Proceedings of the 2001 U.S. Hydrographic Conference, Norfolk, Va.

Robust Automatic Multi-beam Bathymetric Processing

B R Calder and L A Mayer

Center for Coastal and Ocean Mapping
& Joint Hydrography Center
University of New Hampshire
Durham, NH 03824

Abstract

We propose a new methodology for automatically processing high-resolution, high-rate bathymetric data, typically from multi-beam echo sounders. Based on a generic grid of estimation nodes, this method incorporates data from multiple surveys, multiple instruments or multiple passes of the same instrument, and models and maintains depth uncertainty estimates with the data. Consequently, we can extract, at any time, a current estimate of depth, plus a co-located estimate of the uncertainty associated with the depth. The data integration scheme is sufficiently robust to deal with typical survey echo-sounder errors, and is designed so that no data is ever removed; at worst, suspect data is held over until such time that previous estimates have given us sufficient confidence to treat them properly. The algorithm has a low memory footprint, runs faster than data can currently be gathered, and is suitable for real-time use.

We illustrate the algorithm on two survey data sets gathered for entirely different purposes, with different instruments, and in different survey conditions and depth ranges. We have found that with no human intervention, this automatic processing scheme provides depth estimates that are statistically identical to those generated through hand estimates or intuitive binning algorithms such as median and trimmed mean binning.

1 Introduction

Advances in technology and more exacting requirements for hydrographic surveys have led to ever increasing data rates and densities for multi-beam datasets. The time and effort involved in manually processing these datasets is also increasing in proportion; clearly, some automatic approach would be beneficial. However, this leads to a number of problems of reliability, robustness and (domain-specific) safety and liability. In this paper, we propose a new scheme for automatically processing multi-beam bathymetry that attempts to answer these concerns, and specifically the fundamental question of uncertainty in prediction of the depth assigned to a particular point.

There have been a number of approaches to the task of automatically processing high rate bathymetric data. The simplest examples include simple depth and angle gates (i.e., to reject a sounding shallower or deeper than reasonable limits based on a general knowledge of the target area, or from outer beams, where refraction effects are more significant). Slightly more complex examples include filtering based on angles between

Calder B. and Mayer, L.A., 2001, Robust Automatic Multibeam Bathymetric Processing, Proceedings of the 2001 U.S. Hydrographic Conference, Norfolk, VA.

have been proposed. Ware *et al.* [4] suggest a technique to compute mean and standard deviation estimates for entire datasets, and therefore to eliminate some samples based on their deviation from the weighted mean estimated surface. Varma *et al.* [5] use binned data and an efficient database engine to implement much the same scheme. An alternative technique is to consider hypothesis tests based on pseudo-variance estimates, which (combined with a leave-one-out testing scheme) has been used by Eeg [6] to detect spikes in dense MBES data. All of these schemes rely on an estimate of some point statistics of the data in a small area (either geographically, or swath by swath); it is also possible to estimate sounding density and hence attempt to determine modes corresponding to outliers from a histogram. Du *et al.* [7] use this technique to construct an automatic processing scheme that is intended to simulate how a human operator edits data, and is shown to be functionally equivalent to human editing on a small dataset.

Most of the automatic editing schemes that have been proposed operate in either swath mode (i.e., causally as the data is collected), or in spatial mode (i.e., geographically after the data has been geo-referenced). One exception to this is the multi-pass filtering of Lirakis & Bongiovanni [8], which is incorporated into the PFM system [9]. This starts with data in swath mode, processes analogously to many of the filters considered above, and then converts the data into geographic mode for further filtering. This system also adds the concept of a modifiable classification attribute for each sounding, so that a depth can be marked 'Unknown', 'Good' or 'Bad', and many of the tests implemented revolve around transitions between these states.

In our approach to the automatic processing problem, we have implemented a system that constructs a grid of estimation nodes over the area of interest, and attaches to each node at least an estimate of depth and uncertainty at that point (other components could include estimates of sounding density, backscatter strength, etc.) Sounding depths are augmented by an estimate of horizontal and vertical accuracy using the MBES error model of Hare *et al.* [10], and are then offered for incorporation into grid nodes around the nominal location of the sounding. Nodes close to this location accept the depth estimate at almost the computed accuracy; those further away accept the depth estimate but increase the uncertainty according to distance and a user specified dilution function that can be used to balance accuracy with safety. From the nodes' point of view, estimates pushed into their input queue originate from an area around the point of interest, in effect providing an area-based processing scheme.

At a node, soundings are integrated to update the current depth and uncertainty estimates using an optimal Kalman filter. We thereby construct a recursive, continually on-line, estimate of depth and uncertainty at each point of interest, avoiding significant system lag, memory overhead or post-processing steps. To improve robustness, we have also included a re-ordering stage in the processing, prior to incorporation by the Kalman filter. This allows us to delay those observations about which we are uncertain until a later stage of processing, by which point information from the less controversial data allows us to reliably process the remaining outliers. Almost as a by-product, this methodology provides us with an automatic method for robustly cleaning high-rate bathymetric data.

In the remainder of this paper, we outline the theoretical foundation of the methodology and illustrate the operation on two data sets gathered with different instruments, in different depth regimes, and with different processing requirements. One

of the datasets is a massively over-sampled data set [11], which allows us to undertake some statistical comparisons of the surfaces estimated by alternative methods, and we use this to show that the proposed method is consistent with simpler and more intuitive (but also more memory- and processor-hungry) algorithms. We then conclude with some perspectives on future developments of the method.

2 Theory

We begin by revisiting the fundamental question of interest: at any particular point on the earth's surface, what is the best estimate of depth? Almost more importantly, we must also answer the question: how certain are we about this prediction? As with any other data reduction and estimation problem, it is not sufficient simply to report a single point estimate, since this does not provide sufficient information to judge the quality of the value presented. At the very least, we must also provide an estimate of central tendency.

To facilitate the estimation, we envision a system of nodes spread about the true geographical area in which we wish to work (figure 1). We set the estimation in the Bayesian framework, so that we assume that the location of each node is known absolutely, but the position of soundings and their observed depths have uncertainties associated with them. Consequently, we are faced with four problems:

1. How do we take advantage of multiple soundings?
2. How do we impute error estimates in three dimensions to go with each sounding?
3. Given that a sounding will be located arbitrarily with respect to the grid of nodes that we establish, how do we propagate the information inherent in the sounding to the nodes around it?
4. Robustness of estimation.

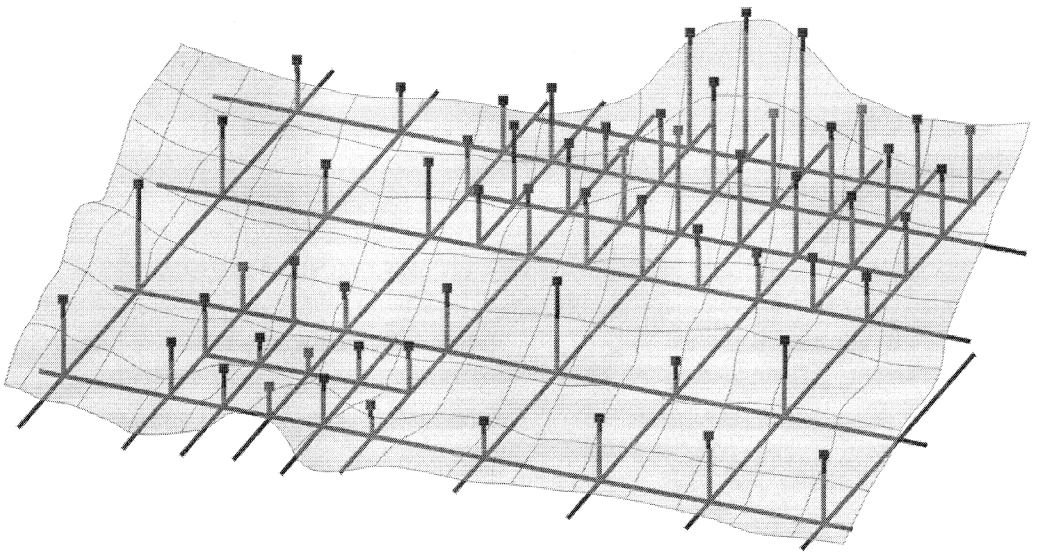

Figure 1: Estimation nodes in the working area. We can in theory arrange the nodes arbitrarily in the space, although in practice a regular grid is used to make the bookkeeping easier. An irregular grid might be used if we know or suspect that a particular area needs a higher sampling rate due to greater slopes or other features.

2.1 Optimal Sounding Integration

An immediate consequence of the transformation from bin based editing and modeling to node based modeling is that we have transformed the estimation problem from one of fitting a nominal plane to one of estimating a constant. That is, at any single location, we expect to see only one depth. Of course, different sonar beams ensonifing overlapping areas of the seafloor will result in different depth estimates, depending on the sub-footprint roughness (or objects), but for the node at an exact point on the surface of interest, there is only one 'true' depth. The problem is that we can only observe this 'true' depth through noisy measurements; our task is to compute a reasonable estimate of the 'true' value from the sequence of soundings.

