WATER-COLUMN VARIABILITY ASSESSMENT FOR UNDERWAY PROFILERS TO IMPROVE EFFICIENCY AND ACCURACY OF MULTIBEAM SURVEYS

Matthew J. Wilson
National Oceanic and Atmospheric Administration, Office of Coast Survey
439 W. York Street, Norfolk, VA 23510
matthew.wilson@noaa.gov

Jonathan Beaudoin
UNH Center for Coastal and Ocean Mapping-Joint Hydrographic Center
24 Colovos Road, Durham, NH 03824

Stephen Smyth
Blue C Designs
21 Dickson Avenue, Halifax, Nova Scotia, Canada B3M 2W4

Abstract

With the advent of underway profilers, sampling the water-column to obtain sound speed corrections is no longer a detriment to hydrographic survey efficiency. Instead, the challenge has become deciding how many casts are necessary to maintain a desired level of multibeam sounding accuracy, while not needlessly overworking the profiler.

Ray tracing uncertainty analysis can determine in hindsight whether a particular sampling interval is adequate or not. Based on this methodology, an algorithm was developed to generate recommended sampling intervals based on successively acquired sound speed profiles, allowing the MVP to run in a "cruise-control" mode where the sampling interval is altered in response to changing oceanographic conditions.

In collaboration with Rolls Royce, the algorithm was implemented in Python and loosely couples with the MVP controller software such that the recommended sampling interval can be adjusted without operator intervention. Integration of the software with the MVP controller was successfully tested aboard the *NOAA Ship Ferdinand R. Hassler* in September of 2012. Initial results from field trials and from analysis of existing data sets are presented.

Introduction

The application of timely measurements of sound speed within a survey area is critical during multibeam surveying; elsewise the refraction of sound in the water-column may be significant enough such that allowable error tolerances are exceeded. The vertical variation of sound speed (referred to as a sound speed profile, or SSP) is applied to the depths reported by a multibeam echo sounder to mitigate the effects of refraction and therefore maintain sounding accuracy requirements. The measurements of sound speed must be taken often enough to capture both spatial and temporal variability of the SSPs; if not, an undersampled water-column will result, and incorrect SSPs will be applied (Beaudoin, 2010a). Refraction-induced depth sounding error resulting from an undersampled water mass costs time and effort to correct during post-processing, and results in decreased sounding coverage and accuracy.

Traditionally, deploying instrumentation over the side of a vessel in order to sample the water-column (herein referred to as a cast) could only be conducted statically—which meant the vessel must be held stationary, and survey operations suspended throughout the duration of the cast. Considering the time taken to conduct the cast (dependent on the water depth), plus the time taken for the vessel to resume the desired survey speed and heading, static sampling of the water-column is generally a detriment to survey efficiency. Thus, the surveyor had to maintain a delicate balance between the loss of survey time taken to acquire the cast, versus the improvement in sounding accuracy (Beaudoin, 2010b).

In the past few decades, underway profilers have been developed, such as the Moving Vessel Profiler (Furlong et al., 1997) from Rolls-Royce, and the Underway CTD (Rudnick and Klinke, 2007) from Oceanscience. Underway profilers such as these allow for the acquisition of SSPs while the vessel is in motion, meaning that a cast can be obtained without suspending survey operations, resulting in a tremendous increase in survey efficiency, as well as a much higher number of casts that can be obtained (Beaudoin et al, 2011).

Surveyors equipped with an underway profiler have a new challenge. Since there is no restriction on the number of casts that can be obtained, surveyors often needlessly oversample the water-column, collecting far more casts than are necessary to maintain desired accuracy levels. This oversampling is not prudent in that it exposes the instrumentation and associated hardware to undue wear and tear, and it also increases the risk of fouling and grounding the towed sensor (Beaudoin, 2010b).

An optimal sampling interval exists somewhere between undersampling and oversampling (Beaudoin, 2010b), such that desired accuracy levels are maintained with a minimal number of acquired casts; however, there is no standard methodology in finding it. There do exist methods to assess full water-column variability (NOAA HSSD, 2012), both prior to and during survey operations, however these are highly subjective (Beaudoin, 2010b), and moreover do not prescribe an optimal sampling rate. Furthermore, there are a multitude of tasks inherent to multibeam surveying, and limited personnel; hence the optimal sampling interval based on observed variability is often determined haphazardly.

With these challenges in mind, a software tool has been developed at the University of New Hampshire's Center for Coastal and Ocean Mapping / Joint Hydrographic Center called "CastTime". The tool will monitor oceanographic variability in real-time based on the acquired sound speed data from underway profilers, and will derive suitable recommendations for optimal water-column sampling intervals accordingly. As a result, CastTime can provide benefits to the NOAA Office of Coast Survey by alleviating the subjectivity associated with sampling interval determination. There is also the potential for an improvement in sounding accuracy, as the constant monitoring of CastTime will help to ensure accuracy requirements are met. Conversely, it will prevent needless overworking of the underway profiler, which can reduce expenditure towards maintenance and replacement parts for the system.

This paper will discuss the methodology utilized in the development of CastTime:

1) constant gradient ray tracing is performed for each acquired SSP,
2) the ray traces are compared, from which the incurred sounding depth bias is derived in a procedure called uncertainty analysis (Beaudoin, 2010a), and
3) an algorithm logic uses the sounding depth bias, the existing sampling interval, and a specified error tolerance to recommend a new sampling interval.

The components listed above were implemented into a Python software program to facilitate data analysis and field testing. The testing was conducted with previously acquired high-density SSP data sets, which were thinned retroactively by means of the algorithm, a standard set interval, and an idealized best solution determined in hindsight. The performance of each thinned set is assessed, analyzed, and compared with regards to sampling efficiency and mitigation of sounding depth bias due to refraction.

In collaboration with Rolls Royce and Blue C Designs, the Python software program was loosely coupled with a customized version of the MVP Controller software, such that the sampling interval recommendations are generated from casts collected real-time, and then sent back to the MVP without operation intervention, to queue the next cast accordingly, resulting in automated, "hands-off" water-column sampling, utilizing algorithm-derived intervals that adapt to changing oceanographic conditions. The Python software program, coupled with the MVP Controller software version, represents the first version of CastTime, and it was successfully tested real-time with the MVP200 onboard the *NOAA Ship Ferdinand R. Hassler* during hydrographic survey operations, and the results analyzed.

This work establishes the software infrastructure for the CastTime tool, and recommendations and lessons learned from the initial trials will be applied toward future versions and continued development. In conjunction with NOAA Hydrographic Systems and Technology Program (HSTP), CastTime will undergo field testing in 2013.

Methodology

Ray Tracing

Using Snell's Law, and assuming the sound speed variation between each measurement in an SSP is linear, the constant gradient method of ray tracing can be used to determine the horizontal and vertical distance traveled, as well as the time elapsed between each measurement (Lurton, 2010). This procedure is repeated in an iterative fashion until a desired endpoint criteria is achieved.

For any given SSP input and multibeam swath width, constant gradient ray tracing calculations are performed in CastTime. The depth and launch angle of the starting point of each of the individual ray traces may be adjusted based on an echo sounder draft and sound speed measured at the echo sounder. Each ray then terminates when the deepest point of the SSP is reached. A separate ray trace is performed for each 1° increment of launch angle over half of the multibeam swath width (the swath is symmetrical so only one side is necessary to compute). Examples are shown in Figure 1a and 1b (with rays in 5° increments for ease of display).

Uncertainty Analysis

The precise horizontal and vertical position of the ray trace as a function of travel time and launch angle, computed for each ingested SSP, allows for quantitative comparison amongst individual profiles in a procedure called uncertainty analysis (Beaudoin, 2010a). This method allows for the estimation of sounding uncertainty due to water-column variability. It is particularly useful in that no actual sounding data is required, which allows for ease of calculations, data management, and processing times. Furthermore, the sounding depth uncertainty is not limited to the seafloor—instead it is quantified over the entire potential sounding space. Thus, the sounding uncertainty at any desired depth, or elapsed travel time, can be recalled. This "look-up" capability is critical in an SSP comparison, when more often than not the profiles being compared are not of the same depth.

