Integration of near-continuous sound speed profile information

John. E. Hughes Clarke (1)
Mike Lamplugh (2)
Edouard Kammerer (1)

(1) Ocean Mapping Group, Dept. Geodesy and Geomatics Engineering, University of New Brunswick
(2) Atlantic Region, Canadian Hydrographic Service
primary contact: jhc@omg.unb.ca

Abstract

Recent advances in underway watermass-profiling instrumentation are now beginning to allow us to view near-continuous cross sections of the sound speed structure of the ocean. The sound speed structure controls both the propagation and refraction of sound from multibeam echo sounders and thus ultimately limits the accuracy of the depth estimates. Prior to the availability of these new instruments, discrete, infrequent (2-12 hourly) profiles were used as an approximation of the true water mass. With this new capability, the inadequacy of the older assumptions can now be quantitatively demonstrated.

The current operational procedure, however, is to utilise these underway profiling instruments only discretely whenever the operator suspects water column changes (normally at 1-2 hour periods). Examples of these data are provided for the August-September time period over Georges Bank during an operational multibeam survey. The instrument has the capability, however, to be deployed repetitively over a 1-2 minute cycle, for as long as is desired. For a small subset of the example data, the system was deployed continuously and revealed that the water column was varying significantly over length scales of as small as a few 100 metres. The quantitative effect associated with these changes is presented showing the scale of sounding error resulting from using increasingly infrequent sampling intervals.

A method is presented in which the swath sonar data collected is post-processed using a model of the time varying water column. To do this requires interpolation between the measured profiles as well as extrapolation of the profiles to the seabed. It is proposed that such an approach should become routine as it represents an improvement in the both the fidelity and the cost-effectiveness of the data.

Introduction

Over the past 10 years shallow water multibeam echo sounders have moved from being novelty instruments to being a standard part of hydrographic survey. Over this decade, we have seen huge steps in the elimination of many of the error sources in these sonars (Hare et al., 1995). As robust bottom detection became the norm, the integration of external sensor information was often seen to be the major culprit. Until 1994 when affordable aided inertial navigation systems (AINS) became available, limitations in motion sensor technology dominated the error budget (Hare et al., 1995). With the common availability of AINS today (Dinn et al. , 1997) , the largest remaining barrier to wide-swath high quality data is proving to be the water column.

Unlike most airborne sensors in which the atmosphere can be treated as a homogenous medium, the high vertical and lateral variation in the physical properties of the ocean must be known sufficiently well to account for ray path refraction. This applies both in the initial assumption about beam steering and the propagation path through the water mass.

Measuring the physical properties that control sound speed in the ocean (salinity, temperature and pressure) is a standard procedure for physical oceanography. These physical oceanographic investigations have clearly demonstrated that over the temporal and spatial scales of common hydrographic surveys (days to weeks and kilometres to 10's of kilometres) the ocean will change significantly. As a result most hydrographic operations take discrete measurements of the ocean at periods of less than 1 day.

Until recently, however, taking sound speed profiles (SSP's) at periods shorter than a several hours was a prohibitively expensive undertaking. Bringing a survey vessel to a halt, lowering a sensor several hundred metres and then taking care of all the data quality assurance and data transfer protocols necessary would commonly involve at least 30 minutes of ship time. Because of this agencies were reluctant to take more frequent observations and thereby implicitly assumed that the space and time variability of the ocean could adequately be described using these sparse observations. As swath sonars have moved to ever wider angular sectors (more sensitive to refraction), and as the other sources of error have been gradually eliminated, it has become clear that this is simply not the case.

Two principal limitations exist:

- the water mass really does change over time scales much shorter than the standard sampling period
- the application of SSP's is almost universally done based on the prior observation only.

This paper presents field results using an underway probe that now provides robust frequent measurements of the spatial and temporal variability in the ocean. The field results are from a multibeam survey done in the mid summer period over Georges Bank, an area of known high spatial water mass variability (Loder and Greenberg, 1986). The paper also presents a new approach to applying these more frequent observations that allows for temporal interpolation of the data after the fact.

Oceanography

Georges Bank is a large shoal on the outer shelf located approximately 200 km southwest of Yarmouth, Nova Scotia. It forms a partial barrier between the Gulf of Maine and the Northwest Atlantic, with depths less than 20 m over the center of the bank. The tidal currents on Georges Bank are dominated by the semi-diurnal component, although there is a significant spring-neap modulation (Bisagni, 1992). The M2 rotary tidal currents have increasing current magnitude associated with decreasing depth (Brown and Moody, 1987), and hence increased stirring over the central shallower regions. In addition to the turbulent tidal mixing, atmospheric cooling and wind mixing keep the entire Bank vertically well-mixed during winter. During late spring, however, increased solar radiation causes thermal stratification over the deeper areas of the Bank where the tidal mixing is less intense.

The transition zone, or tidal mixing front , between the shallower well-mixed area and the deeper stratified region on Georges Bank is a persistent feature in satellite imagery from late spring until early fall (Mayor and Bisagni, 1998). The satellite imagery also shows a persistent front near the shelf break. The historic position of the tidal mixing front as observed from satellites indicates that it migrates back and forth about 8 km over the tidal cycle (Mayor and Bisagni, 1998). The fronts on Georges Bank observed from satellite show excellent agreement with the predicted position of the tidal mixing front based on modeling (Loder and Greenberg, 1986).

The survey data used in this paper is from multibeam sonar surveys collected back and forth across this advecting frontal region during the August-September period in 1999. Similar conditions are common throughout much of the Scotian Shelf and Grand Banks region during the summer period and thus are representative of conditions likely to be encountered by other Canadian survey vessels.