We have chosen to cast this problem as an optimal Bayesian estimation problem. That is, at each node we maintain an estimate of depth and the variance associated with that estimate. As new depth estimates arrive (i.e., as a sounding device goes past the node), we use the current estimate as prior information, and update our current state of knowledge about depth and the prediction uncertainty with the new information inherent in the soundings. The posterior prediction of mean depth and variance are then stored to act as prior information on the next update cycle.

This interpretation of the update scheme is formalized in the Kalman filtering process, which provides the optimality condition for our estimation scheme. In the general case, a Kalman filter can be used to track a multi-variate state vector that evolves dynamically in time, including system noise variance, observation noise, and observation transformations [12, 13]. In this case, we have a very much simpler system since we do not expect *a priori* that the depth at a particular node will change with time (although this may be the case in, e.g., estuarine or littoral environments), so that we can use a static system with no time-dependent dilution of precision. Consequently, we can reduce the problem to the one-dimensional case, resulting in the coupled pair of equations:

$$x(n+1) = x(n) + s(n) \quad s(n) \sim N\left(0, \sigma_s^2\right)$$

$$y(n) = x(n) + e(n) \quad e(n) \sim N\left(0, \sigma^2(n)\right)$$

Here, $x(n)$ is the best estimate depth, and $y(n)$ is the observed depth from the echo-sounder. The $e(n)$ term models the observation noise, which in this case we assume to be zero mean Gaussian with variance $\sigma^2(n)$ for observation n in the sequence at that node, and $s(n)$ models the system noise, expressing the belief that our model may not be exactly correct (we assume that the variance of this term is a constant). The Kalman filter balances the error associated with new data and its current confidence in the estimate held internally. Thus, if the input is of much lower variance than the current estimate, it is given significant weight; conversely, if the input is of much higher variance that the current estimate, it is mostly ignored. This mechanism is responsible for the majority of the properties of the algorithm, and provides a great deal of the robustness (at least in the long term): as the algorithm integrates more and more data, its estimate of the true depth becomes more and more accurate. Consequently, outliers are easily recognized, and are given low weighting as they are incorporated into the surface

estimate. Errors or outliers early in the estimation sequence are still a concern, and approaches to these are covered later.

2.2 Forward Error Prediction

In theory, the process of converting from raw range-angle relationships to sounding depth is a relatively simple process. If we have true geometric angle θ and range r, then the sounding depth given roll R and pitch P is simply:

$$z = r\cos P\cos(\theta + R)$$

Of course, we must correct for refraction, transducer misalignment and reduction to datum, and the equations become significantly more complex in practice. In principle, however, if we have a measure of the errors associated with all of the measurements that we make, we can compute a formal error propagation model and impute the errors associated with the final depths computed. This is the basis of the forward error propagation model of Hare *et al.* [10], and the method implemented in our algorithm.

The model is defined in two parts: those errors that are associated with the transducer itself, and those induced from the auxiliary equipment such as the GPS and Vertical Reference Unit (VRU). For transducer errors, we rely mainly on equipment manufacturers to provide suitable data. In this work, we are dealing with Simrad EM series MBES, for which a relatively well-defined error model is available [14, 15]. In other cases, more empirical models [16] may be used. We also rely on manufacturer information for auxiliary sensor data, and the full list of parameters includes terms for errors in offset correction, instrument errors (attitude, sound speed, etc.) and reduction to datum. A typical vertical error curve (based on an EM3000) is given in figure 2. This model requires a very detailed knowledge of the properties of the measurement devices in use, but the information required is usually available from manufacturer's manuals.

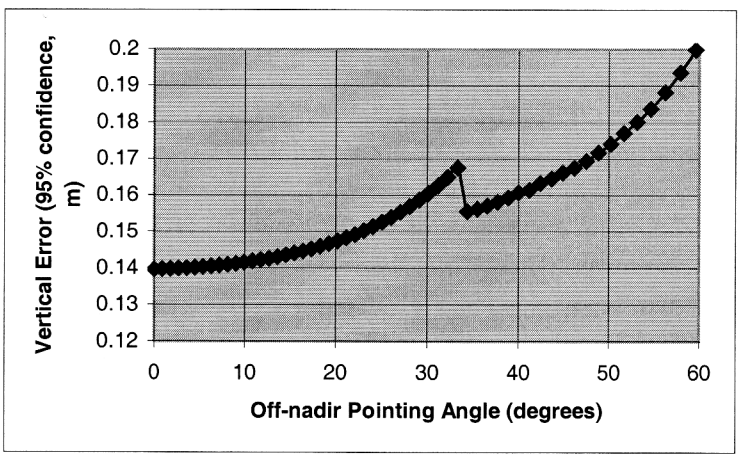

Figure 2: Typical vertical error performance for a Simrad EM3000 (in shallow water mode) and typical auxiliary sensors. Environmental parameters (one standard deviation) are: depth 25m, peak wave height 0.1m, roll and pitch 2°, and sensor errors corresponding to a POS/MV 320. Error magnitude varies significantly with environment and these values are illustrative only. Note that the error is only piecewise monotonically increasing; the small dip at around 33° off nadir is due to the system switching from amplitude detection to phase detection.

2.3 Information Propagation

Having computed horizontal and vertical errors for each sounding, we are now left with the problem of forging a link between where the sounding *is*, and where we need the information (i.e., at each estimation node).

The simplest case is to assume that the seabed, on the scale of the multibeam survey, is effectively flat (figure 3(a)). That is, we make a zero-order prediction of depth, with exponentially increasing uncertainty to account for the fact that our observation does not coincide with the estimation point.

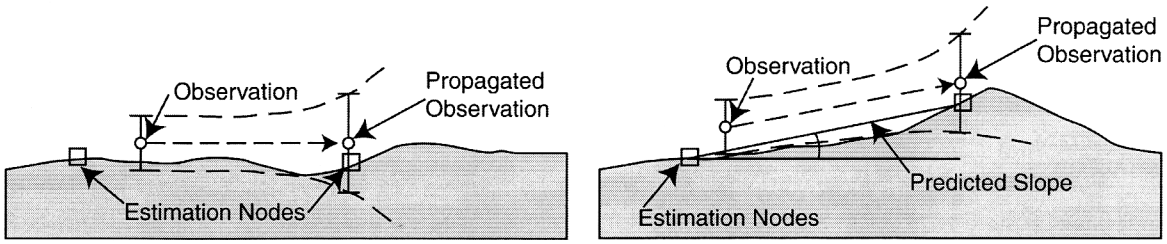

Figure 3: Propagation strategies for referring the sounding data to the estimation nodes. In regions which are essentially flat, we might use a zero-order prediction of depth (left), increasing the uncertainty associated with the data as a function of distance between sounding and estimation node; in regions with significant slope, we might make a first-order prediction of depth (right), with suitably modified confidence bounds.

In some cases, however, we might be working at a larger scale in areas with significant regional slope, for which we need to compensate. In this case, we might make a first-order prediction of slope, and adjust the predicted depth accordingly, always taking care to ensure that the uncertainty propagated covers the predicted depth (figure 3(b)). In particular, this is appropriate where we might be concerned about the effect of shoal soundings in the dataset, or where we have particular hydrographic concerns about safety in particular areas. Indeed, the function used does not have to be stationary, although it would be unusual to have sufficient information to design a suitable non-stationary function in most cases. One possible exception is where a survey is being repeated, and certain problems are known *a priori*. The only difficulty involved in such a scheme is the bookkeeping required.

In the present work, we make a zero-order prediction of depth. The uncertainty exaggeration factor is given by:

$$\hat{\sigma}_P = \hat{\sigma}_V \left(1 + \left[\frac{\left\| \underset{\sim}{x}_i - \underset{\sim}{h}_j \right\| + s_H \hat{\sigma}_H}{\min(\Delta x, \Delta y)} \right]^{\alpha} \right)$$

where $\hat{\sigma}_H$ is the estimated horizontal uncertainty, $\hat{\sigma}_V$ is the estimated vertical uncertainty, $\Delta x, \Delta y$ are the grid spacings in eastings and northings, α is a heuristic, user-defined exponent (typically 2.0), s_H is a factor representing the worst case error that horizontal position error can contribute (typically, s_H = 1.96) and $\underset{\sim}{x}_i$ and $\underset{\sim}{h}_j$ are the location of the sounding and estimation node, respectively. Although this specifies a

continuous increase in uncertainty with distance from the sounding point, we truncate the propagation when the uncertainty reaches an upper limit, σ_{max}. This ensures that we do not continue to propagate information to the extent that our assumptions justifying the zero-order prediction of depth are exceeded, and also that we do not use very poor data in too many nodes. The upper limit is predicted using the appropriate IHO definitions for the survey under consideration [17].

