Cite: Calder, B. and S. Smith, "A Time/Effort Comparison of Automatic and Manual Bathymetric Processing in Real-Time Mode", Proc. US Hydro Conf. (Biloxi, MS), 2003.

A Time/Effort Comparison of Automatic and Manual Bathymetric Processing in Real-Time Mode

Brian Calder Shep Smith
(brc@ccom.unh.edu) (shep.smith@noaa.gov)

Center for Coastal and Ocean Mapping and Joint Hydrographic Center
Chase Ocean Engineering Lab
University of New Hampshire
Durham, NH 03824

Abstract

In September 2002, the NOAA Ship RAINIER hosted an experiment to examine the benefit in processing time from adopting CUBE [Calder & Mayer, 2002] and the Navigation Surface [Smith *et al*., 2002], and to determine whether this methodology could be used in real-time.

Using normal survey methodology, data were collected from multiple platforms in Valdez Narrows, AK (~10x10km, 0-300m depth, slope 40-60°). Seven boat-days of data were collected, containing 22.3×10^6 raw soundings. We obtained a copy of the ship's data after correctors were applied, but before any hand-edits had been done, and then ran the data through CUBE where multibeam data was available, using a TINing algorithm to junction singlebeam, shoreline and other features. Next, we inspected the data using GeoZui3D for visualization and CARIS HIPS for remediation. A final pass through CUBE resulted in an updated 'current survey estimate' to proceed to the next day; after all data were collected, a final sweep through the data completed the editing process. Throughout the process, the ship's survey team recorded time expended on the project, broken into interactive and non-interactive tasks, with finer sub-divisions such as data manipulation time and line-oriented (swath) editing, etc. Since the CUBE effort concentrates on data processing, we recorded only this category.

The experiment showed that the extended algorithm can be used in a real-time mode, and is significantly faster than the traditional processing stream in operator effort to achieve the same result. We found that a tight integration of visualization and remediation tools will be required in order to achieve the most benefit from the method, and that the expected benefit will vary with the quality of the data. The challenging geomorphology precipitated problems that were obviously better solved in line-oriented mode, highlighting the fact that there is no universal solution to the bathymetric modeling problem. Our analysis also shows that, in the traditional processing stream, the most significant time expended during the survey is in swath editing, followed by line planning and quality control. In post-survey, time expenditure is dominated by interactive editing. We note that these results are only a point in a spectrum of survey types, and may not be typical.

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 automatic data processing and manipulation. We are currently testing the CUBE [Calder & Mayer, 2002] and Navigation Surface [Smith *et al*., 2002] algorithms. These techniques are novel in that they operate at the level of 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, resulting in a collection of estimates of 'true depth' and other metrics for QA/QC, 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, from which, combined with hydrographic knowledge and best practice, a 'safe' surface qualified for navigational use can be constructed. 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 tasks that were previously hand-driven in the cartographic realm, with downstream benefits in time and simplicity.

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 re-lock to the true bottom, rather than being confused by the outliers.

As part of the verification process for CUBE and the Navigation Surface, we took part in a survey conducted by the NOAA Ship RAINIER in Alaska at the end of the 2002 field season. The survey was conducted in Valdez Narrows (figure 1a) over five survey days from 13 September 2002. The area is a narrow channel serving Port Valdez and the Valdez Oil Terminal, characterized by a relatively shallow trough oriented roughly northeast/southwest between mountains, with deep, flat areas in Port Valdez and the approaches regions to either end (figure 1b). 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 the 1:20,000 chart insert shown in figure 1a.

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 statistical evidence for how much time was actually expended during the 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. Finally, we conclude with some observations on the effort involved in each method from the operator's point of view, and some lessons learned using the automatic method in the field environment.

Experimental Methods

Data Collection

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 for shallow water around the edge of the survey; RA6 operated an Elac-Nautik 1180 MBES for the deeper areas in the mid-channel, Valdez basin and approaches. Both MBES boats used a POS/MV 320 for attitude and orientation, and a Trimble DSM212 differential receiver for position. Auxiliary information from aerial photography was available for comparison and shoreline work. Tide control was established from Valdez, and differential GPS correctors were derived from the beacons at Potato Point or Cape Hinchinbrook, as appropriate for reception. The depth range was from shoreline to approximately 300 m, with slopes of up to 50-60°.

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 simply 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 remediation. We do 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 cannot split the QA/QC task from the data cleaning task: in CUBE, these are essentially the same process.

Data Processing

For the ship’s crew, the standard NOAA processing chain was followed. Once 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 the data, hence acting as a data interlock. Based on these forms, the XTF data were then converted into HIPS/HDCS 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 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). Single-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 summarized in table 1.

	System		
Day	**RA5 (8101)**	**RA6 (1180)**	**Total**
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. Note that this is raw data volume gathered, rather than data remaining at the end of the survey.

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 MapSheets were used to initialize the per-day set. The data for the day were then run through the CUBE process (figure 3) using the per-day MapSheets. An inspection stage followed, using GeoZui3D [Ware *et al*., 2001] to visualize 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 surface reconstruction was dubious, e.g., figure 4, 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 evidence in the region to the contrary. Spatial (area) mode editing was used in these cases, using subset tiles [Gourlay & DesRoches, 2001] to track the areas of the data that had been inspected. 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 sufficiently 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 an oblique 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 5. 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 8101). In some instances, it is not immediately obvious for a human operator where the true bottom is either.

This problem is line oriented, and best resolved in line mode. We therefore 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 choose to retain more data, paying the cost in extra interactive processing time. A better solution would be arrange for a dynamic angle gate, or other filter, but since we are timing a particular system implementation, we felt that this would confuse matters, and maintained the sub-optimal, but simple, solution outlined above.