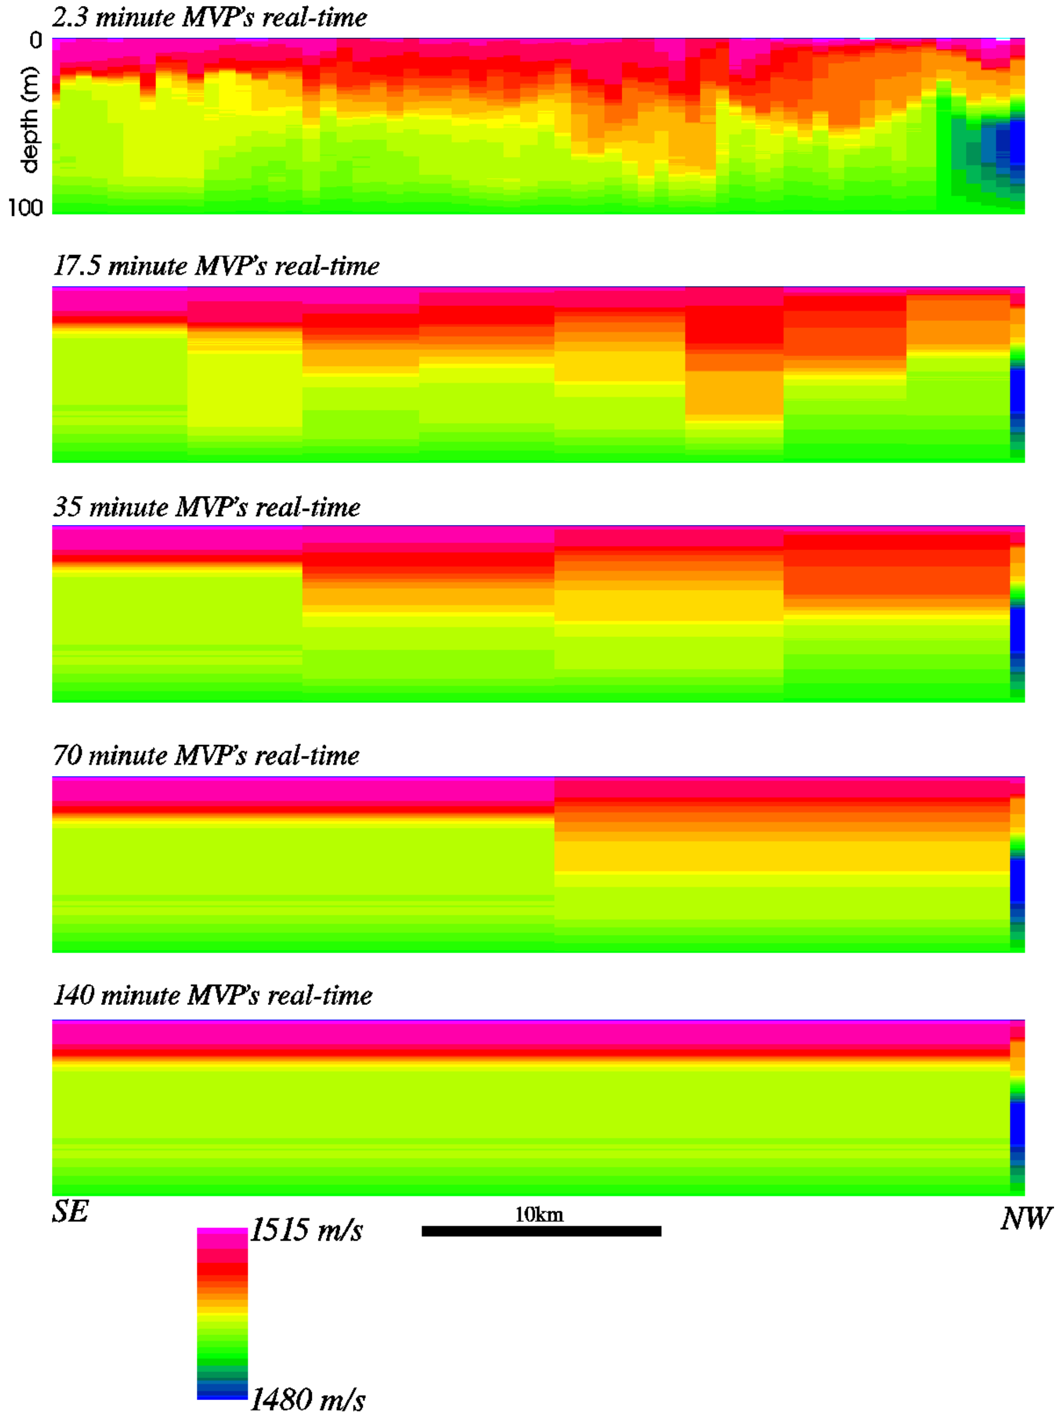

Figure 1: The top profile shows a 45 km long profile of the sound speed variability in the top 100m of water across Georges Bank from SE to NW. The profile was acquired at 10 knots from the MV Anne S. Pierce using a BOT MVP-200 during the simultaneous acquisition of a line of EM1002 multibeam data (examined in this study).

The top profile shows the spacing of the actual measurements (approximately every 140 seconds) whereas the lower profiles show the view of the water column that would be implied if profiles were only collected every 17.5, 35, 70 or 140 minutes. The extrapolation is that implicit in the standard real time application of sound speed profiles, where the last profile measured is used as the only model right up into the collection of a subsequent value.

Instrumentation

The MV Anne S. Pierce is equipped with a Kongsberg Simrad EM1002 multibeam echosounder (Simrad 1999). This system is integrated with an Applanix POS/MV 320 DGPS aided inertial navigation system. The EM1002 transducer is mounted on a retractable ram that pushes it out approximately 80 centimetres from the hull of the vessel. The transducer draft when down is approximately 4.4m. The ram is not equipped for mechanical pitch stabilisation.

The EM1002 forms 111 beams over a user-controlled angular sector. Each beam has an effective beam width of 2.4 degrees fore- aft and ~2 degrees across track. The beam spacing may be equiangular, or equidistant or "in-between". For all operations on Georges Bank, the in-between spacing was used and angular sectors of ~130 degrees were maintained. For the range of water depths considered in this study (75-90m) the ping rate was around 2.5 Hz.

The EM1002 uses three transmit sectors. Inside +/- 50 degrees, 98 kHz is used, and outside this 93 kHz is used. For all operations considered here, a 0.2 ms pulse-length was used. The 111 beams are roll stabilised. The system however, is incapable of pitch stabilisation. A mechanical pitch stabilisation method is available with this system but was not implemented on the Pierce.