2.4 Robustness

All data sets contain some proportion of outliers, typically poor beam solutions due to inaccurate refraction correction during ray solution or acoustic reflectors other than the seafloor. We make the distinction here between outlier-type effects caused by systematic errors in the data (e.g., roll bias, vertical datum correction, etc.) and outliers due to random errors in the system (we may also classify these as artifacts and outliers, depending on the situation). The former are mainly characterized by long time period, and should really be addressed via normal survey QA/QC procedures; the latter must be dealt with either through human editing, or a sufficiently robust estimation scheme.

In essence, we have two basic problems, which revolve around having too much or too little data. That is, if we have only two estimates (or a small number), and they do not agree, how to we determine which is correct? Alternatively, if we have a large number of estimates, but we observe a burst of outliers or a sudden change in value, which estimate is correct? This can be particularly difficult to do if we are processing in real-time, where we cannot easily build context (i.e., base our decision on a large number of samples to be integrated into the node).

Our base problem is one if initialization. Once a good estimate of depth is formed, we can reject outliers through the Kalman weighting scheme. However, we have to initialize the filter somehow, and for lack of better information, we typically choose $(z_0, \sigma_0) = (0, 10^3)$m. Hence, the filter is initially naïve about outliers, and will lock to the first estimate provided – real, or outlier. To protect the filter, we have implemented a median ordered pre-filter queue [18], which effectively permutes the input sequence so that soundings which are suspiciously high or low are delayed until later in the sequence at the node. To avoid significant lag, we limit the size of the queue to 11 soundings; longer queues would provide better protection.

This leaves the problem of reasoning with very little data, in which case outliers have proportionately greater weight. We have resolved this via the method of Eeg [6], which ranks a set of soundings according to likelihood of being an outlier. To avoid gratuitous data loss, this is only applied if the data sequence at a node never fills the median pre-filter queue. Although we do not do this in our experimental implementation, the soundings culled should, of course, be retained for examination by a human operator.

2.5 Overall System

The overall system diagram for the processing chain is shown in figure 4, with detail of the data augmentation scheme in figure 5. We prefer to operate on datagrams as close as possible to those generated by the MBES, since every translation step incurs the risk of loosing some of the input data required for the error model.

Special care must be taken in extracting the estimate grids due to the latency involved in the pre-filter queue; users may either extract the current information from the nodes, or

they may request that the node flush the pre-filter queue into the Kalman filter, and then return the updated estimate. This opens up a non-reversibility in the processing, since a data set processed in two halves and flushed in between would not necessarily yield the same result as one processed in one run.

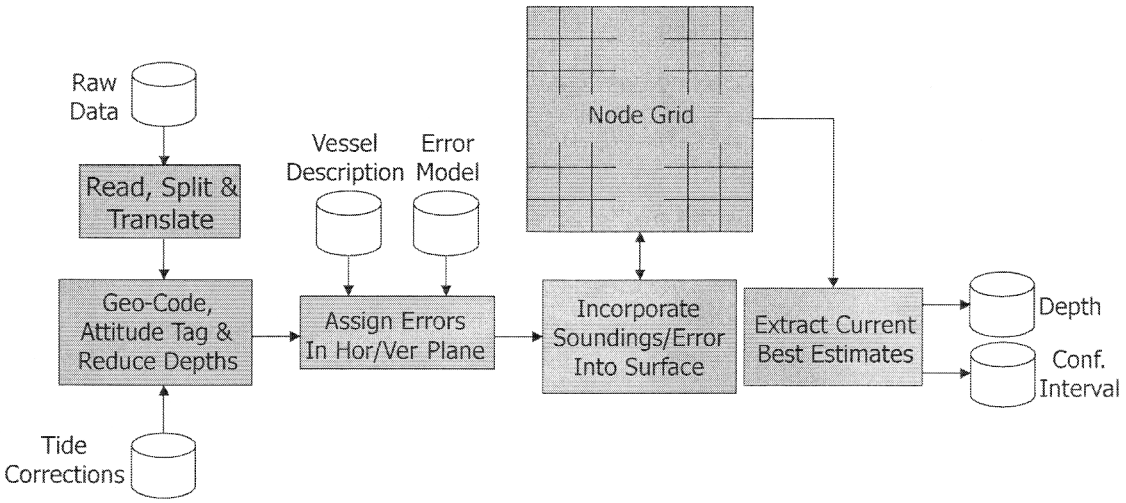

Figure 4: Overall system data path. Conventional corrections are applied to the raw data before errors are assigned in horizontal and vertical planes, and the data are then integrated into the grid of estimation nodes. A separate process extracts current or best estimates of depth and confidence (or uncertainty) as required.

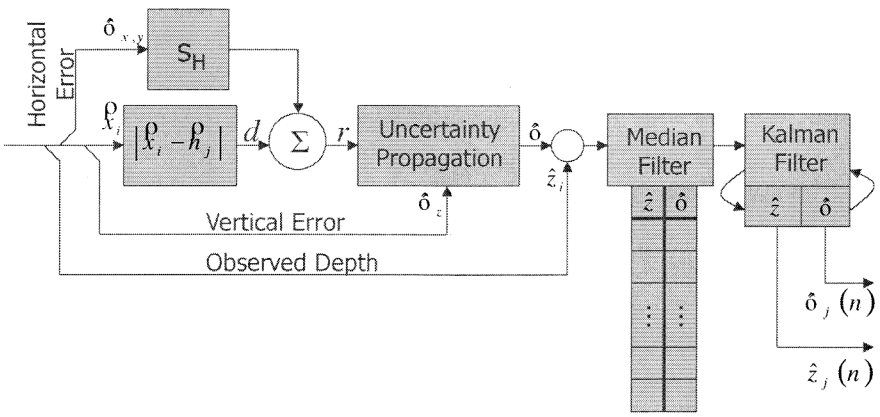

Figure 5: Data Augmentation path (per node). Input depths and errors are used to build a propagated uncertainty depending on worst-case distance from the observation to the estimation node. The data with propagated uncertainties are pre-filtered using a median queue, and are then integrated in the Kalman filter. The outlier rejection stage described in the text is not shown.

3 Experiment

3.1 A 'Reference' Surface: SAX'99

SAX'99 was an ONR sponsored research experiment conducted in the Gulf of Mexico in 1999. The experiment was mainly concerned with acoustic propagation in sediments, and in particular on critical angle and scattering effects [19, 20, 21]. However, as part of

the process, a high-resolution bathymetric survey of the area was carried out using a Simrad EM3000 multi-beam [11]. This data set is unusual in that it covers a very small, relatively flat area of seabed (approximately 1.2×1.2 km at a depth of approximately 20m) many times over, so that we have massively redundant estimates of depths in all areas (figure 6). It thus makes a perfect test bed for new processing algorithms, providing relatively benign data availability but with typical multi-beam processing problems. The massive redundancy also allows us to consider robustness of processing algorithms by selectively reducing the data used and monitoring algorithm performance.

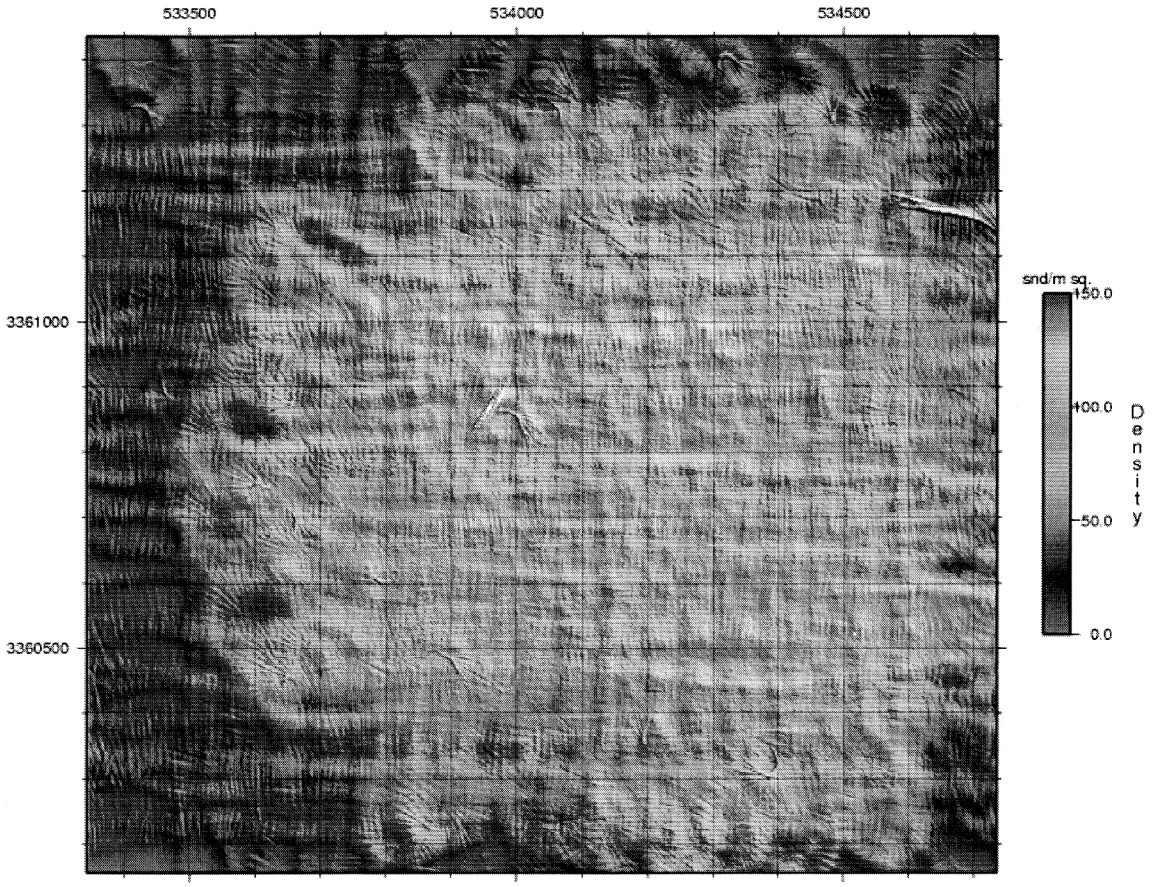

Figure 6: Data density in the SAX'99 high-resolution bathymetric survey center section. Coordinates are UTM (zone 16N) on WGS84; data is colored in soundings/m^2.