After inspection and remediation, the day's data was run through CUBE again to assimilate the new data against the cumulative MapSheets, figure 6. We then constructed a 'current best estimate' surface from the cumulative MapSheets to summarize the state of the survey. Finally, a composite surface was constructed using the CUBE surface (where defined), vertical-beam Echosounder (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 defined 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. This composite surface, at a single resolution of 5m, provided for visualization 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', tides, refraction, merge and pre-filtering as 'preparatory data processing'. All of the target and shoreline processing have been accumulated 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 figure 7 and 8. Times are shown in man-hours assigned against the task. Although, not surprisingly, the line editing (i.e., editing in swath oriented mode) takes up the lion's share 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 9. The vast majority of time is spent in interactive activities at the computer, where processing is limited to operator speed, rather than technology limited. 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 significant 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 10, although it is in the correction direction, with the less dense, lower rate Elac data taking less time to process. We return to this in the discussion.

Finally, figure 11 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. 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 is approximately 12,000 soundings/s, although the rate varies considerably with sounder and depth range. The cost of generating 'current best estimate' grids was relatively constant at approximately 13 min., although this increases slightly over the course of the survey as more of the area becomes active; conversion into HIPS weighted grids for inspection costs another 6 min. per run of CUBE. In total, 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 5, and hence mostly dealing with 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 8101 data.

In summary, the total time to process the survey was 12.8 hrs, with 10.9 hrs (85 %) interactive and 1.9 hrs (15 %) 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 initialization surface as the processing proceeded (described in the following section).

Discussion

A simple comparison of the time taken by manual processing and automatic processing confirms that the automatic method is significantly faster, as expected. A potential concern is that the automatic 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 organization of the computational process. For example, it should be possible to arrange for the CUBE processing to occur on a compute 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 arbitrarily through parallel compute-farm processors, the only real limitation being cost of hardware and complexity of maintenance and scheduling.

The difference in time required to process the data from the different launches (figure 10) is unexpected. If we combine the timing data with the data volume for each system, the average hand-processing rate for the 8101 on launch RA5 is 388 soundings/s, while that for the 1180 on launch RA6 is 47 soundings/s. However, we observe many more problems with the 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, through depth, appears to be noisier than the Reson data. Although this is only to be expected in the deeper 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 automatic 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 initialization surface, which is meant to represent the *a priori* state of information about the survey area before the survey begins. We found that the initialization 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. 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 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 utilized it as the initialization 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 visualize the effect of a CUBE integrated processing system in an approximate manner 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.) Figure 12 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 12 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 grunt-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 things will always break. 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 square up to, however, if we are not to fritter away the gains that we make by implementing new technologies and methodologies.

This field trial of the CUBE/Nav. Surface methods focused our attention on some user interface issues that are important in obtaining the most benefit from the algorithms. 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 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 work too close to the true surface. 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 as we do so because we have to mentally 'fuse' the visualization 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 visualization and editing environments separated. A practical implementation 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 working 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 labor, and a division of control, making it easier to split the task between a number of operators, so reducing the overall 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. 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 very much cheaper to complete a survey before pulling out of an area than to reopen the survey during the next field season (for example). 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 line oriented. In this case, the very significant slope caused a number of false hypotheses to be generated through high-density spatially localized 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% of total man-hours), with line planning and QA/QC activities following behind at 11% and 21%, respectively. Other activities, not counting troubleshooting, amounted to only 23% of the total time, with no category more than 7%. 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%) was taken up with more sounding cleaning, with another significant QC cost (46.5hrs, or 20%). Reporting (13%) and troubleshooting (14%) are also significant. In total, of the 305.06 hours expended so far on the traditional processing path for the survey, 127.97 hrs (42%) is cleaning, and 62.58 hrs (21%) is QA/QC.

The CUBE-based processing path expended 0.81 hrs per run of CUBE over the whole dataset, and 10.9 hrs in interactive processing of data. The vast majority of this time was spent in dealing with the downhill detection problem illustrated in figure 5, 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, 11.8:1 counting just editing time, or 17.6:1 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 visualization 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 visualization and remediation will achieve bigger savings in time and effort than those observed here. Fundamentally, and obviously, the benefit that you achieve depends strongly on the complexity and quality of the underlying data.

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

Acknowledgements

This work would not have been possible without the enthusiastic engagement of the Captain, Officers and Crew of the NOAA Ship RAINIER. We would like to thank CAPT Gardner, LCDR Fletcher, PS Sampadian, CST Rooney and the rest of the survey and deck departments for their commitment, enthusiasm and fortitude in support of this project and the extra work that it entailed. The support of NOAA grant NA97OG0241 is also gratefully acknowledged.

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, B. R. and L. A. Mayer, *Automatic Processing of High-Rate, High-Density Multibeam Echosounder Data*, submitted to Geochem., Geophys., Geosyst. (G3, `g-cubed.org`), DID 2002GC00486, December 2002.

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

Ware, C., M. Plumlee, R. Arsenault, L. A. Mayer, S. M. Smith and D. House, Data Fusion for Interpreting Oceanographic Data, Proc. IEEE Oceans 2001, Honolulu, Hawai’i, November 2001.

Riley, J. L., B. J. Gallagher and LCDR G. T. Noll, “Hydrographic Data Integration with PYDRO”, Proc. 2nd Int. Conf. On High Resolution Survey in Shallow Water, Portsmouth NH, 2001.

Smith, LT S. M., L. Alexander, and CAPT A. Armstrong, *The Navigation Surface: A New Database Approach to Creating Multiple Products from High-Density Surveys*, Int. Hydro. Review, 3(2), pp. 12-26, August 2002.