The Pierce is equipped with a Brooke Ocean Technologies (BOT) Moving Vessel Profiler (MVP-200). This is an underway-profiling instrument developed between BOT and the Ocean Physics branch of the Department of Fisheries and Oceans. Using a heavy streamlined depressor weight, near vertical profiles are obtained by letting the probe fall through the water column. The probe may be instrumented with a sound speed or conductivity/ temperature sensor. The brake for the cable is applied automatically based on the depth (derived from the bridge echo-sounder) and a user-set tolerance limit (usually ~5m above the bottom). Once the brake is applied, the probe very rapidly jumps back about 50% of the way up the water column. After this the probe is winched back in for the next deployment.

The Data Set.

The Pierce –EM1002 survey platform was operated for approximately 90 survey days in the July to November period in 1999. All the field acquisition was focussed on the northern half of the Canadian end of Georges Bank. The survey line presented here was acquired in September. The line is of particular interest as it is one of only two lines that were acquired whilst the MVP-200 was operating in "continuous" mode. In the water depths encountered (75-90m) the probe completes a cycles in ~140 seconds. As a result, at the survey speed employed (~10 knots) these were thus acquired about 700m apart. The line ran from the centre of the bank to the northern flank, thereby crossing the tidal front boundary described by Mayor and Bisagni (1998).

It is interesting to speculate on the stability of the position of this front from line to line. Each survey line was about 2.5 hours long and the front was crossed about 30 minutes before the northern end of the line. The front was thus encountered either 1 hour apart or 4 hours apart. Based on the satellite imagery, this front is known to be advected up to 8000m over a tidal cycle (Mayor and Bisagni, 1998). It would thus be dangerous to try and make assumptions about the spatial stability of the water mass to attempt a static spatial model of the water mass distribution.

Much of the data was collected within the tidally mixed zone in the central part of the bank. Even though the tidal mixing does reduce the stratification, it is clear from the profiles (Figures 1 and 2 top) that significant spatial heterogeneity on a scale of kilometres exists. Even from profile to profile, local perturbations of the thermoclines are seen which may indicate the presence of internal waves. This observation is supported by evidence from the Georges Bank Frontal Study (Loder et al, 1988) During that study, Batfish (another towed undulating physical oceanographic sampling tool) sections across the Bank edge indicated that packets of large-amplitude internal waves were being generated southward from the tidal front during off-bank tidal flow.

Data Handling

All raw Simrad data telegrams were directly imported into the OMG/UNB SwathEd processing software. Sounding data were cleaned, gridded and the backscatter data all mosaicked on a line by line basis, about 2-4 hours behind acquisition. First-pass terrain models and backscatter mosaicks for the entire bank were available at the conclusion of each 14 day survey leg. All data were gridded at resolutions of about 5% of the local water depth. 600 map sheets, covering the entire northern section of the bank were updated throughout the survey. These map sheets were routinely sun-illuminated to investigate small-scale residual artifacts in the bathymetric data.

During the near real-time data processing, it rapidly became apparent that there were significant refraction-related artifacts occurring within the swath data. In response to this, three actions were taken:

1. the frequency of MVP profiles was increased to about every hour
2. a subjective refraction compensation was applied to the data (the “reftool” of SwathEd)
3. the line spacing was closed up to minimise the use of the refraction-compromised outer beam data.

The refraction residuals were noticed to be most pronounced during low sea state periods on the bank. Immediately after the passage of a hurricane, the homogenised water mass allowed much improved control on the refraction problem. Nevertheless, as the water column stratification was reestablished the problem of variable water masses returned.

Given that the data is logged with periodic discrete water column profiles that should allow us to compensate for the changing water masses, there are a number of possible approaches to handling the refraction issue:

The Real Time Approach

On board the vessel in real time, the last water column measured is used by the system to ray trace all the beam solutions. Naturally this model of the water column is most appropriate close to the time and place of that discrete observation. As time progresses and the vessel moves away from that location, the appropriateness of the profile drops. Until, however, a subsequent profile is obtained, there is little choice but to use the most recent.

The shipboard staff were very aware of this and thus whenever a sustained symmetric refraction artifact was noted in the real-time data displays, a new water column profile was obtained. Prior to the availability of the MVP, this required stopping the vessel and lowering a profiler. This time duration of this, together with all the data interfacing issues normally resulted in a 30-45 period in which the vessel was not logging the data. Because of this long time duration, previous surveys tended to delay water column profiles to at least the end of the survey line and usually at periods no closer than 6-12 hours.

With the MVP however, the profile can be obtained without stopping the vessel. It is necessary however to stop the sonar pinging, just for the period during which the serial transfer and data incorporation takes place. This is normally equivalent to about a 1 minute break. If the vessel did not break off the survey line, there would thus be a several hundred metre gap in the data collection. Therefore a compromise was accepted whereby the profile was acquired whilst on-line and logging the bathymetry, but immediately afterwards the vessel would perform a survey speed 360 degree manoeuvre with logging disabled during which the new water column profile was uploaded. At the termination of the turn, the sonar would start pinging slightly before the point at which logging was terminated. In this manner there are no data gaps, and only about a 4 minute loss of survey time is experienced.

