Article

A Time Comparison of Computer-Assisted and Manual Bathymetric Processing

By Brian Calder, Center for Coastal and Ocean Mapping and Joint Hydrographic Center, Chase Ocean Engineering Lab, University of New Hampshire, USA and Shep Smith, NOAA Ship Thomas Jefferson, Marine Operations Center (Atalantic), USA

We describe an experiment designed to determine the time required to process Multibeam Echosounder (MBES) data using the CUBE (Combined Uncertainty and Bathymetry Estimator) [Calder & Mayer, 2003; Calder, 2003] and Navigation Surface [Smith *et al.*, 2002; Smith, 2003] algorithms. We collected data for a small (22.3x10^6 soundings) survey in Valdez Narrows, Alaska, and monitored person-hours expended on processing for a traditional MBES processing stream and the proposed computer-assisted method operating on identical data. The analysis shows that the vast majority of time expended in a traditional processing stream is in subjective hand-editing of data, followed by line planning and quality control, and that the computer-assisted method is significantly faster than the traditional process through its elimination of human interaction time. The potential improvement in editing time is shown to be on the order of 25-37:1 over traditional methods.

Introduction

In order to support improved efficiency of hydrographic data processing, the Center for Coastal and Ocean Mapping & Joint Hydrographic Center have been developing new techniques for auto-matic data processing and manipulation. The results are the CUBE (Combined Uncertainty and Bathymetry Estimator) [Calder & Mayer, 2003; Calder, 2003] and Navigation Surface [Smith *et al.*, 2002; Smith, 2003] algorithms. These techniques are novel in that they generate a surface, meant to represent the best information available about the true depth in the survey area, rather than by selecting individual soundings to represent the summary of the survey.

The two algorithms are complementary. CUBE processes Multibeam Echosounder (MBES) data from raw soundings, using a forward predictive error model [Hare *et al.*, 1995] to quantify the expected variances of each sounding about the true depth. The ultimate goal is to determine the true depth of the water at any point in the survey area; the error estimates allow CUBE to determine how to optimally combine the information inherent in each sounding into a best estimate of this depth, and how much leeway to allow before deciding that a sounding is inconsistent with those previously observed. Each group of self-consistent (but mutually inconsistent) soundings is tracked separately as a depth hypothesis to avoid cross-contaminating the true depth estimate with outliers; after all currently available data is assimilated, the algorithm

uses metrics on the hypotheses to choose the one it considers most likely to be the true depth, and reports it to the user along with the associated metrics. The user's job is then to inspect the algorithm's reconstructions, and determine what went wrong in the cases where the noise level was sufficient for the algorithm's disambiguation engine to choose the wrong hypothesis. The algorithm is very robust to typical MBES noise sources, and hence the area of survey (and volume of data) with which the user has to interact may be significantly reduced since the user only needs to take action in the limited cases where the algorithm fails, and only has to deal with the soundings that are actually causing the observed problems. This can lead to significantly reduced processing time. The algorithmic approach also ensures objectivity in determining which soundings are consistent, in contrast to current subjective methods. The result of the algorithm is a collection of estimates of 'true depth' at points across the survey area along with quality metrics,

which can then be combined to form a surface representation of the depth in the survey area.

The Navigation Surface takes this surface as the foundation of a bathymetric database, along with any other sources of appropriate information, e.g., shoreline depths, point detached positions on obstructions or rocks, other acoustic or non-acoustics depths, etc. From this high-resolution database, combined with hydrographic knowledge and best practice, a 'safe' surface intended for navigational use can be constructed automatically at a scale appropriate for the intended product. Automatic processing of most of the data through CUBE should lead to a significant time advantage for field units; use of surfaces as a database allows the Navigation Surface to carry out automatically many cartographic tasks that were previously hand-driven, with downstream benefits in time, simplicity and applicability of data products.
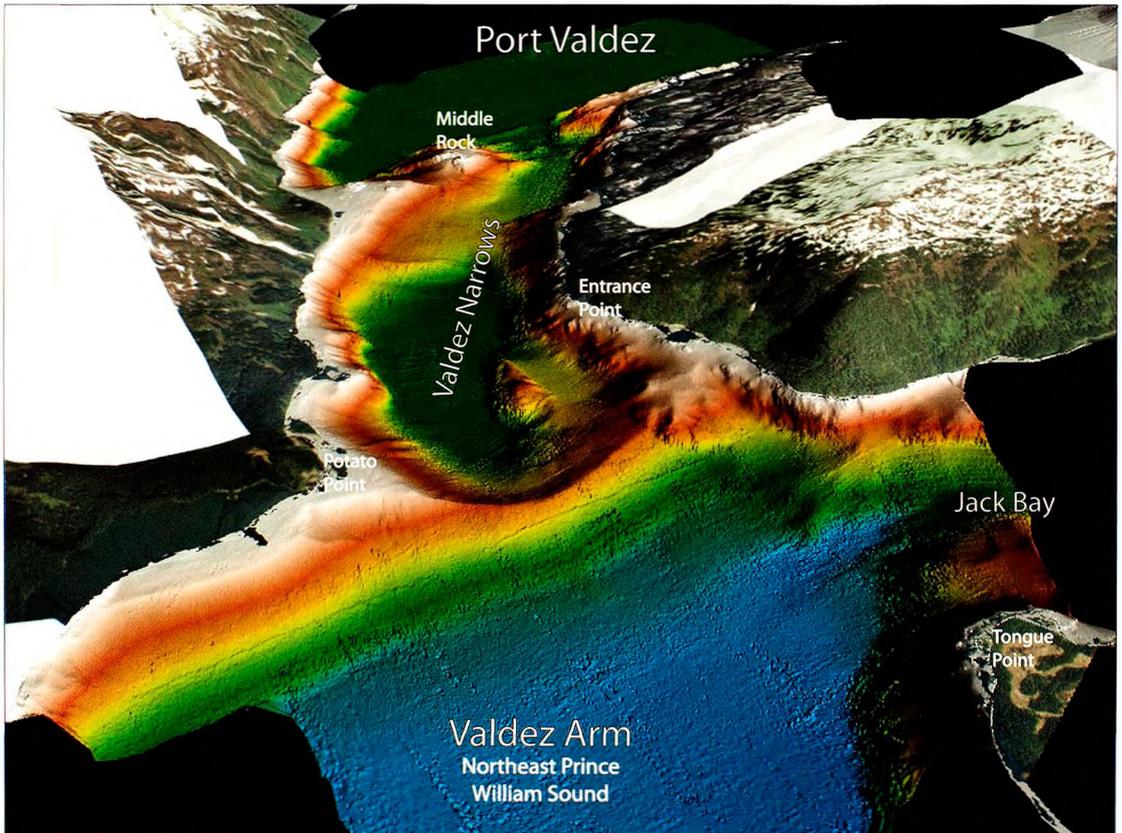

