the majority of the data was processed automatically, and the level of detail in the results is high (Figures 11-12). Preservation of detail is an important concern in automatic processing schemes since an over-zealous procedure could also remove important small features. The dynamic depth gate implemented by the blunder filter bootstrapped by the median depth significantly improves performance in sparse areas for little extra cost, although this affects only deep spikes.

We observed a number of small trackline oriented holidays in the data (Figure 13), which were subsequently tracked to dropped packets in the input data stream (i.e., missing data not recorded by the capture system). Although these holidays are not significant with respect to hydrographic coverage of the area, they illustrate a problem with current data processing methods. There is no way to detect these dropped packets without investigating


Figure 15: Uncertainty colour-coded over bathymetry; view from Great Ledge looking north to Woods Hole passage. The colour-coding is 95 per cent confidence interval predicted from the posterior variance of the depth estimate chosen by the disambiguation engine, with warmer colours indicating higher uncertainty. Prediction variance is a function of the number of soundings assimilated and their component uncertainties. The primary signals evident here are depth range and beam angle, shown in the linear features derived from the line-plan used during the survey

Figure 16: Hypothesis strength colour-coded over reconstructed bathymetry; WHOI dock looking north. Hypothesis strength is a metric indicating how certain the disambiguation engine is about the hypothesis it reported. Green indicates strong evidence for the chosen hypothesis; the scale to red indicates decreasing evidence, implying that there are other plausible solutions

the timestamp on each packet of input data, which is obviously unfeasible, and they are not immediately obvious in points-mode data displays. To demonstrate coverage, only grids at approximately 5m resolution are required, and under any conventional grid construction scheme, these sorts of holidays would not be observed. Here, CUBE has been able to illustrate a potential problem, and provides a way to visualise them so that reasoned quantitative decisions can be made (in this case, to ignore the holidays as hydrographically insignificant).

A use for the number of self-consistent hypotheses is illustrated through the data around the Woods Hole Oceanographic Institution (WHOI) dock. The dock pilings are sufficiently large to return multiple beam hits, and hence CUBE resolves multiple hypotheses, as seen in Figure 14. The obvious geometric arrangement of the multiple hypotheses clearly indicates that these are man-made, although this is not immediately obvious just from the surface, since it is constrained to choose just one hypothesis as 'best'. An objective measure of consistency such as this is a very powerful tool in making decisions about what to keep, and what to ignore. CUBE also provides uncertainty estimates (Figure 15) that provide information about the quality of the chosen hypothesis, and a measure of

'hypothesis strength' (Figure 16) that attempts to measure how sure the algorithm is about the choice of hypothesis that it made. Use of these indicators can further inform processing to best utilise operator time.

To compare the CUBE output with a traditional hydrographic processing chain, we took the preliminary smooth-sheet selected soundings for the survey, and matched them against the CUBE surface, assuming that they are IHO Order 1 accurate (the target for the survey) [IHO, 1998]. For each selected sounding, we found the reconstructed CUBE depth within the horizontal 95 per cent CI for the sounding that minimised the absolute vertical difference between sounding and surface. We then scaled this difference by the vertical 95 per cent CI for the sounding and computed the cumulative probability mass function over the 5902 selected soundings (Figure 17). We observe that just over 95 per cent of the soundings are below the one unit CI limit (135 soundings of 5902, or 2.3 per cent are above) as expected, showing that the CUBE surface agrees very well with the traditional selected sounding approach in this case. The slight bias is probably due to a combination of finite sample effects and the traditional approach of shoal biased selection of soundings.

Figure 17: Cumulative probability mass function for comparison between preliminary smooth-sheet selected soundings and CUBE output surfaces. The horizontal scale is minimum vertical difference between the CUBE surface and the selected sounding assuming that the soundings are IHO Order 1 accurate (the target for the survey). The axis is scaled to the vertical 95 per cent CI for IHO Order 1 survey, so we expect (and observe) 95 per cent of the selected soundings with vertical error less than 1.0

Conclusions

Our current methods of processing Multibeam Echosounder (MBES) data are becoming inadequate as faster and higher resolution systems come online. We have argued that statistical methods of processing data are not only useful, but are in fact required when we consider the properties of MBES data. We have outlined an alternative method for processing such data, which attempts to handle the majority of soundings automatically by focusing on estimation of 'true' depth, rather than selecting 'best' soundings, while building in quantitative estimates of data quality and guideline metrics for QA/QC. We accept that no method will be completely automatic. We have therefore also outlined an inspection and feedback mechanism that attempts to harness the power of automatic methods to bootstrap operator effort. The algorithm can be run in post-processing or real-time mode, and can provide interim results as data is being gathered.