Consider the scenario whereupon two SSPs are acquired and ray traced successively, shown in Figure 1a and 1b, respectively. The SSP is Figure 1a is then outdated, and the SSP in Figure 1b has become a better representation of the water-column. Note that the sound speed in the shallowest layers has decreased, while the depth of the thermocline has risen.

By performing a comparison to the newly acquired SSP in Figure 1b, the outdated SSP of Figure 1a can be assessed of sounding depth bias incurred had it still been used. This is accomplished by first obtaining the total travel times required for each ray to reach the simulated flat seafloor of the newly acquired SSP. Then those travel times are used to access the corresponding horizontal and vertical position of each ray in the outdated SSP, to establish where the echo sounder measured two-way travel time would have been positioned had the outdated SSP been used. Now, for the outdated SSP, the rays no longer terminate at the assumed flat seafloor, but

instead when those travel times are consumed. Plotting the ray traces against one another is shown in Figure 1c.

Note the curvature of the resulting seafloor in Figure 1c, commonly known as a refraction "smile". This serves as a quantitative assessment of the effect of refraction on depth sounding uncertainty. The highest magnitude of vertical uncertainty across the swath will be retained for ensuing algorithm logic (for swath widths 60° or more, this value is almost always associated with the outermost beam of the swath). Also retained is the RMS depth sounding uncertainty of all points in the vertical difference of derived seafloors.
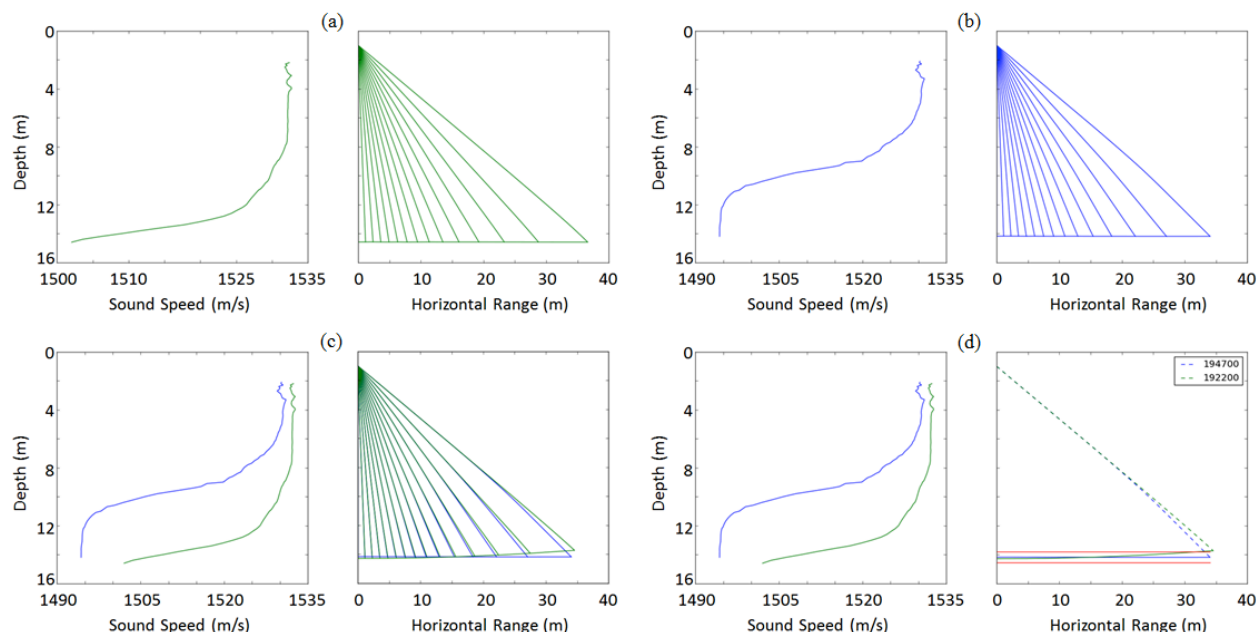

Figure 1. (a) An example of constant gradient ray tracing. An outdated SSP (left) is ray traced (right) using an echo sounder depth of 1.0 m, over a 140° (+/-70°) swath, with launch angles in 5° increments for ease of display. The termination points of each ray simulate a flat seafloor. (b) Like (a), except the newly acquired SSP shows that the water-column has changed considerably, in that the sound speed in the shallowest layers has decreased, while the depth of the thermocline has risen. (c) The SSP's (left) and ray traces (right) of (a) and (b) are overlaid. This time the rays of (a) terminate not at a flat seafloor, but instead when the travel times of (b) are consumed. Comparing the simulated seafloors is an estimation of sounding depth bias. (d) Initial CastTime implementation of the uncertainty analysis, using SSP (a) and (b) as an example. Only the outermost rays are drawn (dashed). Note the SSP timestamps (HHMMSS) in the upper-right hand corner, and the bounds of NOAA allowable error due to refraction (red lines).

For each SSP acquired, this comparison is performed and assessed according to a specific error tolerance. NOAA refraction error tolerances were utilized for the case studies of this report, although this tolerance may be adjusted to meet the requirements of any hydrographic office. The NOAA limit for depth error associated with refraction for multibeam systems is given as 0.30 m plus 0.5% of the water depth (NOAA HSSD, 2012), and the upper and lower bound of this tolerance are represented as red lines in Figure 1d. The SSP timestamps are also given in Figure 1d, and it is shown that 25 minutes have passed between casts. Over these 25 minutes, and across an unspecified horizontal displacement, the limits of NOAA tolerance have been

exceeded in the outermost beams. Note that this issue could be mitigated by reduction of usable multibeam swath width, if this option were preferable to increased sampling frequency.

Algorithm Logic

To examine the temporal behavior of the SSP comparisons, it is helpful to plot the results of the comparisons as a function of elapsed time. For this, a "sawtooth" diagram is utilized, which depicts an assumed linear growth of sounding depth bias (both outer beam error as well as the RMS error across the entire swath) with the time elapsed between SSPs. The SSP comparison from Figure 1d is presented as a sawtooth diagram in Figure 2, and the important parts of the diagram are labeled.
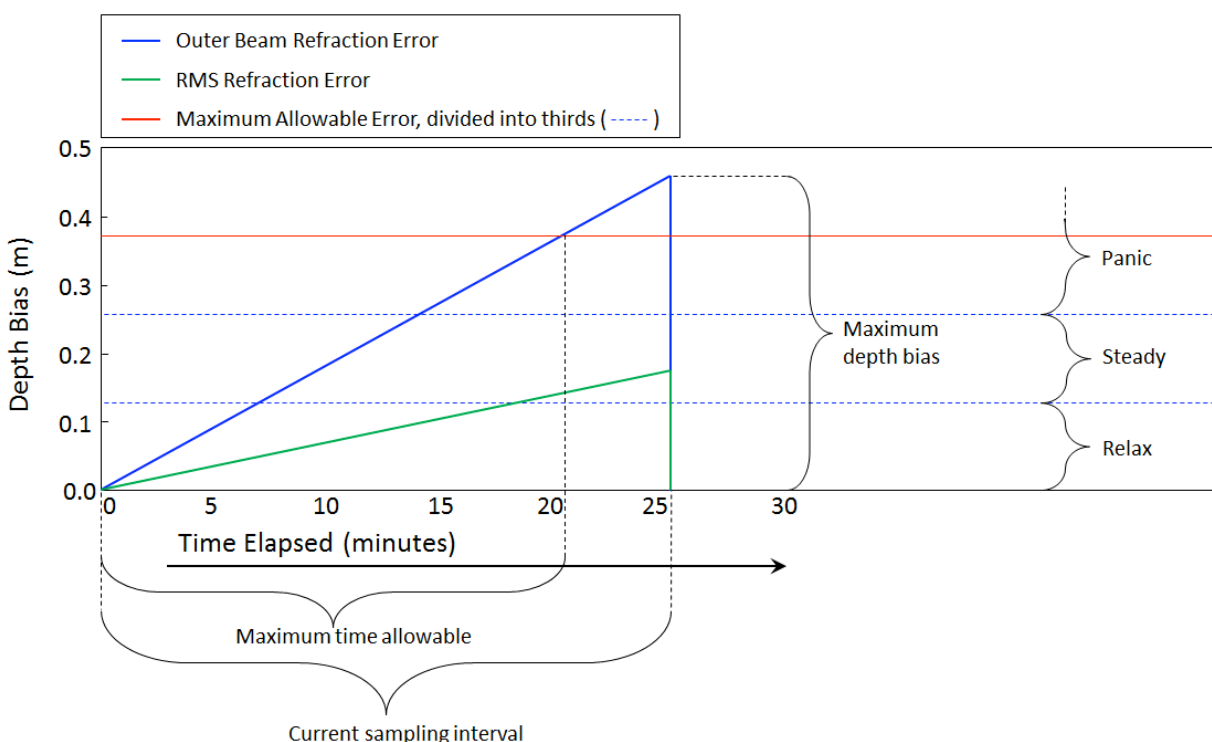