Figure 1: Valdez Narrows, Alaska. Perspective view showing processed bathymetry, topography and features. The topography was generated by TINing the contours in the ENC of the area, and is overlaid with georeferenced orthophoto imagery, which was also used for shoreline. The bathymetric composite includes MBES data, Vertical-Beam Echo Sounders (VBES) data, shoreline positions and points data for targets, combined into a single surface. Note: vertical exaggeration here is 1:1

Use of surfaces as a processing product is a departure from usual hydrographic practice. Their use has significant implications for the hydrographic processing chain as it is currently implemented, and at the same time opens up possibilities for new survey methods that are not supported by current protocols. For example, if the automatic process can deal with the majority of the data, then the operator's primary task is to verify that the algorithm produced the correct result, rather than inspecting every sounding. Or, since the CUBE algorithm can update its estimate of the true depth as new data is gathered, the correct response to a section of data with many observed outliers may be simply to run the MBES over the area again, rather than have an operator painstakingly and subjectively decide where the bottom really is. The additional data effectively improves the signal to noise ratio, hopefully to a level where the algorithm can correctly determine the depth of the true bottom,

rather than being confused by the outliers. Having an inherent measure of the quality of the data (through the simultaneously tracked uncertainty) may also allow us to optimise survey effort through an adaptive survey approach, directing effort to areas which are poorly covered or of more significance, and away from simpler areas where the data may be less dense, but still meets the required standards.

A question of immediate interest is to determine the potential benefits inherent in adopting the new methods. To quantify this advantage, we took part in a survey conducted by the National Oceanographic & Atmospheric Administration (NOAA) Ship *Rainier* in Alaska; the survey took place in Valdez Narrows over five survey days from 13 September 2002. The area is a narrow channel serving Port Valdez and the Valdez Oil Terminal, characterised by a relatively shallow trough oriented roughly Northeast/Southwest between mountains, with deep, flat areas in Port Valdez and the approach regions to either end (Figure 1). The area has typical fjord morphology, with steep rock walls falling rapidly from shoreline to the depths (Figure 2). Survey limits were chosen to approximately match the bounds of a 1:20,000 chart insert depicting Valdez Narrows (National Ocean Service (NOS) Chart 16707).

Our aims for the experiment were threefold. First, we wanted to test a prototype implementation of a real-time integrated CUBE/Navigation Surface and refine a method for its use in a field environment. Second, we wanted to determine how much faster we might expect processing with a CUBE/Navigation Surface system to be. To support this aim, we had to add the third: to gather evidence for how much time was actually used during the traditional processing associated with the survey, and in what categories.

We describe below the data collection for both survey data and time keeping information, and then investigate the distribution of the time expended during the standard processing procedures, contrasting it with that for the automatic method, using some extensions built for this particular project. Then, we conclude with some observations on the effort involved in each method from the operator's point of view, and some lessons learned from using the automatic method in the field environment.

Figure 2: Valdez Narrows, AK. This picture was taken early on the first day of survey, illustrating the rock-faces on the southwest side of the narrows. This type of sheer rock faces, punctuated by waterfalls empting fresh water on top of the channel made for difficult survey conditions

Experimental Methods

Data Collection

The NOAA Ship *Rainier* operates independently in the field for much of the year. In order to define the hydrographic limits for charting in to the shoreline, the ship is host to six survey launches, typically around 30' in length. The ship carries a deep-water MBES system, and four of the launches are equipped with one or more MBES in addition to a VBES. It is also possible to survey with a hull-mounted sidescan sonar on a launch if required. The launches are equipped with motion sensors, GPS receivers and auxiliary equipment as appropriate for the type of sonar in use. The different MBES systems available allow the launches to survey in all water depths that are significant for hydrography.

Data were collected from three launches. *RA2* was used for item investigation, single-beam buffer lines and shoreline verification; *RA5* operated a Reson 8101 MBES (101x1.5° beams over 150° swath, 240 kHz) for shallow water around the edge of the survey; *RA6* operated an Elac-Nautik 1180 MBES (126x1.5° beams over 154° swath, 180 kHz) for the deeper areas in the mid-channel, Valdez basin and approaches. Both MBES boats used a POS/MV 320 for attitude, navigation and orientation, supplied with differential correctors received using a Trimble DSM212 GPS system. Auxiliary information from aerial photography was available for comparison and shoreline work. Tide control was established from Valdez, and differential GPS correctors were received from the beacons at Potato Point or Cape Hinchinbrook, as appropriate for reception. The depths range from shoreline to approximately 300m, with slopes typically 50-60° and occasionally higher (in one or two cases, there appear to be overhangs on rock faces).

The survey proceeded with daily cycle operations where data was collected by the launches during the day and then downloaded to the *Rainier* for processing. Line plans were developed daily, and (after the first day) based on feedback from preliminary processing of the previous day's data. Because of its shallower draft, *RA2* was used to provide reconnaissance for the shallow-water multibeams, establishing where it was safe to take the equipment prior to deployment. Line plans were implemented in Hypack, which was also used for single-beam acquisition and detached positions for rocks and obstructions. MBES data was recorded in XTF format using Triton Elics ISIS.

Once the launch crews started to off-load data, they were requested to fill out a form that detailed the time taken in each stage of the traditional process (appendix A). The form was broken into sections corresponding to the standard processing chain, and into interactive and non-interactive time. This distinction is important because non-interactive time can be improved with better hardware; interactive time can only be improved with more efficient tools, methods and algorithms. The crews were briefed beforehand about the intentions of the timekeeping and the goals of the experiment, and that the information gathered would only be reported in aggregate, and could be entered anonymously if they preferred. Involvement in the timekeeping was voluntary, but unanimously supported. The time sheets were filled in manually, and then transcribed to a spreadsheet for analysis.