Through the data example shown here, we have illustrated the CUBE algorithm. We observe that the algorithm is suitably robust for typical hydrographic systems, and that it handles the majority of data automatically; the algorithm is also sufficiently fast to keep up with data capture rates, even in an experimental research implementation. We have found that the algorithm is not sensitive in its parameters (given calibration of the error model

through installation and patch test measurements), so that it does not need to be returned for each dataset.

We have shown elsewhere [Calder & Mayer, 2001, 2002] that CUBE's estimates are statistically equivalent to more conventional surface estimation techniques, and here that they agree well with a traditionally constructed selected sounding set. We are currently pursuing a project to show hydrographic equivalence (in the sense that the same hydrographic conclusions would be reached using CUBE's results as for a traditional processing scheme). CUBE currently exists as a source code library developed at CCOM-JHC and is available on a non-exclusive royalty-free basis to members of the CCOM-JHC Industrial Consortium. Our intention is for CUBE to be integrated with other processing schemes in order to provide the support services and user interfaces that we have approximated in these experiments, rather than being used on its own.

References

Calder, B. R., and L. A. Mayer (2001), *Robust Automatic Multibeam Bathymetric Processing*, Proc. U.S. Hydro. Conf. 2001, Norfolk, VA, 2001 (reprints: www.thosa.org/us01papers.htm)

Calder, B. R., and L. A. Mayer (2002), *Automatic*

Processing of High-Rate, High-Density Multibeam Echosounder Data, submitted to Geochem., Geophys., Geosyst. (G3, g-cubed.org), DID 2002GC 00486, December 2002

Debese, N. (2001), Use of a Robust Estimator for Automatic Detection of Isolated Errors Appearing in Bathymetry Data, Int. Hydro. Review, 2(2), 32-44

Debese, N. and P. Michaux (2002), Détection Automatique d'Erreurs Ponctuelles Présentes dan les Données Bathymétriques Multifaisceaux Petits Fonds, Proc. Canadian Hydro. Conf. 2002, Toronto, 2002

Du, Z., D. E. Wells, and L. A. Mayer (1996), An Approach to Automatic Detection of Outliers in Multibeam Echosounding Data, The Hydro. Journal, 79, 19-25

Eeg, J. (1995), On the Identification of Spikes in Soundings, Int. Hydro. Review, 72(1), 33-41.

Barnum, S., M. Cisternelli, and R. Brennan (2001), NOAA Ship WHITING Data Acquisition and Processing Report S-B904-WH (Woods Hole, MA; registry number H11077), National Ocean Service, NOAA

Gourley, M., and K. DesRoches (2001) Clever Uses of Tiling in CARIS/HIPS, Proc. 2nd Int. Conf. on High Resolution Survey in Shallow Water, Portsmouth, NH, September 2001

Hare, R., A. Godin and L. A. Mayer (1995), Accuracy Estimation of Canadian Swath (Multibeam) and Sweep (Multitransducer) Sounding Systems, Tech Rep., Canadian Hydrographic Service.

IHO Committee (1998), IHO Standard for Hydrographic Surveys, Int. Hydro. Organization, Special Publication S.44, 4ed

Smith, S., L. Alexander, and A. Armstrong (2002), The Navigation Surface: A New Database Approach to Creating Multiple Products from High-Density Surveys, Int. Hydro. Review, 3(2)

Ware, C., L. Slipp, K. W. Wong, B. Nickerson, D. E. Wells, Y. C. Lee, D. Dodd, and G. Costello (1992), A System for Cleaning High Volume Bathymetry, Int. Hydro. Review, 69(2), 77-94. West, M., and J. Harrison (1997), , 2 ed., Springer-Verlag

Acknowledgements

Biography

Dr. Brian Calder is a Research Assistant Professor at the Center for Coastal and Ocean Mapping & Joint Hydrographic Center at the University of New Hampshire. He graduated MEng and PhD in Electrical & Electronic Engineering from Heriot-Watt University in Edinburgh, Scotland in 1994 and 1997 respectively, but was subsequently seduced into sonar signal processing for reasons that are now obscure. His research interests have previously covered speech recognition, texture analysis, object detection in sidescan sonar, high-resolution sub-bottom profiling, simulation of forward-looking passive infrared images, acoustic modelling and pebble counting. Currently, they revolve around the application of statistical models to the problem of hydrographic data processing.

E-mail: brc@ccom.unh.edu