Figure 2. A sawtooth diagram of a single SSP comparison shows an assumed linear growth of sounding depth bias with time, for both the outermost beam (blue) and RMS of all beams (green). NOAA allowable error (red) has been divided into thirds (dotted blue), the bounds of which correlate with algorithm decisions. The current sampling interval of 25 minutes incurred a maximum depth bias of 0.46 m, which in this case not only enacts the "Panic" condition of the algorithm, it exceeds NOAA requirements altogether. The maximum time allowable is determined by interpolation (or extrapolation) of the outer beam refraction error to the exact moment when the maximum allowable error is reached. The ensuing sampling interval recommendation will be half of the maximum time allowable.

As observed in Figure 2, when the first cast was acquired (at time elapsed=0), the sounding depth bias was set to zero, simply because at that exact moment of acquisition, the collected SSP is assumed to be perfectly representative of the water-column. With each minute of time elapsed, the SSP becomes more and more outdated and the depth bias grows linearly until the maximum outer beam and RMS depth bias is reached. When the next cast is acquired (at time elapsed=25),

the depth bias is again set to zero. Note that the linear growth assumption of the depth bias is chosen for simplicity, and due to the absence of any additional information about changes in the water-column between casts. Also note that all sounding depth biases reported in this paper will be displayed as an absolute value.

Based on the results of the SSP comparison, an algorithm will determine a recommended interval using a simple logic. Note in Figure 2 the measure of "maximum time allowable", which is either a linear interpolation or extrapolation of the maximum depth bias to the point in time when the maximum allowable error is exceeded. This term, along with the current sampling interval, will determine the appropriate reaction based upon the observed maximum depth bias.

In Figure 2, note that the maximum allowable error has been divided into thirds. The recommended sampling interval is dependent upon which third the maximum depth bias falls within, enacting a response condition of either "Relax", "Steady", or "Panic". The response conditions will determine the recommended sampling interval as follows:

- **RELAX:** if the maximum depth bias falls within the lower third of the allowable error, the newly acquired SSP is classified as redundant, i.e. it indicates that the oceanography is unchanged. The current sampling interval is increased by 15%.
- **STEADY:** if the maximum depth bias falls within the middle third of the allowable error, the newly acquired SSP is classified as necessary, i.e. it indicates that the oceanography is changing, but the current sampling interval is adequately capturing the change. The current sampling interval is retained.
- **PANIC:** if the maximum depth bias falls within (or is greater than) the final third of the allowable error, the newly acquired SSP is classified as critical, i.e. the oceanography is changing at rate high enough such that the current sampling interval is insufficient. The current sampling interval is thus reduced sharply, to one half of the maximum time allowable.

The algorithm responses are designed to ensure accuracy requirements are maintained through aggressive reductions in the interval should the allowable tolerance levels be approached, while also minimizing unnecessary casts through a gradual interval relaxation, so long as oceanographic conditions remain unchanging. Examples of each response condition are depicted in Figure 3.

It is the "Steady" condition, the middle ground, which is being sought, wherein the casts reflect enough change in the water-column to be deemed necessary, while also remaining safely inside allowable tolerance levels. In this manner, CastTime can minimize the unnecessary workload to the MVP, while also serving as a constant monitor to changes in the water-column.

The sampling interval can be manually adjusted or a cast can be forced at any time, based upon a surveyor's intuition, or simply when it is convenient, considering vessel movement and other shipboard operations. The manual adjustments, and other user-specifications (described below), should be used to ensure the spatial extents of a survey area are effectively captured. This is necessary, because CastTime—as it currently stands—is based only on time elapsed between successive SSPs, and does not include any spatial component in its algorithm.
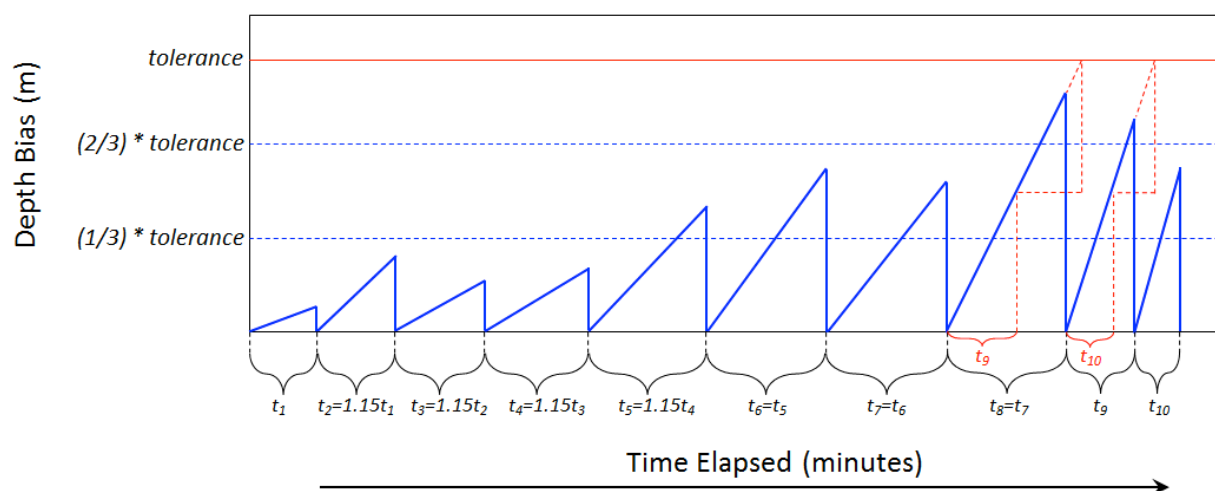
Figure 3. Examples of enacting the "Relax" (t_1–t_5), "Steady" (t_6–t_8), and "Panic" (t_9 – t_{10}) conditions of the algorithm, presented on a sawtooth diagram. The first 5 sampling intervals show a gradual 15% increase, the next 3 remain unchanged, and the final 2 are reduced aggressively to one half of the maximum time allowable for the previous comparison. The final comparison is deemed "Steady", thus t_{11} (not pictured) would equal t_{10}.

A user must specify an initial sampling interval, because the algorithm cannot make a recommendation until at least two casts have been acquired (allowing for the SSP comparison). This should be considered carefully, as it will influence ensuing algorithm decisions. A conservative approach, or if there is reason to suspect high variability, would start with a small sampling interval, e.g. 5 minutes, followed by gradual sampling interval relaxation—if the oceanographic conditions allow. With this approach, the relaxed sampling intervals are effectively "earned". Alternatively, for survey lines fairly short in duration, a good rule of thumb may be to set the initial sampling interval such that a cast is acquired at the start and end of the first line, to maximize the spatial extents of the survey design. For example, if the lines are 10 minutes in duration, the first cast would be conducted as the first survey line is commenced, and then the initial sampling interval set to slightly less than 10 minutes so the next cast is acquired prior to completion of the line. The scenario to avoid would be setting the initial interval to 20 minutes, placing the second cast in a location not far removed from the first, and any existing spatial variability across the survey area would not be detected. This is a danger known as spatial aliasing and should be avoided with an appropriate initial sampling interval (and manual adjustments later as necessary).

The algorithm offers no predictive capability; it is entirely reactive, and based on the SSP comparisons of casts as they are acquired. Thus, an upper bound of the sampling interval must be set to avoid any "surprises", meaning that no matter how consistent or unchanging the water-column appears to be, a cast will still be conducted in accordance with this user-defined upper bound. For example, an upper bound set to 60 minutes will ensure a cast is acquired every hour, regardless of SSP redundancy, which is entirely reasonable for a survey vessel equipped with an underway profiler. Furthermore, if there is suspicion of sudden oceanographic changes (perhaps due to rapid frontal passage or vessel movement into an unsampled area), the sampling interval

should be manually adjusted. In this way, the best practices and intuition developed by prudent surveyors, and simple common sense, are neither lost nor replaced.