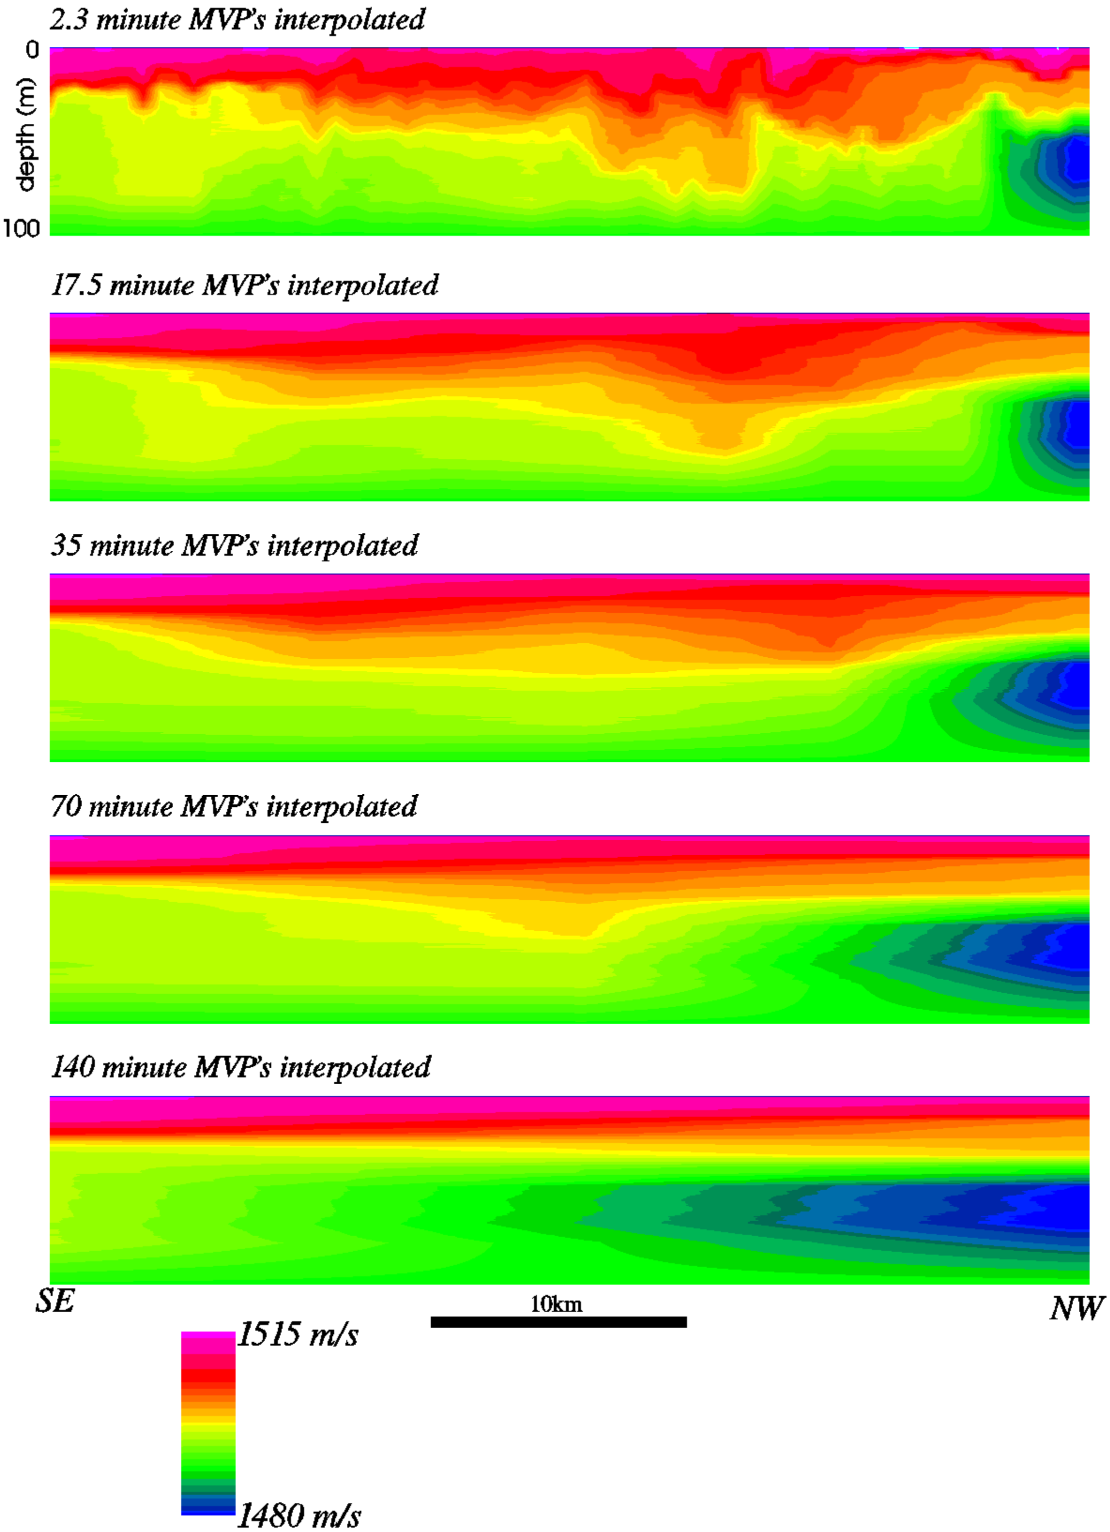

Figure 2: Showing the various spatial models of the water column produced by linear interpolation of the more sparsely sampled profiles. Top shows the results of using the full 2.3 minute sampling period. The profiles below show successively more infrequent sampling (the sample spacing can be inferred from the boundaries seen in Figure 1).

Note the loss of definition of the drop of the thermocline in the centre of the area and the artificial elongation of the incursion of slower speed water (blues) from Georges Basin to the NW. In all cases the profile has been extrapolated to a common 1490 m/s at 100m. This is probably not an appropriate value for the Georges Basin water mass.

The most noticeable limitation of applying the data in real time is that there is an abrupt step in the across track profile (that shows up as a false target on the seafloor) at the time when the new SSP is first applied.

The Interpolated Post-Processing Approach

Naturally during the period immediately preceding the new SSP, the data acquired should really be handled using the new SSP. This is of course is not possible in real time. Nevertheless the sonar saves a number of parameters from which a new refraction solution may be calculated in the luxury of hindsight.

These parameters include:

- the ship-head-relative azimuth and depression angle of each of the 111formed beams. Note that this implies that the sonar has used a valid measure of the surface sound speed and properly accounted for the orientation (roll pitch and heading) at both the time of transmit and the time of receive.
- the sonar-mount-relative steering angles of the transmit and receive beams together with the full asynchronous time series of orientation (acquired at 100Hz) and surface sound speed (acquired at 1Hz). This allows recalculation of the beam steering angles in the event that the initial surface sound speed estimate was invalid.

Given all this data, in post processing a number of different strategies may be applied. The simplest, used herein, is to reprocess each ping using a time-weighted average of the closest profile immediately prior to and after the time of the ping.

Comparison of 2 minute profiles with 17, 35, 70 and 140 minute profiles.

If one intends to do a spatial/temporal interpolation based on discrete profiles, we need to understand the significance of sampling frequency. Naturally the question arises: how low can I set my sample frequency and still have sufficient accuracy? To investigate this the same data is processed using less frequent sampling intervals at the successively longer periods of 17.5, 35, 70 and 140 minutes.