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			

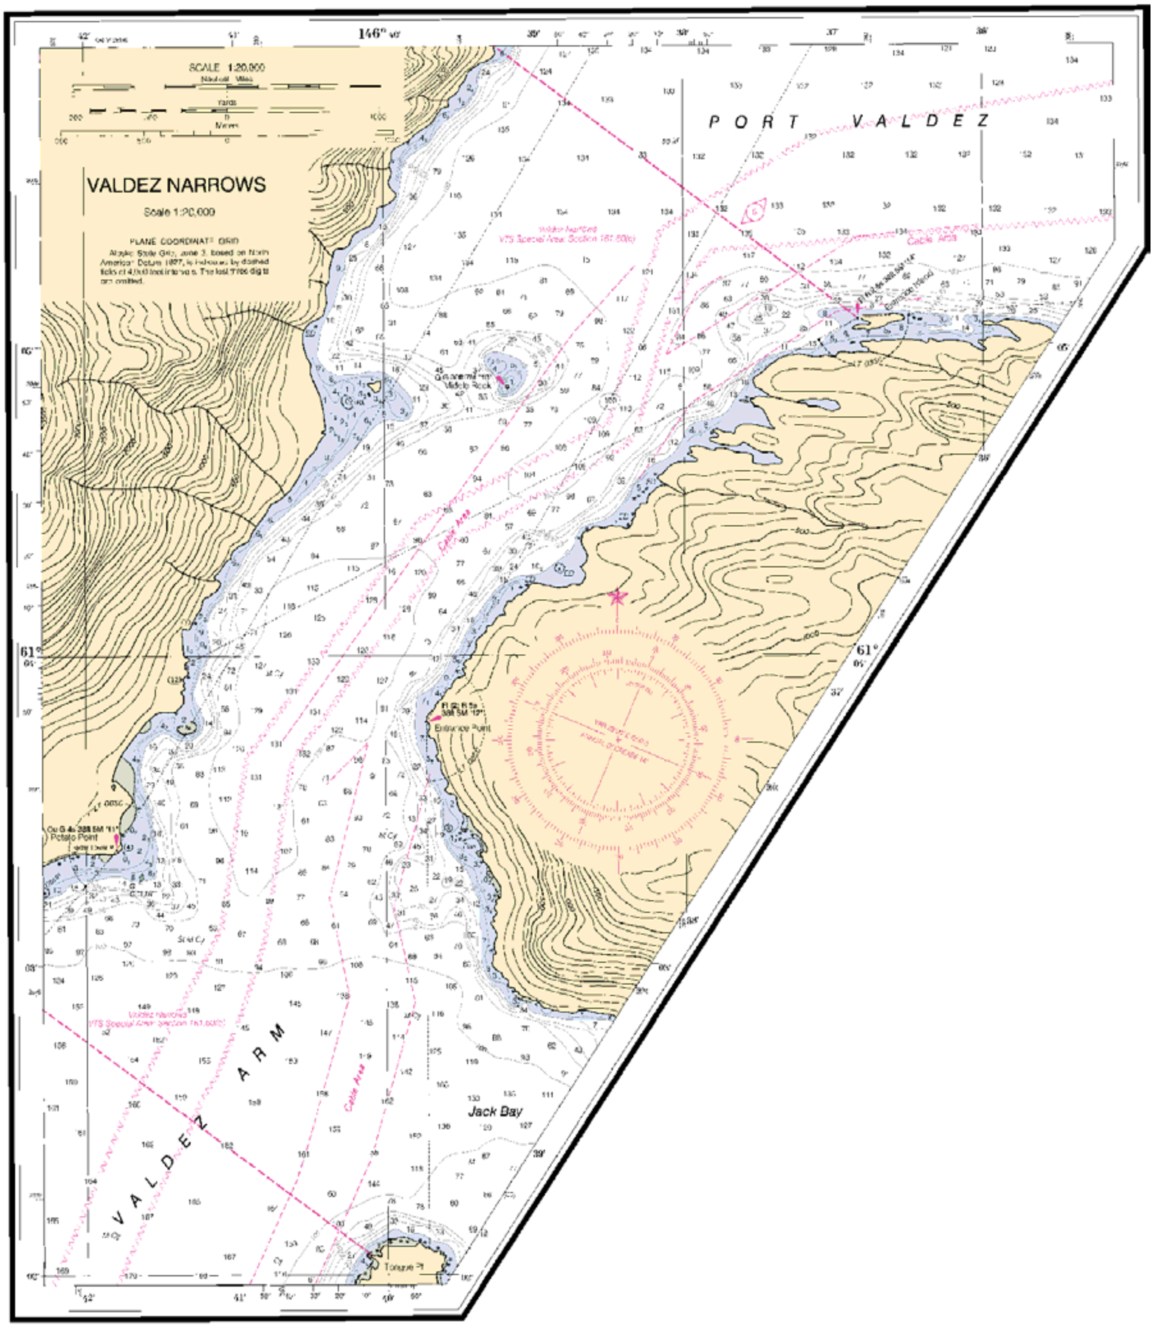

Figure 1a: Valdez Narrows, AK. Chart insert showing the approximate limits of the survey area, which cover the approaches to Port Valdez used to service the Alaskan Oil Pipeline Terminal.

Cite: Calder, B. and S. Smith, "A Time/Effort Comparison of Automatic and Manual Bathymetric Processing in Real-Time Mode", Proc. US Hydro Conf. (Biloxi, MS), 2003.