Note that the algorithm itself is exceedingly simple. This was purposeful—while there are numerous ideas on how to further improve the algorithm (for example, the inclusion of a time series analysis, or a spatial component), it was desired to first start with a very simple and easily understood idea, and incorporate it into a modular, object-oriented software infrastructure. The next step is to test the algorithm in the field and solicit feedback. In this manner, future algorithm enhancements and modifications can incorporate those lessons learned in the field, and furthermore, the software infrastructure already exists to facilitate a seamless implementation of the new and improved versions.

Software Development

CastTime was developed in support of the NOAA Office of Coast Survey, thus its implementation required interaction with the existing hydrographic survey tools and equipment onboard NOAA ships. CastTime therefore is written entirely in Python, a programming language that already has widespread use throughout NOAA, and therefore facilitates potential integration with existing NOAA software. In addition, as an open-source software, Python is easily accessible, and has the flexibility and capabilities needed to accomplish all goals associated with this project. NOAA hydrographic vessels are equipped with the MVP, thus it was logical to tailor much of the development of CastTime towards the MVP system.

The core of CastTime remains the ray tracing, uncertainty analysis, and algorithm logic. Additional Python scripts were composed to tailor the real-time operation of this core towards the MVP framework (though the algorithm could be applied to any sampling equipment). The end result is a complete feedback loop (shown in Figure 4) wherein 1) SSP data is received from the MVP at the conclusion of a cast; 2) the ray tracing, SSP comparison, and algorithm logic are each executed and result in a recommended sampling interval; 3) the MVP is triggered to conduct the cast at the appropriate time; and 4) the cycle repeats. Each of these goals were accomplished with support from both Rolls-Royce and Blue C Designs.

The first part of the CastTime feedback loop involved the ability to receive the SSP data immediately after acquisition, and without user intervention. Network transmission capabilities already existed within the MVP controller software, as part of the Multibeam Interface options within the MVP configuration. The cast data was sent automatically via User Datagram Protocol (UDP) using a binary packing protocol originally implemented for the Naval Oceanographic Office (NAVO) software "NAVO ISS60", in the output file format of asvp (a Kongberg file format in ASCII). The NAVO ISS60 transfer protocol was provided by S. Smyth (personal communication, May 7[th], 2012) of Blue C Designs.
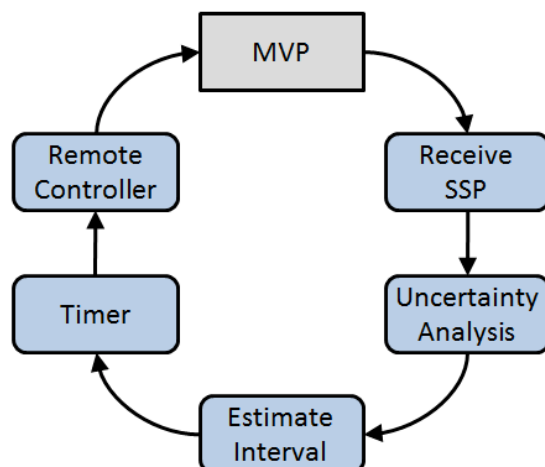
Figure 4. CastTime is a feedback loop between a Python software package and a customized version of the MVP controller software.

After receiving the SSP in Python at the conclusion of a cast, the data is read, stored, and the ray tracing is performed. The uncertainty analysis is performed at the depth of the newly acquired SSP (the final data point of the older SSP is extended, if necessary) and a maximum depth bias is attained. The algorithm logic then results in a recommended sampling interval, which is either accepted or adjusted manually. The chosen sampling interval then initiates a timer, which is viewed as a countdown in Python, thus the surveyor knows exactly how much time is remaining until the next cast. When the countdown is complete, the Python remote controller (see below) sends the start command to the MVP—pending of course the acknowledgement of the surveyor, to ensure vessel readiness and permission from the bridge. After the cast is complete, the data is again transmitted to Python and the cycle repeats.

There was no existing functionality within the MVP controller software to support remote commands from an outside source. Thus, with the support of Rolls-Royce, S. Smyth of Blue C Designs composed version 2.47 of the MVP controller software, and it was released to CCOM-UNH for research and testing purposes on August 31, 2012. The updated software version enables remote control of the MVP from an outside source, allowing for the Python scripts to trigger the MVP to start a cast at the appropriate time, and effectively completes the feedback loop for real-time operation. Implementation of the remote commands feature a binary datagram structure composed by S. Smyth that is packed in Python, and then the datagram is transferred via UDP to the MVP to trigger the cast. The datagram structures of the remote commands were provided by S. Smyth (personal communication, June 21, 2012).

Case Studies

High-Density Data Sets

Prior to a real-time field trial of the algorithm and its implementation within CastTime, it was first desired to test the algorithm in a controlled environment using archived sets of densely-acquired SSPs (acquired with underway profilers with sampling intervals on the order of a few minutes or less). The high rate of SSP data collection allows for near-continuous examination of the water-column sound speed structure beneath the survey vessel.

An example high-density data set is from Portsmouth, New Hampshire, acquired by CCOM-JHC using an MVP30 equipped with an AML Sound Velocity, Pressure, and Temperature (SVPT) sensor. The data was acquired within Portsmouth Harbor and up to 8 km offshore (see the right side of Figure 5) on June 14th and June 21st, 2012. Much of the data offshore was collected as part of survey lines that were generally less than 20 minutes in duration. Transects across the mouth of the harbor are included on the 21st.

As shown on the left side of Figure 5, there was considerably more sound speed variation on the second day. This could be due in part to the different weather on the two days affecting the oceanographic conditions, as the first day was overcast with a light drizzle, whereas the second day was bright and sunny. The higher sound speeds in the uppermost layer on the second day may be a result of surface heating. However, part of the variability observed on the second day is likely due also to the fact that the SSPs were acquired further up into the mouth of Portsmouth Harbor, and are influenced by salinity gradients associated with the Piscataqua River outflow.


Figure 5. SSPs acquired by CCOM-JHC on 6/14/2012 (red) and 6/21/2012 (blue) are plotted on the same axes (left). Cast locations (right) were up to approximately 8 km offshore of Portsmouth Harbor. NOAA chart 13278 is in the background, with soundings in feet.

The SSPs (asvp file format) from both days were ingested in Python, read, stored, and ray traced. Prior to any thinning procedures, the original high-density data set is analyzed—each successive

SSP is compared against the previous using the SSP comparison method as described earlier—and will be used later in the analysis to represent an oversampled water mass.

Thinning and Control Procedures

The algorithm is applied to the high-density data sets retroactively, to simulate the scenario whereupon the surveyor had sampled the water-column in accordance with algorithmically-derived recommendations from the outset. The result is a thinned set of SSPs. This procedure is repeated using standard set intervals commonly used during survey operations, e.g. a 15 or 60 minute fixed interval.

Temporal, linear interpolation between the SSPs of an original, high-density data set results in continuous, uninterrupted data that represents the evolving water-column. This approximation greatly facilitates known interval thinning of the high-density data sets, in that the thinned sets are no longer restricted to the existing SSPs, and SSP samples can instead be derived at any desired time within the data set—or at any time as dictated by a particular water-column sampling interval. Therefore both algorithmic and fixed interval thinning draw upon temporally interpolated SSPs from the original high-density data set as needed to complete each sampling interval. Examples of this temporal interpolation are shown in Figure 6.


Figure 6. Nine SSPs (grey) are temporally interpolated between the old (green) and the new (blue) SSPs. Both SSPs are resampled to common vertical levels of 1 cm intervals (note the depth range of the interpolated SSPs becomes restricted to the shoal and deep limits of the SSPs). For each level, a linear interpolation is performed between sound speed values of the profiles at the desired time.