Throughout the CUBE processing path, we recorded the time taken by the algorithm to process the data, and the time used in interactive inspection and problem remediation. We did not include any component for processes that would be the same in both processing paths (e.g., inspection of attitude or navigation data, target, detached position and shoreline processing), and could not split the QA/QC task from the data cleaning task: for CUBE, these are essentially the same process.

Data Processing

For the ship's crew, the standard NOAA Administration processing chain was followed. Once data was off-loaded, a check-sheet was generated for each launch, listing all of the lines run that day and providing control points to check that each stage of processing was completed. This physical form follows the data and represents authority to process, hence acting as an interlock. Based on these forms, the XTF data were then converted into HIPS/HDCS (CARIS/HIPS [Hydrographic Information Processing System]) format, and the attitude and navigation data were inspected for anomalies. Once cleaned, standard correctors were applied for attitude, tide and refraction, the data were merged and then filtered for a standard angle gate of ±60° and quality

| Day | System | | Total |
| | RA5 (8101) | RA6 (1180) | |
	[soundings]	[soundings]	[soundings]
2002-256	0	362,277	362,277
2002-259	10,217,996	541,466	10,759,462
2002-260	3,447,169	403,761	3,850,930
2002-262	7,192,140	0	7,192,140
2002-264	0	112,682	112,682
Total	20,857,305	1,420,186	22,277,491
%	93.63	6.37	100.00

Table 1: Summary of data collection volume by system and day (in soundings). Note that this is raw data volume gathered, rather than data remaining at the end of the survey

flags as appropriate to the MBES in use. Data were then cleaned in line (swath) mode, before being used for DTM generation (DTMs are used for QA and line planning). Vertical-beam data was converted into HIPS/HDCS format, corrected for tide and sound speed and then merged, inspected and cleaned. Targets for detached positions and shoreline verification were taken into Pydro [Riley et al., 2001] for processing and future correlation with known features. Data volumes and distribution for MBES data are summarised in Table 1.

In order to disturb the processing flow of the ship as little as possible, we arranged with the processing teams to be notified when the MBES data had been merged. We then copied the HDCS data from the ship's server onto a stand-alone disc attached to the commodity PC (Pentium 4 1.6GHz, 1Gbyte RDRAM memory, LaCie FireWire external hard-disc) that we used for processing. Flags in the data, if any, were removed before further processing, except those attached to attitude or navigational data.

To support the day cycle data collection, we maintained two sets of MapSheets (CUBE's internal data structure representing the data), one for 'per-day' work, and one 'cumulative', representing all of the data collected so far. At the start of each day, the cumulative Map-Sheets were used to initialise the per-day set. The data for the day were then run through the CUBE process using the per-day MapSheets. An inspection stage followed, using GeoZui3D [Ware et al., 2001] to visualise CUBE's output surfaces and CARIS/HIPS 5.3β to inspect and

modify the data. Using the surfaces as a guide, we determined areas of data where the depth reconstruction was dubious, e.g., Figure 3, where an outlier point is the only available data due to data sparseness in deep water. In this case, CUBE currently assumes that any data is better than none, and hence reports the spike since there is no sufficiently local evidence to the contrary. Spatial (area) mode editing was used in these cases, using subset tiles [Gourley & DesRoches, 2001] to track the areas of the data that had been inspected. ('Subset mode' is a CARIS/HIPS methodology where all of the data in a specific area are presented to the user in 2D and 3D environments (fully georeferenced) for
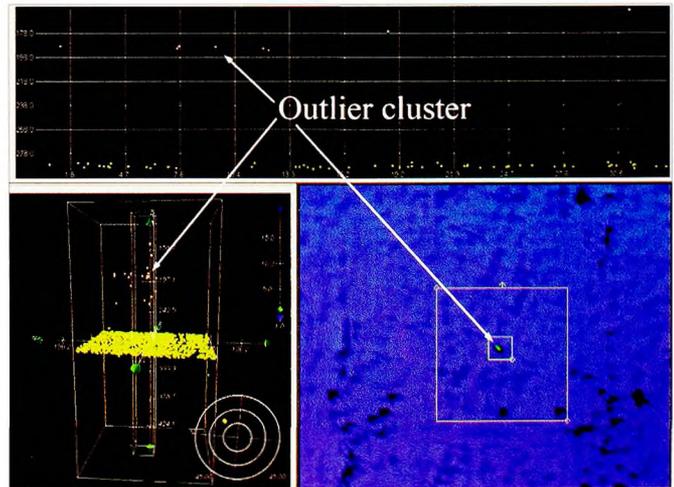

Figure 3: Sparse data in the deep southwest section of the survey area. Data here is sparse, so there are regions where there is not enough evidence from the neighborhood to overcome rogue soundings. In the belief that any data is better than no data, CUBE makes the only available reconstruction
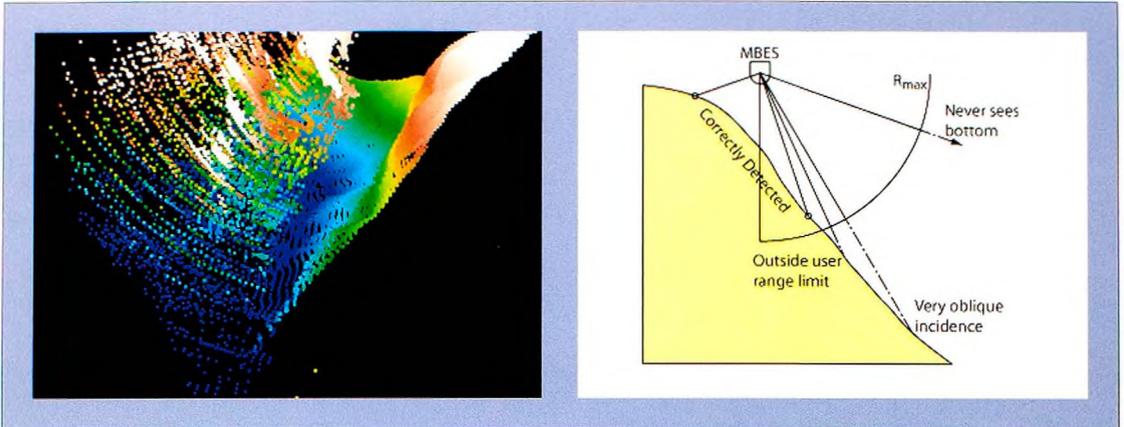
Figure 4: Data from steeply sloping area in northwest of survey area (left) and schematic of data detection difficulties (right). Significant downhill slope and a fixed maximum range at the sonar processor means that a multibeam line at the top of the slope cannot detect the bottom and the top of the slope simultaneously, and suffers from very shallow grazing angles at extreme range thus detecting inconsistent returns. This difficulty puts noise into the dataset where the real data is sparse due to increasing depth