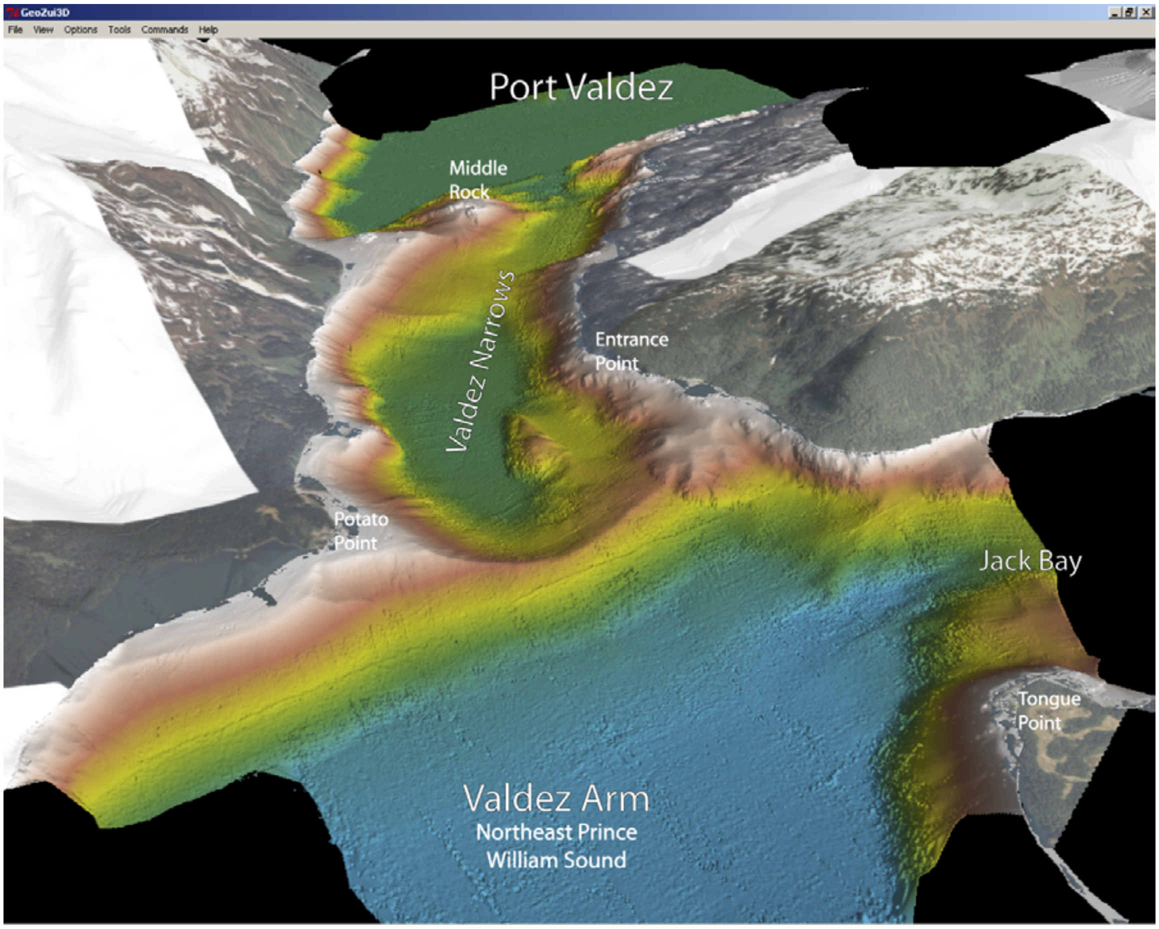

Figure 1b: Valdez Narrows, AK. Perspective view showing processed bathymetry, topography and features. The topography was generated by TINing the contours in the ENC of the area, and are overlaid with georeferenced orthophoto imagery, which was also used for shoreline. The bathymetric composite includes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

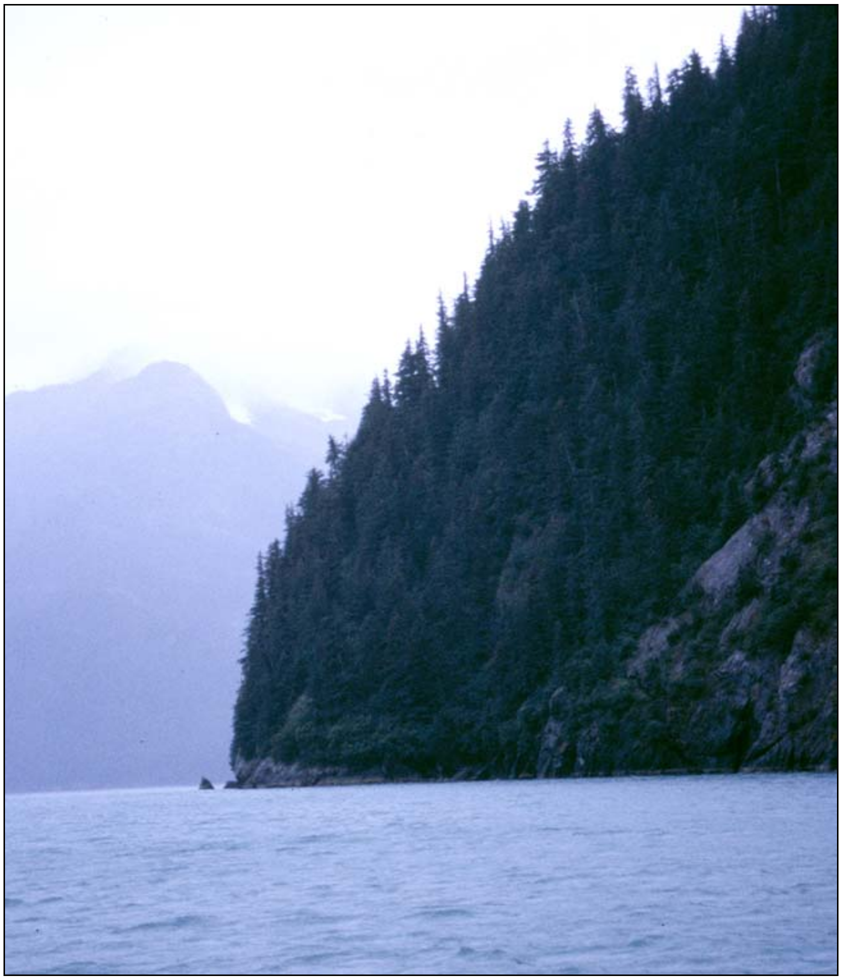

Figure 2: Valdez Narrows, AK. This picture was taken early on the first day of survey, illustrating the rock-faces on the southwest side of the narrows. This type of sheer rock faces, punctuated by waterfalls empting fresh water on top of the channel made for difficult survey conditions.