The results of this approach are illustrated in Figures 3 and 4. In figure 3 one sees graphically the percentage depth error for each of beam for each of the pings over the 2.5 hour survey line (extending over 45 km). For each ping, the depth solution calculated using the interpolated 2.3 minute profiles is subtracted from the profile that would be observed if an approximation of the water column is used either :

- based on the last profile logged (i.e. equivalent to real time) at that chosen sampling period or
- based on a linear interpolation between the two bounding profiles.

For the case of the interpolated 17.5 minute profiles, the errors are kept well in check. There are small perturbations as one crosses the tidal mixing front to the northwest. If only the last sample at 17.5 minutes is used, however, larger errors appear when crossing the tidal fronts and for the most extreme case to the NW, errors over 1% shallower are seen in the outermost beams.

As one drops to a 35 minute sampling period, the interpolated solution has a worst error of about 0.5% due to the fact that the front is artificially extended to the SE. Without interpolation however, the front is not picked up initially causing the same > 1% shallower errors seen with the 17.5 minute profiles.

Extending to a 70 minute sampling period on sees that, for the interpolated case, the artificial extension of the front produces minor (<0.5%) deeper errors. Without interpolation, one sees a problem in the first segment where the dropping of the thermocline is not adequately described (< 0.5% deeper). And again, without a knowledge of crossing the front the > 1% shallower errors are seen.