An example of both adaptive and rigid interval thinning is shown in Figure 7. The interval of the adaptive set is first less than that of the rigid, but a redundant comparison then relaxes the interval beyond that of the rigid. All the while, the rigid intervals never change, regardless of the results of the SSP comparisons. Also shown in Figure 7, is the "control" procedure—the

original, high-density data set is a record of the evolving water-column, and represents "the truth" (or at least the truth to the best that we can measure it), therefore the high-density data set is used to assess the sounding depth bias incurred between sampling intervals of the thinned sets. This is accomplished by performing SSP comparisons for each cast in the original, high-density set to the cast previous in time of the thinned sets. The maximum depth bias associated with each SSP comparison conducted during the control procedure is retained.
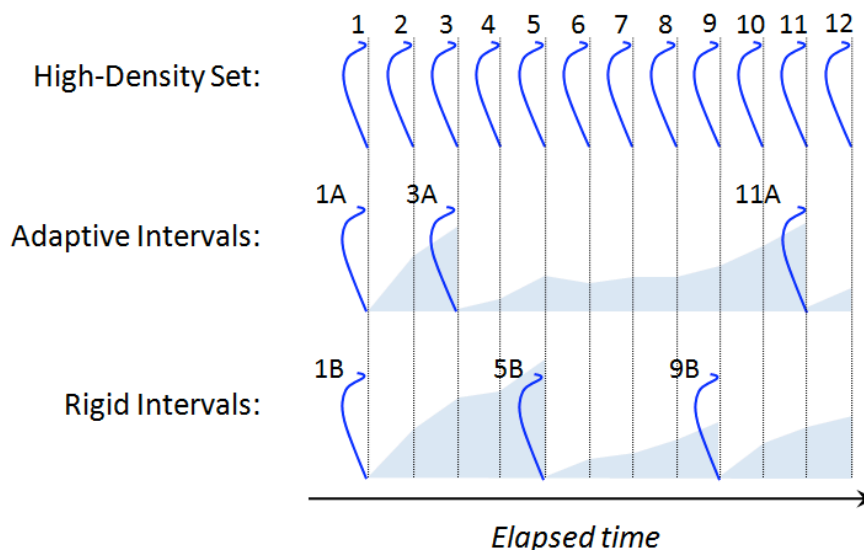

Figure 7. The control procedure of adaptive and rigid interval thinned sets of SSPs, "truthed" by comparison to the original, high-density set. The trend of sounding depth biases resulting from the comparisons of the thinned SSPs versus those of the high-density set are depicted in light blue shading. An SSP number denoted with an "A" or "B" are interpolated SSPs nearest in time to that number of the high-density set. The SSP comparisons are as follows: 1A→2,3; 3A→4-11; 11A→12 (adaptive), and 1B→2-5; 5B→6-9; 9B→10-12 (rigid).

A final thinned set is generated as an idealized, best possible solution, also known as the "Goldilocks" water mass (Beaudoin, 2008). Unlike the adaptive and rigid interval thinning regimes, the best solution draws upon existing SSPs from the high-density set (temporal interpolation is not necessary), and retains only those SSPs necessary to stay within set limits of error tolerance. Thus, the absolute minimum amount of casts are retained to meet uncertainty requirements, and the water mass is sampled to an extent that is "just right" (hence the "Goldilocks" reference). This solution, depicted in Figure 8, can only be determined in hindsight, and will be instrumental in judging the effectiveness of the adaptive and set interval thinning.
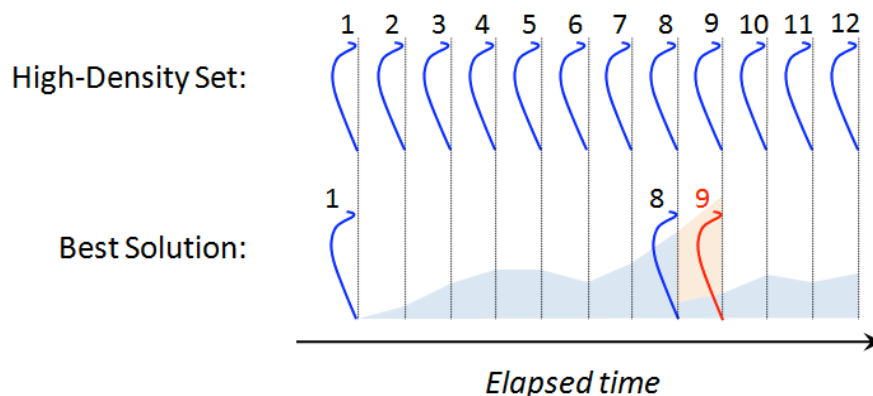
Figure 8. The best solution technique meets uncertainty requirements with the least amount of casts possible. The trend of sounding depth biases resulting from the comparisons of the best solution SSPs versus those of the high-density set are depicted in light blue shading. SSP 1 is retained, and compared to each successive SSP of the high-density set, until uncertainty requirements are exceeded, as they were in the comparison of SSP 1 vs. 9. Thus, SSP 8 is retained, and comparisons continue in this manner. It is possible to have a best solution with only 1 SSP, if all the successive comparisons meet requirements. It is also possible for the best solution to exceed requirements, as would be the case here if SSP 8 vs. 9 were not within requirements. This is unlikely because the SSPs are close in time and space, however it does happen. In this example, only two of the original 12 casts are required to maintain sounding accuracy.

Results

The outer beam sounding depth bias for each of the SSP comparisons that encompass the control procedure can be plotted as a function of time elapsed, on the original sawtooth diagram for each thinned set. In this way, the sounding depth bias incurred during the intervals between SSPs can be viewed, displaying what is essentially a record of variability that was "missed", or not necessary for capture.

Demonstrations of the thinning regimes as applied to the high-density data from Portsmouth are depicted in Figures 9 and 10, for June 14[th] and 21[st], respectively. In addition to the best solution and the algorithm, the data was thinned by way of fixed intervals of 15 and 60 minutes. The initial interval of the algorithm was set at 15 minutes, to match the starting condition of one of the fixed intervals. The upper bound of the algorithm was set to 30 minutes.