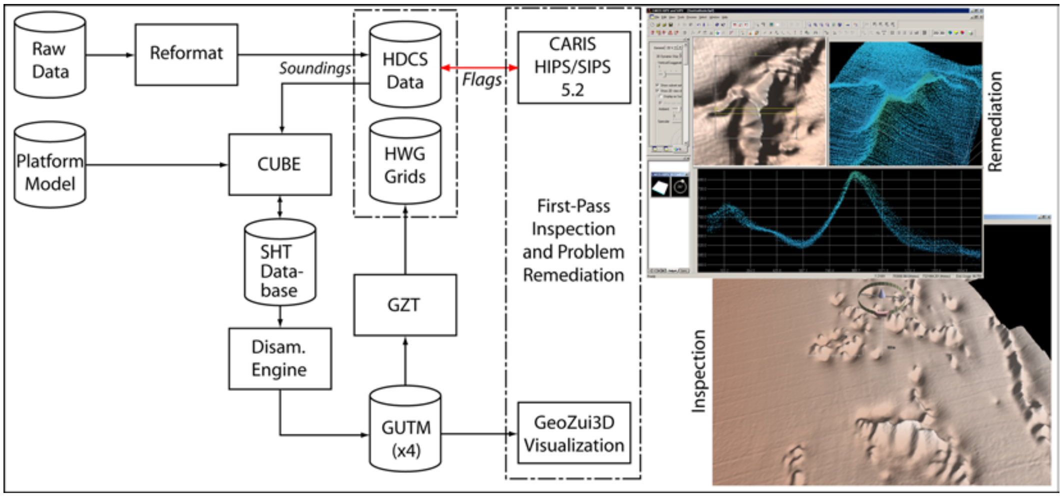

Figure 3: First-pass processing with CUBE to preliminary MapSheets. After CUBE is run, an operator inspects the results, and makes any modifications where the algorithm has failed to correctly determine the proper hypothesis.

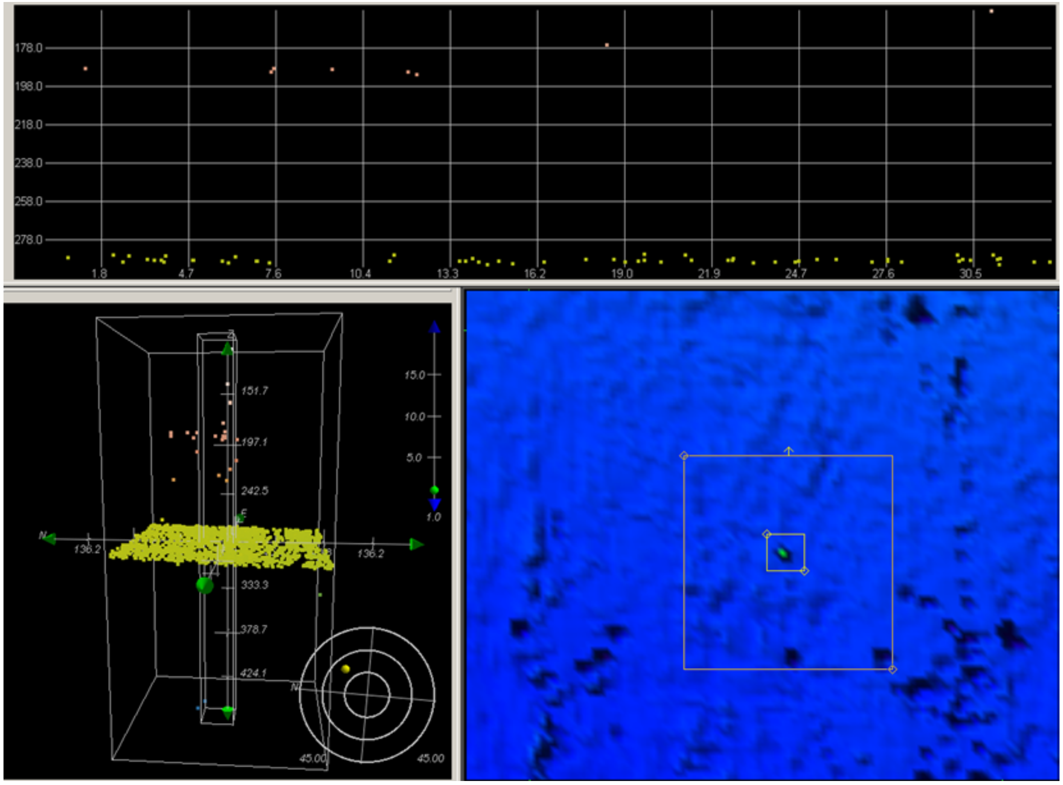

Figure 4: Sparse data in the deep southwest section of the survey area. Data here is sparse, so there are regions where there is not enough evidence from the neighborhood to overcome rogue soundings. In the belief that any data is better than no data, CUBE makes the only available reconstruction.