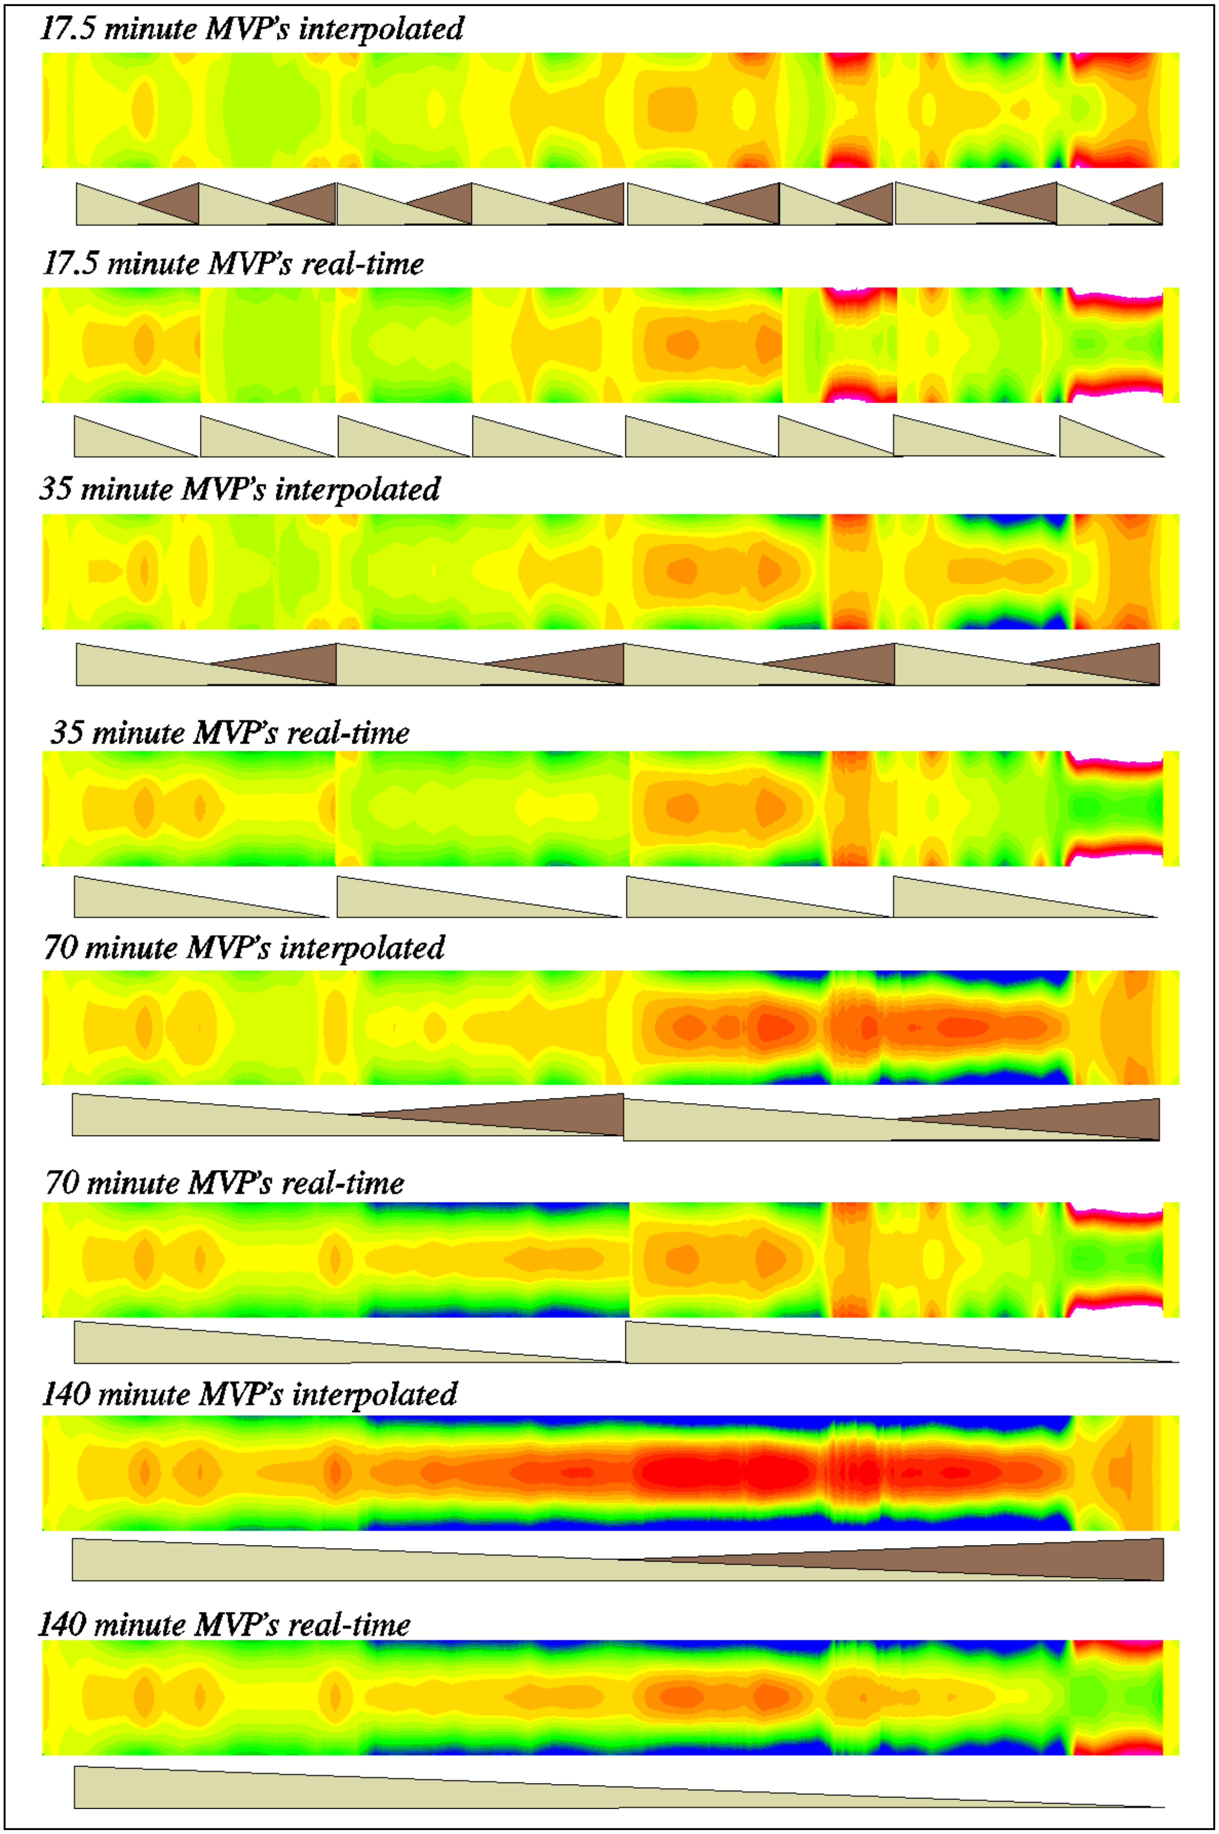

17.5 minute MVP's interpolated

17.5 minute MVP's real-time

35 minute MVP's interpolated

35 minute MVP's real-time

70 minute MVP's interpolated

70 minute MVP's real-time

140 minute MVP's interpolated

140 minute MVP's real-time

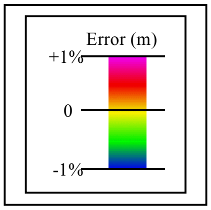

Figure 3 (above and left): Diagrams showing x- axis as distance along the survey line (45 km in total from SE (left) to NW (right)) and y- axis as beam number. The colours indicate the percentage magnitude of the error if only the subsampled approximation of the water column is used (either real time (Figure 1) or interpolated (Figure 2) using sparse SSP's). Pink/White indicate 1% too shallow, dark blue/black indicates 1% too deep. Results for the four considered sampling strategies are shown, 17.5, 35, 70 and 140 minutes.

For the last 140 minute sampling case, it is interesting to see that similar magnitudes problems occur whether the data is linearly interpolated or not. This reflects the fact that the sampling of the water column is occurring at periods far larger than the typical scale of change of the water column. At this point, the extra CPU time involved in the post processing of the interpolated solution actually, by chance, provides a worse solution. In this worst case considered (yet still more frequent than the standard operation policy of most swath survey agencies!) there is nothing to be gained in trying to apply simple spatial/temporal interpolation schemes.

In Figure 4, the average percentage depth error is plotted as a function of beam number for each situation. The average depth error is calculated in two ways. On the left, the standard deviation of the error is shown, whereas on the right the average bias is shown.

Based on an average of 22,000 pings, the average bias due to refraction invariably deviates less than 0.5% for all cases up to the 70 minute sampling period. This time averaging however, disguises, the fact that, unlike random uncorrelated noise in bottom detection or perhaps motion, the magnitude and sign of the refraction errors change very slowly over periods of 5-10 minutes. For these shorter periods the errors may well be significantly larger than the long term average (as the standard deviation (Fig. 4 left) indicates). It is just that over the time scale considered (2.5 hours) the refraction errors range from positive to negative (smiles to frowns) and thus the average profile bias appears little distorted.

Even the standard deviation is misleading as there are periods whilst crossing oceanographic boundaries of several minutes (about a kilometre or more) when the outer beam errors can grow to 2-3% of the total water depth.

Limitations when using underway sound speed profiles

Profiling Frequency

The major limitation will be the upper limit on the frequency of profiles. Whilst these instruments are designed for continuous high speed operations, they have only rarely been used in this mode. This is because of very real concerns about losing a continuously streamed instrument, especially at night in fishing grounds.

As this data set clearly demonstrates, the lower the frequency the poorer the ability to monitor rapidly changing water masses. One possible guide to the right compromise between instrument longevity and data fidelity may be to use the real time surface-sound-speed time-series as an indicator of water mass changes. For example, if the surface sound speed changes by more than 2m/s do another dip, otherwise only do a dip every ~15 minutes.

There is a problem with this approach however. If one dipped infrequently until one came across a frontal region, the linear interpolation scheme that is used in the above examples, would artificially extend the water mass back across the frontal location (e.g., Figure 2 bottom). In this case a weighted linear interpolation (biased toward the prior profile) might have to be used for special case where the probe was triggered due to a surface sound speed anomaly.