inspection and editing; 'subset tiles' are a method of breaking a larger area into manageably sized chunks with a method to colour these tiles in order to indicate whether the hydrographer considers the data to be 'complete', 'partially processed' or 'uninspected'.) We attempted wherever possible to edit only the points that were causing the observed problem, rather than editing all of the data in each subset that we inspected. Our goal in this processing is not to clean the data in the traditional sense, but to improve the signal to noise ratio so that CUBE is no longer confused by the outliers.

It quickly became evident that the MBES systems were having significant difficulties in the regions of high slope. This is primarily a geometric problem: the slope is such that the outer beams on the downhill side either graze the surface at such a shallow angle that the bottom detection is very difficult, or do not receive any return within the range scale required to make the beams on the uphill side of the MBES correctly detect the bottom. Typical data is shown in Figure 4. This poses a difficulty for CUBE (as well as human operators) since the outlier points do not satisfy the normal properties of outliers, which tend to occur at random and moderately sparsely with respect to the true data. Because of the generation mechanism, these outliers appear to cluster strongly in space, and occur where data from downhill passes tends to be sparser (typically, generated by the Elac-Nautik 1180 rather than another pass of the Reson 8101). In many instances, it was

also not immediately obvious for a human operator where the true bottom was.

This problem is line oriented, and best resolved in line mode. We attempted two different remediation methods for the problem, and recorded the processing time used in each one separately. In each case, we started with an original dataset, so that the effects of the different methods were not cumulative. That is, each processing method is independent of the other.

First, we determined by inspection which lines had port side up-hill, and which starboard side, and applied an asymmetric angle gate to the lines of 60° on the uphill side, and 45° on the down-hill, implicitly assuming that slopes are at worst 45°. This method ensures that the majority of the outliers are removed, although we still observe some problems because no MBES can achieve effective bottom detection at a grazing angle of a few degrees. The alternative is to remove more data (say to 30° for a 150° MBES, so that even at 45° slope, the outermost accepted beam is at a grazing angle of 15° as is usual for the outermost beam when fired at a flat seafloor). However, this would entail significant loss of coverage where it would not otherwise be justified, and we chose to retain more data, paying the cost in extra interactive processing time. This simple system illustrates a method using current tools and practices that is closest to the manual-processing path, although it is typically sub-optimal in terms of time expended and coverage achieved.
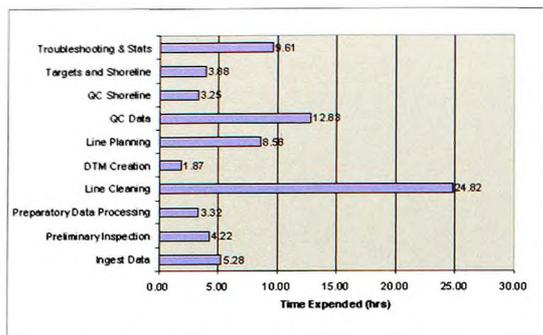
Figure 5: Summary of time expenditure during the survey effort. Some of the categories shown here are aggregated from the detailed data categories actually recorded (see appendix A)
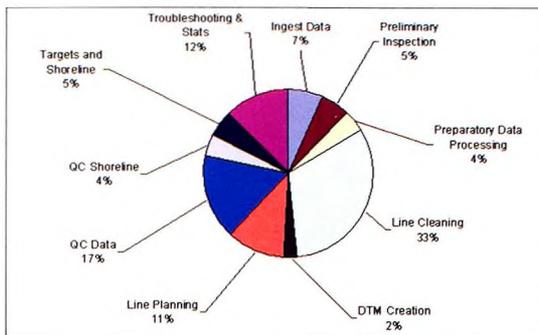

Figure 6: Summary of time expenditure during the survey field program. Majority of data processing time is consumed by line editing, followed by quality assurance/control for data and shoreline features/targets, and line planning

Second, we investigated use of the quality flags supplied with the Reson data (the primary cause of the observed problem, since the Elac-Nautik system was only used for deeper areas). For each sounding reported, the Reson 81-P sonar processor assigns a two-bit quality value indicating whether the data used to compute the sounding passed a 'brightness' and 'co-linearity' test. These test, respectively, that the backscattered energy was sufficient for the return to be real, and that the reconstructed depth does not depart from its neighbouring beams too significantly. Although these are typically only used in the traditional processing scheme to remove soundings that fail both tests, we found that in this case they were a very good indicator of problem data, and re-filtered the data so that only the points that passed both tests were retained. We then ran the pre-flagged data through CUBE, and proceeded with the normal (automatic) processing path.

After inspection and remediation, the day's data was run through CUBE again to assimilate the new data against the cumulative MapSheets. We then constructed a 'current best estimate' surface from the cumulative MapSheets to summarise the state of the survey. Finally, a composite surface was constructed using the Navigation Surface method. The source data consisted of the CUBE surface (where defined), VBES data, point targets with defined depths (e.g., landmass for islands, rocks, obstructions, etc.) and shoreline. The VBES data was gathered directly from the ship's processing stream, rather than being re-processed through CUBE. Point targets were honored in the data wherever they occurred (even if covered by MBES), and a standard GIS (MapInfo) TINing routine was used to form a surface between the sparse VBES

data, and to junction shoreline to VBES, and VBES to MBES [Smith, 2003]. This composite surface, with shoal-preserving down-sampling to a single resolution of 5m, provided for visualisation and overall summary of the progress of the survey.