The experiment data set consists of a total of seven days survey data (one day was spent doing short area tests, where depths were not correctly resolved; this day was removed from subsequent processing in all of the experiments reported). The data were used raw; that is, no editing has been carried out, no beams have been suppressed. Tidal corrections based on measured tides have been applied.

To provide a reference, we computed binned estimates of the bathymetry using 2m bins (projected in UTM coordinates on WGS84) and both simple median and trimmed mean estimation (outliers rejected at 2σ from the mean on each iteration). Before binning, the data were converted into a simpler format to save space and time in

processing. By suitable reduction of data, the whole set was processed in memory in one pass, and the surface estimation was carried out in parallel using four processors of an SGI Origin 2100. Wall-clock time for the process was approximately 10 min. to process 1.2×10^6 pings (153×10^6 soundings), although almost half of this time was spent in processor serial mode, reading the data from disc. The trimmed mean surface is illustrated in figure 7; the median surface is very similar. We observe that there is no practical hydrographic difference between the surfaces, or between the surfaces and the hand-edited equivalent (which took approximately a week to generate). We pursue these comparisons further below.

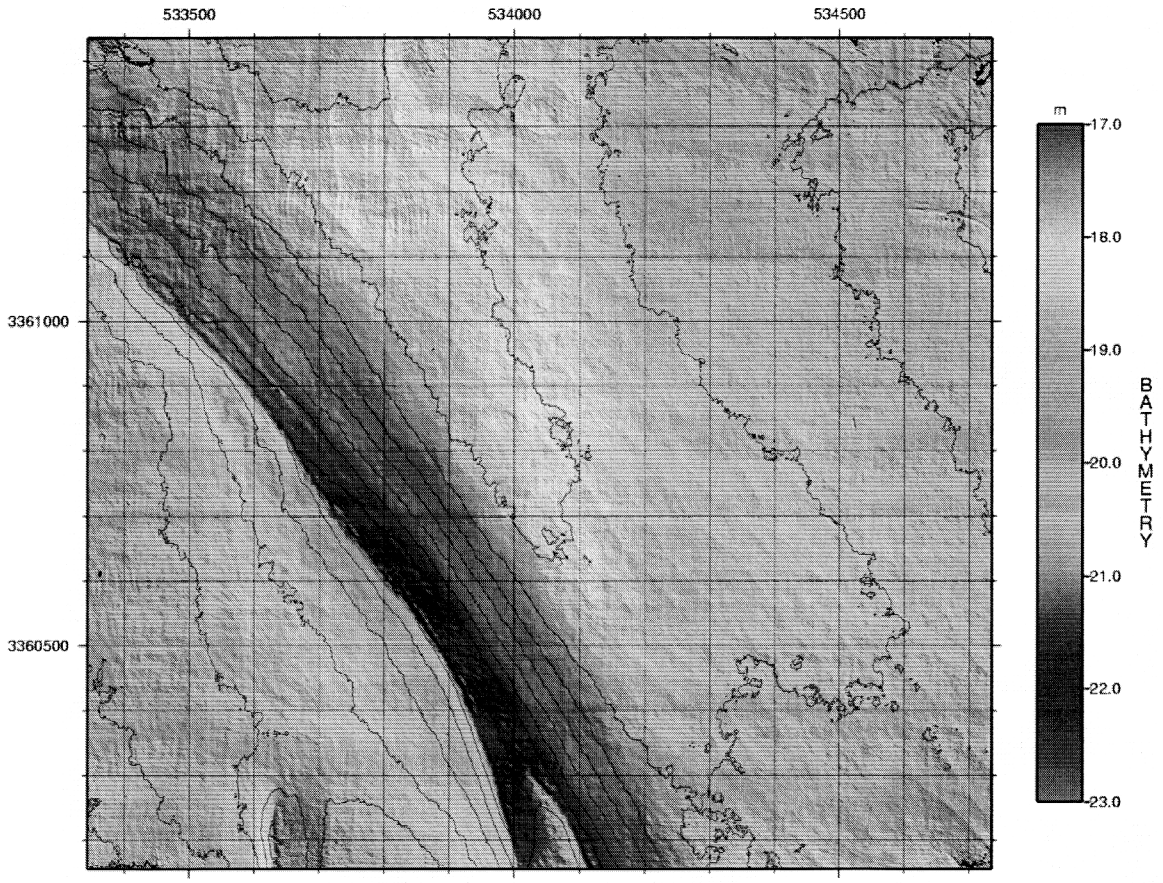

Figure 7: Trimmed mean binned surface estimate for SAX'99. This uses 2m bins in projected coordinates (UTM on WGS84), and represents a total of approximately 153×10^6 soundings.

We next processed the data using the Kalman filtering method. We used raw Simrad telegrams with navigation and attitude data interpolated on-the-fly, the full multi-beam error model, an 11-pt median pre-filter, $\sigma_s = 10^{-3}$ and nodes at 2m spacing. Processing speed is approximately 500 ping/s (approx. 64×10^3 soundings/s for the EM3000); the output depth estimate is illustrated in figure 8, with the associated error uncertainty surface (at 90% confidence) in figure 9. We observe that there are small residual errors in the depths estimated by the recursive estimation, but can attribute these to small vertical data errors not corrected by tides.

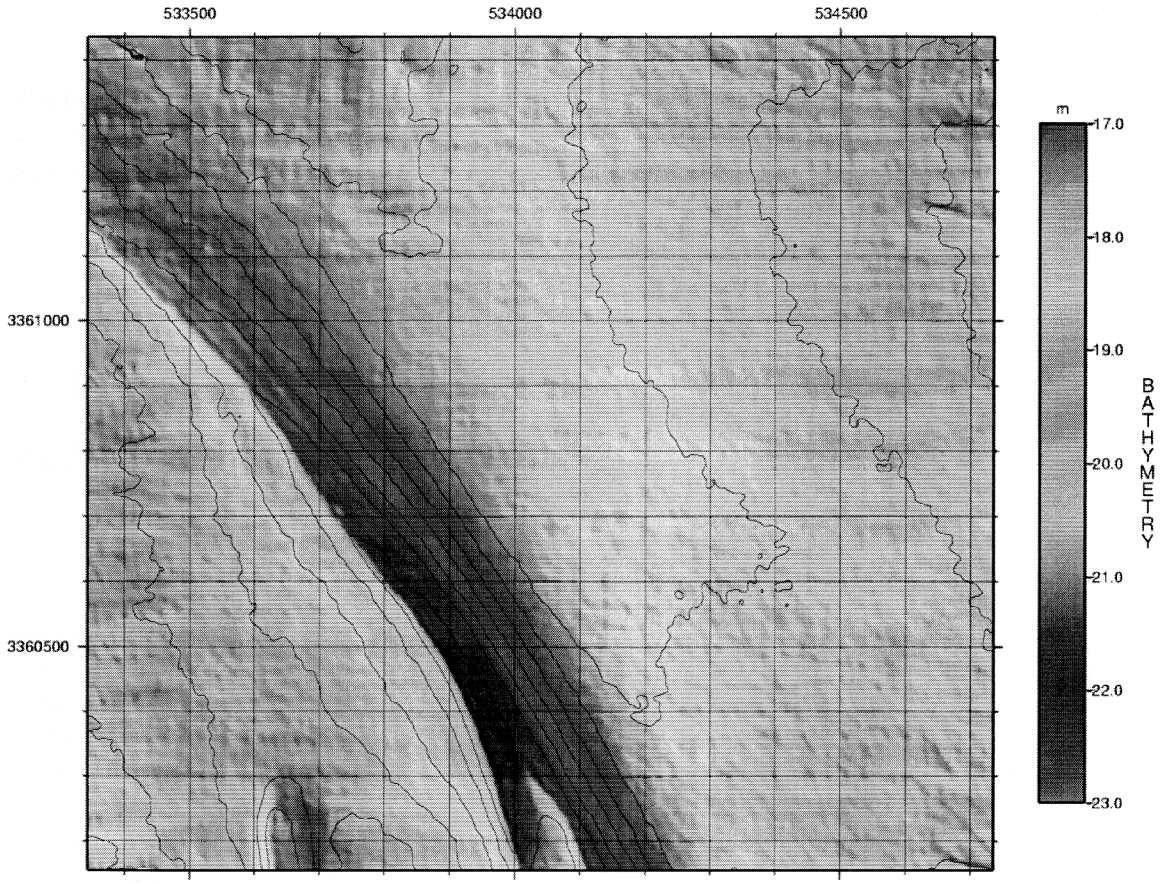

Figure 8: Estimate of depth using the Kalman filtering method. Coordinates are UTM meters on WGS84, with depths in meters. The surface has been shaded from the northeast to highlight the small residual features of the surface.