In both Figures 9 and 10, the high-density set (a) is displayed for comparison purposes, and the thinned sets are the best solution (b), the algorithmically-derived, adaptive sampling intervals (c), and rigid 15 minute (d), and 60 minute (e) sampling intervals. The incurred outer beam sounding depth biases computed from the control comparisons are plotted as a function of time elapsed for each thinned data set (b-e). Like the SSP comparisons of the thinned data sets, note that the incurred outer beam sounding depth biases are also an absolute value.
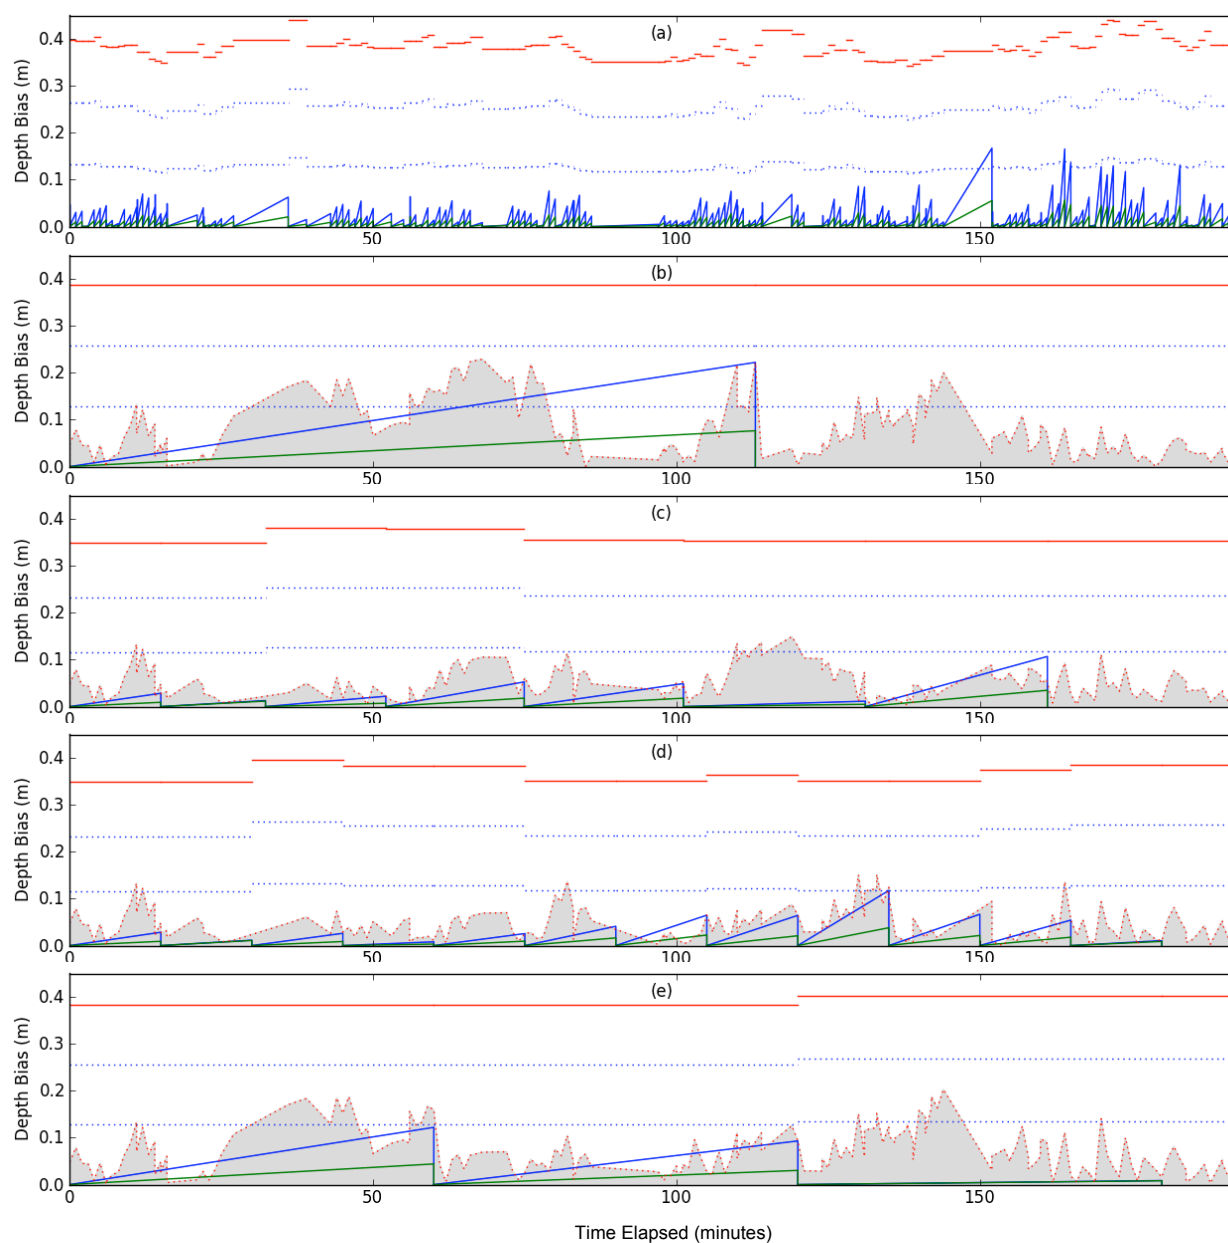
Figure 9. Sawtooth diagrams of SSP comparisons generated from various thinning techniques of 173 high-density SSPs (a) acquired up to 8 km offshore of Portsmouth Harbor (6/14/2012). "Best solution" (b) thinning results in 2 casts, algorithmic thinning (c) results in 8 casts; 15 (d), and 60 (e) minute rigid intervals result in 13 and 4 casts, respectively. Outer beam depth biases from control comparisons with (a) are also plotted (grey shading). All SSP comparisons were generated with a 1.0 m draft over a 140°(+/-70°) swath, and depict both outer beam (blue) and rms (green) refraction error. NOAA maximum allowable error (red) has been divided into thirds (dotted) to denote algorithm bounds.

Figure 10. Sawtooth diagrams of SSP comparisons generated from various thinning techniques of 184 high-density SSPs (a) acquired up to 8 km offshore of Portsmouth Harbor (6/21/2012). "Best solution" (b) thinning results in 22 casts, algorithmic thinning (c) results in 21 casts; 15 (d), and 60 (e) minute rigid intervals result in 12 and 3 casts, respectively. Outer beam depth biases from control comparisons with (a) are also plotted (grey shading). All SSP comparisons were generated with a 1.0 m draft over a 140°(+/-70°) swath, and depict both outer beam (blue) and rms (green) refraction error. NOAA maximum allowable error (red) has been divided into thirds (dotted) to denote algorithm bounds.

As observed in Figure 9, very low variability was observed on the first day, as indicated by the primarily redundant SSP comparisons amongst the high-density data set in (a). This is illustrated further in the best solution (b), wherein it is demonstrated that only 2 casts were actually necessary to avoid depth biases deemed critical. The adaptive intervals (c) relax until the algorithm upper bound is achieved, requiring a total of 8 casts, and with induced depth bias between the intervals primarily redundant. The 15 and 60 minute fixed interval regimes consist of 13 and 4 casts, respectively.

As observed in Figure 10, the second day showed conditions significantly more variable. Even the high-density set (a) cannot always maintain uncertainty requirements, nor can the best solution (b), which required 22 casts. The adaptive intervals (c) show periods of gradual interval relaxation, that were interrupted when critical magnitudes of depth bias were detected at 15, 78, 118, and 143 minutes of time elapsed, and the interval reduced accordingly at each occurrence to maintain uncertainty requirements. Steady conditions were achieved periodically and the sampling interval retained. Note that the initial interval of 15 minutes (chosen to match the 15 minute fixed interval thinning) was a poor selection, considering that time correlated with the transit through Portsmouth Harbor and the salinity gradients therein.

Plotting the control SSP comparisons as in the manner of Figures 9 and 10 is useful for visualization, but a more quantitative assessment of the individual performances of each thinned data set can be achieved by tallying and plotting the magnitudes of the control analysis outer beam depth biases into a histogram. The standard deviation of these values is also computed, and included within the histograms. Algorithm bounds of redundancy and error tolerance are also plotted on the histogram for a nominal depth of 30 m. In addition, the outer beam depth biases associated with the control analysis are also classified according to the algorithm logic discussed previously (i.e. "Relax", "Steady", and "Panic"). These classifications are tallied, normalized, and plotted into bar-stacked histograms for each control analysis. Both types of histograms, for both days of data collection in Portsmouth, are displayed in Figures 11 and 12.

Both histograms of Figure 11 reflect the low variability observed on June 14[th]. The outer beam depth bias standard deviations of each thinned data set are largely irrelevant, because the comparisons of each thinned set were predominantly redundant. During such conditions of low variability, minimization of casts is of primary interest. The 8 casts required by the algorithm was less than the 13 casts required by the 15 minute fixed interval, yet more than the 4 casts required by the 60 minute fixed interval. Note that if the algorithm had been set with a higher initial interval, or upper bound, this would have resulted in fewer casts.

Both histograms of Figure 12 show that the algorithm reacted effectively to the high variability experienced on this day, by tightening the sampling intervals and acquiring more casts than the rigid interval thinned sets. Thus, it is the adaptive sampling regime that most closely resembles the best solution, in terms of the standard deviation of the maximum depth bias, and total number of casts needed.