Results

Manual Processing

We recorded data in detailed categories as shown in appendix A. For display, we have aggregated some of the categories with smaller time expenditure. Data download, conversion and check sheet generation have been aggregated as 'ingestion'; attitude and navigation data editing as 'preliminary inspection'; and tides, refraction, merge and pre-filtering as 'preparatory data processing'. All of the target and shoreline processing have been accu-
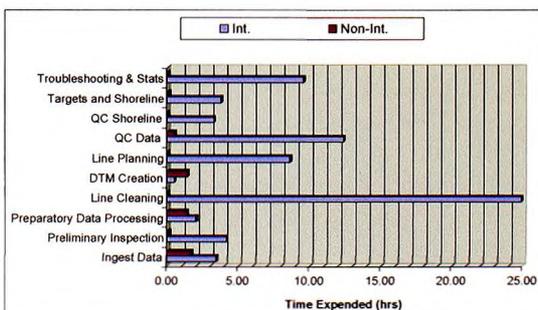

Figure 7: Comparison of Interactive and Non-Interactive time expenditure during the survey fieldwork. The significant lack of non-interactive time (a total of only 7 per cent of the total time expended) implies that our ability to improve the process through hardware alone is limited
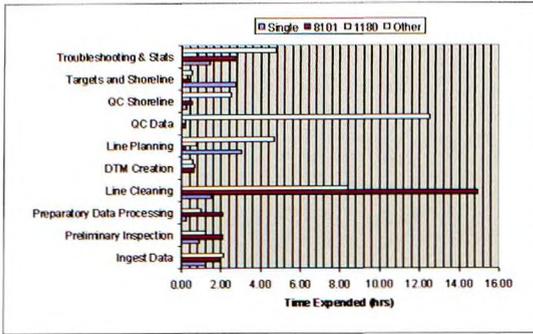
Figure 8: Expenditure of time per survey system. Since there are significant differences in sonar repetition rate and data density, it might be expected that there would be a corresponding difference in time expenditure. In fact only a small difference is observed, although it is in the expected direction
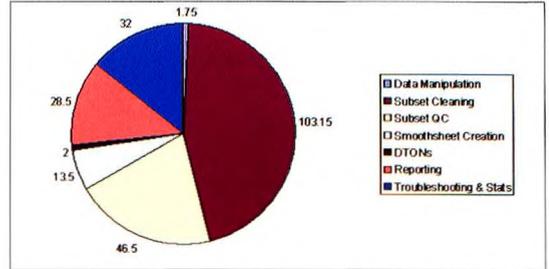

Figure 9: Post-survey expenditure of time (hrs). The vast majority of time is taken up in subset cleaning and QC (a total of 66 per cent). Note that 'DTON' is a 'Danger to Navigation' report (to report a target which is not charted, or is significantly different from what is charted, and may be an immediate danger to surface navigation typical in the area), and 'subset' refers to CARIS/HIPS subset mode (i.e., spatially organised by area) edit and inspection of data

mulated as 'target & shoreline', and all troubleshooting, file management, statistics generation (i.e., for usual purposes, rather than this experiment) as 'troubleshooting & stats'; other categories are reported as they were recorded.

The summary of overall time expenditure during the survey is shown in Figures 5 and 6. Times are shown in person-hours assigned against the task. Although, not surprisingly, the line editing (i.e., editing in swath oriented mode) takes up the majority of the effort, the proportion of time taken by line planning and quality control is also significant. This perhaps corresponds to the mode of operation in this example, which exhibits a classical 'plan, do, review' model, consistent with day-cycle operations. While this mode of operation is limited to a system where platforms gather data during the day and have sufficient down-time to re-cycle the information overnight (rather than running 24hr operations), it is a highly efficient method of making best use of available resources and operating in the field for many months at a time.

The balance between non-interactive and interactive time expended is illustrated in Figure 7. The vast majority of time is spent in interactive activities at the computer, where processing is limited by operator speed, rather than technology. In turn, this means that there is only a small reward in store for improving machine speeds, and that we must consider new techniques for approaching the problem, rather than fine tuning the current approaches.

Since the sonar systems used in the survey produce data in significantly different volumes and at very different rates, it might be expected that this would be reflected in the time taken to process the data. In fact, we find only a small effect, Figure 8, although it is in the correct direction, with the less dense, lower rate Elac data taking less time to process. We return to this in the discussion.

Finally, Figure 9 highlights the data processing subsequent to the survey field program. The single greatest expenditure of time is subset (area) based cleaning, followed by the QC of the data. Combined with the survey line editing, the total expenditure on editing of data is 128 hrs.

Automatic Processing
Automatic processing time was recorded as the time for interactive editing of the data to resolve issues found in the CUBE intermediate depth surfaces. A summary of the time expended is shown in Figure 10. (Note that the two different interactive methods are shown together for efficiency of presentation, but only one was used in each test; the non-interactive times are common between both methods.) We found that a run of CUBE through all of the data took approximately 31 min., with the time taken per day mostly a function of the amount of data gathered as might be expected. On average, this was approximately 12,000 soundings per second (snd s^{-1}), although the rate varies considerably with sounder and depth range and increases proportionately with processor and disc access
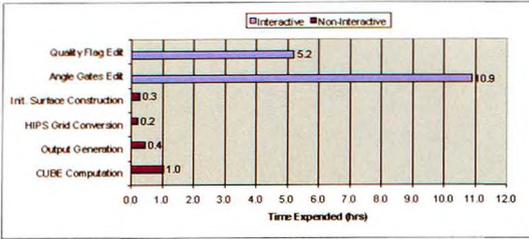
Figure 10: In-survey expenditure of time (hrs.) for CUBE-assisted processing. The experiment considered two different approaches to pre-filtering the data before CUBE-assisted processing and subsequent interactive editing. The interactive time for both are shown here (i.e., 'Quality Flag Edit' for quality-flag based pre-filtering and 'Angle Gates Edit' for asymmetric down-hill angle gate pre-filtering) for efficiency of presentation although only one was used for each run. The non-interactive times were the same in each case, and represent total times (e.g., for two runs of CUBE, etc.) as described in the text

speed. The cost of generating 'current best estimate' grids was relatively constant at approximately 13 min., although this increased slightly over the course of the survey as more of the area became active; conversion into HIPS weighted grids for inspection cost another 6 min. per run of CUBE.