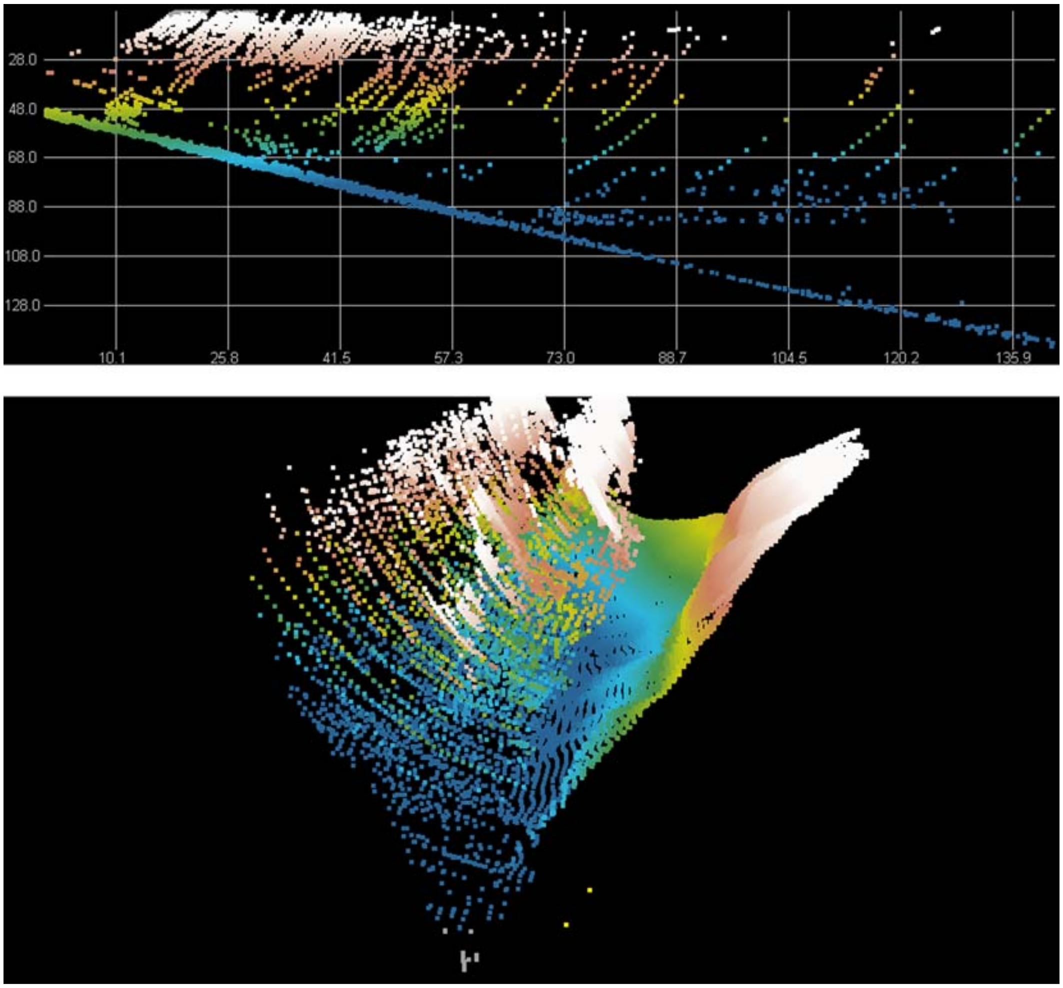

Figure 5: Data from steeply sloping area in northwest of survey area. Significant downhill slope means that a multibeam line at the top of the slope cannot detect the bottom and the top of the slope simultaneously. Detecting inconsistent returns, this difficulty puts noise into the dataset where the real data is sparse due to increasing depth.

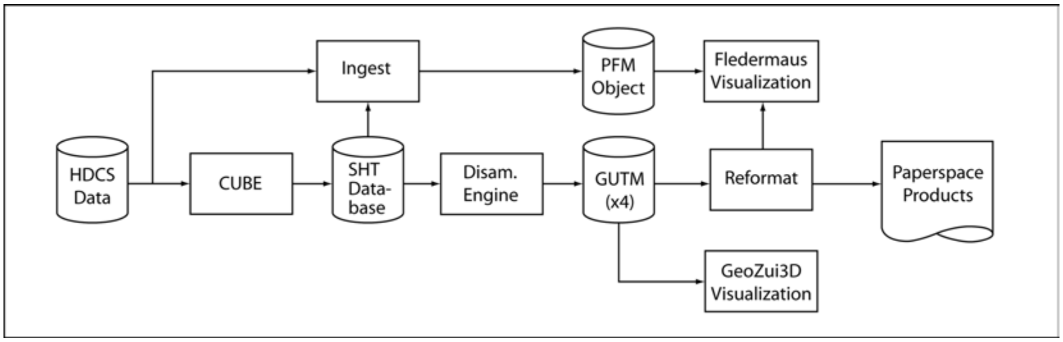

Figure 6: Second-pass processing with CUBE. The data flow is essentially the same as the first, save that the target output is assimilated into the cumulative MapSheets to update them on the current state of the project.

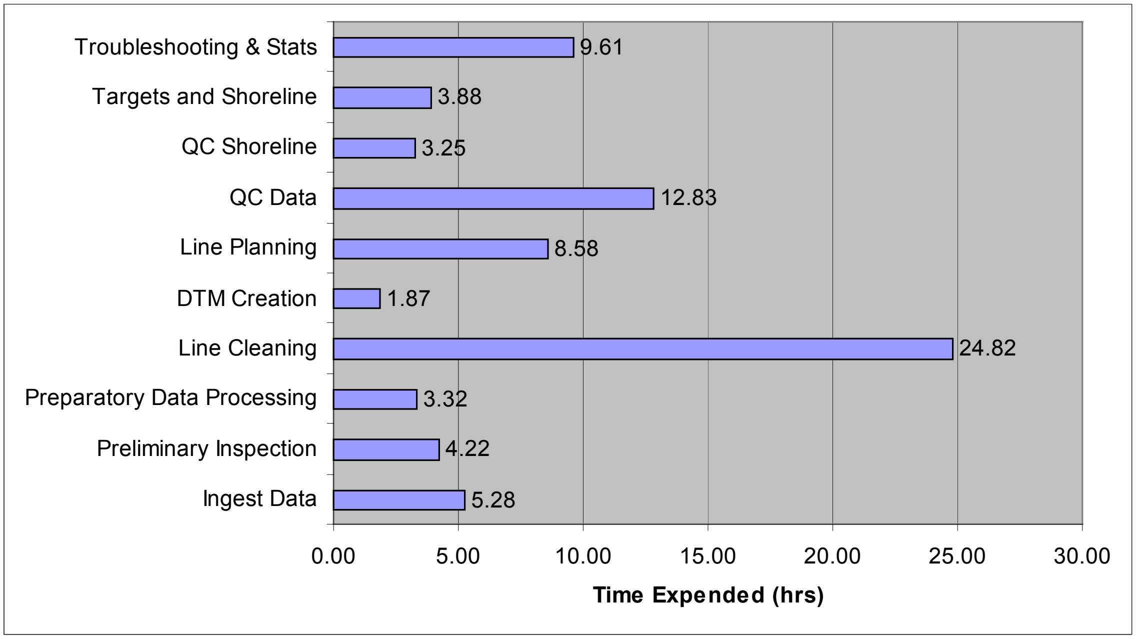

Figure 7: Summary of time expenditure during the survey effort. Some of the categories shown here are aggregated from the detailed data categories actually recorded (see appendix A).