To compare the three surface estimates, we computed point-wise differences of the three resultant DTMs and then computed point histogram estimates of these differences. Figure 10 shows summary histograms for the surfaces generated using all of the available data. There is little difference between the surfaces (the mean absolute difference is on the order of a few centimeters), and the observed differences are hydrographically insignificant at these depths. Since we are happy to accept median estimates as a reasonably robust estimate of depths in the presence of noise, we therefore conclude that the Kalman filtering method generates hydrographically equivalent estimates of depth as more typical processing strategies.

However, the Kalman filtering method also provides an estimate of uncertainty. In this example, the confidence limit estimates are all relatively small due to the abundance of data. The regions of higher uncertainty at the edges of the central survey region are attributable to a relative lack of data, which can be seen by comparison with figure 6. This is more graphically illustrated in figure 11, where we have constructed a perspective plot of the uncertainty surface (vertical exaggeration 500×) and overlaid the tracklines of the survey platform. From this, we can trivially link higher uncertainty with data sparsity.

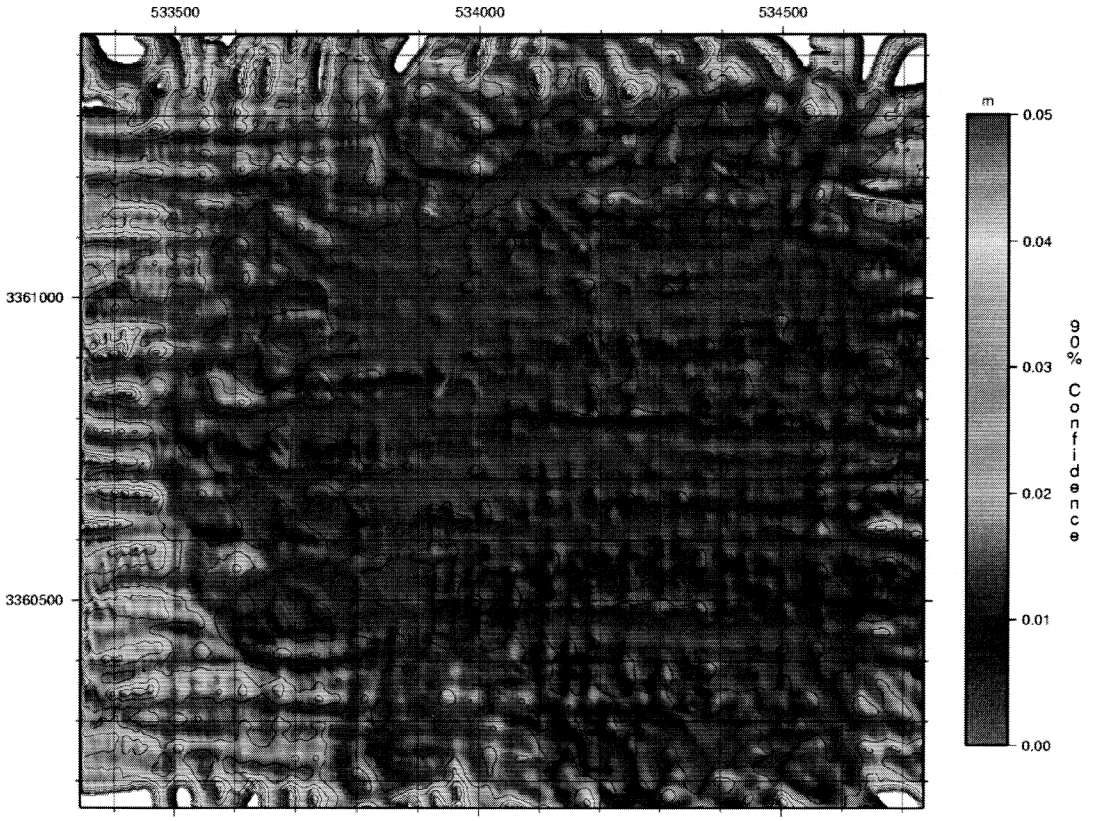

Figure 9: Estimate of depth uncertainty using the Kalman filtering method. Surface represents the 90% confidence interval for estimated depth assuming a Gaussian distribution of errors, and the final estimated variance of the depth estimate. Confidence intervals greater than 5cm have been truncated to provide better dynamic range control.

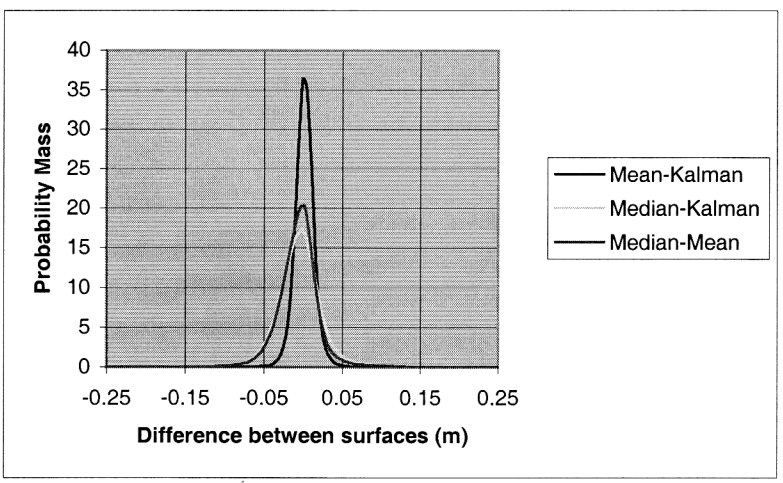

Figure 10: Pointwise difference histograms between surfaces estimated with different algorithms, but using all available data. The mean absolute difference is on the order of a few centimeters in each case, clearly hydrographically and statistically insignificant.

The wealth of data associated with the SAX'99 dataset makes it relatively easy to compute robust estimates with all of the data. We found, however, that as the dataset was progressively subsampled towards more typical survey densities, the recursive estimation algorithm degraded gently, and still remained within nominal '1% of depth' confidence bounds, even at an average data density of 13.7 soundings/m^2 (which is the smallest subset that can be constructed and still retain a minimum overlap of swathes).

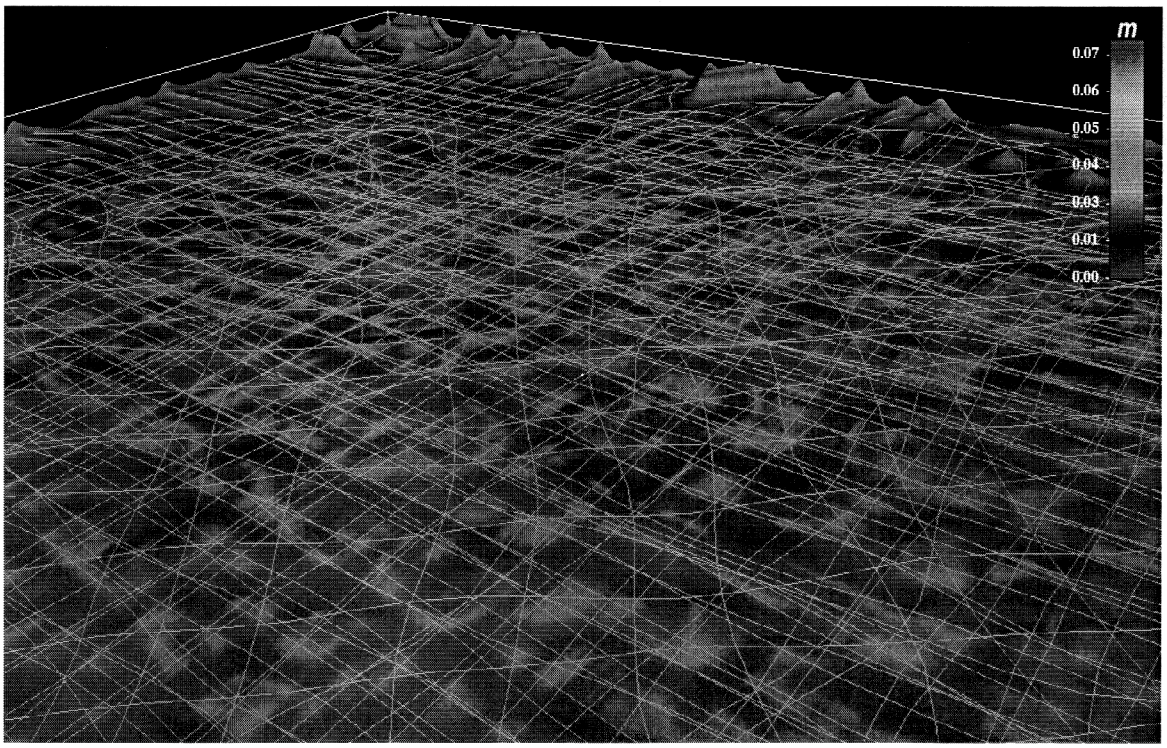

Figure 11: Perspective plot of uncertainty surface (90% confidence interval) generated from full data set with tracklines overlaid. Higher uncertainty is in this case linked almost directly to data sparsity. Vertical exaggeration of the surface is 500×.