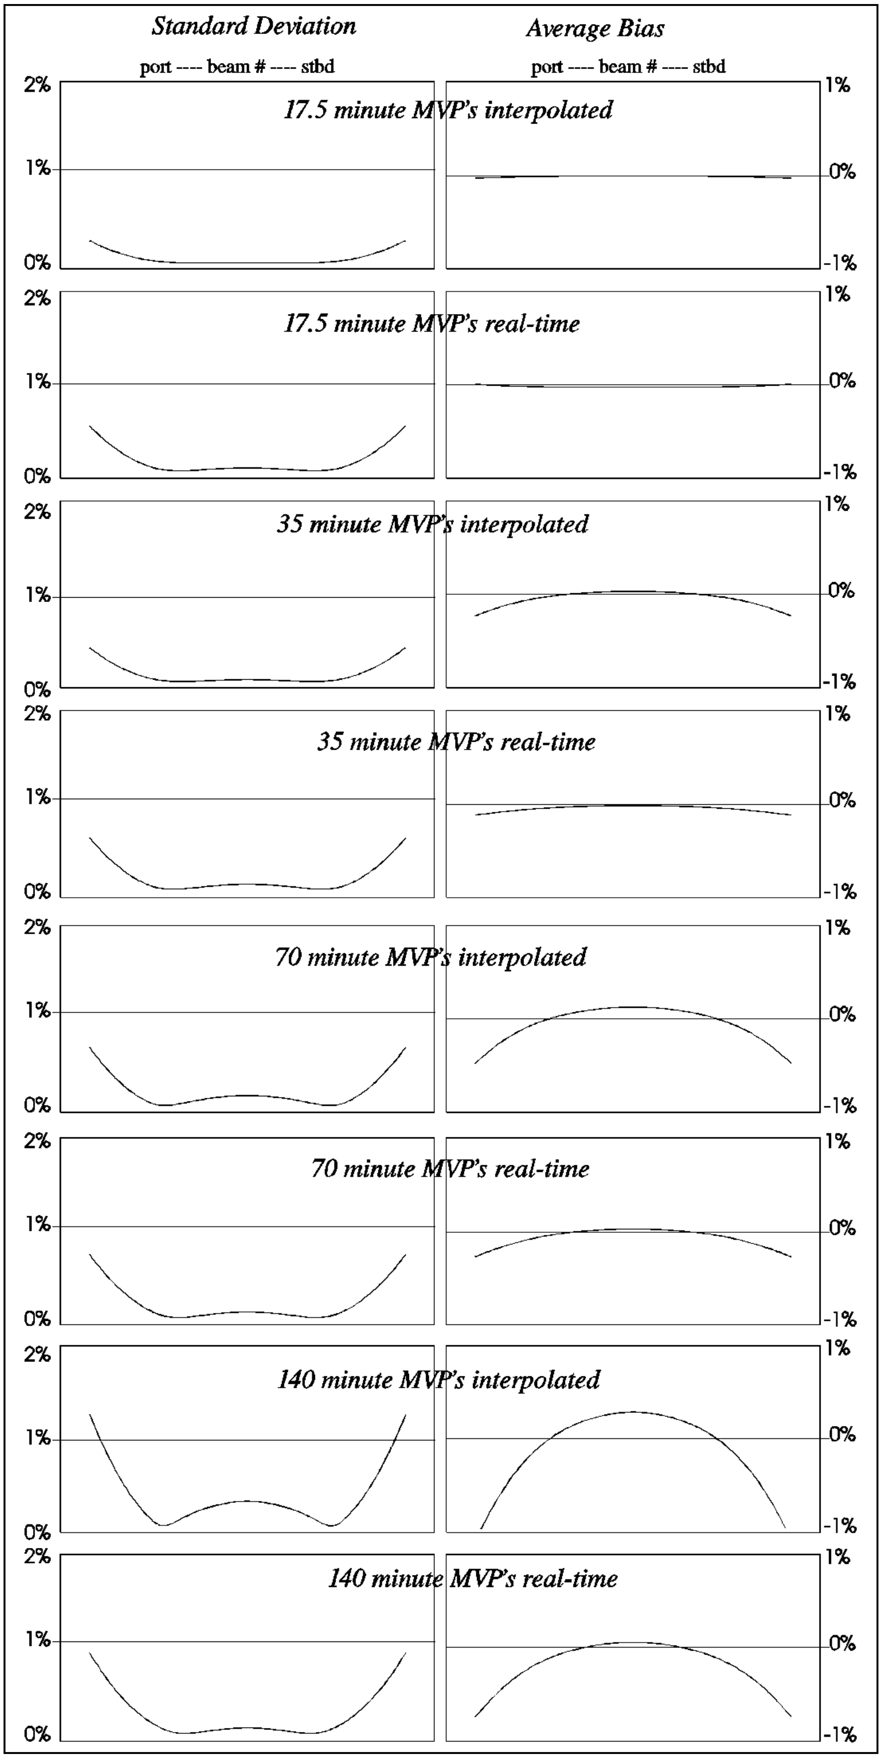

Figure 4: Showing the standard deviation and average bias of the percentage depth errors for the 111 beams of the EM1002 (operating in 130 degree mode).
The average is based on the 22,000 pings collected in the 2.5 hour period used in the previous examples.

One sees a very clear trend toward increasing error with decreasing water column sampling frequency.

Note that for the last (140minute) case, the interpolated time series actually produces worse errors than just the use of the prior water column. This reflects the fact that the sampling is now at far longer distances than the typical changes in the water mass.

Data Corruption
With much more frequent water column profiling and allowing for the fact that the data is at high speed (with thus higher probability of impacting submerged targets) it will be necessary to be able to robustly absorb the expected high occurrence of data quality problems.

Any real SSP tends to have spike or transient anomalies due to a variety of causes including bubbles near the surface, algae, bottom impact, imperfect electrical connection etc. Before such data is entered in the sonar real time software, the hydrographer standardly tries to remove these either through subjective editing. This can be a time consuming process. As however, the frequency of the sampling increases it is increasing valid to superimpose and compare successive profiles. This should allow more robust automated filtering.

Profile Extrapolation
With solitary profiles, the operator would normally try to choose a local deep location so that the profile would be valid for all shallower depths. As we move to underway profiling, we dare not let the probe approach the seabed too closely. The MVP standardly operates to within about 5m of the seabed (it takes an NMEA feed of the depth from the bridge echo sounder). We are thus faced with the problem of extrapolating the 2D cross-section of the water column down to the seabed.

There are a variety of plausible approaches:

1. one could always extend the last value down to the seabed. The danger here is that the basal value is sometimes corrupted.
2. one could extend the basal value with a slow pressure induced increase with depth.
3. one could enter a nominal profile for all profiles to blend with. This has the problem that one may produce a sharp discontinuity at the basal depth whenever an anomalous water mass is encountered.
4. The deeper profile could be based on the closest profile that extends to greater depths.
5. the operator may decide each extrapolation on a case by case basis. This however, will probably be too time consumptive

Proposed Operational Data Handling

The software to integrate the spatially variable water column has only recently been developed. Now that the efficacy of the method has been demonstrated an operational policy needs to be put in place to apply this information in the field.