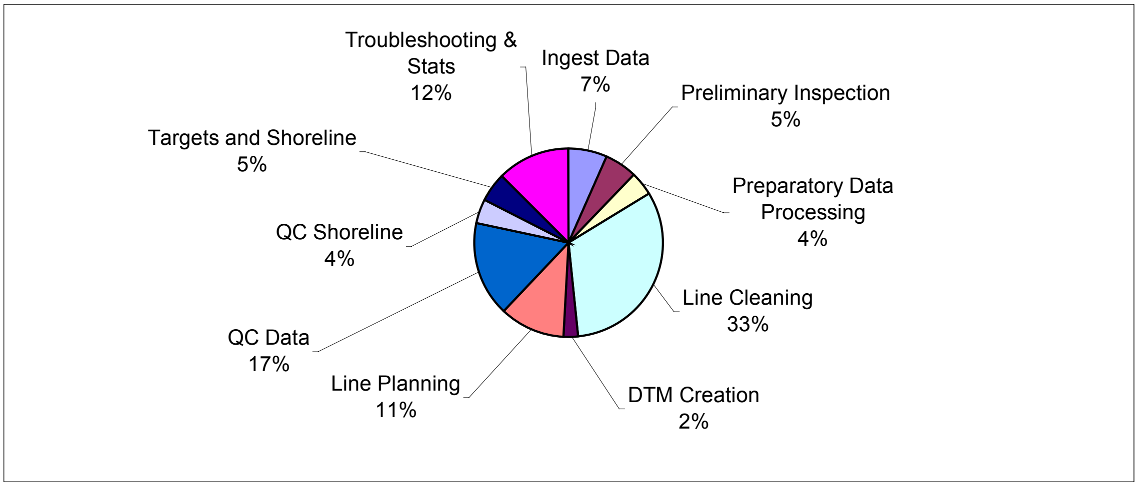

Figure 8: Summary of time expenditure during the survey field program. Majority of data processing time is consumed by line editing, followed by line planning and quality assurance and control for data and shoreline features/targets.

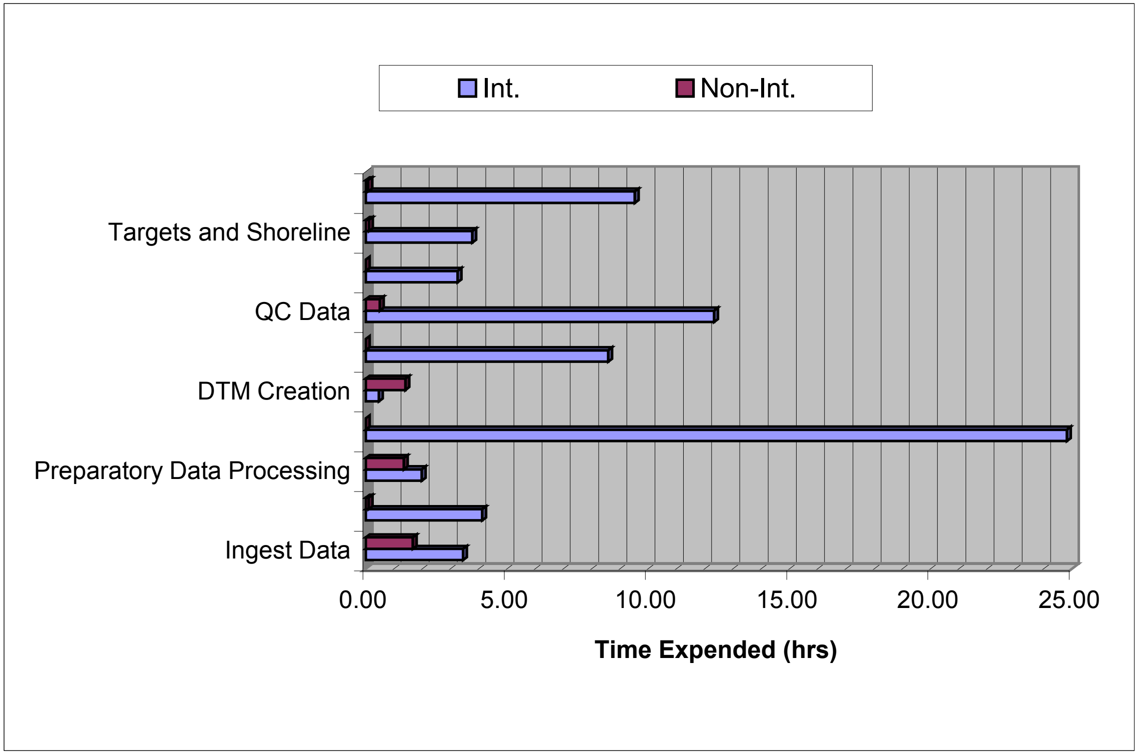

Figure 9: Comparison of Interactive and Non-Interactive time expenditure during the survey fieldwork. The significant lack of non-interactive time (a total of only 7% of the total time expended) implies that our ability to improve the process through hardware alone is limited.