For the simple asymmetric angle filter method, a total of 10.9 hrs were expended in interactive editing, using a mixture of line-oriented (swath) mode and area-oriented mode, as appropriate to solve the problems observed. Most of the time was spent in dealing with the downhill problems illustrated in Figure 4, and hence mostly dealing with Reson 8101-generated data. Filtering by angle took approximately 6 min. per day of survey, although determining which side was downhill took longer, up to 30 min. for a day of Reson 8101 data. The total time to process the survey was 12.8 hrs, with 10.9 hrs (85 per cent) interactive and 1.9 hrs (15 per cent) non-interactive. This total includes: two CUBE runs over the whole area, surface generation and insertion into HIPS weighted grids, interactive editing, and re-generation of the initialisation surface as the processing proceeded (described in the following section).

For the quality-flag method, we found that a total of 5.2 hrs were expended in interactive editing, using primarily area-oriented mode, and some line-oriented editing where required. We found that the quality flags had improved the signal to noise ratio sufficiently in

most areas for the CUBE algorithm to correctly determine the true depth hypothesis in each case, leading to significantly less remedial work after the first pass, even though the level of unflagged noise remained high. With the non-interactive times estimated above, the total time for the processing was 7.1 hrs, with 5.2 hrs (73 per cent) interactive and 1.9 hrs (27 per cent) non-interactive, with the same breakdown of non-interactive time as before.

Discussion

A simple comparison of the time taken by manual processing and computer assisted processing confirms that the assisted method is significantly faster, as expected. A potential concern is that the assisted processing method increases the non-interactive computational load by approximately 0.81hrs per run of CUBE. However, the actual cost of this increase is less than the numbers would suggest, since it can be effectively hidden by careful organisation of the computational process. For example, it should be possible to arrange for the CUBE processing to occur on a computer server tuned for the task, while the user works on another task, or even another survey (it is often the case that a survey party will have multiple surveys, or survey sheets, active simultaneously). Indeed, it is possible to reduce the time required for processing almost arbitrarily through parallel compute-farm processors, the only real limitation being cost of hardware and complexity of maintenance and scheduling. We also note that the set of computer-assisted processing methods show that it is possible to push more of the processing task into the non-interactive category than it is with purely manual methods, exaggerating the benefit that can accrue from this sort of technological improvement. It is typically straightforward to purchase faster computers; it is a more difficult task to purchase faster humans.

The difference in time required to process the data from the different launches (Figure 8) is unexpected. If we combine the timing data with the data volume for each system, the average hand-processing rate for the Reson 8101 on launch RA5 is 388 snd s[-1], while that for the Elac 1180 on launch RA6 is 47 snd s[-1]. However, we observe many more problems with the Reson 8101 due to the extreme slopes in the regions as described previously.

Therefore, there must be some other explanation for the significantly higher effort involved in processing this Elac data. One possible explanation is that the Elac data is sparser and, because it is deeper, appears to be noisier than the Reson data. Although this is only to be expected in the areas where the Elac systems are operated, there is a natural tendency on the part of human operators to 'clean' the 'noise' – even if the data is within specification and consistent – hence expending more time. With the assisted processing described here, only those areas of data that exhibit difficulties are treated, so that effort is focused on only those areas that require work, redressing this balance.

One feature of CUBE is the use of an initialisation surface, which is meant to represent the *a priori* state of information about the survey area before the survey begins. (Note that the result of processing is not directly affected by the initialisation surface; at worst, a bad initialisation surface weakens the strength of the algorithm's robustness, leading to more work for the operator.) We found that the initialisation surface based on prior survey data and in part on the prior ENC for the area was inadequate for many uses. This was mostly because of the shoal biased production process typical in current survey methods, which resulted in differences between the final surface and the TINed version of the selected soundings of over 50-80m in some areas with significant slope. It would be possible to attribute the initial surface with significantly increased *a priori* uncertainty in this case, and let the CUBE algorithm reject this information based on the MBES data. However, this achieves little more than adding more 'noise' to the estimation process, and we have not pursued the idea any further. Moreover, in some cases it might not be possible to obtain any prior data.

We developed one solution to this problem while working on this survey, by constructing a median surface (at a fixed resolution, in this case 15m) to use in place of the prior survey. Due to the extreme amounts of noise, we were obliged to carry out this process in an iterative manner. We first ran CUBE using the prior survey surface, and dealt with the most egregious problems observed. We then constructed the median surface, and corrected problems that were still evident in it. Finally, we junctioned the surface with a mask indicating landmass in the area, so that we had a seamless surface over the whole survey area; a surface

spline [Smith & Wessel, 1990] was used to interpolate over holes. This intermediate surface proved to be sufficiently close to the data to use in the usual mode, and we utilised it as the initialisation surface thereafter. The construction process for the surface took approximately 15 mins. Although it is not common to require this approach, it is suitable for bootstrapping analysis where there is no prior information, for example in an area where there has never been a survey. This approach is most useful in a post-processing mode, but could be adapted for iterative use by working on a day cycle as outlined here.

We can approximate the effect of a CUBE integrated processing system by substituting the editing effort recorded for the traditional process with that found in the computer-assisted process, but keeping the other components of the survey (e.g., reporting, line planning, troubleshooting etc.) constant. Figure 11 illustrates a comparison of the two processing streams indicating the proportion of the time used for each activity. Bearing in mind that the computer assisted process is very much shorter overall, Figure 11 shows that the computer-assisted scheme has much more time spent in 'active' tasks, such as line planning or target and shoreline investigation, and much less in the tedious work of data editing. The higher proportion of 'hydrographic' time suggests that tools to assist in the process of line planning would bring about another significant benefit. The greater proportion of troubleshooting time highlights the difficulties of working in the field for an extended period. It is possible that this time burden cannot be removed, since systems will always fail over time. Indeed, adding another layer of complexity through systems like CUBE might make this worse. It is a significant challenge for software and hardware developers to build systems that will operate correctly under unexpected conditions for extended periods; it is a challenge that we must face, however, if we are not to fritter away the gains that we make by implementing new technologies and methodologies.