3.2 Shelf-Survey: The Ludwig & Walton Pinnacles

During the summer 2000 survey season, the USGS conducted a survey in the Gulf of Mexico to investigate the geology of an area known to contain the remains of coral reefs drowned through sea-level rises [22]. We focus here on the bathymetry of a region known as the Ludwig & Walton Pinnacles [23].

With personnel from the USGS, UNH's Center for Coastal and Ocean Mapping and C&C Technologies, Inc., an eighteen-day survey was conducted from the R/V *Ocean Surveyor*, covering a total of 1650 km^2, in depths from 70-320m. The survey equipment consisted of a Simrad EM1002 intermediate depth MBES, a POS/MV 320 IMU, Dual Trimble 4000 DGPS receivers, backed with a Star-Fix differential station, and a SeaBird 19-02 CTD. Corrections for SVP, ray bending, etc. were carried out at the Simrad topside processor, and data were archived using raw Simrad telegrams. As above, no pre-processing of the data has been done; all beams and all pings are used in the automatic processing. Tidal corrections have been applied using data from the NOAA

tidal station at Panama City, reduced to 89% of measured tides, as recommended by NOAA's Office of Tides. The tide corrections are referred to mean lower low water.

The subsection of data considered consists of four days worth of survey data around the Ludwig & Walton area. Processing was carried out with node spacing of 20m, σ_s = 0m, an 11-pt pre-filter queue, and an outlier culling threshold of $q_{max} = 30$. This section of the survey (figure 13) consists of a relatively constant depth plateau with an offshore shelf break. The area was also found to contain a large collapse feature, typical of the region [24]. Such features are commonly associated with shelf-break slope, or potentially with fresh-water seepage. The data clearly shows the pinnacles reefs expected in the area, and has no significant outlier effects evident. Closer inspection of a perspective rendering of the data (figure 14) reinforce this conclusion, and quite clearly show such features as differential sediment transport regimes on fore and lee sides of the pinnacles. At higher resolution (node spacing 4m, other parameters as before), figure 15, we see significantly more detail of the pinnacle features, as might be expected. In particular, we can see that fine detail of the slump feature is preserved (the small pinnacle-like features on the slump edge at the right of figure 15(b) are approximately 0.5m in height), although we observe few outliers or obvious anomalies. We also observe that the estimates at different resolutions are consistent, as can be seen from figure 16, where contours of the 4m and 20m surfaces are continuous across the join between the two surfaces.

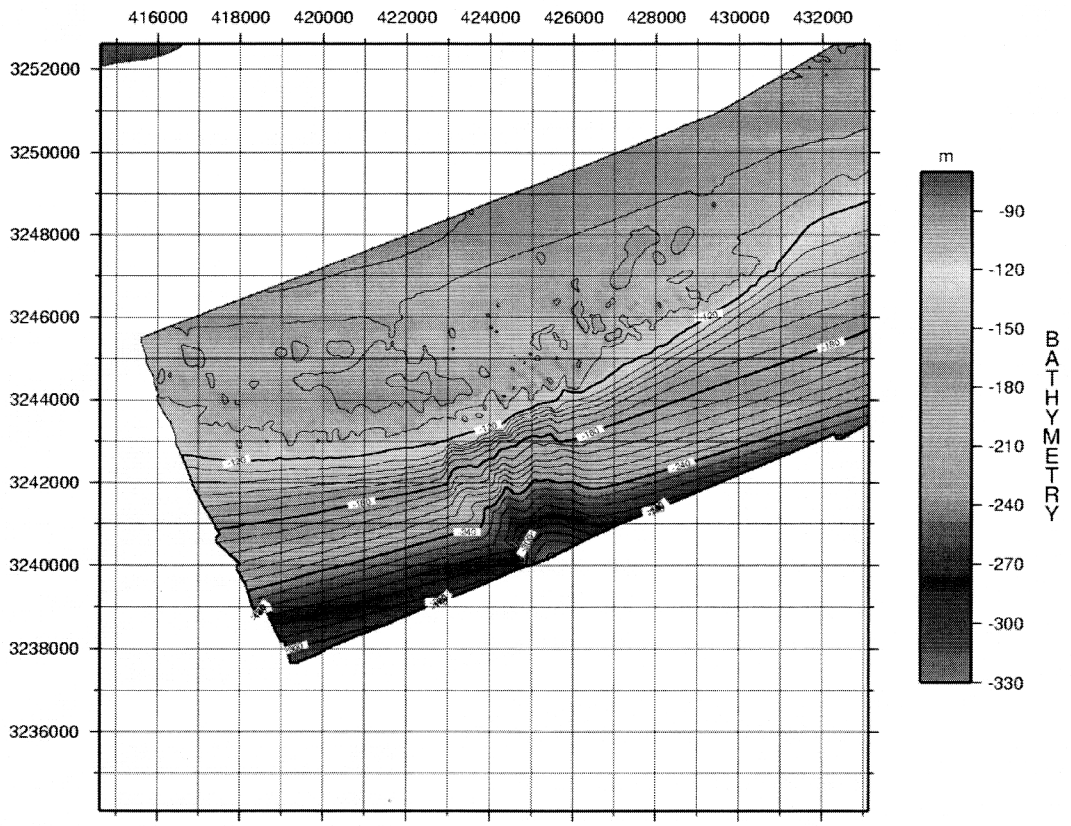

Figure 13: Automatically generated shaded relief DTM of the Ludwig & Walton Pinnacles area, northern Gulf of Mexico. Coordinates are projected meters, using UTM zone 16N on WGS-84. Automatically processed at 20m node spacing, with 'sun' illumination from the northeast.

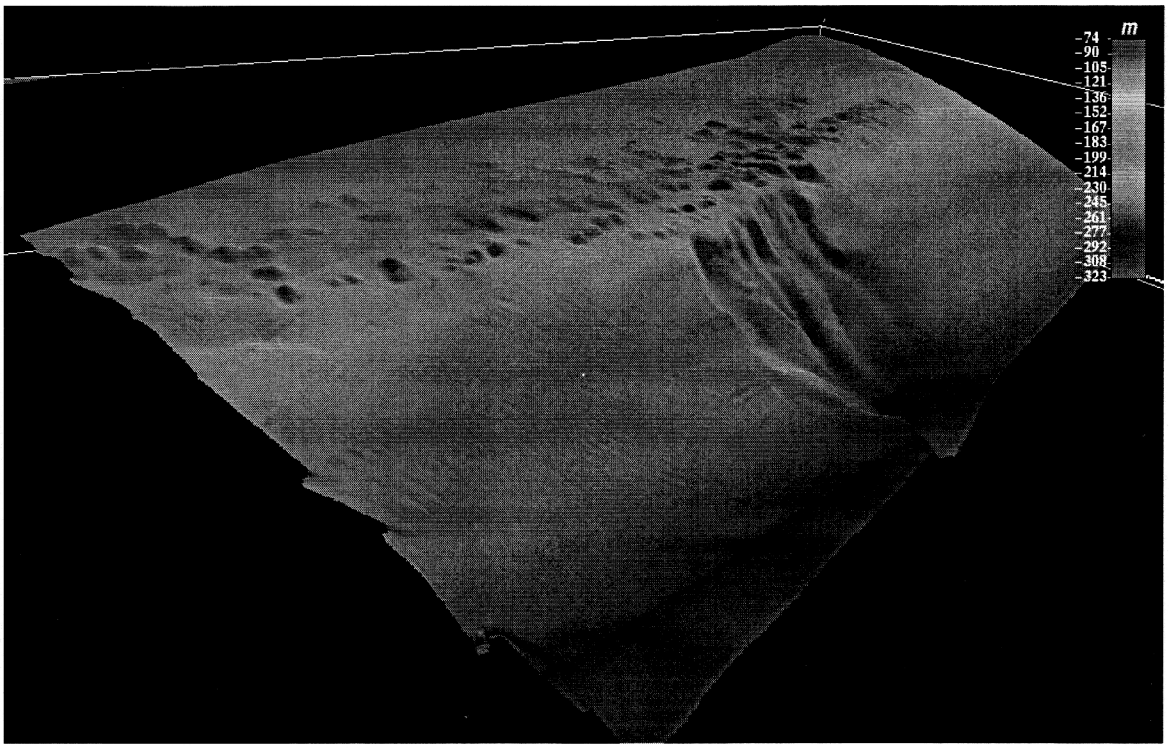

Figure 14: Rendering of the 20m shaded relief DTM from the southwest. Scale is depth in meters below mean lower low water.