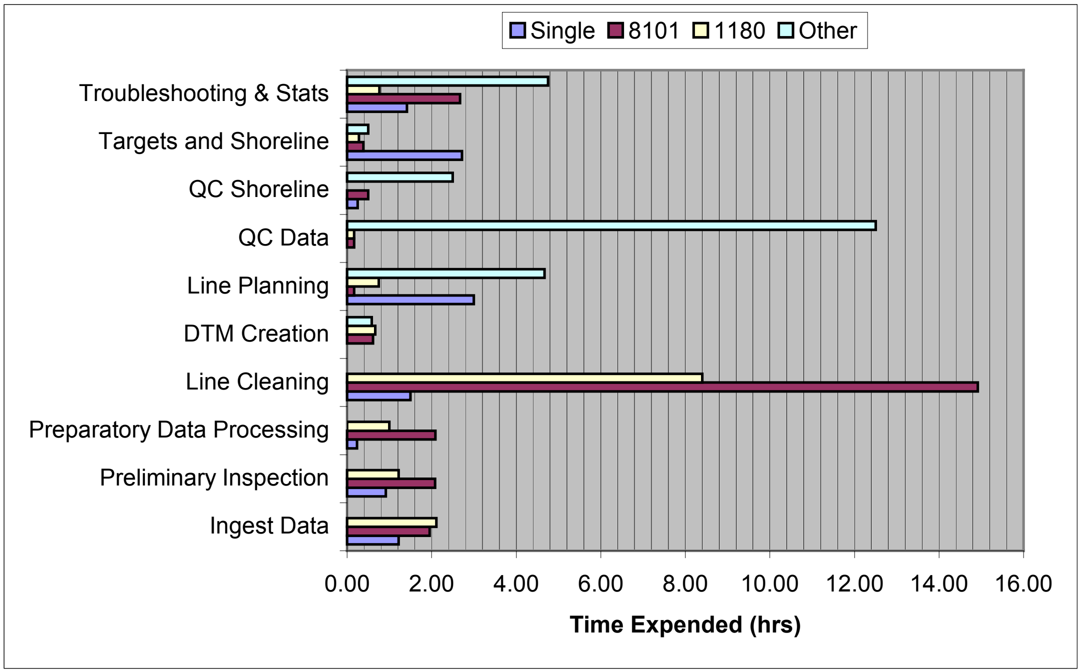

Figure 10: Expenditure of time per survey system. Since there are significant differences in sonar repetition rate and data density, it might be expected that there would be a corresponding difference in time expenditure. In fact only a small difference is observed, although it is in the expected direction.

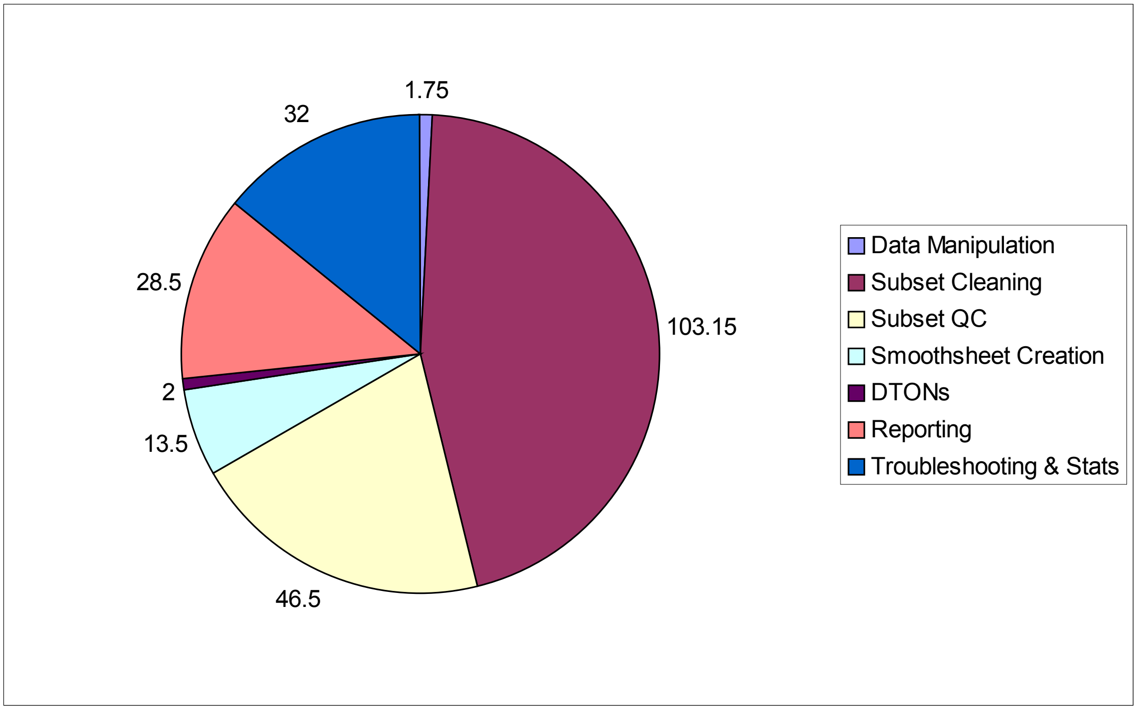

Figure 11: Post-survey expenditure of time (hrs). The vast majority of time is taken up in subset cleaning and QC (a total of 66%).

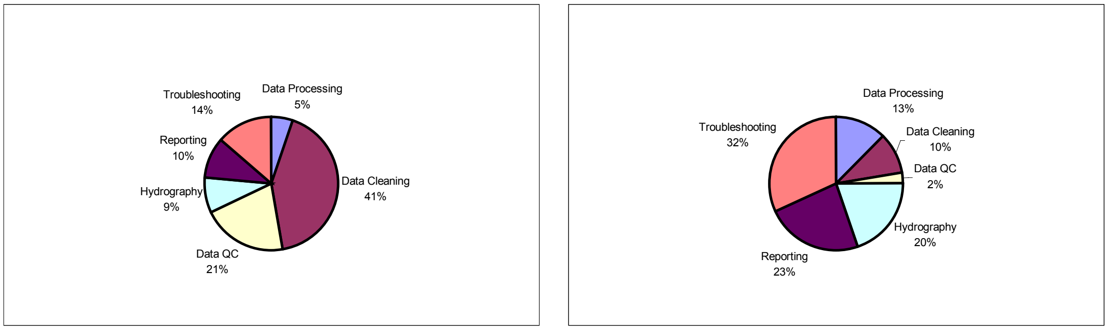

(a): Traditional processing, 305.06 hrs

(b): Computer Assisted processing, 130.56 hrs.

Figure 12: 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 7, 11 and the Results section for corrspondence 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'. 'Toubleshooting' is the sum of the troubleshooting elements from survey and post-survey efforts.