This field trial focused our attention on some user interface issues that are important in maximising the potential benefit of implementing technology such as CUBE and the Navigation Surface. In theory, we do not have to edit every sounding that appears to be an outlier, even where CUBE's disambiguation engine makes the wrong choice of depth

(a): Traditional processing, 305.06 hrs   *(b): Computer Assisted processing, 130.56 hrs*
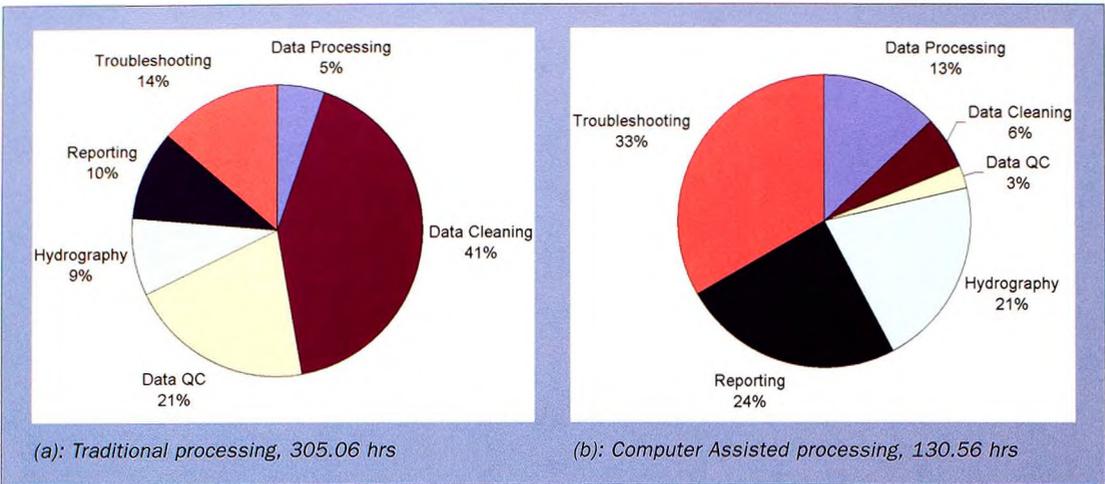
Figure 11: Comparison of possible distribution of time with a computer assisted processing path. The left chart shows the proportion of time spent on the survey using standard methods (amalgamated for clarity); the right chart shows the proportions which might be possible using CUBE instead of the current processing path. The following amalgamations were used; see Figures 5, 9 and the Results section for correspondence to timesheets. 'Data Processing' consists of 'Ingest Data', 'Preliminary Inspection', 'Preparatory Data Processing', 'DTM Creation' and 'Data Manipulation'. 'Data Cleaning' is 'Line Cleaning' and 'Subset Cleaning'. 'Data QC' is 'QC Data', 'QC Shoreline', and 'Subset QC'. 'Hydrography' is 'Line Planning', 'Targets and Shoreline' and 'Smoothsheet Creation'. 'Reporting' is 'DTONs' and 'Reporting'. 'Troubleshooting' is the sum of the troubleshooting elements from survey and post-survey efforts

hypothesis. All we have to do is to improve the signal to noise ratio sufficiently for CUBE to make the correct decisions. This provides a way to maintain the objectivity of the statistical estimates, since we do not have to edit too close to the 'true' data. It is not easy to break the habits engendered by a traditional approach to editing, however, and we found ourselves cleaning all of the data in each area that was investigated. This is partly because we cannot currently see the effects of cleaning outliers, because we have to mentally 'fuse' the visualisation and editing environments. If we could view the CUBE reconstructed surface, data and hypotheses in the same context and be able to do partial re-CUBEing of the data (i.e., only reprocess the data that has been modified), then it would be easier to decide when we have done enough to resolve the problems, and stop the editing process. This would also lead to shorter processing times overall.

We found that one frustrating problem with working on the data 'as needed' was that it was difficult to keep track of which areas had been worked. This was particularly problematic when more than one person was working on the data simultaneously, and was exacerbated by having the visualisation and editing environments separated. A practical imple-

mentation would need to have some way to illustrate which areas had been inspected to avoid repetition of effort. It is possible that this could be combined with the current practice of defining 'subset' areas over the entire survey, and having each one inspected and marked as complete as a way of confirming that the whole survey has been inspected. In this way, we directly shift the focus of processing the survey towards QC inspection, with editing only where required, rather than editing everything and then doing a QC inspection. This division of survey area also naturally leads to a division of labour, and a division of control, making it easier to split the task between a number of operators, so reducing the overall real-time expenditure.

A final implication of the times reported here is that the total computer-assisted interactive processing time was less for the whole survey than was the in-survey preliminary editing in the traditional approach (Figure 12). This means that it should be possible to have a survey ready to leave the ship for reporting and final polishing before the survey vessel leaves the area. This is a very important goal, since it is usually significantly cheaper to complete a survey before pulling out of an area than to reopen the survey during the next field sea-

son. In other circumstances, it may not be possible to return to an area at all. In this case, ensuring that sufficient data of adequate quality is on board before leaving might be the most significant advantage of the type of methods that we have outlined here.

Conclusions

Our experiment shows that the CUBE and Navigation Surface concepts can be applied in real-time mode, and confirms that the automatic processing is significantly faster than the traditional hand-processing methods currently employed.

Our experience in Valdez Narrows shows that not all problems are best solved in spatial mode, since they occur as a function of the data collection process and are intrinsically survey line oriented. In this case, the very significant slope caused a number of false hypotheses to be generated through high-density spatially localised bursts of noise that also happened to be co-located with less-dense data. We found it more efficient to resolve this with filters in line-mode.

In the wider context, we observed that the majority of the effort during the survey (using the traditional approach) was taken up by line-based editing of the data (33 per cent of total person-hours), with line planning and QA/QC activities following behind at 11 per cent and 21 per cent, respectively. Other activities, not counting troubleshooting, amounted to only 23 per cent of the total time, with no category more than 7 per cent. It therefore follows that further development of tools to support the QA/QC procedures and automate the line planning process would be of significant benefit as the survey is being conducted.