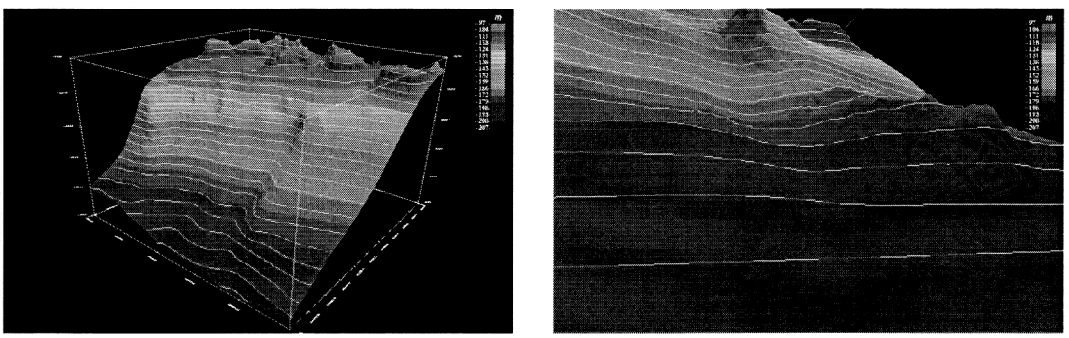

Figure 15: Two perspective views of the high resolution (4m node spacing) bathymetry. Left view is from the southeast, right view is from the southwest. Note particularly the fine detail of the upper escarpment of the slump feature, and of the features on the slump edge (right), which are on the order of 0.5m high.

However, perhaps more interesting than the bathymetry itself is the estimated uncertainty associated with the high resolution surface, figure 17(a), where we see two obvious regimes of behavior, along with what appears to be a roll-induced artifact. The process is more clearly illustrated in figure 17(b), where we have overlaid the predicted 90% confidence interval in meters onto the estimated DTM surface at 4m resolution. From this figure, it is quite clear that the artifacts are trackline oriented, and not depth dependent; further investigation of the orientation shows that the tracklines with higher error were created as the survey platform was heading southwest, and those with lower error when heading northeast.

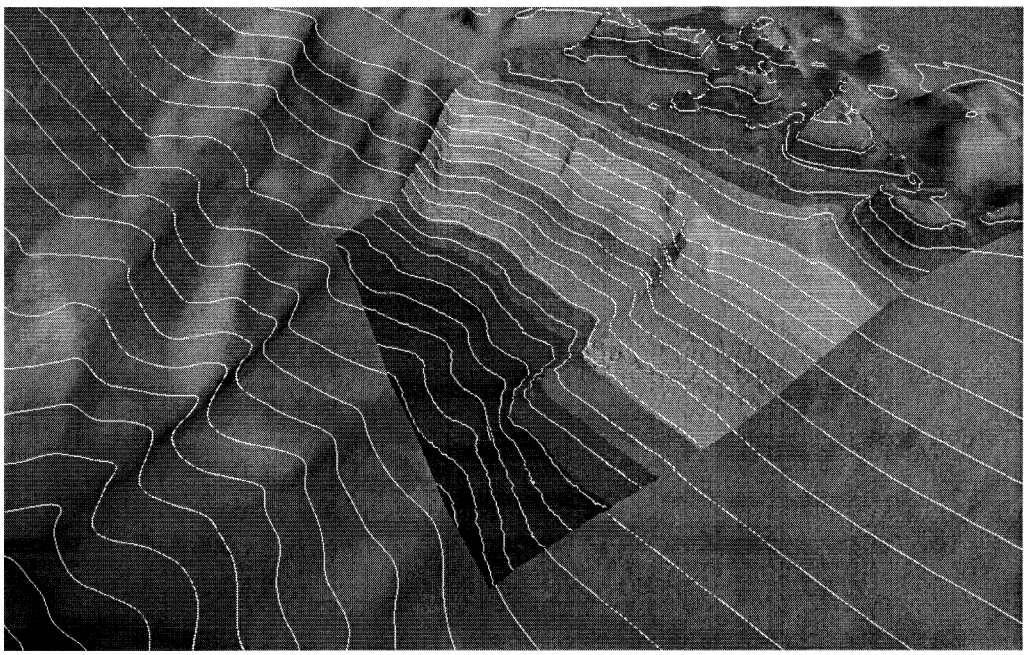

Figure 16: Composite shaded relief DTM constructed by overlaying higher resolution (4m) DTM where available. Contours on the 4m subsection are every 5m; contours on 20m section are every 10m. View from the southeast.

This correlates exactly with the survey logs, which indicate more significant ship motion when running into the swell on southwest legs; significantly increased ship motion always increases the forward predicted error since it acts to magnify any errors we have made in compensating for survey offsets. We note, however, that even in regions of very high error, there are instants where there are troughs of error comparable with those on either side of the trackline. In the future, it might be possible to use such a map of predicted error to find regions, even in bad data, where we could work with greater confidence. Of course, these sorts of estimates could also be used to delineate regions of the survey where further work is required, as well as being a useful clue as to problems observed.

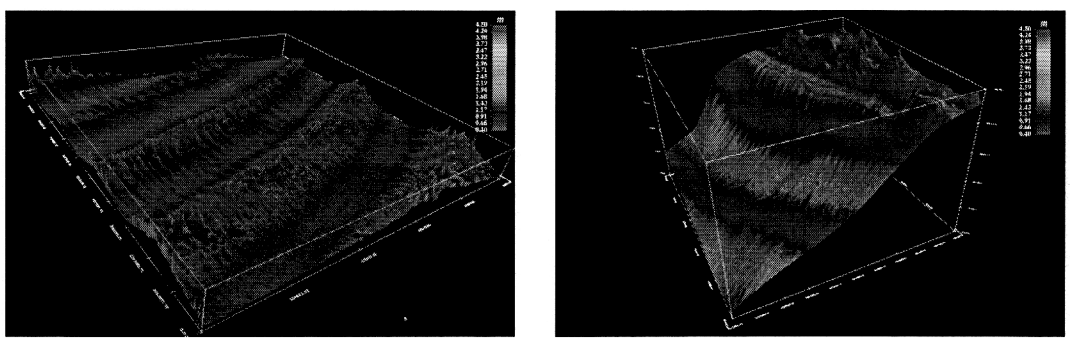

Figure 17: Predicted 90% confidence interval for the high resolution DTM. Left image shows error viewed from the southwest; right image shows the error overlaid on the estimated bathymetry. The regions of higher error are aligned with the survey lines, and are associated with increased ship motion caused by riding into the swell while heading southwest.

4 Discussion

Our results show that it is possible to carry out bathymetric data processing, including automatic cleaning, without having to wait for all of the data to be available, with minimal overhead, and at nominal data capture rates. In addition, the robustness of the algorithm appears to be sufficient to deal with typical stochastic errors inherent in all surveys, and even some non-stochastic errors such as datum shifts, burst-mode sounder failure and badly ray-traced beams.

We observe that, as with all DTM generation tools, our method is fundamentally a weighted griding scheme. However, in its defense, we observe that the method provides a number of advantages that are not present in other schemes (and that we have to have some method of filling in the gaps!) In particular, we provide:

- Individually weighted soundings with non-arbitrary weightings based on forward predicted error,
- Intuitive, but readily modifiable and customizable, model of how data is incorporated into the surface estimate, which is readily adapted depending on survey requirements,
- A continually maintained estimate of estimated bathymetry and associated uncertainty in the estimates, and
- Separation of sounding weighting and sounding incorporation, with data integration to optimally utilize repeated or pseudo-repeated observations.

In practice, the balance between depth prediction and error propagation may depend on the purpose for which the data is collected: in a strict hydrographic processing chain, we might wish to be more conservative and shoal bias all predicted bathymetry, and make the error bounds increase more rapidly. In a geological context, we might need to consider larger scales where our zero-order prediction is no longer valid, and hence incorporate some interaction terms between neighboring estimation nodes to compensate. At present, we specify the propagation terms *a priori* with at best a nod to a Bayesian subjective prior argument for their values. This is a small (but possibly unavoidable) weakness in the algorithm. However, we re-emphasize the point above: it may be a weakness, but at least it is an explicit weakness, and hence will be scrutinized as the method is applied.