Figure 11. A histogram of the magnitudes of the outer beam sounding depth biases (left) and the algorithmic classifications (right) for the control analysis SSP comparisons for each of the various thinning regimes, in addition to the successive SSP comparisons of high-density SSP data acquired up to 8 km offshore of Portsmouth Harbor (6/14/2012). Left side: the standard deviation of the biases for each regime is given in the legend, and algorithm bounds are plotted on the distribution for a nominal depth of 30 m, where redundancy occurs within +/-0.15 m (dotted blue), and error tolerance is exceeded at +/-0.45 m. Right side: the total number of casts contained within each set (listed in parenthesis) is shown as a measure of efficiency.
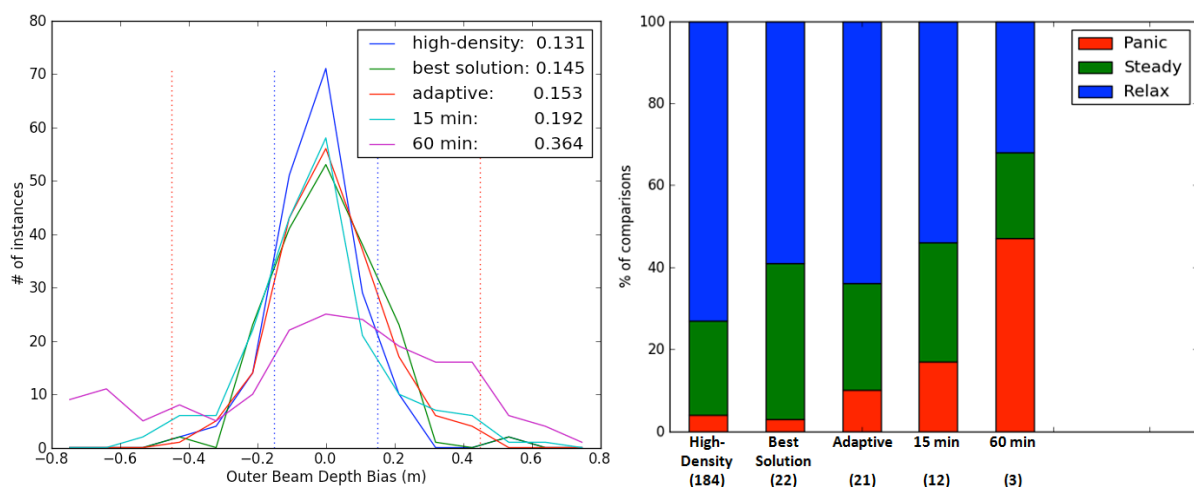

Figure 12. A histogram of the magnitudes of the outer beam sounding depth biases (left) and the algorithmic classifications (right) for the control analysis SSP comparisons for each of the various thinning regimes, in addition to the successive SSP comparisons of high-density SSP data acquired up to 8 km offshore of Portsmouth Harbor (6/21/2012). Left side: the standard deviation of the biases for each regime is given in the legend, and algorithm bounds are plotted on the distribution for a nominal depth of 30 m, where redundancy occurs within +/-0.15 m (dotted blue), and error tolerance is exceeded at +/-0.45 m (dotted red). Right side: the total number of casts contained within each set (listed in parenthesis) is shown as a measure of efficiency.

Note that the initial inputs of the algorithm were set exactly the same for the data sets of both days. Thus, even if the surveyor had no foresight or prior indication as to the expected oceanographic conditions on these days, the algorithm would have nevertheless reacted favorably. This is indicated by the very different results between these two days, which are

18

dictated entirely by the observed oceanographic variability. The adaptive sampling intervals derived in both data sets represent an effective illustration of the potential value of CastTime during survey operations, from standpoints of both efficiency and accuracy.

The thinning and control procedures were applied to several high-density data sets. Most often, the adaptive sampling intervals as dictated by the algorithm most closely resembled the best solution, though not always. There were instances when sudden variability, or variability that is short in duration, was captured more effectively by a rigid interval sampling regime, purely due to chance. This was occasionally observed in data sets short in duration. Data sets encompassing longer periods of acquisition offered more opportunity for the adaptive sampling to capture episodic variability, and then react favorably.

Field Trial

After testing within the controlled environment, the algorithm and its initial Python implementation (which represent the first version of CastTime), were tested in real-time onboard the *NOAA Ship Ferdinand R. Hassler*. In addition to assessing the performance of the algorithm in an actual survey scenario, other goals were to test the integration of the software with shipboard systems, and to solicit feedback from field personnel. The ship was underway from Norfolk, Virginia, on Sept. 6, 2012, working in the offshore approaches to Chesapeake Bay, as shown in Figure 13.


Figure 13. Survey area of the *NOAA Ship Ferdinand R. Hassler*, 9/6 – 9/9/2012, in the approaches to Chesapeake Bay up to approximately 70 km offshore of Cape Henry, with NOAA chart 13003 in the background. Soundings are in fathoms.

CastTime was integrated with the ship's MVP200 system, which features a tow fish, hydraulic winch, and conductor tow cable. The tow fish was equipped with an AML Sound Velocity and Pressure (SVP) sensor. The MVP controller software v2.47 was installed on the shipboard MVP PC and configured with the MVP200. A separate processing computer within the ship's survey lab had a prior installation of Python 2.5 and found to be suitable to execute CastTime.

After the installation and initial tests were performed, CastTime was used during survey operations over an uninterrupted period of more than 16 hours, from midday of Sept. 8[th] until the morning of Sept. 9[th], 2012. The ship's 3.8 m draft of the port and starboard multibeam transducers was set in CastTime. Although the ship planned lines based on usable multibeam swaths of 135°(+/-65°), the value of 140°(+/-70°) used predominantly during data analysis was retained. In this manner, there is higher confidence in the outermost beams, should they decide to be used for data overlap. In addition, there is an extra buffer against exceeding accuracy requirements, and finally, an increased chance of experiencing each of the algorithm's response conditions during the very limited period of underway time for real-time testing.

An initial sampling interval for the algorithm was set at 15 minutes, to correlate with the nominal sampling interval the ship had been using (unless visual inspection of the acquired SSPs indicated higher variability, in which case the interval was decreased to 10 minutes). The survey lines were approximately 30 minutes in duration, thus the upper bound of the algorithm was set to 30 minutes, to ensure a cast was acquired at least once per line.

During the 16 hours of CastTime use, standard survey lines were conducted, and casts acquired according to the algorithmically derived recommendations. Occasionally the interval was adjusted manually, in order to acquire a cast just before or after a turn. Also the interval was sometimes adjusted slightly to ensure spatial extents were adequately characterized. Finally, on two separate survey lines during this period, casts were acquired every 5 minutes. These were not acquired as part of any algorithm decisions. Instead they were control lines to ensure there was no variability occurring within the water-column that was not being detected; essentially a "spot-check" of the adaptive sampling intervals. The SSPs acquired during the real-time testing, and the cast locations, are shown in Figure 14, and the correlating SSP comparisons are shown in Figure 15. Note the higher-density lines in Figure 15, which commence at 0 and 406 minutes of time elapsed, and each result in SSP comparisons deemed redundant.
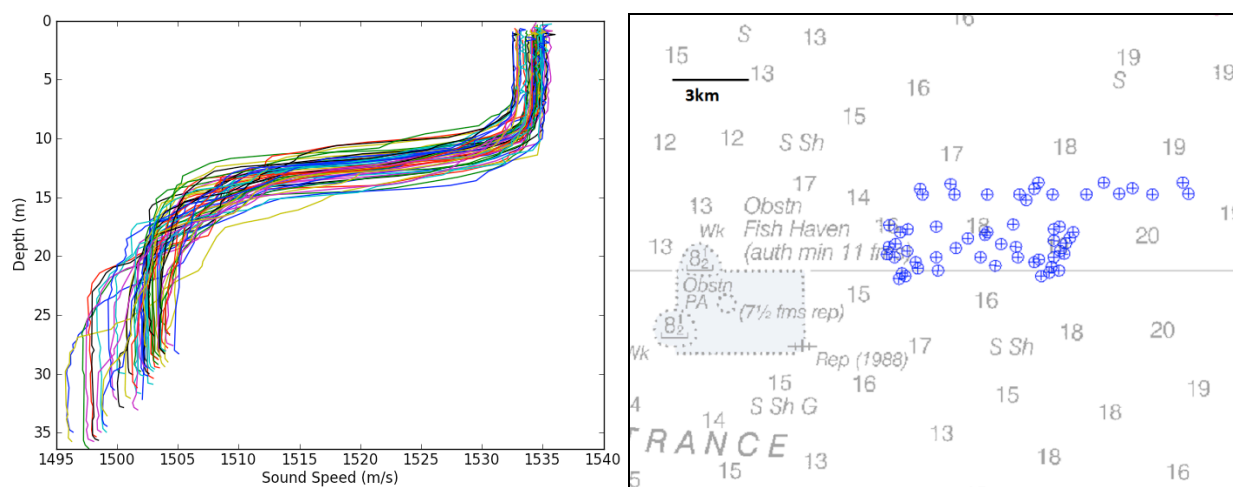

Figure 14. SSPs (left) acquired by *NOAA Ship Ferdinand R. Hassler* on 9/8 − 9/9/2012, and the cast locations (right) with NOAA chart 12200 in the background. Soundings are in fathoms.