It is recommended that the acquisition and storage of the water column data be maintained on a separate logging system. This is because, if the data is intertwined into the sonar data stream, it will be very difficult to isolate the water column information. The bathymetric data is logged only within discrete line based files, whereas the required water column for a line may have been logged during the time window of another line (before or after), or even during the logging gap at the end of a line.

The full resolution of the water column information need NOT be applied in real time, It will be useful to have a reasonable approximation of the water column entered into the real time sonar software. This will allow an approximate monitoring of the scale of the water column problem. But, given the lost time currently required to wait while the SVP is downloaded, and given that the interpolated solution (based on the profile yet to be acquired) is likely to provide a better measure anyway, the is little to be gained by uploading all the profiles immediately.

As long as all the required information is retained, including:

1. full orientation time series.
2. full surface sound velocity time series.
3. full sonar-mount-relative steering angles (and the surface sound speed assumed for these)
4. the exact alignment of the sonar mount

5. full record of changes in mechanical stabilisation of the mount platform
6. full record of the two way travel time estimated for each beam .

-

there is no reason that the refraction solution should not be routinely improved in post- processing.

The water column information should be treated as an external environmental variable much like the tide. Initial data processing (navigation and bathymetric outlier editing) can be adequately completed on the imperfectly refraction corrected data prior to the application of the full interpolated water column model.

A separate verification process will be required to ensure that the water column information is reasonable. For the old infrequent profile acquisition methods, this quality control was done on a profile by profile basis at the time of measurement, This is not however feasible at 1-2minute profile periods. In this case one can take advantage of the expectation that sequential profiles should demonstrate a gradual alteration in the water mass. If the sequence of profiles is presented as a cross-section (e.g. Figure 2 top), any outlying profiles will be immediately obviously based on the profiles obtained a few hundred metres away. Note that this was not possible for the 6-12 hourly profiles in the past (in which it was unreasonable to expect the profile to show any great similarity).

Once this verification process is complete, the full data set of SSP's for a discrete period of data (perhaps 1 24 hour chunk) can be applied in a batch process mode to all the data within that 24 hour period. Whilst this will be a compute intensive process, no human interaction is required after the water column quality assurance step.

An added quality flag should be retained with each swath that records the time to both the prior and next profile used in the water column interpolation. The duration of this time period can be used in estimating the uncertainly of this profile. Recent implementations of the Hare et al (1995) model by Reed et al. (1999) have allowed use of real time environmental data such as:

- Heave magnitude,
- roll/ pitch excursion and
- surface sound speed variability.

To this may now be added the "latency" of the water column information used.

Conclusions

The post-processing of frequent water column information allows a much better control of sounding errors due to the spatial and temporal variations in the water column. The combination of:

- near continuous water column information and
- software that allow seamless space or time variable models of the water column to be applied to the data in post processing

can result in a vastly improved ability to cope with the common oceanographic variability in multibeam sounders.
The biggest limitations currently facing us include:

- long term reliability of these underway sensors.
- editing of imperfect water column data
- interpolation of profiles to the full depth
- adequate monitoring of the surface sound speed variations

The required frequency of profiling ultimately depends on the local oceanography. For the case considered here (a typical summer time mid latitude outer shelf) if the sampling is much sparser than 15 minutes, the ability to cope with frontal regions is severely compromised. If sampling is less frequent than hourly there is little to be gained by doing the post processing based on interpolation as the length scales of the oceanographic features are commonly smaller than the sampling distance.

Many sonar manufacturers are now selling systems with angular sectors over 120 degrees. Until ~1994, motion sensors capable of providing the angular accuracy to match these low grazing angles were not available making their use for hydrographic purposes impractical (Hughes Clarke et al., 1996). Now that the motion sensor issues have been addressed (Dinn et al., 1997) the water column remains the last barrier. Using the commonly accepted practice of SSP's only every few hours, no hydrographic agency would currently seriously consider using angular sectors over 120 degrees. If however, underway vessel profilers can be used, this last impediment to truly wide swath echo-sounding may be removed.

Acknowledgements.

This work would not have been possible without the enthusiastic cooperation of the captain and crew of the MV Anne S. Pierce. The successful integration of the Pierce swath system was primarily the result of the dedicated efforts of Don Dinn. Field processing and first recognition of the scale of the refraction problem was done by Tim Kearns, Andy Craft and Chris Leblanc

References

Bisagni, J.J., 1992: Differences in the annual stratification cycle over short spatial scales on southern Georges Bank, Cont. Shelf Research, 12, 415-435.

Brown, W.S. and J.A. Moody, 1987: Tides. In Georges Bank, R.H. Backus, ed., MIT Press, Cambridge, MA, 100-107.

Dinn, D.F., Furlong,A., Loncarevic,B.D., Penny,M. and Dakin,D.T., 1997, Controlling multibeam sonar errors: Sea Technology, v.38, no.8, p.75-80.

Hare,R., Godin,A. and Mayer,L., 1995, Accuracy estimation of Canadian Swath (multibeam) and Sweep (multi-transducer) sounding systems: Canadian Hydrographic Service, internal report.95pp.

Hughes Clarke, J.E., Mayer, L.A. and Wells, D.E., 1996, Shallow-water imaging multibeam sonars: A new tool for investigating seafloor processes in the coastal zone and on the continental shelf: Marine Geophysical Research, v.18, p.607-629.

Loder, J.W. and D.A. Greenberg, 1986: Predicted positions of tidal fronts in the Gulf of Maine region, Cont. Shelf Research, 6, 397-414.

Loder, J. W., K. F. Drinkwater, E. P. W. Horne and N. S. Oakey. 1988. The Georges Bank Frontal Study: An overview with preliminary results. EOS 69, 1283.

Mavor,P.T. and Bisagni,J.J., 1998, Sea Surface Temperature Fronts in the Georges Bank Region , American Geophysical Union, abstract, 1998 AGU/ASLO Poster: OS32C-01

Reed, B., Hammack, J., Hare, R., Fabre, D., "Horizontal and Vertical Error Estimation in Multibeam Data", Conference Proceedings, Shallow Survey '99, Sydney, Australia, 1999

Simrad, 1999, Simrad EM1002, Multibeam Echo Sounder, Product Description: Kongsberg Simrad.