There is a significant paradigm shift required to consider statistically derived surfaces as the most valid description of the data, rather than the more traditional estimates (e.g., a shoal biased binned estimate). However, we would contend that a significant proportion of the data observed in such conventional estimates represent the upper tails of the sampling distribution of a MBES, and hence are not representative of the actual bathymetry at all. A reliance on such data is probably rooted in the use of physical sounding methods, where if a lead-line showed up shallower within an area that other points, there *must* have been *something* there. With indirect sounding methods, this is not necessarily the case.

We hope that the addition of a predicted error surface to go with the observed depths weakens the arguments against stochastic processing, at least as a preliminary processing stage. We do not believe that any automatic data processing algorithm will compensate

for all errors in data, stochastic or otherwise, and that manual data examination may be required when the uncertainties warrant it. However, our approach should provide reasonable estimates for the depth in most places, and hence reduce the computational burden on the operators somewhat by focusing attention and effort where it is required.

There are a number of directions for further investigation in implementing the algorithm outlined here, particularly with respect to robustness. The problems that we observe with outliers in some datasets result mainly from non-stochastic errors, rather than the random errors we were attempting to model. In effect, we see a failure of our underlying modeling assumptions, rather than failure of the methodology itself. However, this suggests a number of alternatives, such as multiple hypothesis tracking (i.e., where we allow multiple depth estimates, and then resolve the ambiguity in favor of the most likely hypothesis only when a depth estimate is required). Indeed, determining that multiple hypotheses exist could be used as an indicator for model failure, even if the depths implied were unknown. It would also be possible to implement a formal model testing structure [25], and present the posterior probability for each model as another layer of information on which decisions could be based. In some environments, the idea of a static depth estimate at a particular point may not be valid under any circumstances (e.g., where there is significant tidally driven variation, or where datasets are gathered over a significant period of time). In that case, we would have to revert to the full dynamical Kalman model, although specification of model dynamics would be complex. Further input from manufacturers of sonar equipment to define and refine error models for particular multi-beam (and single-beam) equipment would also be beneficial.

Finally, although we have relaxed the constraint of fixed bins, we have not pursued the difficulties and benefits of allowing a more flexible graph of nodes to be specified. In fact, there is no reason for the nodes to be fixed in place (our only assumption is that we know exactly where they are). One possible extension would be to allow nodes to be added or deleted dynamically, to allow us to match the estimated sampling density required for the surface under estimation. Another possibility is to allow the nodes to gravitate to where they are actually required on the surface, rather than having them fixed in place. If the total number of nodes were constrained to be a constant, this would not have a significant computational overhead, although it would require significantly more investment in book-keeping.

5 Conclusions

We have proposed a new methodology for automatic handling of dense bathymetric data. Based on a simple model of the stochastic errors associated with estimates of depth, we constructed an optimal estimator which tracks both an estimate of the depth, and the uncertainty associated with that estimate. The method has a number of advantages to recommend it, in particular that it automatically incorporates estimates of 3D uncertainty, allows for optimal combination of estimates, and provides a methodology for updating older surveys with new data (or using older surveys to constrain interpretation of new data!)

Prosaically, the methodology also allows us to build a low memory overhead, real-time, on-line estimator so that we have, continually available (as data is added), current best estimates of depth and uncertainty at that depth.

We have outlined a pair of robustness modifications to improve performance of the algorithm but note that neither is really the correct solution in this method. Better alternatives would be to exploit a better source model to describe heavy tailed outliers, or to accept a multimodal posterior distribution, or to attempt to track a set of potential depth hypotheses, and then weigh evidence for the most likely, as required. This is the topic of ongoing research.

6 References

[1] M. Gourley and D. Dodd, "HIPS: Hydrographic Information Processing System", *Universal Systems Ltd.*, CARIS White Paper #21.

[2] D. W. Caress and D. N. Chayes, "New software for processing sidescan data from sidescan-capable multibeam sonars", *Proc. IEEE Oceans '95*: 997-1000, 1995.

[3] J. E. Hughes Clarke, L. A. Mayer and D. E. Wells, "Shallow water imaging multi-beam sonars: A new tool for investigating seafloor processes in the coastal zone and on the continental shelf", *Marine Geophys. Res.* 18: 607-629, 1996.

[4] C. Ware, L. Slipp, K. W. Wong, B. Nickerson, D. Wells, Y. C. Lee, D. Dodd and G. Costello, "A System for Cleaning High Volume Bathymetry", *Int. Hydro. Review*, 69(2): 77-94, 1992.

[5] H. P. Varma, M. Boudreau, M. McConnel, M. O'Brien and A. Picott, "Probability of Detecting Errors in Dense Digital Bathymetric Data Sets by Using 3D Graphics Combined with Statistical Techniques", *Lighthouse*, 40: 31-36, 1989.

[6] J. Eeg, "On the Identification of Spikes in Soundings", *Int. Hydrographic Review*, 72(1): 33-41, 1995.

[7] Z. Du, D. Wells and L. Mayer, "An Approach to Automatic Detection of Outliers in Multibeam Echo Sounding Data", *The Hydro. Journal*, 79: 19-25, 1996.

[8] C. B. Lirakis and K. P. Bongiovanni, "Automated Multibeam Data Cleaning and Target Detection", *Proc. IEEE Oceans 2000*: 719-723, 2000.

[9] J. Depner and J. Hammack, "Area based editing and processing of multi-beam data", *Int. Conf. On Shallow Water Survey Technologies*, Sydney, Australia, 1999.

[10] R. Hare, A. Godin and L. Mayer, "Accuracy Estimation of Canadian Swath (Multi-beam) and Sweep (Multi-Transducer) Sounding Systems", *Canadian Hydrographic Service Report Series*, 1995.

[11] R. Flood, V. Ferrini and L. A. Mayer, "Multi-beam Bathymetry and Backscatter at 300kHz in the SAX'99 Study Area, West Florida Shelf", *Eos Trans. AGU, 81(48), Fall Meet. Suppl.*, Abstract OS62A-17, 2000.

[12] P. S. Maybeck, "*Stochastic Models, Estimation and Control*", V.1, Academic Press, 1979.

[13] S. S. Haykin, "*Adaptive Filter Theory*", 3ed., Prentice-Hall, 1995.

[14] E. Hammerstad, "Multi-beam Echo Sounder Accuracy", Kongsberg Simrad AS, *EM Series Technical Note*, 1998.

[15] E. Pøhner, "EM Series Multi-beam Echo Sounder Error Model", *Report from Simrad Subsea AS to FEMME'93*.

[16] J. Eeg, "On the Estimation of Standard Deviations in Multi-beam Soundings", *The Hydro. Journal*, 89: 9-13, 1998.

[17] IHO Committee, "IHO Standard for Hydrographic Surveys", *Special Pub. 44*, 4th Edition, 1996.

[18] T. H. Cormen, C. E. Leiserson and R. L. Rivest, "*Introduction to Algorithms*", MIT Press/McGraw-Hill, 1990.

[19] E. I. Thorsos, K. L. Williams, D. R. Jackson, M. D. Richardson, K. B. Briggs and D. Tang, "An Experiment in High-Frequency Sediment Acoustics: SAX'99", Inst. of Acoustics, Underwater Acoustics Group Conference, 9-12 April 2001, Southampton, UK.

[20] K. L. Williams, M. D. Richardson, K. B. Briggs and D. R. Jackson, "Scattering of High-Frequency Acoustic Energy from Discrete Scatterers on the Seafloor: Glass Spheres and Shells", Inst. of Acoustics, Underwater Acoustics Group Conference, 9-12 April 2001, Southampton, UK.

[21] N. P. Chotiros, D. E. Smith, J. N. Piper, B. K. McCurley, K. Lent, N. Crow, R. Banks and H. Ma, "Acoustic Penetration of a Sandy Sediment", Inst. of Acoustics, Underwater Acoustics Group Conference, 9-12 April 2001, Southampton, UK.

[22] J. V. Gardner, K. J. Sulak, P. Dartnell, L. Hellequin, B. Calder and L. Mayer, "The Bathymetry and Acoustic Backscatter of the Pinnacles Area, Northern Gulf of Mexico", *Cruise Report O-1-00-GM, USGS Open File report 00-350*, http://walrus.wr.usgs.gov.

[23] Ludwig and Walton, "Shelf-edge Calcareous Brominences in North-eastern Gulf of Mexico", *Bull. Amer. Assoc. Petroleum Geologists*, 41: 2054-2101, 1957.

[24] J. V. Gardner, P. Dartnell, K. J. Sulak, B. R. Calder and L. Hellequin (in prep.), "Physiography and Late Quaternary-Holocene Processes of Northeastern Gulf of Mexico Outer Continental Shelf off Mississippi and Alabama", *Gulf of Mexico Science*.

[25] S. Richardson and P. J. Green, "On Bayesian Analysis of Mixtures with an Unknown Number of Components", *J. Roy. Stat. Soc. (B)*, 59(4), 1997.