Figure 15. Sawtooth diagram of SSP comparisons generated from real-time algorithmically-derived sampling intervals onboard *NOAA Ship Ferdinand R. Hassler* during survey operations in the approaches to Chesapeake Bay (Sept 8-9, 2012). Two high-density control lines were conducted at 0 and 406 minutes of time elapsed and showed redundancy. The remaining 44 casts have a mean interval and standard deviation of 22 +/-7 minutes. All SSP comparisons were generated with a 3.8 m draft over a 140°(+/-70°) swath, and depict both the outer beam (blue) and rms (green) refraction error. NOAA maximum allowable error (red) has been divided into thirds (dotted) to denote algorithm bounds.

The sounding depth biases incurred in Figure 15 are primarily a result of the variability in the depth of the thermocline in the water-column, observed in Figure 14. Excluding the redundancy associated with the high-density lines, the SSP comparisons of Figure 15 resulted in "Relax", "Steady", and "Panic" conditions 28, 11, and 4 times, respectively. The "relax" condition was predominant—during times of low variability and stable conditions, the interval would gradually increase, though with occasional comparisons deemed "steady". This trend is observed over large portions of the entire set, from 96-335 (~4 hours), 356-596 (4 hours), 616-838 (3.7 hours), and 887-1031 (2.4 hours) minutes of time elapsed. For those first three periods of time, the upper bound of 30 minutes was achieved. Each of these periods is ended by a panic condition, and subsequent interval reduction. The "panic" condition was enacted at 81, 335, 838, and 875 minutes of time elapsed. The occurrence at 596 minutes was not critical according to the algorithm parameters, but the interval was reduced by a manual adjustment, due to visual inspection of a rising thermocline. Ensuing SSP comparisons after the manual adjustment were primarily redundant, but in this case a conservative approach was taken.

There were no anticipated issues affecting data quality, due to the redundancy noted in the SSP comparisons of the high-density control lines, and ensuing review of the acquired multibeam data during post-processing and quality control procedures confirmed this. Upon visual inspection of the processed soundings and gridded bathymetry, no refraction-induced error was observed, indicating that the 44 SSPs acquired as part of the adaptive sampling regime were sufficient.

Figure 15 has a total of 44 SSPs, with a mean sampling interval and standard deviation of 22 +/-7 minutes. Note the casts acquired as part of the high-density lines were not included in those numbers, as those were acquired as part of this research and not part of standard survey operations. By comparison to the algorithmically-derived set of 44 casts, standard 10 or 15 minute rigid sampling intervals would result in 97 or 65 total casts, respectively, over these 16 hours. Thus, the algorithmic sampling saved the ship 53 or 21 casts, respectively, during this time. There may be times, and certain survey locations, where the algorithm may actually result in a greater number of casts than standard 10 or 15 minute sampling intervals, but generally, working offshore, a scenario of the type exhibited in Figure 4-27 would likely be a common one and result in significantly fewer casts acquired. A cost analysis was not performed as part of this

research, but one can imagine extrapolating these results over an entire field season would result in a significantly reduced workload to the MVP, extended life to the system and its components, as well as reduced expenditure necessary for repair and replacement of those items.

Future Work

The initial version of CastTime (used in the case studies and field trial discussed in this report) was command-line driven, though CastTime has since transitioned to a graphical user interface. Continued feedback from survey personnel will help tailor and customize the interface functionality towards the needs of the user.

There are numerous ideas on how to improve the algorithm. For example, the temporal component of the algorithm can be improved with a true time series analysis of all the casts acquired, in contrast to only the most recent two. In addition, a spatial element of the algorithm could be incorporated, such that CastTime can "learn" from the geographical distribution of the water-column variability exhibited within the casts. Algorithm parameters and settings could be easily adjustable, to match a user's preferences. Error tolerances will also be adjustable, such that CastTime could be used by any hydrographic office, in addition to the NOAA Office of Coast Survey.

The MVP controller software supports the ability to remotely monitor data from the MVP tow fish, winch, and navigation in real-time. With incoming data feeds of this type, there could be other means of water-column variability assessment, in addition to the acquired SSPs. For example, when the tow fish is at the docked position, changes in the surface sound speed could be utilized within the CastTime algorithm, in lieu of visually monitoring a transducer sound speed input, which is a common current practice. Furthermore, in addition to sampling frequency, there can be a protocol to determine an optimal sampling depth, an option of great value for those vessels working in deep waters. Typically, the variability in the uppermost layers of a water-column are more pronounced than at depth, thus it is not always necessary to sample the entire water-column. This is desirable information, as casting the full ocean depth can be time consuming, even for an underway profiler, and also adds significant workload to the tow fish, winch, and cable. With a real-time feed from the tow fish of the oceanographic conditions experienced, a cast termination protocol can be set within the algorithm, such that the cast is terminated when a certain degree of stability is achieved. With a CTD sensor on the tow fish, these capabilities are only improved, as a cast could be terminated based on preset specifications of temperature and/or salinity.

Conclusions

The overall goal of this project was to develop the software infrastructure to enable a proof-of-concept study wherein an adaptive feedback controller adjusts MVP sampling intervals in response to changing oceanographic conditions. The CastTime software can help to remove much of the subjectivity behind the determination of sampling intervals, by providing

quantitative and objective estimates. Furthermore, the case studies and initial field tests suggest that CastTime can greatly minimize sampling redundancy, while also helping to ensure uncertainty requirements are met. The result is potential savings based on a reduced need for repair and replacement parts to an underway profiler. In addition, the automatic comparisons conducted with each acquired SSP may alert a surveyor to water-column variability, ensuring a prompt response and improved overall data accuracy.

A simple control algorithm has been proposed, though it is not yet fully tested nor optimized. The case studies stressed the importance of ensuring a proper initial sampling interval and upper bound of the algorithm. If these parameters are not configured properly, the algorithm can fail. Furthermore, CastTime cannot entirely relieve a surveyor of the duty of water-column variability assessment. The algorithm must be adjusted manually in order to counter the effects of spatial aliasing and sudden changes in the water-column, neither of which the algorithm offers any defense against. It will, however, react to those sudden changes, if they are detected, and it can alert the surveyor to variability that might otherwise go unnoticed. Conversely, it may provide the justification for a surveyor to feel safe in relaxing a sampling interval.

CastTime is not intended to supersede or replace existing tools. A surveyor should always trust their intuition and best practices, as this human element may never be entirely captured within an algorithm. CastTime should be used as an aid to the surveyor in order to make the best possible decisions.

Acknowledgements

References

Beaudoin, J., 2008, Real-time monitoring of uncertainty due to refraction in multibeam echo sounding. Shallow Survey 2008, Portsmouth, NH.

Beaudoin, J., 2010a, Estimation of sounding uncertainty from measurements of water mass variability. PhD. dissertation, Department of Geodesy and Geomatics Engineering, Technical Report No. 271, University of New Brunswick, Fredericton, New Brunswick, Canada, 232p.

Beaudoin, J., 2010b, Real-time monitoring of uncertainty due to refraction in multibeam echo sounding. The Hydrographic Journal, v. 134, p. 3-13.

Beaudoin, J., Floc'h, H., Furlong, A., Lurton, X., and Smyth, S., 2011, Streamline sound speed profile pre-processing: case studies and field trials. U.S. Hydrographic Conference, Tampa, FL, USA, Apr. 25-28. *Conference Proceeding*.

Furlong, A., Beanlands, B., and Chin-Yee, M., 1997, Moving vessel profiler (MVP) real time near vertical data profiles at 12 knots. Proceedings of the Oceans '97 Conference, Halifax, Canada.

Lurton, X, 2010, An introduction to underwater acoustics: principles and applications, 2nd Edition, Springer-Praxis Publishers, Chichester, UK, 680p.

NOAA, 2012, National Ocean Service Hydrographic Surveys Specifications and Deliverables, 181p.

Rudnick, D.L., and Klinke, J., 2007, The underway conductivity-temperature-depth instrument. Journal of Atmospheric and Oceanic Technology, v. 24, p. 1910-1923.