In post-survey work, the vast majority of time (103.5 hrs, or 45 per cent) was taken up with more sounding cleaning, with another significant QC cost (46.5hrs, or 20 per cent). Reporting (13 per cent) and troubleshooting (14 per cent) were also significant. In total, of the 305.06 hours expended so far on the traditional processing path for the survey, 127.97 hrs (42 per cent) was cleaning, and 62.58 hrs (21 per cent) was QA/QC.

The best CUBE-based processing path expended 0.81 hrs per run of CUBE over the whole dataset,

and 5.2 hrs in interactive processing of data. The vast majority of this time was spent in dealing with the downhill detection problem illustrated in Figure 4, a problem we expect only to appear in this type of survey environment. We believe this issue can be solved or ameliorated by automatic filtering processes, which would significantly reduce the amount of interactive time required for computer-assisted processing. The benefit of CUBE is, crudely, 24.7:1 counting just editing time, or 36.9:1 including the QC time, which can be argued to be an integral part of the CUBE inspection process.

We observed that a significant difficulty in processing the data using our prototype integration of CUBE and CARIS/HIPS was transferring the information on problems from the visualisation system, where they are obvious, to the editing system, where they can be resolved. Hence, we expect that better integrated systems with immediate re-CUBE feedback and tight integration of visualisation and remediation will achieve bigger savings in time and effort than those observed here. Fundamentally, and obviously, the benefit that can be achieved depends strongly on the complexity and quality of the underlying data.

Finally, we observe that this survey may be atypical in its noise content (particularly the downhill issues), and hence the time required for processing might not be typical for a survey of this size. Caution should be exercised in drawing wholesale conclusions from the timings presented here.
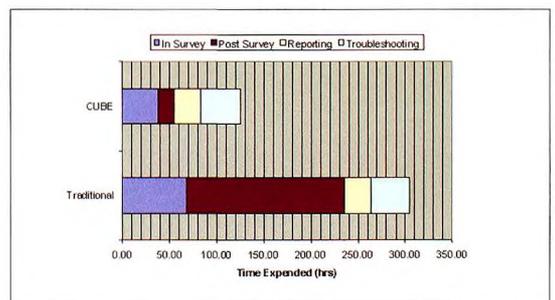

Figure 12: Comparison of total time taken in processing using traditional and CUBE methods. 'In Survey' is time spent processing data during the time that the ship is in the survey area; 'Post Survey' is time after the ship has left. The total time spent processing data using the CUBE method is less than the time expended during survey using the traditional method, suggesting that CUBE can be used to provide verified products while still in the survey area

Acknowledgements

Readers should note that information on data collection and processing procedures reported here as they apply to NOAA standard operating procedures are not intended to be authoritative or exhaustive. The survey (registry number H11182) is described through a Data Acquisition and Processing Report (DAPR) and Descriptive Report (DR), which should be considered as final authority if anything herein differs from them. Use of particular software and hardware in the work described here is not intended as endorsement on the part of the authors, and any trademarks are acknowledged, even if not so marked in the text.

References

Calder, BR (2003). 'Automatic Statistical Processing of Multibeam Echosounder Data', *Int. Hydro. Review*, Vol. 4, No. 1, pp. 53-68

Calder, BR & Mayer, LA (2003). 'Automatic Processing of High-Rate, High-Density Multibeam Echosounder Data', *Geochem., Geophys., Geosyst.* (G3, g-cubed.org), 10.1029/2002GC00486, Vol. 4, No. 6

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

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

Riley, JL, Gallagher, BJ, & Noll, LCDR GT (2001). 'Hydrographic Data Integration with PYDRO', *Proc. 2nd Int. Conf. On High Resolution Survey in Shallow Water*, Portsmouth NH

Smith, LT SM (2003). *The Navigation Surface: A Multipurpose Bathymetric Database,* Masters Thesis, Center for Coastal and Ocean Mapping & Joint Hydrographic Center, University of New Hampshire

Smith, LT SM, Alexander, L, & Armstrong, CAPT AA (2002). 'The Navigation Surface: A New Database Approach to Creating Multiple Products from High-Density Surveys', *Int. Hydro. Review*, Vol. 3, No. 2, pp. 12-26

Smith, WHF, & Wessel, P (1990). 'Gridding with Continuous Curvature Splines in Tension', *Geophysics*, Vol. 55, No. 3, pp. 293-305

Ware, C, Plumlee, M, Arsenault, R, Mayer, LA, Smith, SM, & House, D (2001). 'Data Fusion for Interpreting Oceanographic Data', *Proc. IEEE Oceans 2001,* Honolulu, Hawai'i

Biographies

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

LT Shepard (Shep) Smith is a NOAA Corp officer assigned to the NOAA Ship *Thomas Jefferson*. He has also served aboard the NOAA Ship *Rainier* and S/V *Bay Hydrographer* conducting hydrographic surveys for charting in Alaska and the US east coast. He graduated with a Bachelors degree in Mechanical Engineering from Cornell University in 1993, and with a Masters in Ocean Engineering (Ocean Mapping option) from the NOAA/UNH Center for Coastal and Ocean Mapping/Joint Hydrographic Center at the University of New Hampshire in 2003.

E-mail: brc@ccom.unh.edu and shep.smith@noaa.gov

Appendix A: Timekeeping Forms

This form was used during the Valdez Narrows survey, recorded individually by each operator. Form developed by LT Smith & PS Kim Sampadian.

Daily Processing Log OPR-P151-RA-02 H11182 - Sheet C only			Name:_____ Date/Dn: _____ Launch: _____
Bathymetry	**Time**		**Comments**
	Interactive	Non-Interactive	
Download Data from launch			
Checksheet			
Converting			
H/R/P			
Navigation			
Load Tide			
Load/Correct SVP			
Merge			
Filter			
Line Cleaning			
DTM Creation			
Line Planning for next day operations			
File Management			
QC data			
Features			
Edit target file			
Pydro Insert			
Pydro Edits			
Caris correctors (tide,svp,merge)			
Draw PSS			
Update Shoreline Notes			
Photo download/assigning			
Shoreline Review (Sheet OIC)			
Troubleshooting			
Software			
Hardware			
Data			
Other			
Statistics			